

Associative foundation of causal learning in rats

Cody W. Polack · Bridget L. McConnell ·
Ralph R. Miller

© Psychonomic Society, Inc. 2012

Abstract Are humans unique in their ability to interpret exogenous events as causes? We addressed this question by observing the behavior of rats for indications of causal learning. Within an operant motor–sensory preconditioning paradigm, associative surgical techniques revealed that rats attempted to control an outcome (i.e., a potential effect) by manipulating a potential exogenous cause (i.e., an intervention). Rats were able to generate an innocuous auditory stimulus. This stimulus was then paired with an aversive stimulus. The animals subsequently avoided potential generation of the predictive cue, but not if the aversive stimulus was subsequently devalued or the predictive cue was extinguished (Exp. 1). In Experiment 2, we demonstrated that the aversive stimulus we used was in fact aversive, that it was subject to devaluation, that the cue–aversive stimulus pairings did make the cue a conditioned stimulus, and that the cue was subject to extinction. In Experiments 3 and 4, we established that the decrease in leverpressing observed in Experiment 1 was goal-directed instrumental behavior rather than purely a product of Pavlovian conditioning. To the extent that interventions suggest causal reasoning, it appears that causal reasoning can be based on associations between contiguous exogenous events. Thus, contiguity appears capable of establishing causal relationships between exogenous events. Our results challenge the widely held view that

causal learning is uniquely human, and suggest that causal learning is explicable in an associative framework.

Keywords Causal learning · Operant behavior · Causal intervention · Manipulability theory · Associative learning · Sensory preconditioning

The world is malleable and, to survive and reproduce, organisms must learn how and under what conditions their actions will control exogenous events. Causal learning facilitates successful interactions with the environment, such as throwing a switch to illuminate a dark room or pressing a lever to receive food. Actions are commonly viewed as causes of contingent consequences (i.e., outcomes). Provided that there is concordance between the contingent outcome and the subject's motivational state, the probability of the action is widely viewed as reflecting the extent to which the action is perceived as a cause of the outcome. But the basis for concluding that there is a causal relationship between two exogenous events is less clear. In situations in which two events are merely correlated, performing an action to produce one event when the other event is desired would be wasted energy. Thus, distinguishing a cause–effect relationship from mere correlation is critical. The scientific identification of two events as cause and effect is formalized using well-designed experiments; however, individuals outside the laboratory quickly infer causal relationships without such rigor. Even researchers, at some level, must suspect a causal relationship based on mere correlation in order to decide whether it is worth testing a particular manipulation to see if it will produce an effect. Thus, intervention on a particular variable is indicative of a perceived cause–effect relationship rather than a perceived purely predictive (i.e., correlative) relationship, and the consequences of such an

C. W. Polack · R. R. Miller (✉)
Department of Psychology,
State University of New York at Binghamton,
Binghamton, NY 13902-6000, USA
e-mail: rmiller@binghamton.edu

B. L. McConnell
KU Leuven and Research Foundation–Flanders (FWO),
Leuven, Belgium

intervention are often the basis for effectively strengthening or weakening a learned cause–effect relationship. Our interest here is an organism’s predisposition to seemingly learn causal relationships without scientific methodology, as indicated by their manipulating the surrounding environment. Is higher-order mental processing required, or can basic associative mechanisms alone support apparent causal learning? One strategy to answer this question would be to assess an associative account of an instance of causal learning, and another strategy would be to seek causal learning in a species commonly assumed to be capable of little or no higher-order processing. In other words, can we find behaviors that are indicative of causal learning, but are explicable in terms of simple associations? Here, we conjointly implemented both strategies in an effort to demonstrate behavior indicative of causal learning in rats and to identify potential associative structures underlying the behavior.

While most goal-directed actions (at least among humans) are intuitively viewed as reflecting some basic representation of cause and effect, with the action serving as the cause, how one comes to identify two exogenous events as having a causal relationship is less obvious. Some researchers have argued that contiguity and contingency between a candidate cause and an effect are necessary and sufficient for causal inference (Allan, 1993; Allan & Tangen, 2005), while others have claimed that knowledge of the underlying causal mechanism is additionally required (Waldmann & Hagmayer, 2005). The assessment of causal inference is typically done by verbal report, which assumes veridical introspection. But verbal report is often belied by other measures (e.g., Nisbett & Wilson, 1977). Moreover, causal learning presumably did not evolve in order to fuel discussion, but rather to allow organisms to control outcomes. Hence, actions that reflect causal attribution between an exogenous event and an outcome are more compelling evidence than are verbal reports. These actions may not have any real influence on the outcome, as is the case with superstitious behaviors, but the action would indicate some degree of a perceived causal relationship. Skinner (1948) demonstrated that pigeons develop eccentric behaviors even when the action is not a cause of the outcome, in the mechanistic sense of *cause*. Superstitious behaviors are not exempt from the label of *causal learning* simply because the behavior does not reflect a mechanistic causal relationship; the perception of causality is what we are investigating here. Additionally, nonverbal behavioral assessment of causal judgment allows for questions concerning causal learning to be addressed in nonhuman species as well as humans.

Identifying an external stimulus as a mere predictor, as opposed to a cause, of an event is a daunting task if the organism is not able to intervene in some way. Activation of the mental representation of an outcome by a mere predictor may promote anticipatory behavior (e.g., Pavlov’s dog salivating in anticipation of food), but only a perceived causal

relationship will promote an intervention to alter the occurrence of the cause. Consider the position of a light switch that is correlated with the illumination of a room. There is both contingency and contiguity between the switch position and illumination. But whether there is a perceived causal relationship between them can be determined by a subject’s actions upon entering a dark room. The subject’s manipulation of the light switch (i.e., an intervention) demonstrates behavior indicative of causal learning (Killeen, 1981; Woodward, 2003). A signal that is not perceived as a cause would fail to support manipulation of the signal. For example, people do not request that the weatherman predict good weather because we understand at some level that the weatherman’s forecast is only a signal for, and not a cause of, the weather. As humans, we readily interpret the behavior of other humans as being motivated by the perceived causal relationships between action and consequence, both in everyday life and in laboratory situations (e.g., Chapman & Robbins, 1990; Chatlosh, Neunaber, & Wasserman, 1985; Shanks & Dickinson, 1991). Moreover, we draw on our own experiences with interventions to conclude that interventions are motivated by perceived causal relationships. The extent to which an animal manipulates its environment to achieve some desired end consistent with its motivational state suggests, at some level of analysis, a basic understanding of cause and effect; however, researchers hesitate to attribute causal learning to nonhuman animals. To provide a nonverbal behavioral definition of causal learning, we follow Skinner (1938) and assert that operant conditioning may be, in some part, governed by causal reasoning, and vice versa. Operant conditioning would be taken as reflecting causal learning to the extent that the behavior can be identified as goal directed. That is, when behavior is dependent on the motivational value of the goal, whether it is a food pellet or a certain level of illumination in a room, the actions taken can be interpreted as attempts to achieve that goal (i.e., to cause it to occur). The degree to which an animal’s behavior can be said to be goal directed depends on whether changes in the value of the outcome (e.g., unconditioned stimulus [US] devaluation) produce parallel changes in the operant behavior (e.g., Adams & Dickinson, 1981; Colwill & Rescorla, 1985; Dickinson, 1997; Rescorla, 1987). On the basis of this definition, our present concern is whether evidence of causal reasoning extends beyond an animal’s goal-directed actions to exogenous stimuli.

The immediate question is whether animals can learn a seemingly causal relationship between two contiguous exogenous events and whether that learning can be accounted for in an associative framework. We would conclude that animals do learn such relationships if, given the same contingencies and motivation as humans, they effectively did throw light switches and did not call the weatherman. Alternatively, if these behaviors were observed, one could

conclude that neither humans nor nonhuman animals learn causal relationships. But whatever is concluded, it would not be parsimonious to have two accounts, one for humans and the other for nonhumans, of an intervention observed in both humans and nonhumans. Of course, manipulation of the candidate cause may simply reflect associative chains reaching back from the valued effect. However, if this were true, the same analysis could be equally applicable to goal-directed behavior in humans. We do not view associative accounts of goal-directed behavior as an alternative to causal learning; rather, associations may be the basis of causal learning. The assumption that causal learning requires dedicated reasoning abilities in any organism is likely unnecessary. In humans, the view that causal learning is distinct from associative learning is likely due to the bias of personal introspection when interpreting behavior. By demonstrating behavior analogous to causal learning in rats, we hope to eliminate this bias. Thus, we will continue to refer to *causal learning* in the subsequent text only in the sense that we are looking for behaviors that appear to be like those that have been suggested to require causal knowledge, but we will analyze these behaviors in terms of chains of simple associations between events.

Perceived causes differ from mere signals in that organisms may intervene concerning the occurrence of a cause in order to influence the effect. Therefore, we can present an animal with innocuous exogenous event A, followed by biologically relevant exogenous event B. Then, given the opportunity to manipulate A, subjects should do so if A is perceived to cause B, as long as B has some motivational value. This manipulation should increase the likelihood of A if B is appetitive, and decrease the likelihood of A if B is aversive. If devaluation of the outcome (B) after establishment of an action–outcome associative chain (e.g., Response–Event A–Event B) modifies operant responding (e.g., attenuates some intervention), the action can be regarded as goal-directed behavior, which is conceptually identical to behaviors we operationally define as reflecting causal learning (Dickinson, Campos, Varga, & Balleine, 1996; Parkinson, Roberts, Everitt, & Di Ciano, 2005). Here, we wanted to determine whether causal learning has occurred when the intervention is only indirectly related to the goal. Thus, we can assess whether the relationship between two exogenous stimuli supports goal-directed interventions on the earlier-occurring of the two stimuli. This would establish whether the first occurring stimulus effectively has causal properties. Ultimately, we will suggest that behavior indicative of causal learning may be based on simple associative chains. Alternatively, one might conclude that rats possess higher-order reasoning skills that are typically reserved for explanations of human performance. However, our data suggest that this assumption is not necessary to produce behavior indicative of causal learning in the present task.

