

Contextual control of operant behavior: evidence for hierarchical associations in instrumental learning

Sydney Trask · Mark E. Bouton

Published online: 25 June 2014
© Psychonomic Society, Inc. 2014

Abstract Recent research has suggested that operant responses can be weakened when they are tested in new contexts. The present experiment was therefore designed to test whether animals can learn a context–(R–O) relation. Rats were given training sessions in context A, in which one response (R1; lever pressing or chain pulling) produced one outcome (O1) and another response (R2; chain pulling or lever pressing) produced another outcome (O2) on variable interval reinforcement schedules. These sessions were intermixed with training in context B, where R1 now produced O2 and R2 produced O1. Given the arrangement, it was possible for the animal to learn two distinct R–O associations in each specific context. To test for them, rats were then given aversion conditioning with O2 by pairing its presentation with lithium-chloride-induced illness. Following the aversion conditioning, the rats were given an extinction test with both R1 and R2 available in each context. During testing, rats showed a selective suppression in each context of the response that had been paired with the reinforcer subsequently associated with illness. Rats could not have performed this way without knowledge of the R–O associations in effect in each specific context, lending support to the hypothesis that rats learn context–(R–O) associations. However, despite a complete aversion to O2, responding was not completely suppressed, leaving the possibility open that rats form context–R associations in addition to context–(R–O) associations.

Keywords Context · Operant behavior · Instrumental learning · Occasion setting · Reinforcer devaluation

Recent research suggests that contextual cues can support operant responding (for a review, see Bouton & Todd, 2014). For example, performance after operant extinction, the paradigm in which operant responding declines when the consequence that has reinforced it is omitted, appears to be context dependent. Specifically, if the context is changed after extinction, extinction performance is lost, and the original response is “renewed” (e.g., Bouton, Todd, Vurbic, & Winterbauer, 2011; Nakajima, Urushihara, & Masaki, 2002; Todd, 2013; Todd, Winterbauer, & Bouton, 2012). The renewal effect observed after operant extinction appears to parallel the effect that has been studied extensively after extinction in Pavlovian conditioning (e.g., Bouton & King, 1983; Bouton & Peck, 1989; Thomas, Larsen, & Ayres, 2003).

Related research has also found that operant responding itself can be context dependent. That is, responding is reduced when it is tested in a context that is different from the one in which it was trained (Bouton et al., 2011; Todd, 2013). This result appears to contrast with results commonly observed in Pavlovian conditioning, where the conditioned stimulus often maintains its ability to evoke conditioned responding outside of the original learning context (Bouton & King, 1983; Bouton & Peck, 1989; Brooks & Bouton, 1994; Harris, Jones, Bailey, & Westbrook, 2000; Thomas et al., 2003). However, the effect on operant responding appears to be general. Todd (2013) found that a context switch caused a decrement in operant performance even when the switched-to context had been equivalently associated with the training of a second operant response. Bouton, Todd, and León (2014) reported parallel results with *discriminated* operant performance, in which responding had been reinforced only in the presence of a discriminative stimulus (SD). In those experiments, after behavior was brought under stimulus control, rats were tested for responding in the presence of the SD in the same context and/or a different context.

S. Trask (✉) · M. E. Bouton
Department of Psychology, University of Vermont, Burlington,
VT 05405-0134, USA
e-mail: sydney.trask@uvm.edu

Testing in the different context weakened responding, and the research established that this context switch effect resulted from a failure of the response, and not effectiveness of the SD, to transfer across contexts.

Bouton et al. (2014) suggested that the context switch effect on operant performance can be explained in at least two ways. First, the context might enable the response directly, as if it had entered into a direct association with the response. (The SD would then further modulate the response once it had already been evoked by the context.) In this mechanism, the role of the reinforcer would be to strengthen the context–R association; it would not itself be encoded into the associative structure that underlies performance (Adams, 1980; Dickinson, 1994). Second, like the theoretical action of an SD, the context might also act as an occasion setter that signals or modulates the specific response–outcome (R–O) association that is learned in that context (i.e., context–[R–O]). This mechanism implies a “hierarchical” relation between the context and the R–O association. In this case, the reinforcer would be encoded into the associative structure that underlies performance (e.g., Colwill & Rescorla, 1990b).

