Published in final edited form as:

J Exp Psychol Gen. 2006 February ; 135(1): 92–102.

Reasoning Rats: Forward Blocking in Pavlovian Animal Conditioning Is Sensitive to Constraints of Causal Inference

Tom Beckers

Department of Psychology, Binghamton University, State University of New York and Department of Psychology, University of Leuven, Leuven, Belgium

Ralph R. Miller

Department of Psychology, Binghamton University, State University of New York

Jan De Houwer

Department of Psychology, Ghent University, Ghent, Belgium

Kouji Urushihara

Department of Psychology, Binghamton University, State University of New York

Abstract

Forward blocking is one of the best-documented phenomena in Pavlovian animal conditioning. According to contemporary associative learning theories, forward blocking arises directly from the hardwired basic learning rules that govern the acquisition or expression of associations. Contrary to this view, here the authors demonstrate that blocking in rats is flexible and sensitive to constraints of causal inference, such as violation of additivity and ceiling considerations. This suggests that complex cognitive processes akin to causal inferential reasoning are involved in a well-established Pavlovian animal conditioning phenomenon commonly attributed to the operation of basic associative processes.

Keywords

animal conditioning; cue competition; causal reasoning; associative learning; animal cognition

Ever since the seminal experimental studies with dogs by Pavlov (1927), research into classical conditioning has been flourishing more or less constantly, surviving shifts of paradigm from behaviorism to cognitivism, connectionism, and behavioral neuroscience. Along the way, it has provided indispensable input to fields as diverse as animal cognition, behavior therapy, human psychophysiology, stereotyping in social psychology, back-propagation models of parallel distributed processing, and basic neuroscience.

In Pavlovian conditioning, an initially neutral conditioned stimulus (CS; e.g., a tone) comes to elicit a conditioned response (e.g., freezing) as the result of pairings of the CS with a biologically relevant, unconditioned stimulus (US; e.g., a footshock). Very basic and relatively simple associative principles have often been assumed to account for such Pavlovian

Correspondence concerning this article should be addressed to Tom Beckers, Department of Psychology, University of Leuven, Tiensestraat 102, B-3000, Leuven, Belgium. E-mail: tom.beckers@psy.kuleuven.be.

Kouji Urushihara is now at the Department of Psychology, Osaka Kyoiku University, Osaka, Japan, as a Japan Society for the Promotion of Science Research Fellow.

The research reported in this article was conducted while Tom Beckers was a visiting scholar at Binghamton University, State University of New York and was supported by the Research Foundation—Flanders (FWO—Vlaanderen), of which he is a Postdoctoral Fellow. The research was supported by National Institute of Mental Health Grant 33881 to Ralph R. Miller.

conditioned behavior. Even complex phenomena in Pavlovian conditioning are generally attributed to the operation of these basic associative principles. One such phenomenon and a benchmark for models of Pavlovian conditioning is forward blocking. It was first reported by Kamin (1969), and it has been one of the most intensively studied effects in the field of Pavlovian conditioning ever since. Forward blocking involves the observation that conditioned responding to a target CS X (e.g., a tone) that is paired with the US in compound with another CS A (e.g., a light) is markedly weaker if CS A has itself previously been paired with the same US (i.e., A+ followed by AX+ training, where + represents presentation of the US) than if no elemental A+ training precedes AX+ training. It is as if the presence of the previously conditioned CS A blocks conditioning to CS X.

Forward blocking is generally believed to result directly from the hardwired basic learning rules that govern the acquisition of associative links between CSs and USs. For instance, according to the influential Rescorla–Wagner theory (Rescorla & Wagner, 1972), a given US can support only a limited amount of associative strength. If, when a target CS and a US are paired, another CS (A in our example) is present that has already absorbed most or all of the associative strength supportable by that US, the target CS (CS X in our example) cannot acquire much, if any, associative strength. Metaphorically, in this theoretical framework, conditioning can be described as the encoding of nonredundant predictors of the US. If a reliable predictor is already established and present when the US is presented, valid but redundant new predictors such as CS X are ignored. Another way to frame this is to say that the acquisition of associative strength is determined by the surprise value of the US. CSs that precede the delivery of a surprising US will gain associative strength; CSs that precede the delivery of a nonsurprising US will not. Note that such an account of conditioning effectively assumes that forward blocking reflects a failure to acquire an association between CS X and the US (see Miller & Matzel, 1988, for an alternative view).

Forward blocking can be observed not only in animal conditioning but also in human causal learning and attribution (Dickinson, Shanks, & Evenden, 1984). For example, if human participants first learn that consumption of a particular food item A results in an allergic reaction ($A \rightarrow O$), and then that the combination of that food item with another one (X) also results in an allergic reaction ($AX \rightarrow O$), then X will be judged less allergenic than if no information about A alone was provided (i.e., $AX \rightarrow O$ not preceded by $A \rightarrow O$). This striking functional similarity between human causal learning and animal conditioning has inspired the proposal that associative learning principles derived from animal conditioning research apply to human causal learning as well (Dickinson et al., 1984). According to such logic, when a compound of potential causes A and X is paired with outcome O, less of an associative link between X and the outcome will be established if an associative link between cue A and the outcome is already in place than if no such previous link exists (see De Houwer & Beckers, 2002, for a review).

Recent evidence, however, suggests that in human causal learning, forward blocking is not as stable and ubiquitous a phenomenon as would be expected if the selectivity it reflects were a fundamental, hardwired aspect of human learning. The phenomenon actually appears to a considerable extent to be subject to constraints of causal reasoning. Accordingly, it has been suggested recently that in human causal learning, blocking largely reflects controlled and effortful inferential reasoning processes rather than basic associative processes. A causal reasoning account of forward blocking assumes that people do not fail to acquire the information that the blocked cue X and the outcome co-occurred. Instead, it assumes that people infer by controlled and effortful processing that cue X is not a valid cause of the outcome, the co-occurrence of X and the outcome notwithstanding (De Houwer, Beckers, & Vandorpe, 2005; Lovibond, 2003; Waldmann & Walker, 2005). This inference is based on the assumption that if two effective causes of a given outcome co-occur, a stronger outcome should result than

if only one effective cause is present. In combination with the observation that the actual outcome is as strong in the combined presence of A and X as it was in the presence of A alone, people can logically deduce that potential cause X is not an effective cause of the outcome. Such inferential reasoning may be formally represented by means of the following modus tollens argument¹:

[*if p then q*] If potential causes A and X are both effective causes of a particular outcome, then the outcome should be stronger when both are present than when only one is present.

[*not q*] The outcome is not stronger when A and X are both present than when only A is present.

[*therefore, not p*] Thus, A and X are not both effective causes of the outcome.

Given that A is clearly effective in causing the outcome when presented alone, it follows that X is not an effective cause of the outcome. Hence, people give reduced ratings when asked to judge the causal effectiveness of X. According to this account, if people cannot ascertain the veracity of both premises, the $A \rightarrow O$ trials are not helpful for deducing the causal status of X. Hence, the ratings for X should not differ from what they are when $AX \rightarrow O$ trials are not preceded by $A \rightarrow O$ trials (i.e., blocking should not occur; see below).

