Outcome Additivity and Outcome Maximality Influence Cue Competition in Human Causal Learning

Tom Beckers

University of Leuven and State University of New York at
Binghamton

Jan De Houwer

Ghent University

Oskar Pineño and Ralph R. Miller
State University of New York at Binghamton

Recent research suggests that outcome additivity pretraining modulates blocking in human causal learning. However, the existing evidence confounds outcome additivity and outcome maximality. Here the authors present evidence for the influence of presenting information about outcome maximality (Experiment 1) and outcome additivity (Experiment 2) on subsequent forward blocking. The results of Experiment 3 confirm that, with outcome maximality controlled, outcome additivity affects backward blocking but not release from overshadowing. Finally, the results of Experiment 4 demonstrate that information about outcome additivity has a similar effect on forward blocking if presented after the blocking training instead of before. The results are compatible with the idea that blocking results from inferential processes at the time of testing and not from a failure to acquire associative strength during training.

Cue competition effects are currently among the most intensively studied phenomena in human causal learning (see De Houwer & Beckers, 2002b; Dickinson, 2001, for recent reviews). Cue competition refers to the observation that potential causes of a given effect tend to compete for causal status. That is, the perceived causal efficacy of a given cue (X) in producing a given outcome is determined not only by the cooccurrence of X and the outcome (represented as X+) but also by the degree of contingency between other, competing, cues and the outcome. For instance, in forward blocking, the causal status of X, which in compound with a second cue, A, was repeatedly paired with the outcome, is reduced if A alone was previously repeatedly paired with the outcome (thus, A+ trials followed by AX+ trials; Dick-

inson, Shanks, & Evenden, 1984). The same effect was first demonstrated in Pavlovian conditioning of nonhuman animals (Kamin, 1968). In accordance, when forward blocking was observed in human causal learning, some researchers proposed that associative theories of animal conditioning, such as the highly influential Rescorla-Wagner model (Rescorla & Wagner, 1972), could also adequately account for human causal learning.

Backward blocking, which involves the same procedure as forward blocking but with the trial types in reversed order (i.e., AX+ followed by A+ trials), has been demonstrated in human causal learning as well (Shanks, 1985). Although backward blocking was beyond the scope of the original Rescorla-Wagner model (Rescorla & Wagner, 1972), alternative associative models have been developed that are able to accommodate backward blocking (e.g., Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). The shared notion underlying these associative models is that competition results from an automatic tendency to learn about reliable predictors only and to disregard redundant cues for significant environmental events.

Recently, however, it has been argued that, despite the similarities between Pavlovian conditioning and human causal learning, the occurrence of cue competition effects in human causal learning might not reflect the operation of simple competitive associative principles but instead depend on humans engaging in controlled and effortful inferential reasoning (De Houwer & Beckers, 2003; Lovibond, 2003; Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003). This account assumes that by default, people entertain the assumption that multiple effective causes of a given outcome are additive. Thus, if a reliable cause for a given outcome has been established, adding a second potential cause should result in a more intense outcome if the second cause is effective. From this assumption, the observation that the combined presence of A and X results in a similar outcome as the presence of A on its own leads to the logical conclusion that X is not an effective cause of

Tom Beckers, Department of Psychology, University of Leuven, Leuven, Belgium and Department of Psychology, State University of New York at Binghamton; Jan De Houwer, Department of Psychology, Ghent University, Ghent, Belgium; Oskar Pineño and Ralph R. Miller, Department of Psychology, State University of New York at Binghamton.

Oskar Pineño is now at the Department of Psychology, University of Seville, Seville, Spain.

Tom Beckers is a postdoctoral fellow of the Fund for Scientific Research—Flanders, Belgium. Oskar Pineño is supported by a postdoctoral fellowship from the Spanish Ministry of Education (EX2002–0739). Part of the research reported here was conducted while Tom Beckers was a visiting scholar at the State University of New York at Binghamton, supported by the Fund for Scientific Research—Flanders.

We thank Jessica Fuss and Olga Lipatova for their help in running the experiments reported here, and Olga Lipatova, Jeffrey Amundson, Gonzalo Urcelay, Koji Urushihara, Miguel Vadillo, Daniel Wheeler, and Peter Lovibond for their comments on a draft of this article.

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.ac.be

the outcome. Hence the reduced causal estimate of X that is observed in forward and backward blocking (see De Houwer, Beckers, & Vandorpe, in press, for a review of the role of higher order reasoning processes in the occurrence of cue competition in human causal learning).

The idea that cue competition effects in human causal learning result from inferential reasoning processes sparks a number of unique predictions. One of these is that if people did not entertain the assumption of causal additivity, blocking would not occur. Lovibond et al. (2003) recently set out to test this prediction. They presented their participants with a number of food cues that, alone or in combination, could result in an allergic reaction in a fictitious patient. Afterward, participants had to indicate how likely an allergic reaction was to occur after eating each individual food item. Within this allergy task, they implemented forward- and backward-blocking procedures (i.e., A+ followed or preceded by AX+ trials, among a number of other cues). However, before presenting the blocking trials, they presented a number of other food items, including G and H, which were presented both alone and in compound and which were always followed by an allergic reaction (i.e., G+, H+, and GH+). It is important to note that for half of the participants, the GH compound was followed by the same allergic reaction as were the G and H elements, whereas for the other half of the participants, the GH compound resulted in a *strong allergic reaction* (as opposed to a mere *allergic reaction* on all other trials). In line with an inferential reasoning account, the outcome additivity pretraining markedly modulated the amount of blocking observed: Forward and backward blocking were larger with explicitly additive pretraining than they were with explicitly subadditive pretraining (in fact, backward blocking was not observed at all after subadditive training).

The additive and subadditive groups in the study of Lovibond et al. (2003), however, differed not only with respect to outcome additivity. Participants in the additive group were, both by instruction and by the pretraining treatment, exposed to the very possibility of variability in the intensity of the allergic reaction that a patient could develop. More specifically, during the actual blocking procedure, the outcome presented on both A+ and AX+ trials was clearly submaximal relative to some of the outcomes experienced during pretraining. Participants in the subadditive group were informed only about the possibility of an allergic reaction occurring, without any reference to severity even in the instructions. Therefore, they may have regarded the outcome as the most intense outcome possible or to be reported (or, perhaps more likely, as the only possible nonzero outcome). The exposure to outcome variability in itself may have been necessary and/or sufficient to produce the enhanced forward- and backward-blocking effects in the additively trained group. Indeed, previous research suggests that at least instructed outcome maximality can dramatically modulate forward and backward blocking (De Houwer, Beckers, & Glautier, 2002): If during a blocking procedure, outcomes are consistently accompanied by a verbal label that indicates the nonmaximal status of these outcomes, forward and backward blocking is readily obtained, whereas blocking is not as easily obtained if outcomes are consistently accompanied by a verbal label that indicates that outcome magnitude is maximal. This makes sense from an inferential reasoning perspective because, in order to logically infer that X is not a cause of the outcome, people not only have to assume that effective causes

have additive effects, but they also have to be able to empirically verify that adding X does not increase the outcome over that produced by A. Such verification is effectively prevented if the outcome is already at maximal strength when only A is present (analogous to a ceiling effect preventing a conclusion concerning the effectiveness of a factor in experimental design).

