People want to see information that will help them make valid inferences in human causal learning

STEFAN VANDORPE and JAN DE HOUWER
Universiteit Gent, Ghent, Belgium

According to higher order reasoning accounts of human causal learning (e.g., Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003; Waldmann & Walker, 2005) ceiling effects in forward blocking (i.e., smaller blocking effects when the outcome occurs with a maximal intensity on A+ and AX+ trials) are due to the fact that people are uncertain about the causal status of a blocked cue X in a forward blocking design when the outcome is always fully present on A+ and AX+ trials. This should not be the case for a reduced overshadowing cue Y (B- trials followed by BY+ trials). We tested this hypothesis by asking participants which additional information they preferred to see after seeing all learning trials. Results showed (1) that all participants preferred to see the effect of the blocked cue X over seeing the effect of the reduced overshadowing cue Y (Experiment 1), and (2) that more participants preferred to see the blocked cue X on its own when the outcome on A+ and AX+ trials was fully present than when the outcome on those trials had a submaximal intensity (Experiment 2).

Learning causal relationships is an essential skill in daily life because it allows us to anticipate and control events in the environment. Despite the importance of adequately learning causal relationships, there is still much discussion about whether the underlying processes of causal learning are automatic (i.e., unintentional, uncontrollable, and/or effortless) and stimulus driven (i.e., determined mainly by the experienced events) or whether these processes can be better characterized as controlled (i.e., intentional, controllable, and effortful) and driven not only by experience, but also by language and/or formal reasoning (see, e.g., De Houwer, Beckers, & Vandorpe, 2005; De Houwer, Vandorpe, & Beckers, 2005; Sloman, 1996).

Associative learning models of animal conditioning that have been applied to human causal learning (see, e.g., Denniston, Savastano, & Miller, 2001; Dickinson & Burke, 1996; Miller & Matzel, 1988; Rescorla & Wagner, 1972; Van Hamme & Wasserman, 1994; Wagner, 1981) attribute causal learning to automatic, stimulus-driven processes. According to these models, learning a causal relationship is due to changes in the strength of the association between the representations of cause and effect. The value of a causal judgment is based at least partly on this associative strength. In all existing associative models, controlled reasoning processes play little or no role. Recently, however, various researchers have proposed higher order reasoning accounts (e.g., Lovibond, 2003;

Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003; Waldmann, 2000; Waldmann & Walker, 2005; see also De Houwer & Beckers, 2003, and De Houwer, Beckers, & Vandorpe, 2005, for a review), which all have in common that controlled reasoning processes are assumed to play a crucial role in human causal learning. Lovibond, for instance, argued that "participants may combine the knowledge they have learned about individual cases to generate an inference about a particular cue in the same way as they solve other reasoning tasks" (Lovibond, 2003, p. 98).

Researchers have argued that controlled reasoning processes are particularly important for cue competition. *Cue competition* refers to the fact that a causal judgment about a cue X depends not only on information about the presence of cue X and the presence of the outcome but also on similar information about a cue A with which cue X has co-occurred. *Forward blocking*, for example, refers to the fact that the causal judgment of a cue X will be lower when AX+ trials (cue A and X presented together and followed by the outcome) are preceded by A+ trials (presentation of cue A followed by the outcome) than when only AX+ trials are presented (see, e.g., Dickinson, Shanks, & Evenden, 1984). Another cue competition effect is *reduced overshadowing*. In a reduced overshadowing design, BY+ trials are preceded by B- trials (presentation of cue B not followed by the outcome). Results typically show that causal judgments for cue Y are higher after such training than when only BY+ trials are presented (see, e.g., De Houwer, Beckers, & Glautier, 2002).

In line with the ideas of Lovibond (2003) and Waldmann (2000), De Houwer and Beckers (2003, p. 346) have proposed that forward blocking is due to the fact that participants apply the following deductive rule in order to explain forward blocking:

S.V. is a doctoral researcher at the Special Research Fund of Universiteit Gent. Correspondence should be addressed to S. Vandorpe, Department of Psychology, Universiteit Gent, Henri Dunantlaan 2, B-9000 Ghent, Belgium (e-mail: stefaan.vandorpe@ugent.be).

Note—This article was accepted by the previous editorial team, when Colin M. MacLeod was Editor.

If cue A on its own causes the outcome to occur with a certain intensity and probability, and if cue A and T together cause the outcome to occur with the same intensity and probability, this implies that cue T is not a cause of the outcome.