Support for the involvement of effortful analytical reasoning processes in the occurrence of blocking in human causal learning comes from the observation that blocking in such learning is sensitive to a number of constraints of effortful causal inference. For instance, if forward blocking results from people reasoning from a default assumption of causal additivity (the first premise in the argument above), then blocking should occur only to the extent that people effectively entertain this default assumption. If they do not, the fact that A alone and the compound of A and X result in the same outcome cannot be interpreted as evidence against a causal influence of X because it is possible that X does have a causal impact that cannot be observed because of the lack of additivity. Therefore, if people are trained to not entertain the default assumption of additivity, they should be unsure about the causal status of cue X and blocking should not occur, just as is the case when $AX \rightarrow O$ training is not preceded by $A \rightarrow O$ training. Therefore, it should be possible to reduce forward blocking by explicitly contradicting the assumption of causal additivity. Recent studies have shown just that: When people were shown that two individual food items C and D each resulted in an allergic reaction in a fictitious patient and consumption of the two food items combined resulted in a more intense allergic reaction, subsequent $A \rightarrow O$ training followed by $AX \rightarrow O$ training resulted in robust blocking of X. However, when on the CD compound trials, an allergic reaction occurred that was similar to that which occurred on the elemental C and D trials, subsequent $A \rightarrow O$ training followed by $AX \rightarrow O$ training yielded little if any blocking of X (Beckers, De Houwer, Pineño, & Miller, 2005; Livesey & Boakes, 2004; Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003). When the training with C and D was conducted after the $A \rightarrow O$ and $AX \rightarrow O$ training, a similar effect on blocking was obtained (Beckers et al., 2005).

According to a causal reasoning analysis, blocking should be constrained not only by whether people assume additivity but also by whether people are actually able to empirically verify on AX trials that X does not add to the outcome of A (the second premise in the argument above). Assume that food A results in an allergic reaction with an intensity that corresponds to the maximal intensity that can be measured. If later the combination of food A and food X results in a similar allergic reaction, the causal nonefficacy of X cannot be validly inferred, because of a ceiling effect. Again, under such circumstances, people should be unsure about the causal

¹We do not want to imply that human reasoning is actually based on the formal application of syllogistic logic. Natural logic, inference schemes, mental model construction and validation, constrained causal Bayes net construction, or another form of controlled, analytical processing might be involved.

status of X. Therefore, ratings for X should not differ depending on whether $AX \rightarrow O$ training is preceded by $A \rightarrow O$ training (thus, no blocking should occur) under such conditions. Blocking should be obtained more readily if the outcome that follows A is nonmaximal, because in that case the causal nonefficacy of X can be readily inferred. It was recently demonstrated that blocking is indeed obtained much more readily if submaximal outcomes are presented on $A \rightarrow O$ and $AX \rightarrow O$ trials than if outcomes of a maximal intensity are presented on $A \rightarrow O$ and $AX \rightarrow O$ trials (Beckers et al., 2005; De Houwer, Beckers, & Glautier, 2002).

Other constraints on blocking predicted by a causal reasoning analysis include the idea that blocking should occur only if people have the opportunity to engage in effortful analytical processing. If cognitive capacity (e.g., working memory resources) is lacking, blocking should be less likely to occur. The observation that blocking is greatly reduced if a demanding secondary task is presented simultaneously with the causal learning task, even though this secondary task does not interfere with elementary learning (i.e., it selectively affects blocking), is in line with this constraint (De Houwer & Beckers, 2003; Vandorpe, De Houwer, & Beckers, 2005; Waldmann & Walker, 2005). Note that if blocking were due to the operation of automatic associative processes, one would expect to see an increase, rather than a decrease, in blocking with increasing secondary task load.

The sensitivity of blocking in human causal learning to constraints of causal inference is not anticipated by an associative analysis along the lines of the Rescorla–Wagner model (Rescorla & Wagner, 1972) or any other associative model. Neither pretraining or posttraining with a different set of cues nor the maximal nature of the outcome should have any influence on the fact that X is a redundant cue for the presence of the outcome when $AX \rightarrow O$ trials are preceded by $A \rightarrow O$ training. Therefore, blocking should not be affected by these factors. Similarly, the automatic, bottom-up nature of association formation seems incompatible with the observation that blocking is affected by secondary task load. Accordingly, it might seem that, unlike the suggestion by Dickinson et al. (1984; see also Dickinson, 2001), fundamentally different processes are governing human causal learning and Pavlovian conditioning in animals after all. Causal learning in humans would then at least in part be based on controlled and effortful inferential reasoning processes, perhaps complemented by secondary associative processes,² whereas Pavlovian conditioning in animals would just involve low-level, automatic associative principles or inferential processes much simpler and much less flexible than the ones operating in humans.

Our aim in the present study was to investigate the more challenging idea that in animal Pavlovian conditioning, as in human causal learning, forward blocking to some extent relies on the operation of fairly complex, reasoning-like processes and would therefore be sensitive to modulation by the same constraints of causal inference. Admittedly, this is a provocative proposal, for at least two reasons. First, it assumes much more complex cognitive capacities in lower mammals and birds (the kind of species typically studied in Pavlovian conditioning research) than are generally assumed to exist. Second, such a proposal goes counter to the firm associationist tradition that dominates the animal conditioning field.

We pursued this possibility by testing whether forward blocking in animal conditioning would be similarly sensitive to the constraints on forward blocking that are observed in human causal learning. Specifically, we tested rats' sensitivity to constraints of causal inference in three experiments. In all three experiments, three experimental groups were exposed to a forward

²We are not claiming that forward blocking in human causal learning is always the result of a nonassociative, controlled inferential reasoning process. Forward blocking sometimes occurs in human causal learning in situations that do not meet the requirements for a formal blocking conclusion according to a reasoning analysis (see Beckers et al., 2005; Lovibond et al., 2003). This might indicate that some degree of forward blocking can result from automatic selective associative processes (see the General Discussion).

blocking procedure. An auditory cue (A) was repeatedly paired with a footshock US (A+; elemental training phase). Then A was repeatedly compounded with a click train (X), and the compound of A and X was also followed by footshock (AX+; compound training phase). In each experiment, control groups were exposed to pairings of a different auditory cue (B) with shock in the elemental phase (B+), before receiving AX+ training in the compound phase. All animals were then tested for conditioning to X by measuring suppression of separately trained instrumental responding (lever pressing for water) during test presentations of X. Blocking would be evident in lower conditioned fear responding to X (i.e., less suppression) in the experimental animals than in the control animals. In the first two experiments, we tested the degree to which such a blocking effect would be modulated by pretraining with a different set of cues that explicitly contradicted (Experiment 1) or confirmed (Experiment 2) the general additive nature of cues. In Experiment 3, we assessed the sensitivity of blocking to an experiential outcome intensity ceiling.

In all three experiments, stable instrumental responding by the water-deprived rats was first obtained by making water delivery contingent on lever pressing. Next, the classical conditioning training was performed in the absence of the response lever. After retraining of the instrumental lever-press response, the critical test cue (X) was presented three times while the animals were lever pressing for water, and suppression of lever pressing was recorded.