Even though outcome additivity and outcome maximality effects on cue competition in human causal learning are closely linked according to an inferential reasoning account, the issue of separating their influence has importance from the viewpoint of other accounts of human causal learning. The effect of outcome additivity per se is beyond the scope of contemporary associative models. Theories such as the Rescorla-Wagner model (Rescorla & Wagner, 1972) and related models lack the means to make differential predictions regarding blocking after additive and subadditive training with a set of different cues. The reason for this is that such models do not really allow for learning about one set of cues to transfer to a dissimilar set of cues. However, outcome maximality is clearly within their scope, at least in principle. Indeed, one may readily assume that a submaximal outcome is less salient (has less associability) and/or supports a lower asymptotic value of associative strength than an outcome of maximal intensity. We will return to this issue in the Discussion section of Experiment 1.

A similar argument holds for probabilistic models of causal learning, in particular for the power PC theory (Cheng, 1997). In probabilistic theories, it is assumed that human causal judgments reflect 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: $P(\text{OI}X.A) - P(\text{OI}\neg X.A)$. In a blocking procedure, these probabilities are equal, so the probabilistic contrast amounts to 0. According to the power PC theory, however, in order to yield an estimate of 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 $[1 - P(\text{OI}\neg X.A)]$. As such, the power PC theory predicts that blocking will be sensitive to ceiling information, because a causal estimate cannot be derived if the probability of the outcome is maximal when only A is present. Indeed, if $P(\text{OI}\neg X.A)$ equals 1, then the probabilistic contrast has to be divided by 0, which results in an indeterminate value. In principle, the power PC theory, like other probabilistic models, concerns only outcome probability, not outcome rate or magnitude. Nonetheless, an effect of outcome maximality is at least conceptually consistent with Cheng's arguments. However, a possible effect of outcome additivity does not seem to be within immediate reach of probabilistic models. Indeed, there is no obvious mechanism through which pretraining with one set of cues could affect the way in which causal strength is estimated for another set of cues.

In Experiment 1, we assessed whether the relative strength of the outcome presented during forward-blocking training (i.e., whether the outcome that is presented during forward-blocking training is the strongest outcome that is experienced during the whole of training) can affect the degree of blocking that occurs. In Experiment 2, we then evaluated whether, if we controlled for

outcome maximality, outcome additivity training in itself would still affect forward blocking. In Experiment 3, again controlling for outcome maximality, we examined the effect of outcome additivity training on backward blocking and on release from overshadowing, the latter being a cue competition effect that, according to an inferential reasoning logic, should not be similarly affected by assumptions about outcome additivity. In Experiment 4, we examined whether outcome additivity training given after the blocking training would also affect forward blocking. Such a result would argue against the possibility that differences in elemental versus configural processing would be responsible for the anticipated effect of additivity pretraining on blocking.

Experiment 1

As we previously stated, our aim in Experiment 1 was to investigate whether outcome submaximality (i.e., the mere fact that the outcome presented on A+ and AX+ trials is not the strongest outcome that participants have encountered during the whole of the experiment) might suffice to produce the enhanced forward-blocking effect reported by Lovibond et al. (2003, Experiment 1). We used a food allergy task similar to the one used by Lovibond et al. On arrival, participants were informed that they would assume the role of an allergist and in that capacity review information about a fictitious patient. The first set of records just indicated the extent to which the patient had developed an allergic reaction at various moments in time (outcome preexposure phase). Later records each gave information about one or two food items the patient had eaten, followed again by information about the extent of allergic reaction (blocking phases; see the design in Table 1). At the end, participants were asked to indicate how likely the patient was to develop an allergic reaction on eating each of the individual food items. It is important to note that for half of the participants (maximal condition), all allergic reactions that occurred during the blocking phases were of the same strength as the strongest allergic reaction presented during the preexposure phase (++). For the other half of the participants (submaximal condition), the allergic reactions that occurred during the blocking phases were of the same strength as the moderate allergic reaction presented during the preexposure phase (+). We predicted that forward blocking would be larger in the latter case than it would be in the former case.

Method

Participants. The participants were 93 undergraduate students at the State University of New York at Binghamton (29 men and 64 women, ranging from 17 to 33 years of age) who participated for course credit.

They were randomly assigned to one of the two experimental conditions (45 were assigned to group maximal, and 48 were assigned to group submaximal).

Apparatus. All testing was done on IBM-compatible Pentium PCs. Participants were seated in individual cubicles. A custom-made SuperLab Pro 2.0 (1999) program controlled stimulus presentation and response registration. Participants entered their ratings by means of the numeric keypad on the keyboard.

Stimuli. As food cues, pictures of *cheese, nuts, mushrooms, fish, and strawberries* were used, accompanied below by their written labels. The outcomes consisted of the message *no allergic reaction occurred* printed in green accompanied by a green bar barely rising above 0 (we refer to this outcome henceforth as *no allergic reaction*), the message *an allergic reaction (strength 5) occurred* printed in red accompanied by a red bar rising to 5 (we refer to this outcome henceforth as *moderate allergic reaction*), and the message *an allergic reaction (strength 10) occurred* printed in red accompanied by a red bar rising to 10 (we refer to this outcome henceforth as *strong allergic reaction*).

Design and procedure. After participants indicated their age and gender, they were presented with the following instructions:

Imagine that you are an allergist who tries to discover the cause of allergic reactions in people. You have recently been presented with a new patient. In order to evaluate his condition, you first ask your patient to record his condition at various moments in time. Then you arrange for him to eat various foods and again record his condition. On some sessions, you arrange for him to have two foods at the same time, while on other sessions he has only one food. On each trial, the computer will display a record of your patient's condition. Initial records will only contain information regarding the degree of allergic reactions at different moments in time. A verbal message will be displayed, stating whether an allergic reaction occurred, along with a bar graph representing the strength of allergic reaction. Later records will also include information about the foods your patient has eaten. After you have reviewed all the information about your patient, you will have to judge for each food item separately how likely it is to cause an allergic reaction in your patient. Press the space bar to continue.