Experiment 1

The ability of nonhuman animals to make causal inferences about external stimuli has been a contested issue (e.g., Dwyer, Starns, & Honey, 2009). This series of experiments was designed to provide additional evidence concerning whether rats behave as if they have some sense of causal relations that we can examine within an operant sensory preconditioning experiment modeled after St. Claire-Smith and MacLaren (1983; see Table 1 for our design and the predictions based on the expectation of behavior suggestive of causal learning). Leverpressing was initially elevated for all rats via water reinforcement, and then all subjects learned that leverpressing produced a tone (i.e., a conditioned stimulus [CS]). Subsequently, the CS was either paired (i.e., Condition Pair) or not paired (i.e., Condition Unp) with an aversively loud click train (i.e., a US). Using an unpaired control allowed for equal exposure to the critical stimuli across groups.¹ Orthogonal to the paired/unpaired variable, animals then received devaluation of the US (i.e., Condition Dev), extinction of the CS (i.e., Condition Ext), or merely equivalent context exposure (i.e., Condition Ctx). Following these treatments, animals were allowed to leverpress in a test session without any programmed stimulus presentations or consequences.

The evidence for causal learning here is predicated on the assumption that subjects will intervene to control the occurrence of a cause of an affective event (i.e., the US) but not a mere signal for the same event. However, having an exogenous cue immediately precede a biologically relevant outcome might lead an organism to treat the predictor as if it had causal properties. The present experiment, and the series as a whole, hinges on the assumption that animals will avoid making a response that produces a stimulus that is perceived to be a cause of an aversive event (e.g., Woodward, 2003). If the stimulus is perceived to be only a signal of the US, the animal should not alter its behavior to reduce the frequency of the predictive stimulus, because a mere predictor does not actually affect the likelihood of the outcome's occurrence. A subject's manipulation of a potential cause is an indirect way for the subject to control (in this case, reduce) the occurrence of the putative effect. Manipulating the occurrence of a mere predictor (i.e., not a cause), by definition, does not affect the occurrence of an outcome and should not support a reduction in the response. Importantly, we assumed that

¹ In principle, the use of this control treatment could make the CS a safety signal for the control groups. Causal learning in our basic experimental group could also be manifest if we were to see an increase in leverpressing to generate this putative safety signal. We consider this to be highly unlikely, as two presentations of the CS and of the US would normally be far too few to produce a reliable safety signal. Therefore, we will refrain from considering control of a safety signal in our preparation, although this interpretation would not undermine our discussion of rats manipulating a predictor of the US only if the predictor were viewed as a cause.

Table 1 Design of Experiment 1: Assessing the associative chain in M-SPC

Groups	Shaping	M-SPC	Fear Conditioning	Dev or Ext	Test
	<i>Ctx Shape</i>	<i>Ctx Train</i>	<i>Ctx Train</i>	<i>Ctx Train</i>	Avoidance Expected <i>Ctx Train</i>
Pair-Dev	LP→[H ₂ O+N]	LP→T	2 T→US	24 US	ca
Pair-Ext	LP→[H ₂ O+N]	LP→T	2 T→US	60 T–	ca
Pair-Ctx	LP→[H ₂ O+N]	LP→T	2 T→US	Context	CA
Unp-Dev	LP→[H ₂ O+N]	LP→T	2 T / 2 US	24 US	ca
Unp-Ext	LP→[H ₂ O+N]	LP→T	2 T / 2 US	60 T–	ca
Unp-Ctx	LP→[H ₂ O+N]	LP→T	2 T / 2 US	Context	ca

Pair, paired; Unp, unpaired; Dev, US devaluation; Ext, extinction; Ctx, context; LP, leverpress; M-SPC, operant motor–sensory preconditioning; N, white noise; T, tone; US, clicks; ca and CA, weak and strong conditioned avoidance, respectively. The expected results are based on behavior indicative of causal learning

changes in operant responding consistent with the subject's motivation reflected goal-directed behavior, which we equated with intervention, and therefore constituted behavior indicative of causal knowledge.

The question addressed in Experiment 1 was whether rats would form a chain of associations in which leverpressing initially becomes associated with the CS during leverpress → tone pairings. Then, when the CS and US were paired, a CS → US association would be formed, completing the chain leading from the act of leverpressing to the innocuous tone to the aversive outcome—that is, leverpressing → CS → US. Changes in behavior resulting from this associative structure would be indicative of so-called *causal learning* (i.e., the CS causes the US) if, at test, the rate of leverpressing was reduced to decrease the occurrence of the CS. However, alternative associative accounts might also lead to reduced leverpressing, without appearing to be driven by causal learning. One such possibility is that through conditioning, the CS becomes a second-order aversive reinforcer. This might result in a direct S–R association between the CS and freezing, which is generally considered to be insensitive to manipulations that devalue the US (Holland & Rescorla, 1975). We can assess whether the predicted reduced leverpressing was due to an associative structure indicative of causal learning or to simple second-order reinforcement by using a procedure such as US habituation (e.g., Rescorla, 1973). If the change in instrumental behavior in this preparation were a form of goal-directed behavior, it should be sensitive to the value of the US; US habituation should then result in a decrement in the conditioned suppression of leverpressing. However, if the subjects formed a direct association between the CS and conditioned suppression, we would expect little or no decrement in leverpress suppression as a result of US devaluation. Thus, devaluation of the US allowed us to assess the amount of response suppression that was the result of the CS becoming a second-order reinforcer.

Another alternative to our presumed associative structure underlying causal learning is that CS–US pairings following

the leverpress–CS pairings might have resulted in a direct leverpress–US association, which, in principle, could account for decreased leverpressing at test. Thus, we included a condition in which the CS was extinguished, which allowed for assessment of any direct response–US association that might not have been mediated by the CS. We expected that extinguishing the CS would disrupt the chain of associations presumably underlying causal learning, and would consequently result in a decrement in the conditioned suppression of leverpressing. However, CS extinction should have no effect on leverpressing if the CS did not mediate this change in behavior.

Our central interest was in Group Pair-Ctx, whose behavior would allow us to assess whether leverpressing was suppressed. Without suppression in Group Pair-Ctx, there would be no point in trying to identify the underlying associative structure. Goal-directed behavior, presumably done as an effort to reduce occurrences of the CS, would be indicated as the basis of reduced leverpressing in Group Pair-Ctx if the control groups allowed us to discount a role for the CS having become an aversive reinforcer (i.e., Group Pair-Dev) and a possible role for a direct leverpress–US association that bypassed the CS (i.e., Group Pair-Ext). If the CS–US association were viewed only as the basis of the CS signaling the US instead of serving as a causal agent, the subjects should then not try to manipulate (i.e., reduce) the CS occurrences, just as people do not call the weatherman to prevent rain. A mere predictor does not evoke goal-directed behavior (which in this experiment would be reduced leverpressing) that is designed to modify the predictor.

Method

Subjects

The subjects were 30 male and 30 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony. The mean body weights were 267 g

for males (range 215–339 g) and 196 g for females (range 162–253 g). The subjects were randomly assigned to one of six groups: Pair-Dev, Pair-Ext, Pair-Ctx, Unp-Dev, Unp-Ext, and Unp-Ctx, $n_s = 10$, counterbalanced for sex within groups. The animals were individually housed in a vivarium maintained on a 16/8-h light/dark cycle. The experimental manipulations occurred near the middle portion of the light phase. The animals received free access to Purina Lab Chow, whereas water availability was limited to 20 min per day, following a progressive deprivation schedule initiated 1 week prior to the start of the study.

Apparatus

The apparatus consisted of 12 operant chambers, each measuring $30 \times 30 \times 27$ cm ($l \times w \times h$). The side walls of the chamber were made of stainless steel sheet metal, and the front walls, back walls, and ceilings of the chambers were made of clear Plexiglas. The floor was constructed of 0.3-cm-diameter stainless steel rods, spaced 1.3 cm center to center. Each chamber was housed in a separate environmental isolation chest. This chamber could be dimly illuminated by a houselight (1.12-W, #1820 incandescent bulb) mounted high on a wall of the experimental chamber. Each chamber was equipped with a Med Associates (St. Albans, VT) retractable lever 5 cm wide, which could extend 2 cm into the chamber or could be retracted into a recess in a wall. Each chamber was equipped on the same wall with a solenoid liquid dispenser produced by Lafayette Instruments (Model 80201), which could deliver 0.04 ml of water through the ceiling of a square recess ($3 \times 6 \times 6$ cm, $l \times w \times h$) that was located to the right of each lever. The bottom of the recess, which contained a cup to catch the delivered water, was 3 cm above the floor of the chamber. This constituted Context Train, which was used for all phases of treatment except leverpress shaping.

Ventilation fans in each enclosure provided a constant 76-dB (all auditory measures are based on the C-scale) background noise. Three 45- Ω speakers mounted on the interior right, left, and back sides of each environmental chest were used to deliver a 2-s complex tone (500 and 520 Hz presented simultaneously) that served as the target CS, a 5-s click train (6/s) that served as a US, and a 0.5-s white noise that signaled water delivery, with water and white noise being delivered only during initial shaping of leverpressing. The tone and white noise were each 8 dB above the 76-dB background, whereas the loud click-train US was 30 dB above the background. Each environmental chamber also contained a 100-W bulb, nominal at 120 VAC but driven at 60 VAC, that could emit a 2-s flashing light (0.25 s on, 0.25 s off) that would serve as stimulus L in Experiments 2, 3, and 4.