The present experiment was designed to extend the findings of Bouton et al. (2014) by testing the hierarchical mechanism in more detail. For simplicity, it returned to the free operant, as opposed to discriminated operant, learning paradigm in which reinforcement was contingent upon responding with no presentation of SDs. To experimentally distinguish between the context–R and context–(R–O) mechanisms, we employed a reinforcer devaluation procedure (e.g., Adams & Dickinson, 1981). In a common version of such a procedure, an outcome that has been used to reinforce instrumental behavior is subsequently “devalued” by pairing its presentation with illness. This conditions a taste aversion to the outcome so that the animal begins to reject it when it is made available. In a final phase, the instrumental behavior is tested in extinction—that is, without the newly devalued reinforcer being presented following the response. If the outcome has been encoded as a part of the associative structure that underlies instrumental learning, a postlearning devaluation of the reinforcer should reduce responding. Responding should reflect the current value of the reinforcer previously earned by that response (Colwill & Rescorla, 1985). If, however, the outcome is not encoded into the associative structure (as would be expected in a context–R association), this postlearning change of reinforcer value should not affect instrumental responding, since the context alone would be enough to evoke the response (Adams, 1980).

The experiment employed a within-subjects design (detailed in Table 1) that utilized two contexts (A and B) that were different operant chambers distinguished by their tactile, visual, and olfactory cues. It also utilized two distinct responses (pressing a lever and pulling a chain), fully counterbalanced and represented as R1 and R2, and two

Table 1 Design of the experiment

	Acquisition	Devaluation	Test
Context A	R1–O1, R2–O2	O2–Illness	R1 vs. R2
Context B	R1–O2, R2–O1	O2–Illness	R1 vs. R2

Note. R1 and R2 represent counterbalanced lever or chain. O1 and O2 represent counterbalanced grain pellets or sucrose pellets. Context A and context B represent different chambers in the laboratory, distinguished by tactile, visual, and olfactory cues

reinforcing outcomes (grain pellets and sucrose pellets), fully counterbalanced and represented as O1 and O2. As is summarized in Table 1, all subjects were given training sessions in context A in which R1 responses produced O1 and R2 responses produced O2. These sessions were intermixed with similar sessions in context B, where R1 now produced O2 and R2 produced O1. The design is analogous to one used by Colwill and Rescorla (1990b) to study the associative structures that underlie discriminated operant learning. Given the arrangement, it was possible for the animal to learn two distinct R–O associations in each context (four total context–(R–O) associations). In the next phase, animals were given separate aversion conditioning with O2 (grain pellets for half the animals and sucrose pellets for the other half) by pairing O2 presentations with lithium chloride (LiCl). O2–LiCl pairings were conducted in both contexts A and B until the rats rejected O2. Following aversion conditioning, the animals were given preference tests with both the lever and chain responses simultaneously available in each context. These tests were run in extinction; that is, no outcomes were presented. If the contexts signaled distinct R–O relations, animals should suppress the response associated with the devalued reinforcer in each context—that is, R2 in context A and R1 in context B.

Method

Subjects

The subjects were 16 naïve female Wistar rats purchased from Charles River Laboratories (St. Constance, Quebec). They were between 75 and 90 days old at the start of the experiment and were individually housed in suspended wire mesh cages in a room maintained on a 16:8-h light:dark cycle. Experimentation took place during the light period of the cycle. The rats were food-deprived to 80 % of their initial body weights throughout the experiment.

Apparatus

Two sets of four conditioning chambers housed in separate rooms of the laboratory served as the two contexts (counterbalanced). Each chamber was housed in its own sound attenuation chamber. All boxes were of the same design (Med Associates model ENV-008-VP, St. Albans, VT). They measured 30.5 × 24.1 × 21.0 cm (l × w × h). 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. A lever (Med Associates model ENV-112CM) positioned to the left of the food cup protruded 1.9 cm into the chamber. A response chain (Med Associates model ENV-111C) was suspended from a micro-switch mounted on top (outside) of the ceiling panel. The chain hung 1.9 cm from the front wall, 3 cm to the right of the food cup, and 6.2 cm above the grid floor. Thus, the lever and chain were positioned symmetrically with respect to the food cup. The chambers were illuminated by one 7.5-W incandescent bulb mounted to the ceiling of the sound attenuation chamber, approximately 34.9 cm from the grid floor at the front wall of the chamber. Ventilation fans provided background noise of 65 dBA.