Experiment 1

In the first experiment (see Table 1), we used parameters generally known to result in reliable blocking in our preparation (i.e., 12 elemental training trials and four compound training trials). Pilot work in the Binghamton University, State University of New York (SUNY) laboratory has demonstrated that, in the present preparation, in order to obtain robust blocking, a minimum of 12 elemental trials are required, whereas 8 elemental trials will yield only weak blocking, and 4 elemental trials will yield almost no blocking. We needed a procedure that would result in a strong blocking effect because the aim of this study was to determine if blocking could be reduced by subadditive training.³

Before the actual blocking training, experimental (i.e., blocking treatment) and control animals in the subadditive condition were exposed at the start of the experiment to a demonstration of two effective cues C and D (different from the blocking training cues A, B, and X) that did not have additive outcomes: The combination of the cues C and D resulted in the same shock as did each of the cues in itself (C+/D+/CD+). The experimental and control animals in both other conditions received a similar amount of pre-training, equivalent in terms of the number and intensity of the USs. However, in these conditions, pretraining was irrelevant with respect to additivity. The irrelevant element condition received C+/D+/E+ pretraining, and the irrelevant compound condition received C+/C+/DE+ pretraining. If blocking in Pavlovian conditioning is sensitive to constraints of causal inference, blocking should be reduced by subadditive pretraining, relative to both forms of irrelevant pretraining.

³It is an interesting question why the number of elemental training trials should have such an effect on forward blocking. Perhaps it indicates that associative processes also contribute to blocking. Indeed, one could argue that learning will be closer to asymptote after 12 than after 4 trials, so less associative strength is left for X in the former case than in the latter case (see the General Discussion). Problematic for this account is the observation that a vast number of A+ trials also reduces blocking compared with the number of elemental trials typically used to obtain blocking (see Pineño, Urushihara, Stout, Fuss, & Miller, in press; note, however, that elemental trials were presented after the compound trials in this study). The latter finding is consistent with the extended comparator hypothesis (Denniston, Savastano, & Miller, 2001), according to which blocking results from the comparison of a number of associations at the time of testing instead of from competition at the time of association formation. Another possible explanation is that the number of elemental trials somehow influences assumptions about the maximal level of the US that is possible.

Method

Subjects—The subjects were 36 male and 36 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony at Binghamton University, SUNY, with body weights ranging between 280 and 350 g for males and between 185 and 240 g for females. Subjects were randomly assigned to one of the six groups ($n_s = 12$), counterbalanced within groups for sex. The animals were individually housed in standard hanging stainless-steel wire-mesh cages in a vivarium maintained on a 16/8-hr light/dark cycle. Experimental manipulations occurred near the middle portion of the light phase. The animals were allowed 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. From the time of weaning until the start of the study, all animals were handled for 30 s three times per week.

Apparatus—The apparatus consisted of 12 operant chambers, each measuring $30.5 \times 27.5 \times 27.3$ cm (length $\times$ width $\times$ height). All chambers had clear Plexiglas ceilings and side walls and metal front and back walls. On one metal wall of each chamber, there was an operant lever, and adjacent to it was a niche ($4.5 \times 4.0 \times 4.5$ cm) centered 3.3 cm above the floor. A solenoid valve could deliver 0.04 cc of water into a cup on the bottom of the niche. Chamber floors were 4-mm diameter grids spaced 1.7 cm apart center-to-center, connected with NE-2 neon bulbs, which allowed constant-current footshock to be delivered by means of a high voltage AC circuit in series with a 1.0-M Ω resistor. All chambers were housed in sound- and light-attenuating cubicles. Three 45- Ω speakers were mounted on three different interior walls of each environmental chest. One was used to deliver either a low-frequency tone (300 Hz) or a high-frequency tone (650 Hz) of 8 dB (C scale, SPL) above the ambient background sound of 78 dB, which was produced primarily by a ventilation fan. The second one was used to produce a 6 clicks per second click train 8 dB (C scale, SPL) above background. The third one delivered a white noise 8 dB (C scale, SPL) above background. In addition, a Sonalert (Model H12-02R, Coulbourn Instruments, Allentown, PA) mounted on each environmental chest was able to deliver a high-frequency (1900 Hz) tone 8 dB (C scale, SPL) above background as well as a buzzer able to deliver a buzzing sound 8 dB (C scale, SPL) above background. A light bulb was mounted on the ceiling of the chamber to provide a 1-Hz flashing light stimulus (60-W incandescent bulb nominal at 120VAC driven at 80VAC). The low-frequency tone and the Sonalert served as the A and B stimuli, respectively, counterbalanced within groups. The click train served as target stimulus X. The high frequency tone, the buzzer, and the flashing light were used as the C, D and E stimuli, respectively, counterbalanced within groups. The white noise was used to signal water delivery during lever-press training and testing. Each chamber was illuminated by a dim houselight (1.2 W, #1820). Chamber assignments within the six groups were counterbalanced.

Procedure—The experiment involved 1-hr sessions on each of 18 days. Animals were run in 6 clusters of 12 rats each counterbalanced with respect to group.

Shaping: On Days 1 through 6, lever-press shaping was conducted for 60 min. Animals lever pressed for water in 0.04-cc servings; to make water delivery more conspicuous, a 0.5-s white noise was presented simultaneously with the delivery of water reinforcement. On Days 1 and 2, a fixed-time 120-s (FT-120-s) schedule of noncontingent water delivery operated concurrently with a continuous reinforcement (CRF) schedule for lever pressing. On Day 3, subjects were trained on the CRF schedule alone. On Days 4 – 6, subjects were trained on a variable-interval 20-s (VI-20-s) schedule. After shaping, the levers were removed.

Phase 1: Subadditivity pretraining: On Days 7 through 10, experimental (i.e., blocking treatment) and control animals in the subadditive condition were exposed daily to one pairing of each of two elemental cues (C and D) with the 0.5-s, 0.7-mA footshock US as well as one

pairing of the CD compound with the US. Termination of the US coincided with the termination of the 30-s cue. The animals in the irrelevant element condition were exposed daily to one pairing each of C, D, and E with the US. The animals in the irrelevant compound condition were exposed daily to two pairings of C with the US and one pairing of DE with the US. On Days 7 and 9, the order of trials was C, D, CD (or C, D, E or C, C, DE), and on Days 8 and 10, the order of trials was C, CD, D (or C, E, D or C, DE, C). CSs were presented at Minutes 15, 24, and 45 on Days 7 and 8 and at Minutes 9, 30, and 45 on Days 9 and 10.

Phase 2: Elemental training: On Days 11 through 13, the experimental (i.e., blocking treatment) animals in all conditions were exposed daily to four pairings of A with the US, with A and the US coterminating. The control animals received similar pairings of cue B with the US. The US again consisted of a 0.5-s, 0.7-mA footshock in all groups. Pairings were presented at Minutes 18, 24, 36, and 54 on all days.

Phase 3: Compound training: On Day 14, all animals received four pairings of the compounded presentation of A and X with the US, with parameters similar to those of the previous phase.