The apparatus was changed between leverpress shaping and the operant motor–sensory preconditioning phase in order to minimize instrumental latent inhibition of leverpressing interfering with acquisition of the leverpress–CS association during the subsequent operant motor–sensory preconditioning phase. This manipulation was predicated on prior research that showed latent inhibition to be relatively context specific (e.g., Hall & Honey, 1989). In order to make the context of leverpress shaping distinctive from that of training and testing, the houselight was turned off, a Plexiglas sheet covered the grid floor, an odor cue (methyl salicylate) was present in Context Shape, and the animals received leverpress shaping in a chamber (Context Shape) different from that used during the remainder of the experiment (Context Train).

Procedure

Shaping Over a period of 6 days, all animals were shaped in 60-min daily sessions to leverpress for water on a variable-interval (VI) 40-s schedule in Context Shape. To facilitate leverpress shaping, a 0.5-s white noise presentation coincided with water delivery. On Days 1 and 2 of shaping, water was delivered on a fixed-time 120-s schedule in addition to a continuous-reinforcement schedule for leverpressing. On Day 3, the fixed-time schedule was terminated, and leverpressing was continuously reinforced with water delivery and the white noise. Rats that did not emit at least 50 leverpresses during the session received hand shaping on that day. On Days 4 and 5, the continuous-reinforcement schedule was replaced with a VI 20-s schedule, and on Day 6, the subjects were shifted to the VI 40-s schedule. By the end of Day 6, all subjects met a minimum leverpressing requirement of 50 responses in 1 h.

Operant motor–sensory preconditioning On Day 7, all of the subjects were introduced to Context Train for a single 20-min session. The context shift was intended to encourage new learning during this phase by reducing the effects of operant latent inhibition resulting from prior leverpressing for water. The animals were placed into their respective chambers with the lever extended and received 10 leverpress–tone pairings on a continuous-reinforcement schedule. After an animal completed 10 such pairings, the lever in that chamber was retracted for the remainder of the session to prevent additional, nonreinforced leverpresses from occurring. Additional leverpresses during this phase would have a deleterious effect on our target behavior (baseline leverpressing) as well as adding within- and possibly between-group variance to the number of leverpress–tone pairings during acquisition of the leverpress–tone relationship; therefore, we restricted the animals' opportunity to make

leverpresses beyond the 10 allocated for leverpress–tone training. Although this phase of treatment was expected to reduce overall leverpressing, due to the removal of water reinforcement, the decrement was expected to be minimal, considering the small number of allowed leverpresses and the relatively lean water reinforcement schedule (i.e., VI-40) used during shaping. In any event, the effects of extinction and generalization decrement due to the change in context should have been equivalent across groups.

Fear conditioning On Day 8, in the absence of the lever and water delivery, the animals experienced one 30-min conditioning session in Context Train. There was the potential for the learning that occurred during this phase to become specific to when the lever was absent. However, such a specificity of fear conditioning would have limited the effect of this procedure on leverpressing equally across all paired groups, uniformly reducing the likelihood of our observing differences between Conditions Pair and Unp. The three groups in the paired condition received two presentations of the 2-s tone each, followed immediately by a 5-s US which began upon CS termination. For the paired groups, these presentations began at 5 and 20 min into the session. For the unpaired groups, the US occurred at the same times as in the paired groups, but the CS occurred at 12 and 27 min into the training session.

Devaluation or extinction On Days 9 and 10, animals in Groups Pair-Dev and Unp-Dev received one daily 60-min session of US-alone exposure, whereas Groups Pair-Ext and Unp-Ext received one daily session of CS-alone exposure, and Groups Pair-Ctx and Unp-Ctx experienced daily sessions of equivalent context exposure in Context Train. The levers were retracted throughout this phase for all groups. Subjects in Condition Dev received 12 presentations of the 5-s US per session with pseudorandomized intertrial intervals (ITIs; $M = 5$ min). Subjects in Condition Ext received 30 presentations of the 2-s tone per session with pseudorandomized ITIs ($M = 2$ min). These treatments were demonstrated in the subsequent Experiment 2 to effectively attenuate behavioral control by the US and the CS, respectively.

Testing On Day 11, the levers were inserted at the start of the session, and animals were tested in a 4-min session in Context Train. Responding during the session was recorded in 5-s-long bins. No tones, water, white noise, or clicks were delivered during testing. Avoidance was assessed directly from the number of leverpresses. An elimination criterion was imposed on the data analysis, in which subjects were

eliminated if they had not completed 10 leverpress–CS presentations during the operant motor–sensory preconditioning phase. In practice, none of the rats in Experiment 1 met this elimination criterion.

Results and discussion

The groups' mean leverpresses are indicated in Fig. 1. A 2 (contingency: paired vs. unpaired) \times 3 (treatment: Dev vs. Ext vs. Ctx) factorial ANOVA was conducted on the cumulative leverpresses over the 4-min test session. We found no significant main effect of contingency (i.e., paired vs. unpaired), $p > .17$. This null result is not surprising, given that we expected both the devaluation and extinction treatments to attenuate the impact of paired training. We did find a main effect of treatment (i.e., Dev vs. Ext vs. Ctx), $F(2, 54) = 4.19$, $MSE = 229$, $p < .05$, Cohen's $f = 0.33$, and, critically, an interaction between contingency and treatment, $F(2, 54) = 5.05$, $MSE = 229$, $p < .05$, Cohen's $f = 0.37$. A series of planned contrasts followed, in order to illuminate the source of the interaction. Critically, Group Pair-Ctx reduced leverpressing more than did Group Unp-Ctx, $F(1, 54) = 11.34$, $p < .05$, as a result of the CS–US pairings in the paired subjects relative to the unpaired subjects. If some of this reduction in leverpressing was due to avoidance, it could have been dependent on the CS being perceived as a cause of the aversive US. In light of the lack of a reduction in leverpressing following extinction of the CS (Group Pair-Ext vs. Group Unp-Ext), the reduction in leverpressing observed in Group Pair-Ctx relative to Group Unp-Ctx cannot be explained by aversion to either the CS (i.e., second-order conditioning) or the act of leverpressing (i.e., a direct leverpress–US association). Moreover, the high level of responding following devaluation of the US (Group Pair-Dev vs. Group Unp-Dev) further indicated that the reduced leverpressing observed in Group Pair-Ctx relative to Group Unp-Ctx cannot be explained by the CS having become an independent second-order reinforcer. Neither Group Pair-Dev nor Group Pair-Ext differed significantly

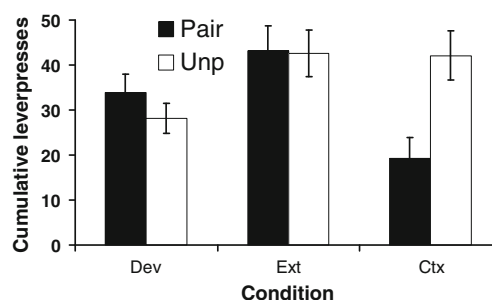


Fig. 1 Experiment 1: Group means representing the average amounts of leverpresses over the 4-min test session for each group. Lower scores represent greater reduction of leverpressing. The black bars represent the paired condition, and the white bars represent the unpaired condition. Error bars represent standard errors.

from their unpaired controls, smallest $p > .39$. Group Pair-Ctx also exhibited more of a reduction in leverpressing than did Groups Pair-Ext, $F(1, 54) = 12.36, p < .05$, and Pair-Dev, $F(1, 54) = 4.65, p < .05$. The lower level of leverpressing in Condition Dev relative to Condition Ext is likely due to some excitatory conditioning to the test context, resulting from many presentations of the US in that context. This is not to say that these associations could not contribute, but only that, with the present parameters, their contributions were negligible.

These statistics support the view that the animals formed a chain of associations in which leverpressing activated a representation of the CS, and critically, the CS activated a representation of the US. Importantly, the lack of any significant avoidance in either Group Pair-Ext or Group Pair-Dev suggests a lack of any appreciable aversion generated directly by the CS or leverpressing in this preparation, leaving only the leverpress \rightarrow CS \rightarrow US associative chain to explain the majority of the reduced leverpressing (i.e., avoidance) observed in Group Pair-Ctx. According to our analysis, the subjects in Group Pair-Ctx avoided making the operant response with the goal of reducing the occurrence of the US.

