Published in final edited form as:
Behav Processes. 2008 March ; 77(3): 299–305. doi:10.1016/j.beproc.2007.07.003.

Pavlovian Backward Conditioned Inhibition in Humans: Summation and Retardation Tests

Gonzalo P. Urcelay, **Olga Perelmuter**, and **Ralph R. Miller**
State University of New York at Binghamton

Abstract

Two experiments using human participants investigated whether a Pavlovian backward inhibitory treatment (nonreinforced trials in phase 1 followed by reinforced trials in phase 2; i.e., AX-followed by A+) produces a stimulus which can pass summation and retardation tests for inhibition. The rationale for conducting these experiments was that previous demonstrations of Pavlovian backward inhibition informed participants about the nature of the outcome before starting the experiment. According to some theoretical views, this is a potential confound. In the present experiments we used a predictive task in which participants had no knowledge about the outcome until phase 2, when reinforcement occurred. The results of Experiment 1 (summation test) and Experiment 2 (retardation test) provide a clear demonstration of backward conditioned inhibition.

Keywords

human contingency learning; Pavlovian inhibition; backwards conditioned inhibition; summation; retardation; two-test strategy

In the last 20 years an associative approach to human contingency learning and causal judgments has generated a great amount of research (see Dickinson, 2001; De Houwer & Beckers, 2002, for reviews). This approach has emerged as a result of a number of reports suggesting that the principles governing Pavlovian conditioning are applicable to what people do when judging relationships among events. A clear example of this is the allergy causal judgment task in which a participant, playing the role of a physician, judges whether a particular food will lead to an allergic reaction in a given patient, based on prior information about the patient. In this example, foods serve as predictors and the allergic reaction serves as the outcome in a way that is analogous to the tones that serve as conditioned stimuli (CSs) and food pellets that serve as unconditioned stimuli (USs) in a typical animal conditioning study. With this preparation, many of the basic phenomena of Pavlovian conditioning initially observed in the animal literature have been replicated in the human literature.

One of the topics that has received attention in both the nonhuman and human traditions is conditioned inhibition. Within the associative framework, a stimulus is conceptually referred to as a conditioned inhibitor when it signals that an otherwise expected outcome will not occur. Operationally, a stimulus (X) acquires inhibitory properties when it has a negative correlation

Mailing Address: Ralph R. Miller Department of Psychology SUNY – Binghamton Binghamton, NY 13902-6000, USA TEL: (607) 777-2291 FAX: (607) 777-4890 E-mail: rmiller@binghamton.edu.

with the outcome or when it is paired (presented in compound without reinforcement) with another stimulus (A) that is otherwise reinforced. This latter type of inhibition is referred to as Pavlovian inhibition and it is usually obtained after interspersed presentations of the excitor (A) followed by reinforcement (+), and presentations of the excitor compounded with the target inhibitor (X) followed by the absence of reinforcement (i.e., A+ / AX−, Pavlov, 1927). However, intermixing reinforced and nonreinforced trials is not the only procedure that leads to behavior indicative of inhibition. Pavlovian inhibition has also been observed when reinforced and nonreinforced trials are presented in separate phases. For example, one might present all reinforced trials in a first phase (A+) and in a second phase present all nonreinforced trials (AX-; Chapman, 1991). This is usually called forward-trained Pavlovian inhibition. Conversely, when these two phases are reversed backward-trained Pavlovian inhibition training is said to occur. For many years, inhibition acquired through the backward inhibition procedure posed a problem for traditional models of learning that could not explain learning about absent cues (Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981). However, recent revisions of these models (Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994) and models that emphasize the use of information at the time of retrieval (Miller & Matzel, 1988) anticipate the occurrence of backward conditioned inhibition.

There are a number of demonstrations of backward inhibition in the human literature. The first demonstration, to our knowledge, is the one provided by Chapman (1991, Experiments 4 and 5). She compared forward and backward Pavlovian inhibition in a symptom-disease task and observed greater forward than backward Pavlovian conditioned inhibition. Nevertheless, the backward inhibition effect was reliable (Experiment 5). Additionally, Larkin, Aitken, and Dickinson (1998; also see Melchers, Lachnit, & Shanks, 2004), using the allergy task, also found stronger Pavlovian forward than backward inhibition, but backward inhibition was evident. Williams and Docking (1995, Experiment 3) used a stock-market predictive task and summation tests, and they too observed Pavlovian backward inhibition. Finally, using the allergy task, Wasserman, Kao, Van Hamme, Katagiri, and Young (1996, Experiment 5) reported an experiment in which they compared forward and backward conditioned inhibition and found similar levels of inhibition across the two procedures.