Reshaping: After the levers were reinstalled, on Days 15 to 17, all animals were retrained to lever press for water in 0.04-cc servings on the VI-20-s schedule. No experimental stimuli were presented.

Test: On Day 18, all subjects were tested on X. The X stimulus was presented three times, for 30 s each, at Minutes 6, 12, and 18 of the session. The number of lever-press responses during X was compared with half the number of lever presses during a 1-min period immediately preceding onset of X.

Results and Discussion—Suppression ratios were calculated for each of the three test presentations of the critical CS X. To this end, we divided the number of responses during the 30-s presentation of X by the sum of this number and the mean number of responses per 30 s during the 1-min pre-CS period (thus, a value of 0 corresponds to complete suppression, and a value of 0.5 corresponds to a complete lack of suppression). Figure 1 shows the mean suppression ratios for experimental and control animals in each condition, averaged over the three presentations of X. Mean suppression ratios over the three presentations of X were analyzed by means of planned comparisons. A first set of three comparisons contrasted responding to X between experimental and control animals in each condition in order to evaluate the degree of blocking in each pretraining condition. Another pair of planned comparisons contrasted the difference in responding to X between experimental and control animals after subadditive pretraining on the one hand and irrelevant element pretraining or irrelevant compound pretraining on the other hand in order to evaluate differences in blocking between subadditive and irrelevant pretraining conditions. Missing data resulting from no responses being emitted during both the pre-CS interval and the CS interval for a given presentation of X were replaced by group means.⁴

As was expected with the present parameters, forward blocking (i.e., less suppression to X in the experimental than in the control animals) was significant with both forms of irrelevant pretraining: $F(1, 66) = 14.01$, $MSE = .009$, $p < .01$, $d = 1.53$ for irrelevant element pretraining, and $F(1, 66) = 36.95$, $MSE = .009$, $p < .01$, $d = 2.48$ for irrelevant compound pretraining. However, the animals that received the explicitly subadditive pretraining did not show any sign

⁴In one set of additional analyses, missing data were replaced by zero values (representing complete suppression); in another set, cases with missing data were excluded. Both sets of additional analyses led to results similar to those in the analyses reported here. The same held true for Experiments 2 and 3.

of a blocking effect, $F(1, 66) < 1$, $MSE = .009$, $d = -0.07$. Further tests confirmed that the difference between the experimental and the control groups (i.e., blocking) was larger in the irrelevant element condition than in the subadditive condition, $F(1, 66) = 7.65$, $MSE = .009$, $p < .01$, $d = 1.60$, and was likewise larger in the irrelevant compound condition than in the subadditive condition, $F(1, 66) = 19.52$, $MSE = .009$, $p < .01$, $d = 2.55$. This result suggests that the assumption of outcome additivity is a basic constraint for blocking to occur, such that blocking does not occur as readily if this assumption is explicitly contradicted.

Experiment 2

In a second experiment, we implemented the inverse manipulation (see Table 2). The blocking procedure now involved parameters that would typically yield weak, if any, blocking (i.e., four elemental training trials followed by eight compound training trials). This time, with a set of cues different from the blocking cues, experimental and control animals in the additive condition were explicitly shown that effective cues had additive outcomes: The combination of two effective cues resulted in a stronger outcome than did either one alone (C+/D+/CD++, with ++ indicating the presentation of a more intense US). Again, the experimental and control animals in both other conditions received a similar amount of pretraining, involving equal amounts and intensities of shocks. However, for these animals, pretraining was irrelevant with respect to additivity. The irrelevant element condition received C+/D+/E++ pretraining, whereas the irrelevant compound condition received C+/C+/DE++ pretraining. If blocking in Pavlovian conditioning is sensitive to constraints of causal inference, blocking might be enhanced by explicitly additive pretraining.

Method

Subjects and Apparatus—Again, 36 male and 36 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony were used, with body weights ranging between 250 and 385 g for males and between 195 and 285 g for females. They were maintained in the same way as the animals used in Experiment 1 and were randomly assigned to one of the six groups ($n_s = 12$), counterbalanced within groups for sex. The apparatus was identical to that used in Experiment 1.

Procedure—The experiment involved 16 days of 1-hr sessions. Animals were run in six squads of 12 rats, counterbalanced with respect to group.

Shaping: On Days 1 through 6, all animals were shaped to lever press for water as in Experiment 1, after which the levers were removed.

Phase 1: Additivity pretraining: On Days 7 through 10, animals in the additive condition were exposed daily to one pairing of each of two elemental cues (C and D) with the 0.7-mA US, as well as one pairing of the CD compound with the 1.0-mA US. (Psychophysically, 1.0 mA is roughly twice as aversive as 0.7 mA). Animals in the irrelevant element condition were exposed daily to one pairing of each of two elemental cues (C and D) with the 0.7-mA US, as well as one pairing of a different elemental cue (E) with the 1.0-mA US. Animals in the irrelevant compound condition were exposed daily to two pairings of an elemental cue (C) with the 0.7-mA US as well as one pairing of the DE compound with the 1.0-mA US. On Days 7 and 9, the order of trials was C, D, CD (or C, D, E or C, C, DE), and on Days 8 and 10, the order of trials was C, CD, D (or C, E, D or C, DE, C). Pairings were presented on Minutes 15, 24, and 45 on Days 7 and 8 and on Minutes 9, 30, and 45 on Days 9 and 10. Cues were all 30 s in duration, USs were all 0.5 s in duration, and the termination of the USs coincided with the termination of the cues.

Phase 2: Elemental training: On Day 11, the experimental animals in all conditions were exposed to four pairings of A with the US. The control animals received similar pairings of cue B with the US. The US consisted of a 0.7-mA footshock in all groups. Pairings were presented on Minutes 18, 24, 36, and 54.

Phase 3: Compound training: On Days 12 and 13, all animals received four pairings of the compounded presentation of A and X with the US, with parameters similar to those of the previous phase.

Reshaping: After the levers were reinstalled, on Days 14 and 15, reshaping was performed as in Experiment 1. No experimental stimuli were presented.

Test: On Day 16, all animals were tested on X as in Experiment 1.

Results and Discussion—Suppression ratios were calculated and analyzed as for Experiment 1 (see Figure 2). Blocking (i.e., less suppression to X in the experimental than in the control animals) was obtained after additive pretraining, $F(1, 66) = 4.09$, $MSE = .014$, $p < .05$, $d = 0.83$, and after irrelevant compound pretraining, $F(1, 66) = 5.24$, $MSE = .014$, $p < .03$, $d = 0.93$, but not after irrelevant element pretraining, $F(1, 66) < 1$, $MSE = .014$, $d = 0.10$. However, the difference in responding between the experimental and control groups (i.e., blocking) was not significantly larger after additive pretraining than after either irrelevant element pretraining, $F(1, 66) = 1.59$, $MSE = .014$, $p = .21$, $d = 0.73$, or irrelevant compound pretraining, $F(1, 66) < 1$, $MSE = .014$, $d = 0.11$. So, whereas explicitly contradicting additivity eliminated an otherwise robust blocking effect in Experiment 1, explicitly confirming additivity did not reliably increase the weak blocking effect obtained here.