On pressing the space bar, the instructions were cleared from the screen, and the experiment began. The experiment consisted of 64 training trials (see design in Table 1) divided into three phases (24 Phase 1 outcome preexposure trials followed by 16 Phase 2 elemental trials followed by 24 Phase 3 elemental and compound trials). The experimental groups differed only in the outcome presented on the outcome-present trials of Phases 2 and 3, which was either the moderate allergic reaction (submaximal condition) or the strong allergic reaction (maximal condition). Each trial started with the message *Press the space bar to load the next record*. On pressing the space bar, this message was replaced by the message *Loading next record . . . please wait*, which stayed on screen for 3 s. Then either the message *[no information available about foods]* (outcome preexposure phase) or one or two food cues (elemental and compound phases) appeared

Table 1
Design Summary of Experiment 1

Group	Phase 1: preexposure	Phase 2: elemental training	Phase 3: compound training
Maximal	-/+ / ++	A++ / Z-	AX++ / KL++ / Z-
Submaximal	-/+ / ++	A+ / Z-	AX+ / KL+ / Z-

Note. A, X, K, and L = cheese, nuts, fish, and mushrooms (partially counterbalanced), Z = strawberries. - indicates that there was no allergic reaction; + indicates that there was a moderate allergic reaction; ++ indicates that there was a strong allergic reaction.

for 5 s, accompanied by one of the outcome messages after 2 s so that the cue(s) and the outcome overlapped for 3 s. Then the screen was cleared, and the next trial started. Total trial duration thus was at least 8 s, with the start of each trial self-paced by the participant. Within phases, trial order was determined randomly for each individual participant. Which food item functioned as which cue was semicounterbalanced across participants and groups. Strawberries always functioned as Cue Z. For one counterbalancing condition within each group, cheese was the blocking cue, A, and nuts the blocked cue, X, with fish and mushrooms as overshadowing control cues, K and L. For a second counterbalancing condition, the blocking and blocked cues were interchanged. For two other counterbalancing conditions within each group, fish and mushrooms functioned as either blocked or blocking cues, with cheese and nuts as overshadowing control cues. After observing all 64 training trials, participants received the following instructions:

Now you have to indicate how likely each of the food items is to cause an allergic reaction in your patient. You will be presented with each of the items, and you have to supply your judgment on a rating scale ranging from 1 (the food item is very unlikely to cause an allergic reaction in the patient) to 9 (the food item is very likely to cause an allergic reaction in the patient), using the keyboard. Press the space bar to continue.

After pressing the space bar, a screen appeared depicting one of the food items along with an anchored rating scale ranging from 1 (*An allergic reaction is not likely*) to 9 (*An allergic reaction is likely*) and the message *Please indicate how likely it is that eating this food item will cause an allergic reaction in the patient*. Participants responded by pressing a numeric key on the keyboard, after which the food item was replaced by another food item until each of the food items was rated. Food items were presented in random order. After all food items were rated, participants were debriefed and dismissed.

Results and Discussion

Mean causal ratings by condition for each of the five experimental cues are depicted in Figure 1. We subjected participants' ratings to a 2×5 analysis of variance (ANOVA), with group (maximal or submaximal) as between-subjects factor and cue (A, X, K, L, or Z) as within-subjects factor. We applied Greenhouse-Geisser corrections where appropriate in this and all following analyses (Greenhouse & Geisser, 1959), and we set α at .05 for all analyses. The ANOVA revealed main effects of group, $F(1, 91) = 17.95$, and cue, $F(3.24, 294.70) = 266.71$, and, more important, a highly reliable Group $\times$ Cue interaction, $F(3.24, 294.70) = 9.43$. Planned comparisons revealed a highly reliable blocking effect (i.e., lower ratings for X than the mean of the ratings for K and L) in group submaximal, $F(1, 91) = 79.78$, as well as a blocking effect in group maximal, $F(1, 91) = 6.76$. However, the blocking effect was larger in the former than it was in the latter group, $F(1, 91) = 18.88$.

The statistical analyses indicate that outcome maximality can have a profound impact on the amount of blocking that is observed. Forward blocking was much stronger when the outcome occurred with submaximal strength during blocking training than when it occurred with maximal strength. This result conceptually replicates and extends the findings of De Houwer et al. (2002). More important, it parallels the results reported by Lovibond et al. (2003), who also obtained greater forward blocking with additive pretraining than with subadditive pretraining but still observed a forward-blocking effect in their

Figure 1. Experiment 1: mean causal ratings for all cues by maximality condition. Error bars represent standard errors.

subadditive condition as well. This parallel suggests that their results may at least in part be due to group differences in perceived outcome maximality rather than to differences in assumed outcome additivity.

It is important to note that unlike additivity pretraining effects per se, effects of outcome maximality are, in principle, within the scope of associative models of human causal learning. In the Rescorla and Wagner (1972) model, it is assumed that the increase of associative strength produced by the pairing of a cue and an outcome is limited by an asymptotic value (represented by λ) specific to that outcome. It is furthermore assumed that this asymptotic value varies directly with outcome intensity so that a more intense outcome supports more associative strength than a less intense outcome. Also, it is assumed that the increase of associative strength on a given trial is a direct function of the associability of the outcome that is presented on that trial (represented by β). Thus, a less intense outcome should result in lower values of both λ and β than a more intense outcome does. On the basis of these assumptions, one can derive predictions from the Rescorla-Wagner model concerning the effect that outcome intensity should have on forward blocking. We have simulated such predictions in Figure 2. Remarkably, the Rescorla-Wagner model (and other associative models like it) predicts a pattern of results exactly opposite the one we observed, that is, if anything, more blocking should be obtained with a more intense outcome than with a less intense outcome, irrespective of whether intensity is assumed to affect asymptote (λ), salience (β), or both.

Figure 2. Predicted associative strength for a blocked cue (X) and for an overshadowing cue (K) as a function of outcome intensity according to the Rescorla-Wagner model (Rescorla & Wagner, 1972), a generic associative model. For the simulation, associative strength for X after 4 A+ followed by 4 AX+ trials was compared with associative strength for K after 4 KL+ trials, assuming $\lambda = 1.0$ and $\beta = .40$ for the maximal outcome and either $\lambda = 1.0$ and $\beta = .20$, $\lambda = .50$ and $\beta = .40$, or $\lambda = .50$ and $\beta = .20$ for the submaximal outcome and assuming equal salience for A, X, K, and L (all set at $\alpha = .90$). Blocking (i.e., the difference between X and K) is markedly smaller in every instantiation of the submaximal condition than in the maximal condition.

Experiment 2

Experiment 1 suggested that a difference in outcome maximality relative to outcomes experienced during pretraining is in itself sufficient to produce a difference in forward blocking. In Experiment 2, we wanted to assess whether additive versus subadditive pretraining would still affect blocking if outcome maximality is controlled. We again used the allergy task of Experiment 1. How-

ever, instead of outcome preexposure, participants received explicit additive or subadditive pretraining in Phase 1 (see the design in Table 2). The additive pretraining involved G+, H+, GH++, I+, and Z- trials (G, H, I, and Z indicating different food cues and + and ++, as before, indicating a moderate or a strong allergic reaction, respectively), and the subadditive pretraining involved G+, H+, GH+, I++, and Z- trials. In both groups, the outcome used during Phase 2 and Phase 3 was of moderate intensity. This way, both groups were equated for the strength of the outcomes presented during pretraining (absent, moderate, and strong allergic reactions were presented equally often in both groups) and for the maximality of the outcome presented during blocking training (which was always submaximal in both groups). If outcome additivity training has an effect on forward blocking independent of differences in outcome maximality, blocking should be obtained in the additive group but not (or markedly less so) in the subadditive group.