All the experiments cited in the previous paragraph provided demonstrations of backward inhibition with humans as participants. However, all these studies made use of preparations that explicitly informed participants about the nature of the outcome prior to the initiation of training. As it is the case with most human contingency learning preparations, before starting the experiment participants were presented with a cover story that gave instructions about the nature of the task and the role participants would assume while participating in the experiment. A typical cover story is the allergy task, in which instructions create an expectation about the occurrence of the outcome (i.e., the allergic reaction) that might be critical for the observation of Pavlovian backward inhibition. Tassoni (1995; Experiment 1) observed lower ratings to a cue that subjects were explicitly informed about its absence relative to a cue that they had not seen on any trial. That is, Tassoni's participants gave lower ratings to a cue when they were explicitly informed it was absent during training (although the participants could see the cue on the screen), relative to a cue for which no information was provided. This observation suggests that in experiments where backward Pavlovian inhibition is observed, the mere mention of the possible occurrence of the outcome will cause participants to believe that nonreinforced cues are preventing the occurrence of the outcome or, in other words, that the target cues and the outcomes are negatively correlated. This suggests that informing subjects about the existence of an outcome creates an expectation of it that would make its absence more noticeable on nonreinforced trials. In situations in which reinforced (A+) and nonreinforced (AX−) trials are interspersed, informing subjects about the occurrence of the outcome previous to the beginning of the experiment does not confound the interpretation because subjects would encounter the outcome anyway soon after the experiment begins.

However, in situations in which all the reinforced trials are presented *after* the nonreinforced trials (i.e., a backward inhibition procedure), subjects presumably revaluate the status of the target cue based on the reinforced trials that occur in phase 2. If subjects were to learn about the outcome before phase 1 trials, then the inhibition observed could be due to heightened expectations of the outcome in phase 1 created by the outcome information, rather than a revaluation process arising from phase 2 reinforced trials.

A second aspect of most of the studies reviewed is that conditioned inhibition was inferred from negative ratings. That is, with the exception of the Williams and Docking (1995) article, most research on conditioned inhibition has been based on participants giving negative ratings to the relationship between the target cue and the outcome. Thus, in those studies the researchers introduced the rating scale and gave explicit instructions about both extremes of the rating scale, emphasizing that if a cue was negatively correlated with the outcome (preventative), then it should be assigned with a negative rating. Williams (1995; 1996) has proposed that a better way to assess inhibitory contingency is to use a transfer test, analogous to the summation test widely used in the animal literature (Rescorla, 1969). In a transfer test participants rate the target cue in compound with a transfer excitor which has been trained independent of the putative inhibitory target cues, and inhibition is inferred if the target cue diminishes the ratings that would be attributed to the excitor alone or to the excitor presented in compound with a novel cue. According to Williams (1995), this procedure avoids giving participants elaborate instructions that might compromise the neutrality of the task. For example, after providing instructions about the anchors of the scale and the information they represent, it is unclear the extent to which any observed inhibition is a result of the contingencies during training as opposed to the instructions concerning the scale. Moreover, the use of a negative scale might encourage participants to use a preventative causal model that they otherwise would not use. Recently, Karazinov and Boakes (2004) have highlighted the same problem and used transfer (summation) tests to assess conditioned inhibition. We agree with Williams (1995; 1996) and Karazinov and Boakes (2004) in thinking that transfer tests provide a less biased assessment of conditioned inhibition, than providing instructions and using negative rating scales. Thus, in the two experiments presented here we assessed Pavlovian backward inhibition using the two-test strategy for assessing conditioned inhibition, which requires both summation and retardation tests (Rescorla, 1969; Hearst, 1972; but see Papini & Bitterman, 1993) that is widely used in the animal literature tradition. This strategy has been proposed in order to control for alternative explanations based on attentional mechanisms. For example, in a summation test, an increase in attention to the target inhibitor might distract subjects from the presentation of the transfer excitor and thus result in low responding. In a retardation test, a decrease in attention to the target inhibitor could retard acquisition of behavioral control and result in retarded responding that is not due to inhibition. However, if under the same training procedures, one conducts both tests for inhibition, these two alternative explanations are rejected on the basis that they are contrary to each other (in one case one has to assume that subjects pay too much attention to the inhibitor [summation] and in the other case one has to assume that subjects paid too little attention to the inhibitor [retardation]). In summary, we used this strategy to make sure that the effect we sought to observe was not the result of too little or too much attention to the target cue.