In one set of boxes, 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) staggered such that odd- and even-numbered grids were mounted in two separate planes, one 0.5 cm above the other. This set of boxes had no distinctive visual cues on the walls or ceilings of the chambers. A dish containing 5 ml of Rite Aid lemon cleaner (Rite Aid Corporation, Harrisburg, PA) was placed outside of each chamber near the front wall.

The second set of boxes was similar to the lemon-scented boxes except for the following features. In each box, 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. The grids of the floor were mounted on the same plane and were spaced 1.6 cm apart (center-to-center). A distinct odor was continuously presented by placing 5 ml of Pine-Sol (Clorox Co., Oakland, CA) in a dish outside the chamber.

The reinforcers were a 45-mg grain-based rodent food pellet (5-TUM: 181156) and a 45-mg sucrose-based food pellet (5-TUT: 1811251, TestDiet, Richmond, IN). Both types of pellet were delivered to the same food cup. The apparatus was controlled by computer equipment located in an adjacent room.

Procedure

Magazine training

On the first day of the experiment, all rats were assigned to a box within each set of chambers. They then received two 30-

min sessions of magazine training in context A, one with each reinforcer type (grain-based food pellet and sucrose-based food pellet). The sessions were separated by approximately 1 h. On the second day, the animals received a single 30-min session of magazine training in context B with each reinforcer type, so that each animal had a total of four magazine training sessions over 2 days. Once all animals were placed in their respective chambers, a 2-min delay was imposed before the start of the session. In each, approximately 60 pellets were delivered freely on a random time 30-s (RT 30-s) schedule. The levers and chains were not present during this training.

Acquisition

On each of the next 2 days, all rats received four 30-min sessions of response training, one for each response in each context. The sessions were separated by approximately 1 h. Only one of the lever or chain manipulanda was available in a particular session. R1 (lever or chain, counterbalanced) was trained in contexts A and B, and R2 (chain or lever, counterbalanced) was also trained in contexts A and B. In context A, animals were reinforced on R1 with O1 (counterbalanced as grain-based pellets and sucrose-based pellets) and R2 with O2 (counterbalanced as sucrose-based pellets and grain-based pellets). In context B, the contingencies were reversed, so that R1 was now reinforced with O2 and R2 was now reinforced with O1. Sessions began as soon as the animal was placed in the chamber. During each session, responding was reinforced on a variable interval (VI) 30-s schedule programmed by initiating pellet availability in a given second with a 1 in 30 probability. No additional response shaping was necessary.

On each of the next 4 days, all rats received two daily 15-min sessions of single response training in each context (four total sessions each day). The sessions were separated by approximately 30 min. During these sessions, either R1 or R2 was available and reinforced according to the same outcome contingencies described above. The order of exposure to each context was double alternated, with each context being equally experienced as the first or second session of the day. The order of response type was also represented equally as the first or second session of the day, alternating every 2 days.

Aversion conditioning

Beginning 24 h following the last acquisition session, animals received aversion conditioning with O2 in contexts A and B with levers and chains removed. On the first day, animals were given 30 O2 pellets in a conditioning chamber on an RT 30-s schedule and then injected with LiCl. For half of the animals, this session was conducted in context A; for the other half, the session was conducted in context B. Once all animals were placed in their respective chambers, a 2-min delay was

imposed before the start of the session. Following this 2-min delay, the RT 30-s schedule began. Once all pellets had been delivered, animals were given a 20-ml/kg IP injection of 0.15 M LiCl, placed in a transportation cart, and returned to the home cage. The following day, animals were placed back into the same chamber and given O1 pellets in an amount equal to the number of O2 pellets consumed the day before with the same procedure. No injection was given following O1 presentations. There were eight such 2-day conditioning cycles, such that each rat had four trials on which O2 was paired with LiCl in each context. Half the animals received the cycles in contexts A and B in an ABABABAB order, and half had the cycles in a BABABABA order.

Test

On the day following the last aversion conditioning cycle, each rat was tested in both contexts (order counterbalanced) with the lever and chain available simultaneously for the first time. All sessions began as soon as the animal was placed in the chamber. Each 10-min test session was separated by approximately 30 min. No reinforcers were delivered during this test.