Method

Participants. The participants were 88 undergraduate students (36 men and 52 women ranging from 17 to 22 years of age) at the State University of New York at Binghamton who participated for course credit. None of them had participated in Experiment 1. They were randomly assigned to one of the two experimental conditions (44 in each group).

Apparatus and stimuli. We used the same apparatus we used in Experiment 1. We used the same food cues we used in Experiment 1 for the blocking phases. In addition, we used pictures of *eggs*, *bacon*, and *toast*, along with their written labels, for Cues G, H, and I of the pretraining phase. The outcomes were the same as they were for Experiment 1.

Design and procedure. Design and procedure were the same as they were in Experiment 1, apart from the following points. Because we presented no outcome-only trials in Experiment 2, we altered the instructions slightly. They now read:

Imagine that you are an allergist who tries to discover the cause of allergic reactions in people. You have recently been presented with a new patient who has a food allergy. In order to evaluate his condition, you arrange for him to eat various foods and record his condition afterward. On some sessions, you arrange for him to have two foods at the same time, while on other sessions he has only one food. On each trial, the computer will display a record of your patient's condition. The records will show what food your patient has eaten, and the degree of allergic reaction after eating this food. A verbal message will be displayed, stating whether an allergic reaction occurred, along with a bar graph representing the strength of allergic reaction. After you have reviewed all the information about your patient, you will have to judge for each food item separately how likely it is to cause an allergic reaction in your patient. Press the space bar to continue.

Instead of an outcome preexposure phase, the experiment started with a pretraining phase in which on each trial, Cue G, Cue H, both G and H, Cue

Table 2
Design Summary of Experiment 2

Group	Phase 1: pretraining	Phase 2: elemental training	Phase 3: compound training
Additive	G+/H+/GH++/I+/Z-	A+/Z-	AX+/KL+/Z-
Subadditive	G+/H+/GH+/I++/Z-	A+/Z-	AX+/KL+/Z-

Note. G, H, and I = bacon, eggs, and toast (counterbalanced); A, X, K, and L = cheese, nuts, fish, and mushrooms (partially counterbalanced); Z = strawberries. - indicates that there was no allergic reaction; + indicates that there was a moderate allergic reaction; ++ indicates that there was a strong allergic reaction.

I, or Cue Z was paired with the appropriate outcome (see the design in Table 2). Each trial type was presented 8 times, resulting in 40 pretraining trials. With the 40 trials of the actual blocking procedure, this resulted in a total of 80 experimental trials. We counterbalanced which food item (eggs, bacon, or toast) served as which pretraining cue across participants orthogonal to the semicounterbalancing of food-cue assignment during elemental and compound phases (resulting in 12 counterbalancing conditions within each experimental group). Within each phase, we randomized the order of trials. After the testing of Cues A, X, K, L, and Z, we also collected ratings for G, H, and I, in random order. The remaining elements of the procedure were as they were in Experiment 1.

Results and Discussion

Mean causal ratings by condition for the five cues of interest are presented in Figure 3. We subjected participants' ratings to a 2×5 ANOVA, with group (additive or subadditive) as a between-subjects factor and cue (A, X, K, L, or Z) as a within-subjects factor. The ANOVA revealed a main effect of cue, $F(4, 344) = 228.70$, and, more important, a highly reliable Group $\times$ Cue interaction, $F(4, 344) = 38.25$. The main effect of group failed to reach significance, $F(1, 86) = 2.66$, $p = .11$. As revealed by planned comparisons, we observed a blocking effect (i.e., ratings for X were lower than the mean of the ratings for K and L) in group additive, $F(1, 86) = 211.12$, and in group subadditive, $F(1, 86) = 4.08$. Most important, the blocking effect was much larger in the former than it was in the latter group, $F(1, 86) = 78.23$.

The statistical analyses clearly show that when outcome maximality is controlled, outcome additivity pretraining still has a profound influence on forward blocking. This result clarifies the

observations by Lovibond et al. (2003) and supports their conclusions (see also the General Discussion section).

Experiment 3

As mentioned in the introduction to this article, human causal learning is not only sensitive to forward-cue competition effects such as forward blocking but also to retrospective revaluation effects such as backward blocking (obtained when A+ training after AX+ training reduces causal ratings for X) and release from overshadowing (obtained when A- training after AX+ training enhances causal ratings for X). Even though retrospective revaluation is at variance with early associative models such as the Rescorla and Wagner (1972) model or the SOP model (Wagner, 1981), it is readily explained by revised associative models in which it is assumed that, on a given trial, a cue that is absent but expected undergoes a change in associative strength that is opposite the change to which it would be subject if it were actually present (Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994). It is also fully compatible with models in which it is assumed that competition occurs during testing rather than during acquisition due to a comparison process in which the associative strength of a cue is assessed relative to the strength of comparator cues (Miller & Matzel, 1988; Miller & Schachtman, 1985). It is interesting to note that in all of these models, it is assumed that the same associative principles underlie backward blocking and release from overshadowing. Accordingly, any manipulation that affects one should also affect the other. To be more specific, if additivity pretraining were to affect backward blocking such that competition was reduced by subadditive pretraining compared with additive pretraining, then such training should affect release from overshadowing in a similar way.

An inferential reasoning account of cue competition makes a different prediction (Lovibond et al., 2003). According to an inferential account, subadditive pretraining reduces blocking relative to additive pretraining because it contradicts the presumed default assumption that a combination of causes should result in a stronger outcome than a single cause in itself, this assumption being crucial for blocking to occur (see the introduction to this article). However, whether cues are assumed to be additive or subadditive is irrelevant in the case of release from overshadowing. If the combination of A and X results in the outcome (AX+), and A in itself is subsequently shown not to result in the outcome (A-), then the causal efficacy of X can be logically inferred from the AX+ trials irrespective of whether causes are assumed to summate linearly or sublinearly (because there is only one potential cause left in the AX compound anyway). As a consequence, according to an inferential reasoning account, additivity pretraining should affect backward blocking but not release from overshadowing, whereas according to associative models that allow for retrospective revaluation, additivity pretraining should, if anything, have similar effects on both. In line with an inferential account, Lovibond et al. (2003) indeed found an effect of pretraining on backward blocking in one experiment but no effect of pretraining on release from overshadowing in a second experiment. However, their evidence was indirect in that the observation that additivity pretraining differently affects backward blocking and release from overshadowing relied on a between-experiments comparison, with a null finding in one experiment but not in the

Figure 3. Experiment 2: mean causal ratings for cues A, X, K, L, and Z by additivity pretraining condition. Error bars represent standard errors.

other. Moreover, it is again unclear to what extent their results reflected differences in outcome maximality rather than outcome additivity. In Experiment 3, we replicated our Experiment 2 with the elemental (A+/Z-) phase of blocking training coming after instead of before the compound (AX+/KL+/Z-) phase. Moreover, we added two groups (one additively pretrained and one subadditively pretrained) in which during the elemental phase, Cue A was not followed by an allergic reaction (A-/Z-; see Table 3) to test whether additivity pretraining would affect release from overshadowing.