The present experiments were conducted to clarify whether Pavlovian backward inhibition is a real phenomenon or if previous demonstrations resulted from potential biases in the procedures used by the researchers. Our intention was to provide a clear demonstration of Pavlovian backward inhibition in a human contingency learning task, not to test the role of instructions in this effect. In the categorization literature, instructions have been observed to facilitate the use of categories and retention of information (Strand, 1975). As applied to the present example, informing the subject about the outcome might facilitate subjects' categorizing trials as reinforced or nonreinforced and using a rule-based strategy to rate cues

at test. Therefore, we chose a predictive task that does not involve providing participants with information about the nature of the outcome. Specifically, participants were given minimal instructions and subsequently run through the different phases of the experiment without any demarcation between phases. Moreover, we assessed inhibition by using the two-test strategy that minimized the use of instructions and eliminated negative rating scales. In Experiment 1, all participants experienced Pavlovian backward inhibition training and were subsequently divided into three groups. Group Excitor was tested on the transfer excitor alone, group Inhibitor was tested on the transfer excitor presented in compound with the target inhibitor, and a third group, a group designed to control for generalization decrement (Gen Dec) at the time of testing, was tested on the transfer excitor in compound with a novel cue. Recently, Karazinov and Boakes (2004) proposed that the optimal control for Pavlovian inhibition in a summation test is a cue that received compound exposure equal to the putative inhibitor during training, but was never presented in compound with the training excitor. We did not use such a control because we reasoned that preexposure would decrease attention to the control cue, and that would favor the observation of inhibition by reducing generalization decrement from the transfer excitor in the control condition. That is, the alternative argument against inhibition with respect to a summation test is that subjects might have paid too much attention to the putative inhibitor and this drew attention from the transfer excitor. According to this argument, if we used a control cue that received preexposure (and hence decreased attention), we would favor the observation of inhibition. To further control for this attentional alternative explanation, Experiment 2 used the same training parameters as Experiment 1. But instead of testing immediately after Pavlovian backward inhibition training, the target inhibitor was reinforced and retardation in the acquisition of excitatory control of behavior was analyzed by comparing the target inhibitor to a novel cue that underwent the same reinforcement regimen. If subjects were paying too much attention to the inhibitor, then acquisition should be *faster* rather than retarded.

Experiment 1: Summation test

Method

Participants—The participants were 60 female and male undergraduate students at the State University of New York at Binghamton who participated in partial fulfillment of a course requirement. After reading and signing an informed consent form, they were assigned to one of three main treatment groups, counterbalanced as closely as possible for gender: 21 participants were in group Excitor, 21 in group Inhibitor, and 18 in group Gen Dec.

Apparatus and Stimuli—The experiments were conducted on six IBM compatible PCs. Inputs were made through a standard keyboard. All visual stimuli were displayed on 14-inch VGA monitor screens. SuperLab Pro 2.0 (Cedrus) was the software application used to present the stimuli and to record responses. A picture of a red inverted-U or a blue square, (the training excitor A and the transfer excitor B, counterbalanced within groups) and a green triangle or a yellow circle (X or Y, counterbalanced within groups), each measuring 4 cm × 4 cm, were used as the cues. A humorous picture of a cross-eyed human baby measuring 12.5 cm × 12.5 cm was used as the outcome (O). The stimuli were presented on a black background. The black background alone was displayed during intertrial intervals (ITIs).

Procedure (see Table 1)—At the beginning of the experiment, each participant was presented with the following set of instructions on the computer screen:

> This is an investigation of how people learn. It is not a test of your personal abilities or skills. Your name will not be linked with any of the data. You are about to witness several intermittent events on the computer monitor. These events will last for about

> five minutes. During this time, please do not look away from the monitor, as you might miss seeing an event. At any time you may be asked a question concerning what you have seen. The question will have a rating scale below it.
>
> <<Press the SPACE BAR to continue>>

After pressing the space bar, the participants saw a sequence of fifteen nonreinforced trials of a compound of cues AX (side by side); this constituted the compound conditioning phase. On these trials, the compounds of shapes were presented for 2 s followed by an ITI. The interstimulus interval (measured from cue termination to cue onset) was 3, 7, or 11 s, varied randomly (i.e., mean ITI = 7 s). The elemental training phase began 7 s after the nonreinforced phase. During this phase, participants were presented with the A or B stimulus followed by the outcome in a pseudorandom sequence (the sequence was A→O, B→O, B→O, A→O, B→O, A→O, A→O, B→O, A→O, B→O). Outcome duration was 1 s, and its presentation began immediately after cue termination (0 s delay). The interstimulus interval during this phase (from outcome offset to cue onset) was the same as in the previous phase. Stimulus identity, duration, and intertrial intervals were borrowed from a previous study in our laboratory that made use of the same task (Stout, Amundson, & Miller, 2005). The test phase occurred 7 s after the last elemental training trial. Before the test screen, participants read the following instructions:

> You will now be shown a colored shape and will be asked for your estimate of the probability that the picture of the baby will appear next.
>
> <<Press the SPACE BAR to continue>>

At test the target shape (i.e., the transfer excitor B) or compound of shapes (i.e., the transfer excitor B in compound with either X or Y) were presented in the same screen position as before with the following question below:

> What is the probability that the picture of the baby will follow next?
>
> Using a scale from 0% to 100%, VERY CAREFULLY enter an estimate and then press the SPACE BAR to record your estimate. It is not necessary to enter the '%' symbol.