It should be noted that, in principle, second-order reinforcement might or might not be diminished by subsequent US devaluation. However, an account based on the CS becoming a second-order reinforcer that is independent of the US can be rejected because the reduced leverpressing was not observed following US devaluation. “Second-order reinforcement dependent on a CS–US association” is alternative terminology for goal-oriented behavior, which we take as being indicative of causal learning, provided that the subject intervenes to alter the frequency of occurrence of the second-order reinforcer (i.e., the tone). However, within this framework, it is still possible that the CS could have been viewed as merely signaling the US rather than causing it. The former view is congruent with a purely Pavlovian conditioning account, in which conditioned freezing competes with leverpressing and does not reflect goal-oriented behavior. This response competition account is based on Pavlovian conditioning in which the leverpress \rightarrow CS \rightarrow US associative chain in Experiment 1 mediated an expectation of the aversive outcome that triggered Pavlovian freezing. Importantly, this framework precludes any role of causal learning, which we have characterized as attempts to control the outcome via operant goal-directed behavior in the form of reduced leverpressing. Both accounts rely on the leverpress \rightarrow CS \rightarrow US associative chain identified in Experiment 1, but only goal-directed behavior (which, in this case, is suppression of operant responding) could be viewed as an intervention. Ultimately, a conclusion concerning whether a response competition or causal-learning account applies to the observations of Experiment 1

depends on whether responding was motivated by Pavlovian or operant mechanisms, respectively. Experiments 3 and 4 were conducted to differentiate between the response competition and causal accounts of the reduced leverpressing observed in Experiment 1. It should be noted that an account based on response competition here is different from that previously proposed (Dwyer et al., 2009) to account for prior claims of causal learning in rats (Blaisdell, Sawa, Leising, & Waldmann, 2006; Leising, Wong, Waldmann, & Blaisdell, 2008). In that case, two active responses (nose poking and leverpressing) competed with one another (see the General Discussion for further details). The response competition account proposed here is that the passive response of freezing competed with the active response of leverpressing.

In summary, Experiment 1 suggested that, if the reduced leverpressing seen in Group Pair-Ctx in Experiment 1 can be identified as goal-directed behavior, rats not only perceived a causal relationship between an intervention (leverpressing) and an outcome (the tone, in this case), but were also capable of behaving as if they perceived a causal relationship between two exogenous events (the tonal CS and the aversive US). The present conclusions (in conjunction with the results of Exps. 3 and 4) indicate that causal perception can arise from purely associative structures based on the chaining of related events.

Experiment 2

Experiment 1 was the central experiment in this series. Experiment 2 was merely a direct assessment of the efficacy of the CS–US conditioning, CS extinction, and US devaluation treatments used in Experiment 1. To do this, we assessed the aversive qualities of the tone (CS) and the loud clicks (US). We sought to demonstrate that the US was initially aversive and was effectively devalued via repeated presentations (i.e., habituation). Additionally, we sought to demonstrate that the CS was initially neutral and that CS–US pairings produced a reliable conditioned response that could then be extinguished by CS-alone presentations. US devaluation of the clicks and experimental extinction of the tone were conducted in manners identical to those of Experiment 1, and the operant motor–sensory preconditioning procedure of Experiment 1 was used to assess stimulus control (see Table 2 for the present design). With Group US-Dev, we tested the effectiveness of repeated presentations of the US in attenuating unconditioned responding (UR) to the US (measured by suppression of ongoing leverpress responding in the presence of the US). Group US-Ctx received no additional US presentations during the devaluation or extinction phase, thereby providing a measure of an intact UR to the clicks that could be compared to the UR following US devaluation (Group US-

Table 2 Design of Experiment 2: Assessing the effectiveness of the critical manipulations used in Experiment 1

Groups	Shaping	Sham M-SPC	Fear Conditioning	Dev or Ext	Test Expected Response
	<i>Ctx Shape</i>	<i>Ctx Train</i>	<i>Ctx Train</i>	<i>Ctx Train</i>	<i>Ctx Train</i>
US-Dev	LP→[H ₂ O+N]	LP→L	2 T→US	24 US	US→ur
US-Ctrl	LP→[H ₂ O+N]	LP→L	2 T→US	24 US	no stimulus→baseline
US-Ctx	LP→[H ₂ O+N]	LP→L	2 T→US	Context	US→UR
T-Ext	LP→[H ₂ O+N]	LP→L	2 T→US	60 T-	T→cr
T-Ctrl	LP→[H ₂ O+N]	LP→L	2 T / 2 US	Context	T→cr
T-Ctx	LP→[H ₂ O+N]	LP→L	2 T→US	Context	T→CR

US, unconditioned stimulus (clicks); Dev, US devaluation; Ctx, context; Ext, extinction; LP, leverpress; N, white noise; M-SPC, operant motor-sensory preconditioning; L, flashing light; T, tone; cr and CR, weak and strong conditioned suppression, respectively; ur and UR, weak and strong unconditioned suppression, respectively; →, followed by; /, unpaired with; baseline, suppression of leverpressing with no stimulus presentation. Note that functionally Experiment 2 consisted of two independent experiments, the US condition which assessed suppression to the US with and without US devaluation and baseline leverpressing, and the tone (T) condition which assessed suppression to the tone as a CS with and without extinction and baseline suppression to the tone when it had not been paired with the US

Dev). To assess the completeness of our devaluation treatment, we also compared the UR in Group US-Dev to the performance of a group that leverpressed in the absence of any stimulus presentations during test (Group US-Ctrl).

We additionally tested the effectiveness of nonreinforced presentations of the CS (i.e., tone), following CS-US pairings, in producing behavioral extinction of conditioned suppression of leverpress responding. In a manner analogous to the devaluation comparisons, following CS-US pairings, we compared conditioned suppression in Group T-Ext, which received nonreinforced presentations of the CS, to that in Group T-Ctx, which did not receive extinction treatment, to assess whether extinction was effective relative to unaltered conditioned responding to the CS. Additionally, we were interested in the potential of the CS to elicit suppression of leverpressing without prior CS-US pairings. Therefore, we included Group T-Ctrl, which received presentations of the CS unpaired with the US. During the test phase, the CS was presented during leverpressing, as in the other CS test groups. This allowed us to compare the level of unconditioned suppression to the CS with the amount of conditioned suppression that was acquired as a result of the CS-US pairings.

At test, the subjects in Group US-Dev (which received US devaluation treatment) were expected to show attenuated suppression to the US relative to Group US-Ctx, which did not experience devaluation of the US. Weak suppression was also expected in Group US-Ctrl, which was tested in the absence of any overt stimulus, thereby allowing assessment of baseline levels of leverpressing. Within Condition T, attenuated conditioned suppression to the tonal CS was expected in Group T-Ext, which received extinction training of the CS, relative to the subjects in Group T-Ctx, which did not receive extinction training. Weak conditioned suppression was also predicted in Group T-Ctrl, which was tested in the presence of the CS when the nominal CS was an associatively neutral stimulus.

Method

Subjects and apparatus

The subjects were 36 male and 36 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony. The mean body weights were 318 g for males (range 232–362 g) and 224 g for females (range 184–281 g). The subjects were randomly assigned to one of six groups: US-Dev, US-Ctrl, US-Ctx, T-Ext, T-Ctrl, and T-Ctx, all $n_s = 12$, counterbalanced for sex within groups. The same apparatus, stimuli, and context manipulations described in Experiment 1 were used.

Procedure

Shaping On Days 1–6, the rats were shaped in Context Shape to leverpress for water (concurrent with a 0.5-s white noise) on a VI 40-s schedule using procedures identical to those in Experiment 1. Shaping provided the relatively high rates of leverpressing that were needed to implement operant motor-sensory preconditioning.

Operant motor-sensory preconditioning (sham treatment) On Day 7, all subjects were introduced to Context Train with the levers extended for a single 20-min session. This phase was identical to that in Experiment 1, with one notable exception: The animals received 10 leverpress-flashing-light (2-s duration) pairings on a continuous-reinforcement schedule rather than leverpress-tone pairings, as we were seeking to directly assess responding to the CS and US in the present experiment. Following the 10th presentation of the flashing light, the lever was retracted to prevent additional, nonreinforced leverpresses from occurring.

Fear conditioning On Day 8, in the absence of the lever and water delivery, animals experienced one 30-min conditioning session in Context Train. Five of the six groups received two presentations of a 2-s tone, each followed immediately by a 5-s loud click train, which served as the US. For the paired groups, these presentations began at 5 and 20 min into the session. For the unpaired group (T-Ctrl), the US occurred at the same time as in the paired groups, but the CS occurred at 12 and 27 min into the training session.

Devaluation or extinction On Days 9 and 10, the animals in Groups US-Dev, US-Ctrl, and T-Ext received exposure to the US alone or the CS alone in Context Train, while the US-Ctx, T-Ctrl, and T-Ctx groups experienced equivalent context exposure during daily 60-min sessions. Groups US-Dev and US-Ctrl received 12 daily 5-s US presentations with pseudorandomized ITIs of 3, 4, 5, 6, and 7 min. These parameters had previously been demonstrated in our laboratory to produce an attenuated fear response to the aversive clicks in a lick suppression preparation (Laborda & Miller, 2011). The rats in Group T-Ext received 30 presentations of the 2-s tone (CS) over the course of each session, with ITIs of 1, 2, and 3 min. These procedures were identical to those of Experiment 1.

Testing On Day 11, the levers were reinserted into Context Train for a 6.5-min session. At the start of the session, all animals were allowed to leverpress in the absence of any nominal stimulus (i.e., water, white noise, flashing light, or tone). At 2.5 and 5.5 min into the test session, all groups except for Group US-Ctrl were presented with a stimulus. The groups in Condition T received 1-min presentations of the CS, while Groups US-Ctrl and US-Dev received 1-min presentations of the loud clicks (US). Leverpress responding was recorded in 30-s-long bins. A suppression ratio was computed using the equation $A / (A + B)$, where A was the response rate during stimulus presentation (or equivalent time bins for Group US-Ctrl) and B was the baseline response rate during the 90 s before the stimulus presentation. The B rate was calculated using a longer period of time in order to have a more reliable baseline for the suppression ratio. An elimination criterion based on completion of 10 leverpress–stimulus presentations during the operant motor–sensory preconditioning phase was applied to the data analysis throughout this series. This was done because Experiments 1, 3, and 4 relied on training that occurred during the operant motor–sensory preconditioning phase, and the number of trials that occurred during that phase was partially dependent on the rats' behavior. For consistency, underperforming rats in the present experiment were similarly excluded, even though this phase was irrelevant. Three of the subjects were excluded on the basis of this criterion, two from Group US-Dev and one from Group T-Ctrl.