Reacquisition testing

On the next and final day of the experiment, all rats received four 15-min sessions of reinforced single-response training, following the procedure used in acquisition. This test was conducted to see whether the O2 pellet would still serve as a reinforcer and support operant responding. All sessions began as soon as the animal was placed in the chamber. Context order was counterbalanced and was orthogonal to the test order during the previous extinction test. Thus, half of the rats that had previously been tested first in context A were tested first in A, and half were tested first in B. Similarly, half the animals tested first in B during the extinction test were tested first in A, and half were tested first in B. During these sessions, either the lever or the chain was available and reinforced on the same schedule and contingencies that were used in acquisition. In context A, R1 was again rewarded with O1, and R2 was rewarded with O2. In context B, R1 was rewarded with O2, and R2 was rewarded with O1.

Data treatment

The data were subjected to either *t*-tests or analyses of variance (ANOVAs) as appropriate. For all statistical tests, the rejection criterion was set to $p < .05$. During the course of the experiment, 1 animal died unexpectedly and was therefore excluded from all analyses.

Results

Acquisition

The results of the two 30-min and four 15-min training sessions in each context are shown in Fig. 1. A 2 (context) \times 2 (response) \times 6 (session) ANOVA found a significant main effect of context, $F(1, 14) = 22.76, p < .001$, with overall responding being higher in context B than in context A. There is no clear explanation of this difference. There was no significant main effect of response, $F < 1$. A significant main effect of session was found, $F(5, 11) = 47.10, p < .001$, but there was no significant three-way interaction, $F < 1$. There was no effect of either outcome type or response type. For example, on the last day of acquisition, a 2 (context) \times 2 (response) \times 2 (O1 identity) ANOVA revealed no significant main effects of context, $F(1, 13) = 1.22, p > .05$, response, $F < 1$, or O1 identity, $F(1, 13) = 2.61, p > .05$. There was also no significant three-way interaction, $F(1, 13) = 3.63, p > .05$. Similarly, a 2 (context) \times 2 (response) \times 2 (R1 identity) ANOVA was run to assess responding on the final day of instrumental training. This revealed no significant main effects of context, $F(1, 13) = 1.13, p > .05$, response, or R1 identity, $F_s < 1$. There was also no significant three-way interaction, $F < 1$.

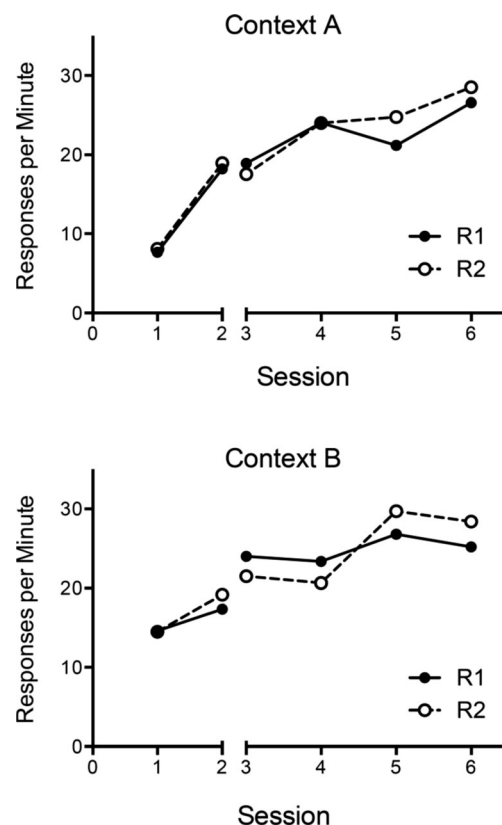


Fig. 1 Acquisition of R1 and R2 responding in the two contexts during the instrumental training phase

Aversion conditioning

The number of O2 pellets consumed on each trial of the aversion conditioning treatment is shown in Fig. 2. The animals gradually decreased the number of pellets they consumed. Interestingly, the rats showed very little change in the aversion from the first conditioning trial to the second conditioning trial, which occurred in different conditioning chambers. A paired-samples *t*-test confirmed that the mean number of pellets consumed on the first trial did not differ from that consumed on the second trial, $t(14) = 1.87, p > .05$. This could be due to a lack of generalization of the aversion between conditioning chambers. Another possibility is that O2 had received extensive exposure during acquisition and the animals, therefore, may have had a great deal of latent inhibition to overcome in conditioning the taste aversion, making conditioning slower. On the final trial, animals consumed a mean of .27 of the 30 pellets delivered. A one-sample *t*-test confirmed that this number was not significantly different from 0, $t(14) = 1.30, p > .05$. Thus, the aversion to O2 appeared to be complete.