After entering the rating and pressing the space bar, a debriefing screen was presented.

Data Analysis—Ratings were analyzed with analysis of variance (ANOVA) and subsequent planned comparisons. Simple *t*-test were used when appropriate. Effect sizes (Cohen's *f*)were calculated with the algorithm provided by Myers and Well (2003; p 210).

Results and Discussion

As it can be observed in Table 1, participants tested on the compound containing the transfer excitor and the target inhibitor (group Inhibitor) exhibited lower ratings than participants who rated the transfer excitor presented in compound with a novel stimulus (group Gen Dec). Participants tested on the transfer excitor alone (group Excitor) showed the highest ratings of the three groups.

These impressions were confirmed by a one-way analysis of variance (ANOVA) with group (Excitor vs. Inhibitor vs. Gen Dec) as a factor $F(2, 57) = 31.36$, $p < .001$, $MSE = 802.4$, Cohen's $f = 1.0$. Visual inspection of the results indicates a large generalization decrement effect, in that ratings to the transfer excitor presented in compound with a novel cue were less than half of those for the excitor alone. Therefore, the critical comparison was between the Inhibitor and the Gen Dec groups. A planned comparison using the overall error term of the ANOVA was

conducted to compare the ratings of group Inhibitor with those of group Gen Dec. This comparison indicated that the mean ratings differed, $F(1, 57) = 5.71$, $p < .05$, Cohen's $f = 0.34$.

The present experiment used a transfer (summation) test and demonstrated that without any knowledge about the existence of the outcome, Pavlovian backward conditioned inhibition can still be observed. One might argue that the Pavlovian backward inhibition effect results from X distracting subjects from B and this being more pronounced after AX- trials, that is, too much attention to X. For this reason, summation tests alone for conditioned inhibition are viewed as inadequate. The gold standard for assessing inhibition is to use both a summation test and a retardation test (Hearst, 1972; Rescorla, 1969).

Experiment 2: Retardation test

The second experiment, therefore, was designed to determine if the same training parameters would result in retarded emergence of behavioral control. To answer this question, we used two groups of students that experienced the same number and type of trials as those in Experiment 1, with the only difference being that, after phase 2 of training, participants experienced in phase 3 one reinforced trial each with the target cue, a novel (control) cue, and the cue that had only been trained during phase 2 (i.e., B). The rationale for reinforcing B was to maintain continuity with phase 2, that is, we wanted to avoid a disruptive transition between phases 2 and 3. Subsequently, retardation of behavioral control was assessed.

Method

Participants—The participants were 36 female and male undergraduate students at the State University of New York at Binghamton who participated in partial fulfillment of a course requirement. After reading and signing an informed consent form, they were randomly assigned to one of two treatment groups, counterbalanced as close as possible for gender (ns = 18).

Apparatus and stimuli—The test apparatus was the same as in Experiment 1, with the addition of another cue, a purple inverted letter Y, which was counterbalanced within groups in serving as stimuli A and B. This stimulus was added because this study was part of a larger experiment that does not merit inclusion in the present manuscript. We also changed the rating scale to simplify the task and data collection. Therefore, in the present experiment participants used a rating scale that ranged from 1 to 9, thereby allowing a rating to reflect a single key stroke (see Procedure below). The anchors of the scale were illustrated with the legends ‘Least probable’ and ‘Most probable’ to simplify the rating procedure without suggesting to the participants negative contingencies or conditioned inhibition.

Procedure (See Table 2)—The procedure was similar to the one used in Experiment 1. During phases 1 and 2, participants received the same treatments as in Experiment 1. During phase 3, all participants experienced one reinforced trial of X, Y, and B, with the order counterbalanced within groups.

At test, the target shape was presented. Group Inhibitor was tested on X and group Control was tested on Y. However, to allow within-participants comparisons all participants were tested on the remaining cue after providing the first rating. Because we changed the rating scale, we also provided slightly different instructions. Each shape was presented in the same screen position with the following questions and a 1 to 9 scale below it:

> What is the probability that the picture of the baby will follow next?
>
> 1 — 2 — 3 — 4 — 5 — 6 — 7 — 8 — 9 Referring to the scale above VERY CAREFULLY enter an estimate using the appropriate number key on the keyboard.