Results and discussion

Figure 2 depicts the groups' mean suppression ratios. Only data from the first test were used for analysis, due to a general lack of baseline leverpressing during the 1.5 min immediately prior to the beginning of the second test (i.e., after 4 min of nonreinforced leverpressing). A one-way analysis of variance (ANOVA) conducted on the pre-CS scores found no significant effect of condition in baseline performance, $p > .10$. Due to the lack of variance in the suppression ratios of Group US-Ctx, this group was omitted from the ANOVA conducted on the suppression ratios. The ANOVA conducted on the suppression ratios of the remaining groups found an effect of condition, $F(4, 55) = 6.18$, $MSE = .05$, $p < .05$, Cohen's $f = 0.61$. Figure 2 suggests that the loud clicks were an effective US in our preparation, as indexed by the complete suppression in Group US-Ctx relative to Group US-Ctrl, which was tested in the absence of any stimuli and showed essentially no suppression (suppression ratio near 0.5). To assess the effectiveness of the US devaluation treatment, we conducted Mann–Whitney U tests because of the absence of variance in Group US-Ctx. A comparison between Group US-Ctx and US-Ctrl confirmed suppression in the presence of the US, $U = 0$, $z = 4.44$, $p < .05$, whereas the comparison between Groups US-Ctx and US-Dev indicated strong attenuation of this suppression as a result of our devaluation procedure, $U = 6$, $z = 4.16$, $p < .05$. Twenty-four presentations of the US appear to have reduced the US's aversive properties, as evidenced by reduced unconditioned suppression in Group US-Dev relative to Group US-Ctx. Moreover, the level of suppression in Group US-Dev was similar to the level observed in Group US-Ctrl, indicating that the devaluation treatment was nearly complete. This provides evidence that the clicks were initially aversive and that the US devaluation treatment rendered them less so, to the point that they failed to evoke a level of responding that differed significantly from baseline.

In Condition T, strong suppression in the presence of the CS was observed in Group T-Ctx, which received pairings

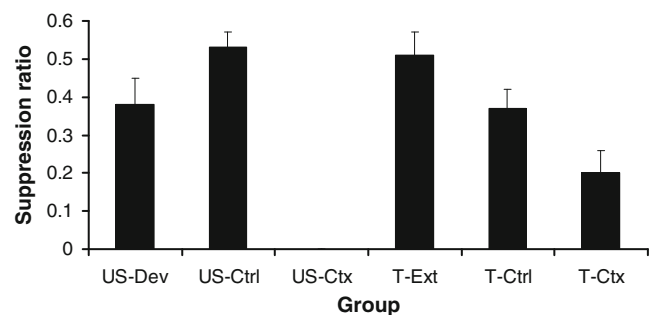


Fig. 2 Mean suppression ratios from Experiment 2. Lower scores represent more suppression (i.e., a greater fear response), and error bars represent standard errors.

of the CS and the US, relative to Group T-Ctrl, for which the CS was associatively neutral, $F(1, 55) = 4.15, p < .05$. This indicates that the CS–click pairings made the CS an excitatory CS. However, 60 presentations of the CS in the absence of the clicks following excitatory conditioning rendered the CS unable to evoke a conditioned response. This is evident in the comparison of Group T-Ext relative to Group T-Ctx, $F(1, 55) = 16.41, p < .05$. Thus, we demonstrated that our devaluation and extinction treatments were effective at attenuating the fear response to the US and the CS, respectively.

Experiment 3

In order to determine whether the reduced leverpressing observed in Experiment 1 arose from goal-directed behavior or conditioned freezing, we used Pavlovian conditioned inhibition as a probe in Experiment 3. It has been observed that Pavlovian inhibitors act to attenuate conditioned responding to CSs that were made excitatory by being paired with a different US of the same valence as the excitator used in inhibition training (Dickinson & Dearing, 1979). This conclusion has been supported by demonstrations in which a Pavlovian conditioned inhibitor for footshock has attenuated responding to a conditioned excitator for an aversive noise (Nieto, 1984). The operant analog to a Pavlovian conditioned inhibitor, a *negative discriminative stimulus*, has been found to be more specialized in its influence. Negative discriminative stimuli have been shown to primarily influence the specific response–reinforcer association with which it was trained, rather than behaviors motivated by a different reinforcer (Bonardi & Hall, 1994). This suggests that, although a Pavlovian inhibitor for a footshock US should transfer to a Pavlovian excitator for a different aversive stimulus (e.g., the clicks used in the present series) due to the inhibitor’s potential to attenuate fear in general, this inhibitor should have little effect on instrumental behavior motivated by aversive auditory reinforcement, such as our loud clicks. Thus, a conditioned inhibitor for footshock should not suppress the avoidance response motivated by goal-directed behavior in Experiment 1 (Group Pair-Ctx), unless that response was due to a Pavlovian fear response. (This is not to say that goal-directed behavior could not be inhibited by a signal trained to inhibit the specific response–outcome relationship within our operant motor–sensory preconditioning procedure; such a signal might be considered a causal preventor.) In other words, if the training in Experiment 1 merely established a Pavlovian predictive relationship, as opposed to a causal relationship, between the CS and the US, a conditioned inhibitor for footshock should reduce conditioned suppression to an associatively activated representation of our auditory US, in the same manner that it

should reduce conditioned suppression to a Pavlovian excitator for the footshock US. But, if the CS was perceived as a cause of the US, the reduction in leverpressing should be less affected by the conditioned inhibitor, because interventions are driven by goal-directed operant learning.

In Experiment 3, as indicated in Table 3, all subjects were given training that established B as a Pavlovian conditioned inhibitor of footshock (US1). All subjects also received standard operant motor–sensory preconditioning training with CS X. Following this phase, Condition Operant (Op) received CS X–US2 pairings in which the US was a loud click train, just as in Experiment 1. Condition Pavlovian (Pav) received CS Y–US2 pairings, to establish Y as a simple Pavlovian excitator. At test, half of the groups, orthogonal to the Op/Pav factor, were tested on Y in compound with neutral stimulus Z, and the other half were tested on Y in compound with conditioned inhibitor B. If the results of Experiment 1 were due to causal learning, we should see moderate reduced responding in Condition Op, indicative of avoidance of leverpressing, and little effect of the inhibitor, whereas the conditioned inhibitor would be expected to reduce conditioned suppression in Condition Pav if conditioned-inhibition training was successful.

Method

Subjects

The subjects were 24 male and 24 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony. The mean body weights were 284 g for males (range 207–336 g) and 200 g for females (range 160–231 g). The subjects were randomly assigned to one of four groups: Op-Inh, Op-Neut, Pav-Inh, and Pav-Neut, $n_s = 12$, counterbalanced for sex within groups. The rats were housed and water deprived as in Experiment 1.

Apparatus

The apparatus consisted of the 12 chambers used in the prior experiments. Additionally, a 1900-Hz SonAlert and a buzzer mounted in each environment chest could each deliver a signal 8 dB above background that would serve as cues B and Z, counterbalanced within groups. The SonAlert, buzzer, and flashing light were all presented for 30-s, except during testing, in which cues were presented throughout the entire test session to better approximate the test conditions of Experiment 1. A 0.5-s duration, 0.8-mA constant current footshock, which served as US1, could be delivered through the grid floor of the chamber during conditioned inhibition training.

The apparatus was altered as in the prior experiments to create two contexts, indicated as Context 1 (Ctx 1) and

Table 3 Design of Experiment 3: Addressing response competition with a conditioned inhibitor

Groups	Accl Ctx 2	CI Train Ctx 1	Shaping Ctx 1	M-SPC Ctx 2	Fear Conditioning Ctx 2	Test Ctx 2	Expected Results
Op-Inh	2 B / 2 Z	21 L-US1 / 60 LB-	LP→[H ₂ O+HL off]	10 LP→X	2 X-US2 / 2 Y	YB	Cr
Op-Neut	2 B / 2 Z	21 L-US1 / 60 LB-	LP→[H ₂ O+HL off]	10 LP→X	2 X-US2 / 2 Y	YZ	Cr
Pav-Inh	2 B / 2 Z	21 L-US1 / 60 LB-	LP→[H ₂ O+HL off]	10 LP→X	2 Y-US2 / 2 X	YB	cr
Pav-Neut	2 B / 2 Z	21 L-US1 / 60 LB-	LP→[H ₂ O+HL off]	10 LP→X	2 Y-US2 / 2 X	YZ	Cr

Op, operant condition; Pav, Pavlovian condition; Inh, inhibition condition; Neut, neutral condition; Ctx, context; B and Z, SonAlert and buzzer (counterbalanced); L, flashing light; US1, 0.8-mA footshock; LP, leverpress; HL, houselight; Accl, acclimation; M-SPC, operant motor-sensory preconditioning; X and Y, tone and white noise, counterbalanced; US2, loud clicks; Cr and cr, moderate and weak suppression, respectively. The expected results are based on behavior in the Op condition being instrumental.

Context 2 (Ctx 2), to encourage new learning during the operant motor-sensory preconditioning phase. Ctx 1 differed in a number of features from Context Shape used in the prior experiments (i.e., the houselight was on instead of off, and a grid floor was used rather than a Plexiglas floor). These changes were necessary in order to use the white noise stimulus as an additional cue in later phases of the present experiment and to allow conditioned inhibition training with footshock to occur in the same context as leverpress shaping. The animals received inhibition training and leverpress shaping in Ctx 1, whereas all other treatments and testing took place in Ctx 2.