At least two reasons can account for this finding. First, it could be that our manipulation simply lacked sufficient power to detect a difference between the groups. The fact that the blocking effect was at least as large in the irrelevant compound pretraining condition as in the additive pretraining condition might seem to argue against this. However, in both groups, animals experienced that two cues produced a stronger outcome than just one cue during pretraining. Although the animals in the irrelevant compound condition never experienced the elements of the compound in isolation, they might in principle have formed the expectation that two cues should result in a stronger outcome than one. Even though such an explanation goes beyond our original predictions, it would be generally consistent with a symbolic account. Note, however, that even a planned comparison in which the blocking effect in the additive and irrelevant compound conditions was contrasted with the blocking effect in the irrelevant element condition did not reach significance, $F(1, 66) = 2.45$, $MSE = .014$, $p = .12$, $d = 0.78$. Also, the size of the blocking effect in the former two conditions was markedly smaller than that of the blocking effect obtained in the irrelevant pretraining conditions of Experiment 1.

Second, and perhaps more interesting, it could be that additivity is somehow the default assumption for the animals. Thus, explicitly demonstrating additivity (Experiment 2) might have a limited impact on blocking because it only confirms an assumption already in place, whereas explicitly violating additivity by demonstrating subadditivity (Experiment 1) affects blocking more substantially because it effectively disconfirms the default assumption. Such an interpretation is consistent with the fact that without any pretraining, blocking tends to be a fairly reliable phenomenon in animal Pavlovian conditioning (given that appropriate parameters are used; see the introduction to Experiment 1). It is also consistent with the occurrence of phenomena such as overexpectation, in which reinforcement of a compound of two elements that have both previously been reinforced elementally reduces conditioned responding to the elements (e.g., Kremer, 1978). It might also explain why many contemporary associative models, in which additivity is a hardwired feature of associative learning (rather

than a flexible assumption from which blocking is deduced), have been so successful in explaining the bulk of conditioning phenomena (see De Houwer, Vanderpe, & Beckers, 2005, and Lovibond et al., 2003, for related arguments).

It might look as if there is a discrepancy between the amplification of blocking by additive pretraining that is observed in human causal learning and the relative lack of such amplification we seemed to observe here. However, additive pretraining has been shown to increase blocking in human causal learning studies only in comparison to a nonadditive condition (Lovibond et al., 2003; see also Beckers et al., 2005). Lovibond et al. (2003) provided no direct evidence that additive pretraining enhances blocking relative to a neutral control condition, and it is probable from the available evidence that in human causal learning, additivity training only helps when additivity for some reason cannot be assumed or cannot be verified (e.g., because of a ceiling; see Beckers et al., 2005; De Houwer et al., 2002).

Experiment 3

In the third experiment, blocking training was again administered with parameters expected to result in weak blocking at most (four elemental training trials and four compound training trials). However, before the blocking training, all animals received a number of unsignaled footshocks of two different intensities. For experimental and control animals in the maximal condition, subsequent blocking training involved the stronger of the two shocks. In two other conditions, equated with the first condition on either absolute shock intensities during preexposure (submaximal low condition) or absolute shock intensity during blocking training (submaximal high condition), blocking training involved the weaker of the two preexposure shocks (see Table 3). In the case in which blocking training was conducted with the strongest of the two preexposure shocks, an inference concerning whether the presence of X increases the outcome relative to the presence of A alone would be occluded by a possible ceiling effect: The outcome on A+ and AX+ trials would be the most intense outcome ever experienced by the animal, so it is unclear whether the animal would be able to anticipate anything more intense. Such a ceiling effect would obviously be absent in the case in which the weaker of the two preexposure shocks was used during blocking training: If presentation of the compound of A and X is followed by the same submaximal shock as the presentation of A alone, then it is logically plausible to infer that X does not add anything to the outcome produced by A. Alternatively stated, if blocking is modulated by the magnitude of the shock used during the blocking training relative to prior shock that the subject has experienced (i.e., blocking would be increased by using a definitely submaximal shock during blocking treatment), this would again suggest that rats are sensitive to formal constraints of causal reasoning. In order to minimize any US preexposure deficit in subsequent conditioning, a clear context switch was made between the preexposure phase and the actual blocking treatment (see *Procedure* section).

Method

Subjects and Apparatus—Again, 36 male and 36 female experimentally naive, Sprague-Dawley-descended rats obtained from our own breeding colony were used. Bodyweight ranges were 250–310 g for males and 180–255 g for females. As before, animals were randomly assigned to one of the six groups ($n = 12$), counterbalanced within groups for sex. The subjects and the apparatus otherwise conformed to those in Experiments 1 and 2.

Procedure—The experiment involved 13 days of 1-hr sessions. Animals were run in six squads of 12 rats. Conditions were run interspersed.

Shaping: On Days 1 through 6, animals were shaped to lever press for water as before. After shaping, levers were removed.

Phase 1: US preexposure: On Days 7 and 8, the experimental and control animals in all groups were exposed to two shocks of low intensity and two shocks of high intensity (each 0.5 s in duration) on each day. For animals in the submaximal high condition, the low-intensity shock was 0.75 mA and the high-intensity shock was 1.0 mA. For all other animals, the low-intensity shock was 0.4 mA and the high-intensity shock was 0.75 mA. (Psychophysically, the difference between 0.4 mA and 0.75 mA is roughly equivalent to the difference between 0.75 mA and 1.0 mA.) All animals received the high-intensity shocks at Minutes 12 and 36 and the low-intensity shocks at Minutes 30 and 48 on each day. During Phase 1, in order to minimize any US preexposure deficit in subsequent conditioning, a context different from the other phases was installed by presenting a distinctive olfactory cue (produced by two drops of 98% methyl salicylate, a mint odor, on the top surface of a wooden cube placed inside the sound-attenuating environmental isolation chest but outside of the experimental chamber) and a distinctive visual cue (houselight turned off).

Phase 2: Elemental training: On Day 9, the experimental animals in all conditions were exposed to four pairings of A with the US, with the termination of the 0.5-s US coinciding with the termination of the 30-s A stimulus. The control animals received similar pairings of cue B with the US. For the submaximal low condition, the US consisted of a 0.4-mA footshock. For both other conditions, the US consisted of a 0.75-mA footshock. The pairings were presented at Minutes 18, 24, 36, and 54 of the session.

Phase 3: Compound training: On Day 10, all animals received four pairings of the compounded presentation of A and X with the US, with parameters similar to those of the previous phase.

Reshaping: After the levers were reinstalled, on Days 11 and 12, all animals were retrained to lever press for water as before. No experimental stimuli were presented.

Tests: On Day 13, all animals were tested on X as in the previous experiments.