Anchors at 1 and 9 read “Least probable” and “Most probable,” respectively. After entering the rating, the next rating screen appeared. Finally, when the participants had entered the second rating, the debriefing screen was presented.

Results and Discussion

As can be observed in Table 2, the ratings for the target inhibitor (X) were substantially lower than those for the novel cue (Y), which did not receive any inhibition treatment. Because we tested all subjects with both cues but with the order counterbalanced, we conducted a 2 × 2 (Order [XY vs. YX] x Test [Test 1 vs. Test 2]) mixed ANOVA. This analysis revealed only an interaction,$F(1, 34) = 6.14$, $p < .02$, $MSE = 6.83$, Cohen's $f = 0.37$, with no significant main effects (smallest $p = .12$), which is indicative of higher ratings of Y than X regardless of test order. This was further corroborated by a within-participants test, comparing all participants pooled across test order, which found the mean for Y to be greater than for X, $t(35) = 2.42$, $p < .05$, Cohen's $f = 0.36$. Thus, we conclude that the target inhibitor showed slower emergence of behavioral control than a novel cue.

General Discussion

Experiments 1 and 2 demonstrated Pavlovian backward conditioned inhibition as assessed by summation (Experiment 1) and retardation (Experiment 2) tests in a human predictive learning task. In Experiment 1 inhibition was inferred from a comparison of the ratings of a transfer excitor alone, with the ratings given to the transfer excitor presented in compound with a novel cue (group Gen Dec), and the ratings to the transfer excitor presented in compound with the target inhibitor (group Inhibitor). Ratings to the latter compound were lower than those given to the compound of the transfer excitor and a novel cue. In Experiment 2 we used the same parameters but tested to see whether acquisition of behavioral control by the target inhibitor would be retarded compared to acquisition of behavioral control by a novel cue. Both between- and within-participants comparisons demonstrated that acquisition was retarded in the target inhibitor. It is important to note that our retardation test did not directly assess retarded emergence of behavioral control *during* reinforced trials but rather afterwards. This is the reason why we chose to administer only *one* reinforced trial during retardation training, before subsequently assessing the relationship between the cues and outcomes. This limitation stems from the fact that if we wanted to assess behavioral control during retardation training we would have had to inform subjects about the nature of outcome prior to the retardation test pairings. Nevertheless, the current approach resembles that used in lick-suppression experiments with rats as subjects, in which only a few *offline* pairings (typically three or four trials without any data collection) reveal in subsequent tests retarded emergence of behavioral control (e.g., Amundson, Wheeler, & Miller, 2005).

Backward Pavlovian inhibition was initially problematic for some models of learning that did not allow for learning about absent cues (e. g., Rescorla & Wagner, 1972; Wagner, 1981). That is, according to these models ratings for X should not change as a result of reinforcing A because on those trials X was not presented and therefore should not undergo any associative change. The observation of Pavlovian backward inhibition (Chapman, 1991) as well as backward blocking (Shanks, 1985) encouraged researchers to revise these influential models to account for these observations. Van Hamme and Wasserman (1994) modified the Rescorla-Wagner model to allow for learning about absent cues by proposing that, for a cue that is expected but not presented on a given trial, the alpha parameter for a cue that is expected but not presented becomes negative. Consequently the associative strength of the absent cue changes in the opposite direction of the cue if it had been presented. According to this model, in the case of Pavlovian backward conditioned inhibition, during phase 1 the target inhibitor and the training excitor are presented together and consequently an association between them is formed. During phase 2 the excitor used in phase 1 is presented alone followed by reinforcement. Therefore,

the target inhibitor acquires negative associative strength because it is expected but not presented. Although the psychological correlates of assigning a negative value to a learning parameter are not known, the model can account for the present observations.

Another model for which the initial observations of backward blocking and Pavlovian backward inhibition posed a problem is Wagner's SOP (1981). However, Dickinson and Burke (1996; also see Aitken & Dickinson, 2005) proposed a revision of the model that accounts for these observations. The original SOP stated that elements that compose a stimulus representation can be represented in one of three states in memory: the high activity state (A1), the low activity state (A2), and the inactive state (I). According to Wagner, when a stimulus is presented, its representation passes from the inactive state into state A1, and after the stimulus' presentation is completed the representation decays through state A2 back to the inactive state. The critical feature of SOP is that an excitatory association develops between two stimuli that are simultaneously in state A1, whereas a cue becomes an inhibitor when the cue is represented in A1 and the outcome is represented in A2. Dickinson and Burke (1996) modified SOP (MSOP) by adding two more assumptions: 1) they proposed that when the representations of two stimuli are in state A2, then an excitatory association will develop between them, 2) when an outcome is in state A1 and a cue is in state A2, an inhibitory association will develop. (Wagner originally stated that inhibitory associations develop when the cue is in state A1 and the outcome in state A2, but not the reverse.) Applied to backward conditioned inhibition, MSOP poses that during phase 1 a within-compound association between A and X is established. During phase 2, presentation of A (followed by the outcome) retrieves X into state A2; in other words, during the reinforced trials of phase 2 the outcome is represented in A1 and the putative inhibitor X is associatively retrieved into A2 by the excitatory cue A (by virtue of the within-compound association that was established between them during phase 1) and as a consequence an inhibitory association between X and the outcome is established. Thus, MSOP can also account for the present results.