Procedure

Acclimation to B and Z On Day 1, all rats received two 30-min sessions in Ctx 2 with the lever retracted. During the first session, two 30-s presentations of the buzzer occurred at 10 and 23 min into the session. During the second session, two 30-s presentations of the SonAlert occurred at 10 and 23 min into the session. These cue presentations occurred during separate sessions because of equipment limitations prohibiting presentation of both cues during a single session. These presentations were intended to decrease any potential unconditioned responding to cues B and Z in the test context. Additionally, elemental preexposure to these cues was intended to discourage their being perceived as a single configured stimulus with the excitator during conditioned-inhibition training.

Conditioned inhibition On Days 2–4, all rats received 90-min daily sessions in Ctx 1 with the lever retracted. The subjects experienced seven reinforced presentations of the flashing light and 20 nonreinforced presentations of the flashing light concurrent with cue B in each session. The mean intertrial interval was 200-s from stimulus onset to onset. All reinforced trials coterminated with a 0.8-mA constant-current 0.5-s footshock.

Shaping On Days 5–10, in Ctx 1, leverpress shaping occurred. The shaping procedure was identical to that used in the prior experiments, except that a 0.5-s interruption of the houselight rather than the 0.5-s white noise cue signaled water delivery. Leverpress shaping followed conditioned inhibition training in order to allow shaping to reduce any deficit in responding produced by the footshock administered during inhibition training.

Operant motor-sensory preconditioning On Day 11, in Ctx 2, operant motor-sensory preconditioning training was conducted as in Experiment 2. The only difference was that the animals received 10 presentations of cue X, which was now either a tone or white noise, counterbalanced with Y within groups.

Fear conditioning On Day 12, in Ctx 2, animals experienced one 30-min session in the absence of the lever. The two groups in the Pav condition received two 2-s presentations of cue Y followed immediately by a 5-s loud click train, which began upon CS termination and served as an aversive US (as revealed by Exp. 2). Additionally, rats in the Pav condition received two 2-s nonreinforced presentations of cue X. The two groups in the Op condition received similar treatment, except that cue X presentations were paired with the US2 (completing the operant motor-sensory preconditioning procedure, as in Exp. 1), whereas cue Y presentations were nonreinforced. The paired presentations began 5 and 20 min into the session, and the unpaired cue presentations began 12 and 27 min into the session.

Testing On Day 13, all of the rats were tested in Ctx 2. To approximate the testing situation in Experiment 1, the test stimuli were presented immediately upon placing each rat in the chamber. All rats received stimulus Y concurrent with either conditioned inhibitor B, for those in the inhibitor condition, or neutral stimulus Z, for those in the neutral condition. The presence of Y should have had little effect on conditioned suppression in Condition Op because, in these groups, Y was irrelevant with regards to the

presentation of any US. However, in Condition Pav, Y was predictive of US2, and therefore should have produced conditioned suppression. The question of interest in Condition Op was whether the conditioned inhibitor (B) would affect leverpressing. Y was presented at test simply to match test conditions with Condition Pav. Importantly, any effect produced by the presence of Y in Condition Op should have been equal within this condition. Thus, Condition Op should give us a measure of avoidance of leverpressing based on training similar to that used in Experiment 1, whereas in Condition Pav, responding should depend on the stimulus presented at test producing a fear response, rather than avoidance of the leverpress. The test compounds were presented for the duration of the test session. The rats were able to leverpress throughout the 4-min session without water delivery, and their responses were recorded in 5-s time bins throughout the test session. As in Experiment 1, an elimination criterion was imposed on the data analysis, in which subjects were eliminated if they had not completed 10 leverpress–stimulus X presentations during the operant motor–sensory preconditioning phase. Five of the subjects were eliminated on the basis of this criterion, three from Group Pav-Neut and one each from Groups Op-Inh and Op-Neut.

Results and discussion

Figure 3 depicts the leverpressing group means at test. We expected the Pavlovian conditioned inhibitor for footshock (B) to negatively summate with a Pavlovian conditioned excitator for the US (Y in the Pav condition) when they were presented in compound at test, resulting in an attenuated fear response to the excitator. The central question of this experiment was whether the Pavlovian conditioned inhibitor would also attenuate the reduction of leverpressing resulting from an operant motor–sensory preconditioning procedure similar to that used in Experiment 1. A decrease in the

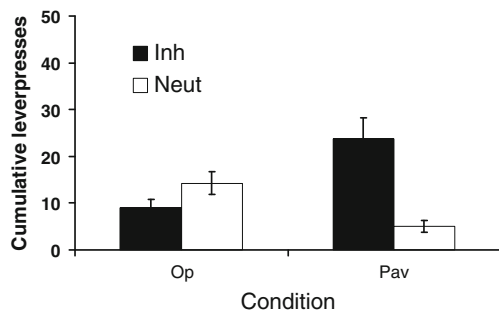


Fig. 3 Experiment 3: Group means representing the average amounts of leverpresses over the 4-min test session for each group. Lower scores represent greater reduction of leverpressing. The black bars represent the inhibitor test condition, and the white bars represent the neutral test condition. Error bars represent standard errors.

previously observed reduction of responding (i.e., an increase in leverpressing) would imply that the results of Experiment 1 could be explained by response competition, which would not support an interpretation based on causal attribution to an exogenous cue. However, if rats reduced leverpressing because leverpressing was perceived as producing a cause of the US, one would expect little effect of the Pavlovian inhibitor in the operant condition.

A 2 (test condition: Pav vs. Op) \times 2 (inhibitor: Inh vs. Neut) factorial ANOVA was conducted on the cumulative leverpresses emitted in the first 4 min of the test session. An interaction, $F(1, 39) = 16.99$, $MSE = 97.09$, $p < .05$, Cohen's $f = 0.61$, and a marginal main effect for the presence of the inhibitor, $F(1, 39) = 3.71$, $p < .07$, Cohen's $f = 0.25$, but no main effect of test type, $F < 1$, were detected. It should be noted that a slight crossover occurred, indicating that the interaction was not due to a ceiling effect or scaling differences across the range of data. Planned contrasts were conducted in order to interpret the source of the interaction. Group Pav-Inh emitted more leverpresses than Group Pav-Neut, indicating that the inhibitory training produced a cue that negatively summated with excitors for aversive USs in general, $F(1, 39) = 17.70$, $p < .05$. There was no significant difference detected between Groups Op-Inh and Op-Neut, indicating that the Pavlovian inhibitor did not negatively summate with other factors that reduced leverpressing in the operant condition, $p > .10$.

A comparison between Groups Op-Inh and Pav-Inh was conducted to determine whether the basic finding observed in Experiment 1 was replicated. If the conditioned inhibitor fully attenuated fear produced by the Pavlovian excitatory cue Y, Group Pav-Inh would be analogous to Group Unp-Ctx in Experiment 1, whereas Group Op-Inh would exhibit reduced leverpressing similar to that produced by operant motor–sensory preconditioning in Group Pair-Ctx in Experiment 1. This provides a conservative assessment of whether we actually observed deliberate avoidance relative to an unpaired control. This assessment is conservative because, to the extent that the inhibitor did not fully attenuate fear, the difference between the two groups should have been reduced. The difference between Groups Op-Inh and Pav-Inh was significant, $F(1, 39) = 13.49$, $p < .05$, thereby supporting the basic finding of Experiment 1. However, direct comparisons across Conditions Op versus Pav should be considered cautiously, considering the differences in the training procedures. Nevertheless, the central finding here is that these two conditions yielded different behaviors in the presence of the inhibitor, on the basis of direct comparisons with their appropriate control groups. Thus, our critical finding was the interaction, which is based on the comparison of the differences within particular conditions, not on direct comparisons between individual groups across conditions.

This experiment was based on the assumption that a conditioned inhibitor attenuates a generalized fear response

to cues but does not as readily attenuate instrumental behavior motivated by a US different from that used in inhibition training. The observed interaction supports the conclusion that the leverpress reduction in Experiment 1 was due to instrumental intervention (avoidance) rather than to a conditioned fear response, such as freezing, competing with leverpressing. In other words, rats avoided performing actions that had previously been learned to produce the CS because the CS was perceived as a cause of the US.

Experiment 4

Experiment 4 complemented Experiment 3 in addressing whether the reduced leverpressing observed in Experiment 1 was due to instrumental intervention or competition with leverpressing by a Pavlovian fear response. Differentiation between these two possibilities in Experiment 4 was based on evidence in the human causal learning literature that causes are more sensitive to cue competition than are Pavlovian signals (e.g., overshadowing, Pineño, Denniston, Beckers, Matute, & Miller, 2005; and blocking, De Houwer, Beckers, & Glautier, 2002). In Experiment 3, we found that a Pavlovian inhibitor for footshock had a greater effect on a Pavlovian signal for an aversive noise than on an intervention to avoid the aversive clicks. In Experiment 4, we asked whether an overshadowing treatment had a greater effect on an intervention to avoid aversive clicks than on a Pavlovian signal (see Table 4). If this expectation were to be confirmed in Experiment 4, Experiments 3 and 4 would conjointly constitute a double dissociation of Pavlovian fear responding and instrumental intervention, and thus would be consistent with causal learning occurring in our operant motor–sensory preconditioning preparation.

Method

Subjects

The subjects were 24 male and 24 female experimentally naive, Sprague-Dawley-descended rats obtained from our

own breeding colony. The mean body weights were 234 g for males (range 173–320 g) and 198 g for females (range 160–230 g). The subjects were randomly assigned to one of four groups: Op-OV, Op-No, Pav-OV, and Pav-No, $n_s = 12$, counterbalanced for sex within groups. The rats were housed as in the prior experiments.