Method

Participants, stimuli, and apparatus. The participants were 96 undergraduate students (27 men and 69 women, ranging from 17 to 57 years of age) at the State University of New York at Binghamton who participated for course credit. None of them had participated in Experiments 1 and 2. Participants were randomly assigned to the four treatment groups (24 in each group). We used the same stimuli and apparatus we used for Experiment 2.

Design and procedure. The design and procedure were the same as in Experiment 2, apart from the following changes. The elemental phase trials were presented after instead of before the compound phase trials. Moreover, for the participants in the additive and subadditive release from overshadowing groups, unlike in the backward-blocking groups, presentations of the A cue were accompanied by the allergy absent outcome during the elemental phase (see Table 3). Also, the first sentence of the instructions for the rating phase was slightly altered in line with the instructions used by Lovibond et al. (2003). It now read (additions in bold here only): *Now you have to indicate how likely each of the food items is to cause an allergic reaction of any strength in your patient.* Likewise, rating screens now read: *Please indicate how likely it is that eating this food item will cause an allergic reaction of any strength in the patient.*

Results and Discussion

Mean causal ratings by condition for the five cues of interest are presented in Figure 4. We subjected the ratings to a $2 \times 2 \times 5$ ANOVA, with pretraining (additive or subadditive) and treatment (backward blocking or release from overshadowing) as between-subjects factors and cue (A, X, K, L, or Z) as a within-subjects factor. The ANOVA revealed main effects of treatment, $F(1, 92) = 6.06$, and cue, $F(3.22, 295.85) = 101.90$, as well as a Treatment $\times$ Cue interaction, $F(3.22, 295.85) = 83.95$. Most important, however, this Treatment $\times$ Cue interaction was qualified by a higher order Pretraining $\times$ Treatment $\times$ Cue interaction, $F(3.22, 295.85) = 3.37$. Planned comparisons contrasting ratings for X with the mean of ratings for K and L revealed backward

blocking, $F(1, 92) = 13.12$, but only with additive pretraining, $F(1, 92) = 20.33$, and not with subadditive pretraining, $F(1, 92) < 1$. Accordingly, backward blocking differed between additive and subadditive pretraining, $F(1, 92) = 7.59$. Release from overshadowing (again probed by contrasts comparing ratings for X with the mean of ratings for K and L) was also significant, $F(1, 92) = 30.36$. However, in contrast with backward blocking, it did not differ between additive and subadditive conditions, $F(1, 92) < 1$: Release from overshadowing was reliable both with additive pretraining, $F(1, 92) = 20.33$, and with subadditive pretraining, $F(1, 92) = 10.78$. It is important to note that the effect of additivity training on the difference between X on the one hand and the mean of K and L on the other hand was different for backward blocking and release from overshadowing, $F(1, 92) = 6.56$.

In Experiment 3, subadditive pretraining reduced backward blocking relative to additive pretraining (effectively abolishing backward blocking altogether), while at the same time leaving release from overshadowing unaffected. It is important to note that the effect of additivity pretraining on backward blocking was different from the (not statistically detectable) effect of such training on release from overshadowing. The fact that backward blocking is modulated by type of pretraining mirrors the significant modulation of forward blocking observed in Experiment 2 and lends support to the findings reported by Lovibond et al. (2003, Experiment 1). A remarkable difference between Experiments 2 and 3 nevertheless lies in the fact that in Experiment 3, backward blocking was not observed at all after subadditive pretraining, whereas in Experiment 2, in the subadditive condition a small but reliable forward-blocking effect was still obtained, although it was significantly smaller than it was in the additive condition. This lends support to the view that the residual forward blocking that was obtained after subadditive pretraining in Experiment 2 might have been due to low-level attentional processes, because such attentional processes are assumed to not be involved in backward blocking. Indeed, most attentional theories (e.g., Mackintosh, 1975) assume an asymmetry between forward blocking (in which selective attention can be involved) and backward blocking (which cannot rely on selective attention with respect to X during acquisition, but see Kruschke & Blair, 2000, who argue that differences in attention at test contribute to backward blocking).

More important, however, is the observation of a robust modulation of backward blocking by additivity training and the lack of such modulation of release from overshadowing. Whereas it is difficult to see how associative theories that can account for retrospective revaluation could account for effects of additivity

Table 3
Design Summary of Experiment 3

Group	Phase 1: pretraining	Phase 2: compound training	Phase 3: elemental training
Backward blocking additive	G+/H+/GH+/I+/Z-	AX+/KL+/Z-	A+/Z-
Backward blocking subadditive	G+/H+/GH+/I+/Z-	AX+/KL+/Z-	A+/Z-
Release from overshadowing additive	G+/H+/GH+/I+/Z-	AX+/KL+/Z-	A-/Z-
Release from overshadowing subadditive	G+/H+/GH+/I+/Z-	AX+/KL+/Z-	A-/Z-

Note. G, H, and I = bacon, eggs, and toast (counterbalanced); A, X, K, and L = cheese, nuts, fish, and mushrooms (partially counterbalanced); Z = strawberries. - indicates that there was no allergic reaction; + indicates that there was a moderate allergic reaction; ++ indicates that there was a strong allergic reaction.

Figure 4. Experiment 3: mean causal ratings for cues A, X, K, L, and Z by additivity pretraining condition and revaluation condition. Error bars represent standard errors.

pretraining on backward blocking to begin with, one should at least assume that, if it has such an effect, it would have similar effects on other forms of retrospective revaluation that are assumed to rely on the same associative principles.

As we stated before, contemporary associative theories at first sight seem to have little means to account for effects of additive versus subadditive pretraining with one set of cues on subsequent competition between a set of different cues. However, Livesey and Boakes (2004) have recently proposed an associative account in which differences in degree of elemental or configural processing are deemed responsible for the effects of additivity pretraining on forward and backward blocking. Our aim in Experiment 4 was to test the plausibility of their account.

Experiment 4

According to Lovibond et al. (2003; see also Mitchell & Lovibond, 2002), effects of additivity training on forward and backward blocking are indicative of the role of inferential reasoning processes in cue competition. However, Livesey and Boakes (2004) have recently argued that additivity training might, albeit indirectly, act to influence competition by associative means. Additivity implies that the outcome that occurs following a compound can be reduced to the sum of the outcomes of the elements that make up that compound. As a result, they argued, people would be encouraged to process the cues that are presented to them elementally. In contrast, if the outcome of a compound cannot be

reduced to the sum of the outcomes of its constituent elements, people would be inclined to process the cues that are presented to them configurally. That is, they would tend to treat the compound of two cues as a separate, novel cue that has no bearing on the causal status of its constituent elements.