The comparator hypothesis (Miller & Matzel, 1988; Miller & Schachtman, 1985), a response rule for the expression of associations, also anticipates the observation of backward Pavlovian conditioned inhibition. The comparator hypothesis explicitly denies the existence of negative associations but states that behavior indicative of conditioned inhibition is the result of competition at the time of testing between associations of null or positive (but never negative) value. As applied to backward Pavlovian inhibition, the comparator hypothesis states that when participants are presented with the target inhibitor, they indirectly retrieve a representation of the outcome mediated by the training excitor (because these two cues were presented together during phase 1, and the training excitor was consistently followed by the outcome during phase 2) and compare this outcome representation with that activated by the target inhibitor, which is null because the inhibitor was not reinforced during training. If the indirectly activated representation exceeds the strength of the representation activated directly by the target inhibitor, then behavior indicative of inhibition is anticipated. Thus, the present results are anticipated by a number of contemporary learning theories. However, the present experiments do not allow us to discriminate between these diverse accounts of Pavlovian backwards inhibition.

The present experiments were conducted to overcome two limitations of previous demonstrations of Pavlovian backward conditioned inhibition. One of these limitations is that all previous demonstrations (Chapman, 1991; Larkin, Aitken, & Dickinson, 1998; Melchers, Lachnit, & Shanks, 2004; Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996; Williams & Docking, 1995) provided information about the outcome before the beginning of the experiment. This could have had two consequences. First, the information concerning the outcome could be considered a presentation of the outcome itself (i.e., an unsignaled outcome presentation), which could result in contextual conditioning. Conditioning of the experimental

context could potentially be responsible for the observation of Pavlovian backward conditioned inhibition in the same way that it is widely viewed as the excitatory basis for inhibition produced by explicitly unpaired presentations of a cue and an outcome.

The second consequence of providing information about the outcome is more cognitive. It has been demonstrated that reliable associative learning can occur as a result of providing verbal instructions about cues and outcomes, even in the absence of the target cues after providing these verbal instructions (Cook & Harris, 1937; Tassoni, 1995). Thus, it is not surprising that alerting participants to the nature of the outcome may have favored the observation of inhibition (Chapman, 1991; Larkin, Aitken, & Dickinson, 1998; Melchers, Lachnit, & Shanks, 2004; Wasserman, Kao, Van Hamme, Katagiri, & Young, 1996; Williams & Docking, 1995). A second general limitation of previous demonstrations (with the exception of Williams & Docking, 1995) was the use of rating scales that allowed negative ratings. Although a rating scale with positive and negative ends has some parallel with the approach-withdrawal tests used in the animal literature (Wasserman, Franklin, & Hearst, 1974; Hearst & Franklin, 1977), the mere presence of a negative end-value in the scale might have encouraged participants to judge the relationship between the target and the outcome as being negative in instances where they otherwise would have not rated them as being negative. Thus, the use of transfer tests such as the one employed in Experiment 1 seems to be a more conservative approach to the study of inhibition in human contingency learning.

The present results parallel previous observations of posttraining inflation effects observed in the animal literature. Using rats as subjects, Amundson, Wheeler, and Miller (2005) reported enhanced conditioned inhibition after inflating the training excitor (providing extra reinforced trials), which is operationally similar to what we presented during phase 2 of training although they provided both nonreinforced and reinforced trials during phase 1. Similarly, Baker, Murphy, and Metha (2003) observed behavior indicative of inhibition to a noise after presenting 48 noise-alone trials followed by 48 shock-alone trials in the same context. These blocks of unpaired CS and US are another example of inflating the companion cue (the context, in this case) and observing evidence of inhibitory behavior. However, the same argument we made concerning the prior human studies could be leveled against the Baker et al. summation test because they exposed subjects to the US (i.e., they trained the transfer excitor) before administering the nonreinforced presentations of the CS. The parallels between the results of the experiments reported here and those observed with nonhuman animals suggest that the same mechanisms may be operating across mammalian species.