Apparatus

The materials were the same as in Experiment 3; however, the flashing-light stimulus (0.25-s on, 0.25-s off) was 2 s in duration, and neither the buzzer nor the SonAlert was used. The Ctx 1 and Ctx 2 features were consistent with those in Experiment 3.

Procedure

Shaping On Days 1–6, leverpress shaping occurred using a procedure identical to that in Experiment 3.

Operant motor–sensory preconditioning On Day 7, motor–sensory preconditioning was conducted using the procedure of Experiment 3.

Fear conditioning On Day 8, the same fear-conditioning procedure described in Experiment 3 was used for rats in Groups Op-No and Pav-No. Groups Op-OV and Pav-OV received similar training, with the only difference being that the flashing-light stimulus, serving as a potential overshadowing (OV) cue, was presented simultaneous with the reinforced cue, which was X for Group Op-OV and Y for Group Pav-OV.

Testing On Day 9, in Ctx 2, rats were tested using the test procedure described in Experiment 3. As in Experiment 3, an elimination criterion was imposed on the data analysis, in which subjects were eliminated if they had not completed 10 leverpress–stimulus X presentations during the operant motor–sensory preconditioning phase. Three rats were

Table 4 Design of Experiment 4: Addressing response competition with cue competition

Groups	Shaping Ctx 1	M-SPC Ctx 2	Fear Conditioning Ctx 2	Test Ctx 2	Expected Results
Op-OV	LP→[H ₂ O+HL off]	10 LP→X	2 XL-US / 2 Y	Y	cr
Op-No	LP→[H ₂ O+HL off]	10 LP→X	2 X-US / 2 Y	Y	CR
Pav-OV	LP→[H ₂ O+HL off]	10 LP→X	2 YL-US / 2 X	Y	Cr
Pav-No	LP→[H ₂ O+HL off]	10 LP→X	2 Y-US / 2 X	Y	CR

Op, operant condition; Pav, Pavlovian condition; OV, overshadowing condition; No, no-overshadowing condition; Ctx, context; LP, leverpressing; HL, houselight; M-SPC, operant motor–sensory preconditioning; X and Y, tone and white noise, counterbalanced; L, flashing light; US, clicks; cr, Cr, and CR, weak, moderate, and strong suppression, respectively. The expected results are based on behavior in the Op condition being instrumental.

eliminated on the basis of this criterion, one from each of the following groups: Op-No, Pav-OV, and Pav-No.

Results and discussion

We expected the overshadowing treatment to have a larger impact on reduced leverpressing in the operant motor–sensory preconditioning condition than the Pavlovian fear condition if leverpress reduction in the operant condition was due to causal learning. In fact, the overshadowing treatment only attenuated leverpress reduction in the operant condition (see Fig. 4). A 2 (test condition: Op vs. Pav) \times 2 (training condition: OV vs. No) factorial ANOVA was conducted on the cumulative leverpresses emitted during the 4-min test. An interaction was observed, $F(1, 41) = 21.44$, $MSE = 127.09$, $p < .05$, Cohen's $f = 0.67$. This interaction had a slight crossover, suggesting that the interaction was due neither to a ceiling effect nor to differences in scaling across the range of observed leverpressing. Main effects of test condition, $F(1, 41) = 19.23$, $p < .05$, Cohen's $f = 0.64$, and training condition, $F(1, 41) = 17.16$, $p < .05$, Cohen's $f = 0.60$, were also detected. Planned contrasts were conducted in order to identify the source of the interaction. Consistent with the overshadowing treatment having a stronger effect on the operant motor–sensory preconditioning procedure than on the Pavlovian condition, Group Op-OV emitted more leverpresses than did Group Op-No, $F(1, 41) = 40.27$, $p < .05$. In contrast, no difference was observed between Groups Pav-OV and Pav-No, $F < 1$; both groups exhibited low, but not complete, suppression. Rather unexpectedly, there was no overshadowing effect evident in the Pavlovian condition. The present procedures and parameters may not have been sufficiently sensitive to produce appreciable overshadowing of Pavlovian cues. In prior studies assessing cue competition using leverpress suppression, rats were tested under conditions in which they had been trained to stable baselines, and suppression ratios were computed;

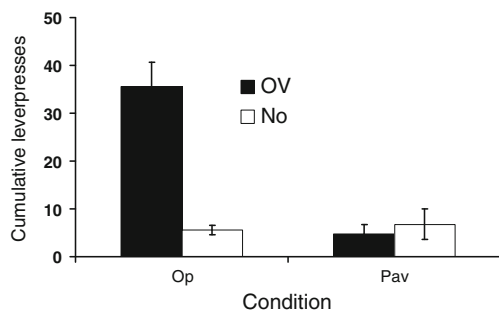


Fig. 4 Experiment 4: Group means representing the average amounts of leverpresses over the 4-min test session for each group. Lower scores represent greater reduction of leverpressing. The black bars represent the overshadowing condition, and the white bars represent the no-overshadowing condition. Error bars represent standard errors.

both of these procedures reduce within-group variability. However, for similarity with Experiment 2, our test situation contained no pre-CS interval and consequently did not allow for the calculation of suppression ratios (i.e., the stimulus was presented at the onset of the test session). That our test detected overshadowing in the operant test condition indicates that the operant test was more sensitive to cue competition than was the Pavlovian test. This suggests that the leverpress reduction in the operant condition was not driven by Pavlovian fear competing with leverpressing. The following statistics support this conclusion. As in Experiment 3, the critical finding was the interaction, not the individual comparisons.

Experiment 4 demonstrated that the operant motor–sensory preconditioning condition is more sensitive to overshadowing than is Pavlovian conditioned suppression, and Experiment 3 demonstrated that the Pavlovian condition was more sensitive to attenuation by a conditioned inhibitor than was the operant condition. Taken together, Experiments 3 and 4 provide a crossed double dissociation between reduced leverpressing produced by the operant motor–sensory preconditioning procedure and that produced by Pavlovian fear. These experiments collectively indicate that the leverpress reduction in Experiment 1 was not due to response competition between conditioned freezing and leverpressing, but instead to instrumental intervention to reduce the frequency of CS occurrences. This testifies to the rats having viewed the CS as a cause of the US. Thus, leverpress suppression was seemingly produced by a chain of associated events due to the same mechanisms that guide reinforced actions. The relationship between two exogenous cues, the CS and the US in Experiment 1, appears to be the same as that involved with throwing a switch to turn on a light or pressing a lever to obtain a food pellet. Thus, rats appear to sometimes treat exogenous cues as if they have the same causal properties as their own actions.

Conclusions

In the present series, we isolated the specific associations involved in an operant motor–sensory preconditioning paradigm and assessed whether this associative structure was based on goal-oriented behavior or on a Pavlovian fear response. Experiment 1 demonstrated that operant motor–sensory preconditioning produced an avoidance response dependent on a particular associative structure, leverpressing \rightarrow CS \rightarrow US, and that extinction of the CS or devaluation of the US attenuated this response. Experiment 2 simply confirmed that our US was able to suppress leverpressing and that US devaluation eliminated this suppression, as well as that the CS–US pairings enabled the CS to suppress leverpressing, whereas extinction of the CS eliminated this suppression. In Experiment 3, we compared the

influence of a conditioned inhibitor on conditioned responding following operant motor–sensory preconditioning relative to Pavlovian conditioned responding, and showed that the operant training condition was less sensitive to negative summation. Finally, in Experiment 4 we found that conditioned responding in the operant training condition was more sensitive to the influence of cue competition than was the Pavlovian training condition. Thus, the associative structure identified in Experiment 1 can be identified as producing a goal-directed instrumental response and was not produced by a conditioned fear response that competed with leverpressing.

An alternative approach to the question of apparent causal knowledge in rats is based on the conclusion that rats perform in accord with causal Bayes net models (Blaisdell et al., 2006). In Blaisdell et al.'s preparation, rats observed a light that predicted the occurrences of food or a tone on separate trials. With few light–food and light–tone trials, integration of these memories led to the expectation of food (indicated by an increase in nose poking into a food hopper) when the tone was presented, presumably because occurrence of the tone implied that the light had occurred, even if it was unobserved. (It is worth noting that with many light–tone trials, an inhibitory relationship between tone and food might have been established; Stout, Escobar, & Miller, 2004.) When the rats were later able to intervene by turning on the tone via leverpressing, there was less expectancy of food than among controls that simply had the tone presented on a yoked schedule. Presumably, leverpressing, having produced the tone, interfered with the diagnostic information provided by the tone (i.e., that the light, and hence the food, had occurred). That is, leverpressing had displaced the light as a cause of the tone. These results led the authors to conclude that the animals initially had inferred the light to be a cause of the tone. Blaisdell et al. rejected traditional associative models of learning because these models did not anticipate that an action–tone pairing would interfere with retrieval of an independently trained light–tone memory. Moreover, associative models that did anticipate such interference (e.g., Anderson, 2003; Miller & Escobar, 2002) failed to anticipate within-trial interference (i.e., interference on the first leverpress–tone trial, which was observed) without assuming real-time information processing (i.e., learning during, as well as at the termination of, a trial). But surely animals are real-time information processors, so these prior observations might be explained by associative interference theory, without assuming any causal attribution to the light. Further experiments demonstrated that presenting an exogenous stimulus in place of the leverpress did not produce the same result (Leising et al., 2008). Consequently, Leising et al. concluded that something is special about an animal's intervention. However, interventions can be viewed as being highly salient, making them more likely to interfere with

prior learning than are exogenous stimuli; although this says something about the special quality of actions, it raises doubt as to whether this implies a causal relationship, as opposed to mere signaling by exogenous stimuli. An additional concern that has been raised regarding this paradigm is that response competition between leverpressing and nose poking, rather than the causal inference that had been theorized, might have produced the observed reduction in nose poking (Dwyer et al., 2009).