Evidence from predictive learning studies suggests that configural processing of compound cues has a detrimental effect on blocking (Williams, Sagness, & McPhee, 1994). The idea is that if cues are processed elementally, the target cue fails to acquire associative strength during the compound phase of blocking training (in the case of forward blocking) or loses associative strength during the elemental phase (in the case of backward blocking) because of the associative strength acquired by the blocking cue, A. However, if cues are processed configurally, causal ratings for the blocked cue, X, are based on the associative strength acquired by the AX compound, and causal ratings for K and L are based on the associative strength acquired by the KL compound. Because the associative strength acquired by A is irrelevant for the associative strength acquired by the AX compound (which is processed as a novel cue, just like the KL compound), causal ratings for X will be similar to the causal ratings for K and L; hence, no blocking will be obtained. As previously stated, according to Livesey and Boakes (2004), additivity training exerts its effect on blocking by influencing the degree of elemental versus configural processing. They provided some support for this claim by showing that independent manipulations that affect the degree of configuring (e.g., variations in spatial separation between cues) also affect the degree to which additivity training modulates blocking. It is nevertheless unclear whether an account in terms of elemental versus configural processing could provide a full explanation for the effects of additivity pretraining on subsequent blocking, because such an account would seem to anticipate similar effects of additivity pretraining on other cue competition effects such as release from overshadowing. Experiment 3 above, as well as Experiment 2 of Lovibond et al. (2003), suggests, however, that additivity pretraining has markedly different effects on blocking and release from overshadowing (also see De Houwer et al., 2002, Experiment 3).

If configural versus elemental processing were nevertheless somehow crucially involved in the effect that additivity training has on blocking (rather than outcome additivity and configural vs. elemental processing each having an independent influence on cue competition), then it should make a rather dramatic difference whether information about outcome additivity is presented before or after the actual blocking training. That is, information about outcome additivity should have an impact on blocking only if it is presented beforehand, given that a configural or elemental mode of processing has to be induced before the actual blocking training is carried out for it to be able to influence how associative strength is acquired. In contrast, according to an inferential account of blocking (and also according to performance-focused associative models of causal learning, e.g., Miller & Matzel, 1988), order of presentation of the various types of information should not matter much. Irrespective of whether information about outcome additivity is presented before or after the actual blocking training, as long as all information is properly retained, it allows the participant to evaluate whether the target cue, X, should (have) augment(ed) the outcome produced by the blocking cue, A, if it were an effective cause of the outcome. Therefore, in Experiment 4, we replicated Experiment 2, but now reversing the order of the actual blocking

training and the additivity training, effectively turning the additivity manipulation in a posttraining instead of a pretraining procedure (see Table 4).

Method

Participants, stimuli, and apparatus. The participants were 72 undergraduate students (10 men and 62 women, ranging from 18 to 40 years of age) at the University of Leuven, Leuven, Belgium, who participated for course credit. None of them had participated in related experiments before. Participants were randomly assigned to the two treatment groups (36 in each group). We used the same stimuli and apparatus we used for Experiments 2 and 3.

Design and procedure. Design and procedure were as they were in Experiment 2, apart from the fact that the actual blocking training (elemental phase and compound phase) was presented before the additive or subadditive training phase. We maintained the slight change in instructions for the rating phase introduced in Experiment 3.

Results and Discussion

Mean causal ratings by condition for the five cues of interest are presented in Figure 5. We subjected the ratings to a 2×5 ANOVA, with posttraining (additive or subadditive) as a between-subjects factor and cue (A, X, K, L, or Z) as a within-subjects factor. The ANOVA revealed a main effect of cue, $F(3.47, 242.85) = 156.68$, as well as a tendency toward a Posttraining $\times$ Cue interaction, $F(3.47, 242.85) = 2.13$, $p = .09$. Specific contrasts comparing ratings for X with the mean of ratings for K and L revealed forward blocking overall, $F(1, 70) = 61.44$. It is important to note that blocking was highly reliable with additive posttraining, $F(1, 70) = 51.06$, as well as with subadditive posttraining, $F(1, 70) = 15.52$, but was more so with the former than with the latter type of posttraining, $F(1, 70) = 5.14$.

In Experiment 4, subadditive posttraining reduced forward blocking relative to additive posttraining, much like pretraining did in Experiment 2. The fact that a reliable blocking effect was still obtained in the subadditive group also parallels the results of Experiment 2. In the present case, it may reflect an involvement of automatic selective attentional processes in forward blocking, some participants failing to reevaluate conclusions that they may have drawn along the way during the actual blocking training (probably assuming additivity by default) in light of the information presented subsequently or a combination of both. The idea that some participants might fail to update their judgments on the basis of the posttraining information would also explain why the effect of additivity posttraining obtained here, even though clearly significant, seems smaller than the effect of additivity pretraining observed in Experiment 2.

The fact that posttraining of outcome additivity reliably influenced forward blocking in Experiment 4 raises serious doubts about the suggestion that differences in elemental versus configural processing induced by additive or subadditive training, respectively, were mainly responsible for the effects of additivity pretraining observed by Lovibond et al. (2003) and in the present Experiments 2 and 3. Instead, such training seems to modulate how information about cue-outcome pairings acquired during training is combined at the time of retrieval to give rise to differences in causal judgments about the target cue, X.

General Discussion

In four experiments, we investigated the effects of perceived outcome maximality and assumed outcome additivity on cue competition in human causal learning. Experiment 1 suggested that differences in outcome maximality (i.e., whether the outcome presented during blocking training was the most intense outcome ever presented during the whole of the experiment) have a profound influence on forward blocking. Experiment 2 demonstrated that, if outcome submaximality is kept constant across conditions, differences in outcome additivity pretraining also modulate forward blocking. Experiment 3 confirmed that outcome additivity pretraining has a similar effect on backward blocking but not on release from overshadowing. Experiment 4 showed that posttraining of additivity has a similar effect on forward blocking as had pretraining in Experiment 2.

Effects of perceived outcome maximality and assumptions about outcome additivity are closely related, according to an inferential reasoning account of cue competition in human causal learning. Explicitly demonstrating subadditivity of outcomes serves to disconfirm the assumption of participants that, when combined, effective causes should produce their designated outcome with greater intensity or probability than when presented in isolation, an assumption that underlies blocking according to an inferential reasoning account. In contrast, explicitly demonstrating additivity of outcomes should, if anything, serve to confirm this assumption. Explicitly demonstrating outcome submaximality, then, should increase the certainty with which people can actually verify whether the above-mentioned assumption is met. If the outcome presented on AX+ trials is of the same submaximal strength as the outcome presented on A+ trials, then it should be clear that the assumption that A and X are both effective causes is false. Hence, blocking should result. If the outcome presented on AX+ and A+ trials is always of maximal extent, then the participant should realize that a potential additive effect of X on the outcome produced by A cannot be observed because of a ceiling effect.