It is noteworthy that there seems to be an asymmetry in the effects of posttraining inflation and deflation effects. Surely it is harder to observe effects of posttraining inflation than posttraining deflation (e.g., Grahame, Barnet, & Miller, 1992). The reason why inflation effects are harder to observe than deflation effects might be that, while inflation effects increase the associative strength of the companion stimulus, they also decrease the association between the companion and the target cue (i.e., the within-compound association). Thus, one of these processes favors the observation of revaluation while the other does not. In contrast with inflation effects, deflation effects are more readily observed because both the association between the companion cue and the target cue and the association between the companion cue and the outcome are extinguished during the revaluation process. This might explain why an inflation effect is harder to obtain than a deflation effect (which itself, is parameter dependent, e.g., Holland, 1999).

In summary, the two experiments reported here provide a clear demonstration of the existence of backward Pavlovian inhibition. Within the present experiments we intentionally used a task that did not provide any information to the participants about the nature of the outcome. Furthermore, instead of using a negative scale to infer conditioned inhibition, we made use of

transfer tests (i.e., summation tests) and retardation tests, which are clearly less biased tests to evidence inhibitory phenomena.

Acknowledgements

The authors thank Jeffrey C. Amundson, Jonah Grossman, Rachael Hessner, Olga Lipatova, Alyssa Orinstein and Daniel Wheeler for their comments on an earlier version of this manuscript. Inquiries concerning all aspects of this research should be addressed to Ralph R. Miller, Department of Psychology, SUNY- Binghamton, Binghamton, NY 13902-6000, USA; e-mail: rmiller@binghamton.edu.

References

Amundson JC, Wheeler DS, Miller RR. Enhancement of Pavlovian conditioned inhibition achieved by posttraining inflation of the training excitor. Learning & Motivation 2005;36:331–352.

Aitken MR, Dickinson A. Simulations of a modified SOP model applied to retrospective revaluation in human causal learning. Learning & Behavior 2005;33:147–159. [PubMed: 16075835]

Baker AG, Murphy RA, Mehta R. Learned irrelevance and retrospective correlation learning. Quarterly Journal of Experimental Psychology 2003;56B:90–101. [PubMed: 12623540]

Chapman GB. Trial order affects cue interaction in contingency judgment. Journal of Experimental Psychology: Learning, Memory, and Cognition 1991;17:837–854.

Cook SW, Harris RE. The verbal conditioning of the galvanic skin response. Journal of Experimental Psychology 1937;21:202–210.

De Houwer J, Beckers T. A review of recent developments in research and theories on human contingency learning. Quarterly Journal of Experimental Psychology 2002;55B:289–310. [PubMed: 12350283]

Dickinson A. The 28th Bartlett memorial lecture. Causal learning: An associative analysis. Quarterly Journal of Experimental Psychology 2001;54B:3–25. [PubMed: 11216300]

Dickinson A, Burke J. Within-compound associations mediate the retrospective revaluation of causality judgments. Quarterly Journal of Experimental Psychology 1996;49B:60–80. [PubMed: 8901386]

Grahame NJ, Barnet RC, Miller RR. Pavlovian inhibition cannot be obtained by posttraining A-US pairings: Further evidence for the empirical asymmetry of the comparator hypothesis. Bulletin of the Psychonomic Society 1992;30:339–402.

Hearst, E. Some persistent problems in the analysis of conditioned inhibition. In: Boakes, RA.; Halliday, MS., editors. Inhibition and learning. Academic Press; London: 1972. p. 5-39.

Hearst E, Franklin SR. Positive and negative relations between a signal and food: Approach-withdrawal behavior to the signal. Journal of Experimental Psychology: Animal Behavior Processes 1977;3:37–52.

Holland PC. Overshadowing and blocking as acquisition deficits: No recovery after extinction of overshadowing or blocking cues. Quarterly Journal of Experimental Psychology 1999;52B:307–334. [PubMed: 10605392]

Karazinov DM, Boakes RA. Learning about cues that prevent an outcome: Conditioned inhibition and differential inhibition in human predictive learning. Quarterly Journal of Experimental Psychology 2004;57B:153–178. [PubMed: 15204114]

Larkin MJW, Aitken MRF, Dickinson A. Retrospective revaluation of causal judgments under positive and negative contingencies. Journal of Experimental Psychology: Learning, Memory, and Cognition 1998;24:1331–1352.

Mackintosh NJ. A theory of attention: Variations in the associability of stimuli with reinforcement. Psychological Review 1975;82:276–298.

Melchers KG, Lachnit H, Shanks DR. Within-compound associations in retrospective revaluation and in direct learning: A challenge for comparator theory. The Quarterly Journal of Experimental Psychology 2004;57B:25–53. [PubMed: 14690848]

Miller, RR.; Matzel, LD. The comparator hypothesis: A response rule for the expression of associations. In: Bower, GH., editor. The psychology of learning and motivation. 22. Academic Press; San Diego, CA: 1988. p. 51-92.