Given our working definition of behavior indicative of causal learning being demonstrated by goal-directed behavior, we agree with Blaisdell et al. (2006) that causal learning (or behavior that may appear to reflect causal knowledge) can, in principle, be shown in rats. It is surely possible that the subjects in Blaisdell et al.'s experiment did indeed acquire causal knowledge between leverpressing and the tone. But this conclusion is questionable due to certain assumptions and design features of that experiment. Thus, the present series was designed to determine whether rats exhibit behavior indicative of causal learning between exogenous events using a novel procedure: Rather than assuming competition between independently trained actions and cues, the present set of experiments capitalized on the cause–effect relationship implied by an animal's intervention.

In the present set of experiments, rats that received operant motor–sensory preconditioning did not avoid leverpressing because they were appreciably afraid of the CS per se. Instead, they seemingly had learned that the CS would cause the US (and that leverpressing produced the CS). The data suggest that the CS was not merely a signal for the US, because a signal would not support intervention—in this case, avoidance of leverpressing. We conclude that causal learning by nonhumans is not limited to goal-directed actions, but extends to exogenous events. CS and US events were treated as cause and effect, respectively, after exposure to two pairings. Seemingly, the present apparent causal learning was based on a simple chain of associations, as illustrated by Experiment 1: leverpressing → CS → US. Thus, predictive relationships appear to be able to support causal attribution when no competing information is provided. Any account that anticipates predictive relationships potentially describes causal learning, evidenced through an attempt by the subject to control the predicted outcome, given the opportunity to manipulate the antecedent event.

In Experiments 3 and 4, we found that the predictive relationship was perceived as causal only when the animal was able to act on the predictive cue. That is, cues X and Y had equivalent pairings with the US, but only X had been the consequence of an action by the subjects. This implies that the ability to intervene might not only be necessary to assess causal learning, but also facilitates its development.

The pursuit of exactly how causal learning is established is beyond the scope of our present discussion. To the extent that

nonverbal behavioral assessment of causal learning is telling, our data suggest that rats can form causal relationships between exogenous events. The causal attribution evidenced in these experiments rests on an associative framework, implying that contiguity and contingency are sufficient preconditions for causal inference for rats, and therefore likely are for humans as well. Although the language of this article has been framed in terms of rats being able to causally reason about their environment, the real point is that causal reasoning does not need to be cast as a product of higher-level processing. What looks like causal reasoning can be simultaneously described in terms of the learning of operant contingencies. Experiment 1 demonstrated that these operant contingencies are not necessarily based on directly reinforced behavior; an exogenous stimulus can serve as a mediator for a response–outcome association. Through a double dissociation, Experiments 3 and 4 supported a causal-intervention interpretation of the behavior observed in Experiment 1, as opposed to a Pavlovian, fear-induced freezing account. Critically, behavior suggestive of causal learning appears to parallel that of goal-directed operant conditioning. There may be intense argument over the definition of *causal*, but the behavior is the same as goal-directed behavior. Our conclusions that rats are capable of behaviors that parallel causal learning in humans, and that causal learning in general may result from integrated chains of associations, will surely be met with some skepticism. Ultimately, these assertions should only be taken as suggestive, and such a complex question could only be resolved by further studies with convergent or opposing results.

A number of reviewers have indicated that invoking causal learning to explain our present data is not only unnecessary, but unhelpful. We could not agree more. The purpose of this exercise was to evaluate what behaviors indicate causal knowledge (i.e., interventions via goal-directed actions). What is causal knowledge? Probably, it is nothing more than action motivated by associations. Alternatively, if we conclude that causal learning is more than mere associative learning in human preparations, we may want to attribute to our rats qualitatively similar cognitive abilities.

Author Note This research was supported by National Institute of Mental Health Grant 33881. We thank Aaron Blaisdell, Mario Laborda, Helena Matute, Gonzalo Miguez, James Witnauer, Jason King, Lee Vilinsky, and Ben Williams for comments on earlier versions of the manuscript. *Open Access* This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Adams, C. D., & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology*, *33B*, 109–122.
- Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin*, *114*, 435–448.
- Allan, L. G., & Tangen, J. M. (2005). Judging relationships between events: How do we do it? *Canadian Journal of Experimental Psychology*, *59*, 22–27.
- Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, *49*, 415–445. doi:10.1016/j.jml.2003.08.006
- Blaisdell, A. P., Sawa, K., Leising, K. J., & Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, *311*, 1020–1022. doi:10.1126/science.1121872
- Bonardi, C., & Hall, G. (1994). Discriminative inhibition is specific to the response–reinforcer association but not to the discriminative stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, *20*, 278–291. doi:10.1037/0097-7403.20.3.278
- Chapman, G. B., & Robbins, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, *18*, 537–545.
- Chatlosh, D. L., Neunaber, D. J., & Wasserman, E. A. (1985). Response–outcome contingency: Behavioral and judgmental effects of appetitive and aversive outcomes with college students. *Learning and Motivation*, *16*, 1–34. doi:10.1016/0023-9690(85)90002-5
- 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. doi:10.1037/0097-7403.11.1.120
- De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, *55A*, 965–985.
- Dickinson, A. (1997). Bolles’s psychological syllogism. In M. E. Bouton & M. Fanselow (Eds.), *Learning, motivation, and cognition: The functional behaviorism of Robert C. Bolles* (pp. 345–367). Washington, DC: American Psychological Association.
- Dickinson, A., Campos, J., Varga, Z. I., & Balleine, B. (1996). Bidirectional instrumental conditioning. *Quarterly Journal of Experimental Psychology*, *49B*, 289–306.
- Dickinson, A., & Dearing, M. F. (1979). Appetitive–aversive interactions and inhibitory processes. In A. Dickinson & R. Boakes (Eds.), *Mechanisms of learning and memory* (pp. 203–232). Hillsdale: Erlbaum.
- Dwyer, D. M., Starns, J., & Honey, R. C. (2009). “Causal reasoning” in rats: A reappraisal. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 578–586. doi:10.1037/a0015007
- Hall, G., & Honey, R. C. (1989). Contextual effects in conditioning, latent inhibition, and habituation: Associative and retrieval functions of contextual cues. *Journal of Experimental Psychology: Animal Behavior Processes*, *15*, 232–241.
- Holland, P. C., & Rescorla, R. A. (1975). The effect of two ways of devaluing the unconditioned stimulus after first- and second-order appetitive conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *1*, 355–363. doi:10.1037/0097-7403.1.4.355
- Killeen, P. R. (1981). Learning as causal inference. In M. L. Commons & J. A. Nevin (Eds.), *Quantitative analysis of behavior: Vol. 1. Discriminative properties of reinforcement schedules* (pp. 89–112). Cambridge: Ballinger.
- Laborda, M. A., & Miller, R. R. (2011). S–R associations, their extinction, and recovery in an animal model of anxiety: A new associative account of phobias without recall of original trauma. *Behavior Therapy*, *42*, 153–169.
- Leising, K. J., Wong, J., Waldmann, M. R., & Blaisdell, A. P. (2008). The special status of actions in causal reasoning in rats. *Journal of Experimental Psychology: General*, *137*, 514–527. doi:10.1037/0096-3445.137.3.514
- Miller, R. R., & Escobar, M. (2002). Associative interference between cues and between outcomes presented together and

- presented apart: An integration. *Behavioural Processes*, 57, 163–185.
- Nieto, J. (1984). Transfer of conditioned inhibition across different aversive reinforcers in the rat. *Learning and Motivation*, 15, 37–57.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231–259. doi:10.1037/0033-295X.84.3.231
- Parkinson, J. A., Roberts, A. C., Everitt, B. J., & Di Ciano, P. (2005). Acquisition of instrumental conditioned reinforcement is resistant to the devaluation of the unconditioned stimulus. *Quarterly Journal of Experimental Psychology*, 58B, 19–30. doi:10.1080/02724990444000023
- Pineño, O., Denniston, J., Beckers, T., Matute, H., & Miller, R. R. (2005). Contrasting predictive and causal values of predictors and of causes. *Learning & Behavior*, 33, 184–196.
- Rescorla, R. A. (1973). Effects of US habituation following conditioning. *Journal of Comparative and Physiological Psychology*, 82, 137–143.
- Rescorla, R. A. (1987). A Pavlovian analysis of goal-directed behavior. *American Psychologist*, 42, 119–129.
- Shanks, D. R., & Dickinson, A. (1991). Instrumental judgment and performance under variations in action–outcome contingency and contiguity. *Memory & Cognition*, 19, 353–360.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York: Appleton-Century.
- Skinner, B. F. (1948). “Superstition” in the pigeon. *Journal of Experimental Psychology*, 38, 168–172.
- St Claire-Smith, R., & MacLaren, D. (1983). Response preconditioning effects. *Journal of Experimental Psychology: Animal Behavior Processes*, 9, 41–48.
- Stout, S., Escobar, M., & Miller, R. R. (2004). Trial number and compound stimuli temporal relationship as joint determinants of second-order conditioning and conditioned inhibition. *Learning & Behavior*, 32, 230–239. doi:10.3758/BF03196024
- Waldmann, M. R., & Hagmayer, Y. (2005). Seeing versus doing: Two modes of accessing causal knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 216–227. doi:10.1037/0278-7393.31.2.216
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.