Results and Discussion—Suppression ratios were calculated and analyzed as before (see Figure 3). Irrespective of absolute shock intensity, blocking (i.e., reduced suppression to X in the experimental group relative to the control group) was reliable in animals trained with the weaker of the two shock intensities experienced during preexposure: $F(1, 66) = 11.67$, $MSE = .006$, $p < .01$, $d = 1.39$ for the submaximal high condition, and $F(1, 66) = 11.36$, $MSE = .006$, $p < .01$, $d = 1.38$ for the submaximal low condition. However, blocking was completely absent in animals receiving blocking training with the stronger of both preexposure shocks, $F(1, 66) < 1$, $MSE = .006$, $d = 0.01$. Accordingly, the difference in responding between the experimental and control groups was significantly larger with submaximal shock intensity than with maximal shock intensity: $F(1, 66) = 5.75$, $MSE = .006$, $p < .02$, $d = 1.38$ for the difference in blocking between the maximal and the submaximal high conditions, and $F(1, 66) = 5.59$, $MSE = .006$, $p < .03$, $d = 1.37$ for the difference between the maximal and the submaximal low conditions. This demonstrates that under circumstances that do not ordinarily yield blocking, explicit experience with more intense outcomes can unveil blocking in animal conditioning, in line with a causal reasoning analysis.

General Discussion—Experiment 1 demonstrated that under circumstances that result in robust blocking after irrelevant pretraining treatment, explicit exposure to subadditivity during pretraining effectively abolishes blocking. Experiment 2 suggested that the reverse does not hold: When parameters are used that yield weak, if any, blocking after irrelevant pretraining, exposure to additivity during pretraining does not demonstrably enhance blocking. Experiment 3 demonstrated that blocking is obtained more readily if the blocking procedure involves a

shock intensity that is submaximal than if the blocking procedure involves the most intense shock intensity that animals have ever experienced, independent of the absolute intensity of the shock.

This sensitivity to cue additivity and outcome intensity ceiling is remarkably parallel to what has recently been documented for human causal learning and human electrodermal conditioning (Beckers et al., 2005; Lovibond et al., 2003; Mitchell & Lovibond, 2002). In human causal learning, these results have been used to argue that blocking does not (or at least does not always) result from automatic selective associative processing. It seems that the present results call for a similar conclusion regarding blocking in Pavlovian animal conditioning. Indeed, none of the currently existing theories of associative learning is able to account for our results.

Existing associative theories are actually silent when it comes to effects of additivity versus subadditivity pretraining because they do not allow abstract learning about the rules that apply to one set of cues to transfer to an entirely different set of cues. In the context of human causal learning studies, Livesey and Boakes (2004) proposed a way out for associative accounts of blocking by suggesting that subadditive pretraining encourages a configural mode of processing. That is, because the effect produced by the compound of two cues cannot be reduced to the sum of the effects produced by its constituent elements in the case of subadditive pretraining, people would be inclined to process that compound and all subsequent compounds they encounter as different, configural cues instead of as combinations of two elements. As a result, the compound of A and X would be processed as an entirely new cue, so that $A \rightarrow O$ training does not affect learning on the $AX \rightarrow O$ trials, and responding to X should not be expected to be different after $A \rightarrow O/AX \rightarrow O$ training than after $B \rightarrow O/AX \rightarrow O$ training. There is indeed evidence that manipulations that encourage a configural mode of processing reduce blocking in humans (Williams, Sagness, & McPhee, 1994). However, configuring does not provide a satisfactory account for the effect that subadditive training has on blocking in human causal learning because subadditive training has similar effects when presented after the actual blocking procedure and when presented before, which is not compatible with the idea that the effect of subadditive training is due to changes in the way compounds are encoded (Beckers et al., 2005). Moreover, encouraging configuring in rats does not seem to have effects similar to those that it has in humans; in fact, configural training does not seem to have an effect on later learning in rats at all (Williams & Braker, 2002). In line with this latter evidence, in Experiment 1 subadditive training did not diminish generalization from $AX+$ compound training to responding to X at test in the subadditive control group compared with the irrelevant control groups (see Figure 1). If subadditive training influenced the degree of configuring, one would expect to see less conditioned suppression to X in the subadditive control group, relative to both other control groups, unlike what we observed.

Obviously, an explanation in terms of configuring would also not apply for the results of Experiment 3, in which no cues were presented during pretraining at all. Effects of outcome ceiling might in principle be dealt with in associative terms by assuming that relative outcome magnitude (rather than absolute outcome magnitude) affects outcome ceiling and asymptotic associative strength. However, simulations show that an associative model that does so would actually predict the opposite pattern of results from the one obtained here, that is, stronger blocking with increasing outcome magnitude. According to the Rescorla and Wagner (1972) model, on each trial on which a US is presented, the associative strengths (V) of the CSs that are present are updated according to the following equation:

$$\Delta V_j = \alpha \beta (\lambda - \sum V_j).$$

In this equation, α and β represent the salience or associability of CS j and the US, respectively. The maximal associative strength that is supportable by the US is represented by λ . Finally,

ΣV_j represents the summed existing associative strengths of all CSs that are present on a given presentation of the US. According to this model, blocking occurs because, during A+ training, A acquires a certain degree of associative strength. On subsequent AX+ trials, both A and X acquire only limited associative strength because of the associative strength already acquired by A decreasing the common error term ($\lambda - \Sigma V_j$). Conditioned responding is assumed to be a direct function of associative strength. In order to simulate effects of outcome intensity on blocking, one can vary either US salience (β), the asymptotic associative strength supported by the US (λ), or both. Figure 4 depicts simulated associative strength for cue X in the experimental condition and in the control condition of a blocking procedure. For the simulations, associative strength for X after four A+ trials followed by four AX+ trials was compared with associative strength for X after four B+ followed by four AX+ trials, assuming $\lambda = 1.0$ and $\beta = 0.6$ for the maximal condition and either $\lambda = 1$ and $\beta = 0.3$, $\lambda = 0.5$ and $\beta = 0.6$, or $\lambda = 0.5$ and $\beta = 0.3$ for the submaximal conditions, assuming equal salience for A, B, and X (all set at $\alpha = .6$; the simulations yield identical patterns for any nonzero value of α). Blocking (i.e., the difference in the associative strength of X between the experimental and control groups) is markedly smaller in every instantiation of the submaximal condition than in the maximal condition. Other existing associative models either yield the same result or predict no difference between maximal and submaximal conditions.

In contrast, sensitivity to cue additivity and outcome ceiling is entirely expected from a causal reasoning perspective. In order to conclude after $A \rightarrow O/AX \rightarrow O$ training that X is not a valid cause of the outcome, one has to assume that if both were effective, A and X would summate. Moreover, when this assumption is made, in order to empirically verify that X does not add to the outcome produced by A, the outcome should not be at ceiling if A is presented alone. Accordingly, the observation that human causal learning and human electrodermal conditioning are sensitive to effects of cue additivity and outcome ceiling, as well as a number of other recent findings, have been taken to indicate that blocking in human learning reflects the operation of symbolic causal reasoning processes (De Houwer, Beckers, & Vandorpe, 2005; Lovibond, 2003; Waldmann & Walker, 2005). The fact that blocking in rat conditioning is sensitive to pretrained cue additivity and experiential outcome ceiling strongly suggests that parallel processes are operating in rodents. Obviously, this conclusion endows relatively simple mammals such as rats with remarkable cognitive abilities, rather than confining these capabilities to higher mammals such as humans and (some) other apes. It also suggests that a comparative approach to inference making and reasoning is a potentially worthwhile endeavor.