Table 4
Design Summary of Experiment 4

Group	Phase 1: elemental training	Phase 2: compound training	Phase 3: posttraining
Additive	A+/Z–	AX+/KL+/Z–	G+/H+/GH++/I+/Z–
Subadditive	A+/Z–	AX+/KL+/Z–	G+/H+/GH+/I++/Z–

Note. G, H, and I = bacon, eggs, and toast (counterbalanced); A, X, K, and L = cheese, nuts, fish, and mushrooms (partially counterbalanced); Z = strawberries. – indicates that there was no allergic reaction; + indicates that there was a moderate allergic reaction; ++ indicates that there was a strong allergic reaction.

Figure 5. Experiment 4: mean causal ratings for cues A, X, K, L, and Z by additivity pretraining condition. Error bars represent standard errors.

Their close relatedness in light of an inferential account notwithstanding, it is important to disentangle the potential contribution of perceived outcome maximality and assumed outcome additivity to cue competition because they might have quite different implications for other potential accounts of cue competition. For instance, acquisition-focused associative models such as the Rescorla and Wagner (1972) model might, in principle, accommodate outcome maximality effects on blocking by means of variations in the salience of the outcome and the asymptotic strength that it supports. However, as demonstrated in the *Results and Discussion* section of Experiment 1, contemporary associative learning models actually predict the opposite pattern of results than the one observed here; that is, they predict more rather than less blocking with maximal outcomes compared with submaximal outcomes (see Figure 2). It is not clear at present how acquisition-focused associative models could be adapted to correct this without fundamentally changing their mode of operation.

Effects of differences in additivity pretraining, on the other hand, are not immediately within the scope of elemental models of associative learning such as the Rescorla and Wagner model (1972). The additional assumption that different kinds of pretraining would engage different learning mechanisms (i.e., a more elemental or a more configural mode of association formation) at first seems to go some way toward providing a semiassociative account of additivity pretraining effects. However, this account fails to explain why such pretraining does not have a similar effect on release from overshadowing (which should be similarly affected by configuring as would be backward blocking; Experi-

ment 3). Moreover, it is incompatible with the fact that not only pretraining but also additive versus subadditive posttraining affects blocking (see Experiment 4).

Unlike associative models, Cheng's (1997) extension of probabilistic contrast models is at least conceptually compatible with the observed effect of outcome maximality on blocking. Effects of outcome additivity pretraining, however, seem to be beyond the scope of these models, at least in their present form.

There is a final alternative account for the results obtained in Experiments 2–4 that we have not mentioned yet. Information about cue additivity or subadditivity may differ in the kind of rules they allow the participants to extract. For instance, the observation that G alone, and the compound of G and H are each followed by the same outcome allows the participant to establish the rule that if an elemental cue is followed by a certain outcome, and a compound of that cue with a second cue is followed by the same outcome, then the second cue in itself will be followed by the same outcome as well. If this rule is applied to the information supplied during the actual blocking procedure (A is followed by a certain outcome and the compound of A and X is followed by the same outcome), it results in the conclusion that X will be followed by the same outcome, and blocking is not to be expected (actually, from a strict application of this rule, equal ratings for A and X would be expected). In the case in which the compound of G and H is followed by a stronger outcome than G or H individually, a similar rule cannot be established. Therefore, differences in the rules that participants extract during additivity or subadditivity training and apply to the information supplied during the actual blocking training might contribute to the differences in blocking that are observed (for other evidence of rule governed processing in human contingency learning, see Shanks & Darby, 1998). Although this would still suggest the involvement of rather complex and effortful cognitive processes, it is less obvious whether this should be considered genuine counterfactual inferential reasoning.

How this simpler form of rule learning could account for the results obtained in Experiment 1 is not clear, however. As such, an account in terms of inferential reasoning seems to be more comprehensive both in terms of the present results and in terms of other evidence for the involvement of deliberate, effortful reasoning processes in human causal learning that has been accumulating recently (see De Houwer et al., in press, for an in-depth review). For instance, blocking seems to critically depend on the availability of working memory resources, as manipulated by the difficulty of a secondary task (De Houwer & Beckers, 2003) and is modulated by the specific causal scheme linking cues and outcomes (e.g., De Houwer et al., 2002; Waldmann, 2000; Waldmann & Holyoak, 1992). When given the opportunity to receive additional information about particular cues in a causal learning task, participants prefer information about cues that are most informative from a causal reasoning point of view (Vandorpe & De Houwer, 2004). If people are given verbal information about the presence or absence of an alternative cause after training, they are able to retrospectively adjust their causal judgment accordingly (De Houwer, 2002). Cue competition effects seem to appear only in participants that are afterward able to report appropriate inferential reasoning (Vandorpe, De Houwer, & Beckers, in press). These and other findings (e.g., De Houwer & Beckers, 2002a, 2002c), to-

gether with the present findings, are most easily accommodated by an inferential reasoning account.

Still, small but reliable forward-blocking effects were also obtained under circumstances that should not allow for blocking to occur according to an inferential reasoning account, particularly in the maximal outcome intensity condition of Experiment 1 and after subadditive pretraining in Experiments 2 and 4. The observation of a small forward-blocking effect in Experiment 1 is not very informative, actually, because the ceiling that was imposed on outcome intensity was deliberately experiential and not instructed. The fact that the outcome that is presented during the actual blocking training is not the strongest one that has been experienced overall suffices for participants to be assured of outcome submaximality in the submaximal condition, but the opposite does not necessarily hold for the maximal condition. More remarkable is the residual forward-blocking effect observed after subadditive pre- or posttraining. The residual effect after subadditive posttraining might in part or in whole be due to a failure to revise causal judgments once they are established in some participants. However, a similar argument does not hold for the residual forward-blocking effect after subadditive pretraining. An associative processing failure due to acquired inattention seems the most plausible explanation here (see the *Results and Discussion* section of Experiment 3).

As previously stated, none of the currently available associative models is able to account for the present results. Models in which blocking is viewed as an acquisition failure (forward blocking) or a loss of associative strength (backward blocking) are especially fundamentally incompatible with the flexibility in cue competition that is demonstrated here, particularly in Experiment 4. Performance-focused models such as the comparator hypothesis (Miller & Matzel, 1988; Miller & Schachtman, 1985) can in principle be more easily adapted to fit the present data. The main reason for this is that in such models, it is assumed that, in a blocking procedure, information about the blocked cue is effectively acquired even though not expressed and is thus available to the cognitive system (Miller & Escobar, 2001). The reason that the associative knowledge about the blocked cue is not expressed is that it is downplayed by stronger, competing associations at the time of judgment. In other words, the blocked cue has positive absolute associative strength but a low relative associative strength. Because conditioned performance (i.e., causal ratings) reflects relative and not absolute associative strength, the blocked cue elicits a low causal rating, despite the fact that an association between that cue and the outcome has been acquired. If we supplement the comparator hypothesis with an additional process that modulates the degree of comparison, so that causal ratings would sometimes reflect absolute associative strength instead of relative associative strength, then modulation of blocking becomes a principled possibility (see Pineño, Denniston, Beckers, Matute, & Miller, in press, for a similar proposal). However, it remains unclear what this supplemental process should look like, and, in particular, how it could give rise to a differential sensitivity to pretraining conditions of blocking and release from overshadowing. Given also the evidence cited above, it seems implausible that it could function entirely bottom-up (i.e., purely on the basis of stimulus contingencies). Nevertheless, such a model might provide the starting ground on which to study how inferential and associative processes in human causal learning might interact or how

seemingly inferential processes might someday be explained from a more molecular view.