Test

Results of the first 3 min of the devaluation test are shown in Fig. 3. The results are also shown over all minutes of the test, collapsing over context, in Fig. 4. Even though O1 and O2 were never presented during testing, responding in each context was sensitive to the current value of the reinforcer previously earned by each response in each context. A 2 (context) \times 2 (response) ANOVA over the first 3 min of the test session revealed no main effect of either context or response, $F_s < 1$. A significant context \times response interaction was found, $F(1, 14) = 4.88, p < .05$, suggesting that responding on each manipulandum was dependent on the testing context. Responding was higher on R1 than on R2 in context A. The opposite was true in context B, with responding being lower on R1 than on R2. A 2 (context) \times 2 (response) ANOVA over the entire 10-min test session revealed no main effects of

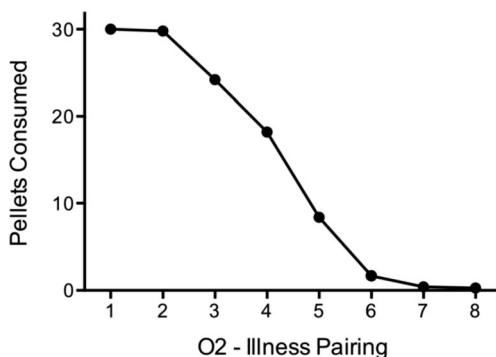


Fig. 2 Number of O2 pellets consumed across aversion conditioning trials. Adjacent trials were conducted in different contexts (A or B)




Fig. 3 R1 and R2 responding in the first 3 min of the test session

either context or response, $F_s < 1$, and also did not reveal a significant context \times response interaction, $F(1, 14) = 3.18, p > .05$. However, when one statistical outlier ($Z = 3.08$; Field, 2005) was excluded from the test data, the interaction was highly significant, $F(1, 13) = 11.67, p < .01$. Removal of the outlier did not affect the statistical pattern of the previous acquisition or aversion conditioning results.

Reacquisition

Responding during the reacquisition sessions, when R1 and R2 were paired again with their usual pellets in each context, is shown in Figs. 5 and 6. Responding in each context was sensitive to the value of the reinforcer earned by each response. A 2 (context) \times 2 (response) ANOVA revealed no main effect of response, $F < 1$, but there was a significant main effect of context, $F(1, 14) = 5.64, p < .05$, with responding being slightly higher in context B. Importantly, there was a significant context \times response interaction, $F(1, 14) = 102.58, p < .001$, meaning that responding on each response was dependent on the reacquisition context. As during the extinction test, in context A responding was higher on R1 than on R2, and in context B responding was higher on R2 than on R1. Responding for the devalued reinforcer during reacquisition did not differ from responding during the extinction test, $t(14) = 1.49, p > .05$. Additionally, animals consumed an average of only 1.26 % of the O2 pellets and all of the O1 pellets earned during reacquisition. Clearly, the aversion to O2 was still strong.




Fig. 4 R1 and R2 responding collapsed across context and response identity throughout the test session




Fig. 5 Responding for the devalued and still valued pellets separated by context and throughout the reacquisition test

In order to further examine the role of the completeness of the aversion to O₂, a final analysis excluded animals that consumed any devalued pellets either on the last devaluation cycle or during the reacquisition test. When just including the animals that completely rejected O₂ ($n = 13$), there was still an average of 4.74 responses per minute for the devalued pellet during the reacquisition test. A one-sample t -test confirmed that this was still significantly different from 0 responses per minute, the expected amount of responding if the behavior was completely mediated by a representation of the outcome, $t(12) = 4.70$, $p < .01$.

Discussion

In this experiment, animals showed selective suppression of the unique response that had produced the devalued outcome in each of the two contexts. That is, there was a preference during testing for R1 over R2 in context A and for R2 over R1 in context B. Such results demonstrate that the rats had knowledge of the R–O relationships that had been in effect in each of the contexts during acquisition. They extend the findings of Bouton et al. (2014) and confirm that animals can form associations between the context and the response–outcome association in a free operant paradigm. They are also consistent with results reported by Colwill and Rescorla (1990b), who studied the control of instrumental responding by 30-s light and noise discriminative stimuli in a




Fig. 6 Responding for the devalued and still valued pellets collapsed across context and response identity throughout the reacquisition test

discriminated operant procedure. In that experiment, like the present one, rats suppressed responding of a specific response in a specific stimulus depending on the status of the outcome with which it had been associated during that stimulus during training. Analogous to the rats in the Colwill and Rescorla (1990b) experiment, the rats in the present experiment were specifically trained with different R–O relations in the two contexts. Such a procedure could have encouraged hierarchical learning. The present data thus provide an “in principle” demonstration that hierarchical context–(R–O) control can be learned.