One way of formalizing the idea that Pavlovian conditioning to some extent reflects causal reasoning is by assuming that Pavlovian conditioning relies on Bayesian structure learning. According to structure learning algorithms, an abstract representation of the causal structure of a set of variables (termed a causal Bayes net) is inferred from the available data by comparing the likelihood of a number of possible causal structures on the basis of the conditional probabilities among the variables (Pearl, 2000). Applied to a blocking procedure, these structure learning algorithms can be used to decide whether a structure in which CS X is a cause of the US or a structure in which CS X is not a cause of the US is more compatible with the presented contingency information (see Gopnik et al., 2004, for a detailed Bayes nets account of blocking). According to recent extensions of the Bayes nets formalism, a priori assumptions can limit the space of causal structures that are considered, for example, by imposing constraints on how candidate causal structures are parameterized (Tenenbaum & Griffiths, 2003; see also Sobel, Tenenbaum, & Gopnik, 2004). This way, prior experience can result in a switch from estimating the likelihood of candidate causal models that implement a linear (i.e., additive) integration rule (which seems to be the default in human causal learning and perhaps in Pavlovian conditioning as well) to considering causal models that implement a nonlinear (i.e., subadditive) integration rule (see Lagnado, Waldmann, Hagmayer, & Sloman, in press, and Waldmann, 1996, for earlier proposals from a causal model theory perspective).

on how prior assumptions about integration rules can affect causal learning). Human causal induction experiments indeed suggest that the context in which causal induction is performed affects people's a priori assumptions about the parameterization of candidate causal structures (Griffiths, 2005) in a way that is consistent with the modified causal Bayes nets formalism proposed by Tenenbaum and Griffiths (2003).

One specific instance of Bayes net theory that is at least conceptually compatible with an effect of outcome maximality on blocking is Cheng's (1997) power probabilistic contrast (PC) theory, which is an extension of probabilistic contrast models (e.g., Cheng & Holyoak, 1995). In these models, it is assumed that human causal judgment reflects the outcome of probabilistic contrasts, in which the probability of the outcome given a certain cue is compared with the probability of the outcome given the absence of the cue while controlling for the presence or the absence of other cues. In the case of a blocking procedure, this implies comparing the probability of the outcome in the presence and in the absence of the blocked cue, X, while keeping the presence of the blocking cue, A, constant: $P(O|X.A) - P(O|\neg X.A)$. In a blocking procedure, these probabilities are equal, so the probabilistic contrast computes to zero and a low causal estimate of X results. However, according to the power PC theory, in order to provide an estimate of generative causal power, the probabilistic contrast has to be normalized for the base rate of the outcome in the absence of X by dividing the above contrast by $[1 - P(O|\neg X.A)]$. As such, the power PC theory predicts that blocking will be sensitive to ceiling information because a generative causal estimate cannot be derived if the probability of the outcome is maximal when only A is present. Indeed, if $P(O|\neg X.A)$ equals one, then the probabilistic contrast has to be divided by zero, which results in an indeterminate value. Under these circumstances, people will be uncertain about the causal status of X, just like they will be if no A 3 O trials are presented and $P(O|\neg X.A)$ cannot be estimated at all. As a result, blocking should not occur. However, the power PC theory, like other probabilistic models, concerns only outcome probability, not outcome rate or magnitude. Therefore, the effect of outcome maximality that we obtained here is only conceptually consistent with Cheng's arguments.

Ultimately, we cannot categorically exclude the possibility that somehow an associative model could be devised that is able to account for the present results. It is likely that such a model would then also be able to account for other, seemingly exclusively symbolic aspects of human causal reasoning. Minimally, the present results suggest that human and nonhuman animal cognition are less different than might be deduced from apparent demonstrations of higher order, symbolic processing in humans that are not easily reduced to subsymbolic processes.

Finally, even if animal (including human) learning cannot be entirely reduced to subsymbolic processes, associative learning probably plays a crucial role in animal conditioning and learning all the same. For one thing, irrespective of the fact that blocking appears to be greatly enhanced when conditions allow for a valid blocking inference according to a causal reasoning analysis, blocking can still occur—when using appropriate parameters—under circumstances that do not seem to allow for such inference, both in human causal learning (e.g., Beckers et al., 2005; Lovibond et al., 2003) and in animal Pavlovian conditioning (e.g., Experiment 1, irrelevant conditions). Exactly how automatic associative processes and controlled symbolic processes combine and interplay to give rise to human causal judgments and animal conditioned responses promises to stay an intriguing question for some time to come. One distinct possibility is that association formation most often relies on fairly simple, noncompetitive processes rather than being driven by complex error-reduction functions. The associative system might simply be *gluing together* everything that cooccurs (provided that it is attended). More complex, inferential processes such as are instantiated in the theory-based causal Bayes nets formalism (Tenenbaum & Griffiths, 2003) might then operate on this knowledge to give rise to conditioned performance that can be quite sophisticated. Presumably,

these inferential processes can have differing degrees of complexity depending on situational requirements and constraints.

Acknowledgements

We thank Jim Esposito and Mike Karmin for their help in running the experiments and Frank Baeyens, Marc Buehner, Tom Griffiths, Oskar Pineño, and Daniel Wheeler for valuable suggestions and discussions during the execution and write-up of this work.

References

- Beckers T, De Houwer J, Pineño O, Miller RR. Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2005;31:238–249.
- Cheng PW. From covariation to causation: A causal power theory. *Psychological Review* 1997;104:367–405.
- Cheng, PW.; Holyoak, KJ. Complex adaptive systems as intuitive statisticians: Causality, contingency and prediction. In: Meyer, JA.; Roitblat, H., editors. *Comparative approaches to cognition*. MIT Press; Cambridge, MA: 1995. p. 34-45.
- De Houwer J, Beckers T. A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology* 2002;55B:289–310.
- De Houwer J, Beckers T. Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology* 2003;56B:345–357.
- De Houwer J, Beckers T, Glautier S. Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology* 2002;55A:965–985.
- De Houwer J, Beckers T, Vandorpe S. Evidence for the role of higher-order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior* 2005;33:239–249. [PubMed: 16075842]
- De Houwer, J.; Vandorpe, S.; Beckers, T. On the role of controlled cognitive processes in human associative learning. In: Wills, AJ., editor. *New directions in human associative learning*. Erlbaum; Mahwah, NJ: 2005. p. 41-63.
- Denniston, JC.; Savastano, HI.; Miller, RR. The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In: Mowrer, RR.; Klein, SB., editors. *Handbook of contemporary learning theories*. Erlbaum; Mahwah, NJ: 2001. p. 65-117.
- Dickinson A. The 28th Bartlett memorial lecture: Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology* 2001;54B:3–25.
- Dickinson A, Shanks DR, Evenden J. Judgement of act–outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology* 1984;36A:29–50.
- Gopnik A, Glymour C, Sobel DM, Schulz LE, Kushnir T, Danks D. A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review* 2004;111:1–30.
- Griffiths, TL. *Causes, coincidences, and theories*. Stanford University; Stanford, CA: 2005. Unpublished doctoral dissertation
- Kamin, LJ. Predictability, surprise, attention, and conditioning. In: Campbell, BA.; Church, RM., editors. *Punishment and aversive behavior*. Appleton-Century-Crofts; New York: 1969. p. 279-296.
- Kremer EF. The Rescorla–Wagner model: Losses in associative strength in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes* 1978;4:22–36. [PubMed: 621465]
- Lagnado, DA.; Waldmann, MR.; Hagmayer, Y.; Sloman, SA. Beyond covariation: Cues to causal structure. In: Gopnik, A.; Schulz, L., editors. *Causal learning: Psychology, philosophy, and computation*. Oxford University Press; Oxford, England: in press