Miller, RR.; Schachtman, TR. Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In: Miller, RR.; Spear, NE., editors. Information processing in animals: Conditioned inhibition. Erlbaum; Hillsdale, NJ: 1985. p. 51-88.

Myers, JL.; Well, AD. Research design and statistical analysis. 2nd ed.. Lawrence Erlbaum Associates Publishers; Mahwah, NJ, US: 2003.

Papini MR, Bitterman ME. The two test strategy in the study of inhibitory conditioning. Journal of Experimental Psychology: Animal Behavior Processes 1993;19:342–352. [PubMed: 8228833]

Pavlov, IP. Conditioned reflexes. Anrep, GV., editor. Oxford University Press; London: 1927. Trans.

Pearce JM, Hall G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not unconditioned stimuli. Psychological Review 1980;82:532–552. [PubMed: 7443916]

Rescorla RA. Pavlovian conditioned inhibition. Psychological Bulletin 1969;72:77–94.

Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. Classical conditioning II: Current research and theory. Appleton-Century-Crofts; New York: 1972. p. 64-99.

Tassoni CJ. The least mean squares network with information coding: A model of cue learning. Journal of Experimental Psychology: Learning, Memory, and Cognition 1995;21:193–204.

Shanks DR. Forward and backward blocking in human contingency judgement. Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology 1985;37B:1–21.

Stout SC, Amundson JC, Miller RR. Trial order and retention interval in human predictive judgment. Memory & Cognition 2005;33:1368–1376.

Strand BZ. Effects of instructions for category organization on long-term retention. Journal of Experimental Psychology: Human Learning and Memory 1975;1:780–786.

Van Hamme LJ, Wasserman EA. Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. Learning and Motivation 1994;25:127–151.

Wagner, AR. SOP: A model of automatic memory processing in animal behavior. In: Spear, NE.; Miller, RR., editors. Information processing in animals: Memory mechanisms. Erlbaum; Hillsdale, NJ: 1981. p. 5-47.

Wasserman EA, Franklin SR, Hearst E. Pavlovian appetitive contingencies and approach vs. withdrawal to conditioned stimuli in pigeons. Journal of Comparative and Physiological Psychology 1974;86:616–627. [PubMed: 4823236]

Wasserman, EA.; Kao, S.; Van Hamme, LJ.; Katagiri, M.; Young, ME. Causation and association. In: Shanks, DR.; Holyoak, KJ.; Medin, DL., editors. The psychology of learning and motivation. 34. Academic Press; New York: 1996. p. 208-264.

Williams DA. Forms of inhibition in animal and human learning. Journal of Experimental Psychology: Animal Behavior Processes 1995;21:129–142. [PubMed: 7738496]

Williams, DA. A comparative analysis of negative contingency learning in humans and nonhumans. In: Shanks, DR.; Holyoak, KJ.; Medin, DL., editors. The psychology of learning and motivation. 34. Academic Press; New York: 1996. p. 89-131.

Williams DA, Docking GL. Associative and normative accounts of negative transfer. Quarterly Journal of Experimental Psychology 1995;48A:976–998.

Table 1

Design of Experiment 1: Summation Test

Group	Phase 1	Phase 2	Test	Mean	*SEM*
Excitor			B	83.3	6.52
Inhibitor	15 AX−	5 A+ / 5 B+	BX	13.7	6.98
Gen Dec			BY	35.4	4.91

Note. A, B, X, and Y denote geometrical figures used as stimuli. Numbers next to pairings denote number of trials. "+" = outcome. "−" = no outcome. Excitor = group tested on the transfer excitor alone. Inhibitor = group tested on the transfer excitor in compound with the target inhibitor. Gen Dec = Group tested on the transfer excitor in compound with an irrelevant cue. Mean = rating on a scale of 0 to 100. SEM = standard error of the mean.

Table 2

Design of Experiment 2: Retardation Test

Condition	Phase 1	Phase 2	Phase 3	Test	Mean	*SEM*
Inhibitor	15 AX−	5 A+ / 5 B+	1 B+ / 1 Y+ / 1 X+	X	3.69	0.50
Control				Y	5.22	0.50

Note. A, B, X, and Y denote geometrical figures used as stimuli. Numbers next to pairings denote number of trials. "+" = outcome. "–" = no outcome. Inhibitor = target inhibitor after one reinforced trial. Control = cue with no inhibitory history after one reinforced trial. Mean = mean rating of all participants for each cue (X = Inhibitor; Y = Control) on a scale of 1 to 9. SEM = standard error of the mean.