It is worth noting that a considerable amount of responding remained on the response that had led to the devalued reinforcer. Traditionally, such residual responding is seen as evidence that the stimulus—or by inference, the context—is sufficient to evoke responding regardless of the current value the outcome delivered for the response (Adams, 1980; Colwill & Rescorla, 1986). An alternative explanation might be that the aversion conditioned to O₂ was partial or incomplete. However, a considerable amount of residual responding was still seen when we included only the animals that *completely* rejected O₂ consumption during (1) the final devaluation cycle and (2) the reacquisition test in the data analysis. Furthermore, responding during the reacquisition test did not differ from that in extinction, suggesting that the outcome was not promoting behavior above the level of responding that occurred without it. Thus, the results suggest that the aversion to O₂ was sufficiently strong that we have to consider the possibility that some responding was not completely controlled by an R–O association. One possibility, of course, is a simple context–R association. Note that this would mean that context–R associations are formed *in addition* to context–(R–O) associations. While the decrement seen in responding for the devalued outcome, relative to the valued outcome, must be due to some knowledge or representation of that outcome (through an R–O association), responding that remains for devalued outcome despite a complete rejection of that outcome could be attributable to a context–R association (Colwill & Rescorla, 1986).

Recent work from our laboratory has more explicitly suggested a role for a context–R association in simple operant learning. For example, Thraikill and Bouton (2014) subjected animals to either a small or a large amount of instrumental training in which lever pressing was associated with a single outcome. Subsequently, half the animals in each group received presentations of the outcome paired with LiCl, and the other half received pairings with saline. The animals given a small amount of instrumental training, but not a large amount, showed a reinforcer devaluation effect, consistent with the idea that the large amount of training resulted in a response (a “habit”) that was not mediated by an R–O association. Nevertheless, when responding was tested both in the original training context and in a context in which no instrumental

training had taken place, a significant context switch effect was evident in all groups. This result implies that the context can influence the response directly, and not (as in the present experiment) by modulation of an R–O association. Thus, the context might also control instrumental responding via a simple context–R association.

It is also worth noting that on its own, residual responding after reinforcer devaluation remains an incomplete and indirect way to evaluate context–R associations. Aside from the aforementioned evidence that habit might be sensitive to context change (Thrailkill & Bouton, 2014), there are no direct methods to test for context–R associations. Several alternative factors could also account for residual responding. First, certain procedures tend to promote more residual responding than do others. For example, Colwill and Rescorla (1990a) conducted an experiment in which the reinforcer was delivered directly into the mouth of the animal following the response through surgical implantation of a fistula. They found that following a devaluation procedure, animals that received the outcome intraorally showed zero residual responding for that reinforcer during an extinction test, as well as during a reacquisition test. However, animals that received reward using the standard magazine delivery method continued to respond during extinction, as well as during the reacquisition test. Colwill and Rescorla (1990a) suggested that magazine delivery of the reinforcer may create weaker evidence of an R–O association because there are other actions (e.g., approach of the food cup) that occur before the animal obtains the reinforcer. Thus, the animal is free to perform the target response during testing and then suppress food cup approach in order to reject an expected outcome it finds aversive. Intraoral delivery eliminates the approach response. Alternatively, intraoral administration of the outcome may promote better generalization between the outcome presented during aversion conditioning and the one presented during instrumental training.

A second possibility to account for the present residual responding is that the training procedure could have promoted a certain amount of generalization between the two contexts. During acquisition, animals had equal experience with the R1–O1 contingency as they had with the R1–O2 contingency (the same was true of the R2–O2 and R2–O1 contingencies), creating the opportunity for each response to become equally associated with the two outcomes. Residual responding could still occur on the response that led to the devalued outcome as a function of how much its association with the *nondevalued* outcome generalized from the other context. On this view, responding could still be governed by an R–O association, and the residual responding during the preference and reacquisition tests results simply from contextual generalization. Additionally, the animals could have generalized to some extent between the two responses. The present data are unable

to conclusively distinguish between these possible explanations of residual responding.