- Livesey EJ, Boakes RA. Outcome additivity, elemental processing and blocking in human causality judgements. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology* 2004;57B:361–379.
- Lovibond PF. Causal beliefs and conditioned responses: Retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 2003;29:97–106.
- Lovibond PF, Been S-L, Mitchell CJ, Bouton ME, Frohardt R. Forward and backward blocking of causal judgement is enhanced by additivity of effect magnitude. *Memory & Cognition* 2003;31:133–142.
- 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.
- Mitchell CJ, Lovibond PF. Backward and forward blocking in human autonomic conditioning requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology* 2002;55B:311–329.
- Pavlov, IP. *Conditioned reflexes*. Oxford University Press; Oxford, England: 1927.
- Pearl, J. *Causality: Models, reasoning, and inference*. Oxford University Press; New York: 2000.
- Pineño O, Urushihara K, Stout S, Fuss J, Miller RR. When more is less: Extending training of the blocking association following compound training attenuates the blocking effect. *Learning & Behavior*. in press
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical conditioning II: Current research and theory*. Appleton-Century-Crofts; New York: 1972. p. 64-99.
- Sobel DM, Tenenbaum JB, Gopnik A. Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science* 2004;28:303–333.
- Tenenbaum, JB.; Griffiths, TL. Theory-based causal inference. In: Becker, S.; Thrun, S.; Obermayer, K., editors. *Advances in neural information processing systems*. 15. MIT Press; Cambridge, MA: 2003. p. 35-42.
- Vandorpe S, De Houwer J, Beckers T. Further evidence for the role of inferential reasoning in forward blocking. *Memory & Cognition* 2005;33:1047–1056.
- Waldmann, MR. Knowledge-based causal induction. In: Shanks, DR.; Holyoak, KJ.; Medin, DL., editors. *The psychology of learning and motivation: Vol. 34. Causal learning*. Academic Press; San Diego, CA: 1996. p. 47-88.
- Waldmann MR, Walker JM. Competence and performance in causal learning. *Learning & Behavior* 2005;33:211–229. [PubMed: 16075840]
- Williams DA, Braker DS. Input coding in animal and human learning. *Behavioural Processes* 2002;57:149–161. [PubMed: 11947995]
- Williams DA, Sagness KE, McPhee JE. Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 1994;20:694–709.

Figure 1.
Experiment 1: Mean suppression ratios for experimental and control groups by pretraining condition. Error bars represent standard errors of the means.

Figure 2. Experiment 2: Mean suppression ratios for experimental and control groups by pretraining condition. Error bars represent standard errors of the means.

Figure 3.

Experiment 3: Mean suppression ratios for experimental and control groups by outcome intensity condition. Error bars represent standard errors of the means.

Figure 4.

Simulated associative strengths for experimental and control animals in a maximal condition and in three instantiations of a submaximal condition. Note that suppression ratios are inversely related to associative strength. See text for details.

Table 1

Design of Experiment 1

Condition and group	Phase 1: Pretraining	Phase 2: Elemental	Phase 3: Compound
Subadditive			
Experimental	4 C+/4 D+/4 CD+	12 A+	4 AX+
Control	4 C+/4 D+/4 CD+	12 B+	4 AX+
Irrelevant element			
Experimental	4 C+/4 D+/4 E+	12 A+	4 AX+
Control	4 C+/4 D+/4 E+	12 B+	4 AX+
Irrelevant compound			
Experimental	4 C+/4 C+/4 DE+	12 A+	4 AX+
Control	4 C+/4 C+/4 DE+	12 B+	4 AX+

Note. The + represents a 0.7-mA footshock 0.5 s in duration; C, D, and E represent a 30-s high tone (650 Hz, 8 dB above background), 30-s buzzer (8 dB above background), and 30-s flashing light (0.5 s on, 0.5 s off), counterbalanced; A and B represent a 30-s Sonalert (1900 Hz, 8 dB above background) and 30-s low tone (320 Hz, 8 dB above background), counterbalanced; X represents a 30-s click train (6/s, 8 dB above background). Slashes separate interspersed trials.

Table 2

Design of Experiment 2

Condition and group	Phase 1: Pretraining	Phase 2: Elemental	Phase 3: Compound
Additive			
Experimental	4 C+/4 D+/4 CD++	4 A+	8 AX+
Control	4 C+/4 D+/4 CD++	4 B+	8 AX+
Irrelevant element			
Experimental	4 C+/4 D+/4 E++	4 A+	8 AX+
Control	4 C+/4 D+/4 E++	4 B+	8 AX+
Irrelevant compound			
Experimental	4 C+/4 C+/4 DE++	4 A+	8 AX+
Control	4 C+/4 C+/4 DE++	4 B+	8 AX+

Note. The + and ++ represent 0.7-mA and 1.0-mA footshocks, respectively, both 0.5 s in duration; C, D, and E represent a 30-s high tone (650 Hz, 8 dB above background), a 30-s buzzer (8 dB above background), and a 30-s flashing light (0.5 s on, 0.5 s off), respectively, counterbalanced; A and B represent a 30-s Sonalert (1900 Hz, 8 dB above background) and a 30-s low tone (320 Hz, 8 dB above background), respectively, counterbalanced; X represents a 30-s click train (6/s, 8 dB above background). Slashes separate interspersed trials.

Table 3

Design of Experiment 3

Condition and group	Phase 1: Preexposure	Phase 2: Elemental	Phase 3: Compound
Maximal			
Experimental	4 ++/4 +	4 A++	4 AX++
Control	4 ++/4 +	4 B++	4 AX++
Submaximal high			
Experimental	4 ++++/4 ++	4 A++	4 AX++
Control	4 ++++/4 ++	4 B++	4 AX++
Submaximal low			
Experimental	4 ++/4 +	4 A+	4 AX+
Control	4 ++/4 +	4 B+	4 AX+

Note. The +, ++, and ++++ represent 0.4-mA, 0.75-mA, and 1.0-mA footshocks, respectively, each 0.5 s in duration; A and B represent a 30-s Sonalert (1900 Hz, 8 dB above background) and a 30-s low tone (320 Hz, 8 dB above background), respectively, counterbalanced; X represents a 30-s click train (6/s, 8 dB above background). Slashes separate interspersed trials.