A similar argument can be developed for probabilistic contrast models. In these models, it is assumed that causal judgment is based on a comparison of conditional probabilities in which the presence or absence of competing cues is kept constant. Blocking arises because the conditional probability of the outcome (i.e., given the presence of the blocking cue, A) is equal in the presence and in the absence of the blocked cue, X, so that the probabilistic contrast equals zero. However, a probabilistic contrast calculated over nonconditional probabilities yields a different outcome: Because the probability of the outcome is greater when the blocked cue, X, is present than when no cue is present, a positive probabilistic contrast is obtained. So, by supplementing probabilistic contrast models with a mechanism that allows a shift between conditional and nonconditional probabilities as the input for the calculation of probabilistic contrasts, modulation of blocking could be obtained. The assumption that subadditive pre- or posttraining for some reason encourages a shift toward nonconditional probabilities would then suffice to explain why such training has a detrimental effect on blocking, as observed in Experiments 2–4. Also note that shifting between conditional and nonconditional probabilities would not affect release from overshadowing, because the conditional and the nonconditional probability of the outcome in the absence of the critical cue, X, are both equal to zero. Therefore, a modified probabilistic contrast model would be able to account for the differential sensitivity to additivity training of blocking and release from overshadowing. However, again it remains to be spelled out what such a shifting mechanism should look like and why it would be differentially triggered by additive and subadditive pre- or posttraining.

One could argue that the prominent involvement of inferential reasoning processes in human causal learning might explain why blocking seems a more fragile phenomenon in human learning than it does in Pavlovian conditioning. Indeed, associative models were initially developed to account for Pavlovian conditioning phenomena. Therefore, the fact that in their current form they fall short in explaining the flexibility and modularity of cue competition in human causal learning suggests that fundamentally different processes are involved in human causal learning (in which inferential reasoning would play a major role) and in Pavlovian conditioning (which would mainly rely on low-level, automatic associative processes). However, recent evidence points out that similar effects of additive versus subadditive pretraining and/or outcome maximality can be found not only in human causal learning but also in human electrodermal Pavlovian conditioning (Mitchell & Lovibond, 2002) and even in Pavlovian conditioning in rats (Beckers, Miller, De Houwer, & Urushihara, 2005). The remarkable similarity in findings between these different preparations and species suggests that, to a considerable extent, similar processes are at work in human causal learning and in Pavlovian conditioning in human and nonhuman animals after all. The fact that these processes seem to some extent akin to inferential reasoning processes suggests that future research might dramatically change our outlook on what Pavlovian conditioning is about.

References

- Beckers, T., Miller, R. R., De Houwer, J., & Urushihara, K. (2005). *Reasoning rats: Forward blocking in Pavlovian animal conditioning is sensitive to constraints of causal inference*. Manuscript submitted for publication.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104, 367–405.
- De Houwer, J. (2002). Forward blocking depends on retrospective inferences about the presence of the blocked cue during the elemental phase. *Memory & Cognition*, 30, 24–33.
- De Houwer, J., & Beckers, T. (2002a). Higher-order retrospective revaluation in human causal learning. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 55B, 137–151.
- De Houwer, J., & Beckers, T. (2002b). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 55B, 289–310.
- De Houwer, J., & Beckers, T. (2002c). Second-order backward blocking and unovershadowing in human causal learning. *Experimental Psychology*, 49, 27–33.
- De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 56B, 345–357.
- De Houwer, J., Beckers, T., & Glautier, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 55A, 965–985.
- De Houwer, J., Beckers, T., & Vandorpe, S. (in press). Evidence for the role of higher-order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior*.
- Dickinson, A. (2001). The 28th Bartlett memorial lecture. Causal learning: An associative analysis. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 54B, 3–25.
- Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 49B, 60–80.
- Dickinson, A., Shanks, D. R., & Evenden, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 36A, 29–50.
- Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, 32, 95–112.
- Kamin, L. J. (1968). "Attention-like" processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimulation* (pp. 9–31). Miami, FL: University of Miami Press.
- Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review*, 7, 636–645.
- Livesey, E. J., & Boakes, R. A. (2004). Outcome additivity, elemental processing and blocking in human causality judgements. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 57B, 361–379.
- Lovibond, P. F. (2003). Causal beliefs and conditioned responses: Retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 97–106.
- Lovibond, P. F., Been, S.-L., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgement is enhanced by additivity of effect magnitude. *Memory & Cognition*, 31, 133–142.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276–298.
- Miller, R. R., & Escobar, M. (2001). Contrasting acquisition-focused and performance-focused models of acquired behavior. *Current Directions in Psychological Science*, 10, 141–145.
- Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 22 (pp. 51–92). San Diego, CA: Academic Press.
- Miller, R. R., & Schachtman, T. R. (1985). Conditioning context as an associative baseline: Implications for response generation and the nature of conditioned inhibition. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 51–88). Hillsdale, NJ: Erlbaum.
- Mitchell, C. J., & Lovibond, P. F. (2002). Backward and forward blocking in human autonomic conditioning requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 55B, 311–329.
- Pineño, O., Denniston, J. C., Beckers, T., Matute, H., & Miller, R. R. (in press). Contrasting predictive and causal values of predictors and of causes. *Learning & Behavior*.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Shanks, D. R. (1985). Forward and backward blocking in human contingency judgment. *Quarterly Journal of Experimental Psychology: Comparative & Physiological Psychology*, 37B, 1–21.
- Shanks, D. R., & Darby, R. J. (1998). Feature- and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 24, 405–415.
- SuperLab Pro (Version 2.0, 1999). [Computer software]. San Pedro, CA: Cedrus Corporation.
- Vandorpe, S., & De Houwer, J. (2004). *People ask for information that will help them make valid inferences in human causal learning*. Manuscript submitted for publication.
- Vandorpe, S., De Houwer, J., & Beckers, T. (in press). Further evidence for the role of inferential reasoning in forward blocking. *Memory & Cognition*.
- Van Hamme, L. J., & Wasserman, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning and Motivation*, 25, 127–151.
- Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behavior. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum.
- Waldmann, M. E. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 53–76.
- Waldmann, M. E., & Holyoak, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, 121, 222–236.
- Williams, D. A., Sagness, K. E., & McPhee, J. E. (1994). Configural and elemental strategies in predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 694–709.

Received July 13, 2004

Revision received September 10, 2004

Accepted September 16, 2004 ■