In summary, the present experiment found that in distinct contexts, animals selectively suppressed the response that was previously paired with a reinforcer that had been devalued by its association with LiCl. The results thus confirm that distinct context–(R–O) associations *can* be learned during free operant training and suggest that these associations can be important in guiding preference for instrumental behavior. The results do not preclude the possibility that other types of associations (such as context–R associations) may be formed in conjunction with context–(R–O) associations and account for the portion of behavior that is not mediated by a representation of the outcome. Although such a possibility is consistent with levels of responding observed here during extinction tests, there are other possible accounts of that result.

Author Note This research was supported by NIH Grant RO1 DA033123 to M.E.B. The experiment was included as part of a thesis submitted by S.T. in partial fulfillment of the requirements of a master's degree at the University of Vermont. We thank Eric Thrailkill and Scott Schepers for their comments on the manuscript.

References

- Adams, C. D. (1980). Post-conditioning devaluation of an instrumental reinforcer has no effect on extinction performance. *Quarterly Journal of Experimental Psychology*, *32*, 447–458.
- Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology*, *33B*, 109–121.
- Bouton, M. E., & King, D. A. (1983). Contextual control of the extinction of conditioned fear: Tests for the associative value of the context. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 248–265.
- Bouton, M. E., & Peck, C. A. (1989). Context effects on conditioning, extinction, and reinstatement in an appetitive conditioning preparation. *Animal Learning & Behavior*, *17*, 188–198.
- Bouton, M. E., & Todd, T. P. (2014). A fundamental role for context in instrumental learning and extinction. *Behavioural Processes*, *104*, 13–19.
- Bouton, M. E., Todd, T. P., & León, S. P. (2014). Contextual control of discriminated operant behavior. *Journal of Experimental Psychology: Animal Learning and Cognition*, *40*, 92–105.
- Bouton, M. E., Todd, T. P., Vurbic, D., & Winterbauer, N. E. (2011). Renewal after the extinction of free operant behavior. *Learning & Behavior*, *39*, 57–67.
- Brooks, D. C., & Bouton, M. E. (1994). A retrieval cue attenuates response recovery (renewal) caused by a return to the conditioning context. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 366–379.
- Colwill, R. M., & Rescorla, R. A. (1985). Postconditioning devaluation of a reinforcer affects instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes*, *11*, 120–132.
- Colwill, R. M., & Rescorla, R. A. (1986). Associative structures in instrumental learning. *Psychology of Learning and Motivation*, *20*, 55–104.

- Colwill, R. M., & Rescorla, R. A. (1990a). Effect of reinforcer devaluation on discriminative control of instrumental behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*, 40–47.
- Colwill, R. M., & Rescorla, R. A. (1990b). Evidence for the hierarchical structure of instrumental learning. *Animal Learning & Behavior*, *18*, 71–82.
- Dickinson, A. (1994). Instrumental conditioning. In N. J. Mackintosh (Ed.), *Animal learning and cognition*. New York, NY: Academic Press.
- Field, A. (2005). *Discovering statistics using SPSS*. Thousand Oaks, CA: Sage Publications.
- Harris, J. A., Jones, M. L., Bailey, G. K., & Westbrook, R. F. (2000). Contextual control over conditioned responding in an extinction paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*, *26*, 174–185.
- Nakajima, S., Urushihara, K., & Masaki, T. (2002). Renewal of operant performance formerly eliminated by omission or noncontingency training upon return to the acquisition context. *Learning and Motivation*, *33*, 510–525.
- Thomas, B. L., Larsen, N., & Ayres, J. B. (2003). Role of context similarity in ABA, ABC, and AAB renewal paradigms: Implications for theories of renewal and for treating human phobias. *Learning and Motivation*, *34*, 410–436.
- Thraillkill, E. A., & Bouton, M.E. (2014). Contextual control of instrumental actions and habits. Manuscript submitted for publication.
- Todd, T. P. (2013). Mechanisms of renewal after the extinction of instrumental behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*, 193–207.
- Todd, T. P., Winterbauer, N. E., & Bouton, M. E. (2012). Contextual control of appetite. Renewal of inhibited food-seeking behavior in sated rats after extinction. *Appetite*, *58*(2), 484–489.