Likewise, to explain reduced overshadowing, De Houwer, Beckers, and Vandorpe (2005) have proposed that participants might apply the following rule:

If cue A on its own does not cause or predict the outcome but cue A and T together do cause or predict the outcome, this implies that cue T is a cause or predictor of the outcome.

The deductive rule that is assumed to underlie blocking is, however, valid only if one can verify whether X adds anything to the effect of A. For instance, if the outcome on A+ trials always occurs with a maximal intensity, one cannot verify whether cue X increases the probability or intensity of the outcome during the AX+ trials. It could be that X is a cause, but that its effect cannot be seen because A already causes the outcome to a maximal extent. Therefore, the causal status of cue X should be unsure. Importantly, the causal status of a reduced overshadowing cue Y should always be sure. On the basis of B- and BY+ trials, one can infer that Y has to be the cause of the outcome given that B is not a cause of the outcome, irrespective of whether the outcome occurs with a maximal or submaximal intensity on the BY+ trials. De Houwer, Beckers, and Glautier (2002; see Beckers, De Houwer, Pineño, & Miller, 2005, and Wu & Cheng, 1999, for related results) indeed found evidence for ceiling effects in forward blocking. In Experiment 1 of De Houwer et al. (2002), blocking was smaller (and not significantly different from zero) when the outcome on A+ trials and AX+ trials always occurred with a maximal intensity of 10 out of 10 than when the outcome always occurred with a submaximal intensity of 10 out of 20. In a similar vein, Vandorpe and De Houwer (2005) found that reduced overshadowing was much stronger than forward blocking when the outcome was merely present on A+, AX+, and BY+ trials.

The explanation of higher order reasoning accounts for ceiling effects in forward blocking is that participants do not have the necessary information to make valid inferences about the causal status of the blocked cue X (e.g., De Houwer & Beckers, 2003; Lovibond et al., 2003; Waldmann & Walker, 2005). The present experiments were designed to provide a direct test of this explanation. In two experiments, we asked participants after the learning stage which additional information they would prefer to see. We predicted that participants would prefer to see information that allowed them to make an inference about cues for which insufficient information was provided until then. In other words, the likelihood of whether participants would want to see additional information regarding the causal status of a cue was taken as an index of the uncertainty regarding the causal status of the cue. The learning stage of Experiment 1 involved both a blocking design (A+, AX+) and a reduced overshadowing design (B-, BY+). After all trials were presented, we asked participants whether they would prefer to see the effect of the blocked cue X when it was presented on its own or the effect of a reduced

overshadowing cue Y when it was presented on its own. Because ceiling effects applied to the latter but not the former, a valid inference could be made about the causal status of Y but not of X. Participants should therefore prefer to see the effect of the blocked cue X. In Experiment 2, we investigated whether the participants' preference to see the effect of the blocked cue X would be reduced if the outcome was submaximal on A+ and AX+ trials. In both experiments, we used the food allergy paradigm wherein cues are foods and the outcome is an allergic reaction.

EXPERIMENT 1

In Experiment 1, we presented A+ and B- trials in the first learning stage and AX+ and BY+ trials in the second learning stage. After presentation of all trials, we asked participants whether they would prefer to see what would happen when cue X was presented on its own or what would happen when cue Y was presented on its own. Because the causal status of the blocked cue X could not be inferred with certainty whereas the causal status of cue Y could be, the higher order reasoning account predicted that participants would prefer to see cue X on its own.

Method

Participants. Sixteen first-year psychology students at Universiteit Gent participated for course credit. All were native Dutch speakers.

Design, Stimuli, and Materials. During the first learning stage, six A+ and six B- trials were presented. During the second learning stage, AX+ and BY+ trials were presented six times each. The sequence of trials was randomized for each participant and learning stage separately. Names of the following four foods were used for the different cues (translated from Dutch): mushrooms, kiwi, fish, and potatoes. The cues were presented as colored pictures of foods (15 cm high, 10 cm wide) against a white background, with the name of the food under the picture in a black color.

The task was presented on a Pentium IV PC with a 17-in. screen, and implemented using a custom-made Inquisit program. Four different allocations of the foods to the different cues were used, counterbalanced across participants. Each food was allocated once to cue A, B, X, and Y.

Procedure. After reading the learning instructions (see the Appendix), the participants pressed a key to start the learning stage. Each learning trial started with the presentation of one or two foods in the center of the screen. After 2,000 msec, either the message "allergic reaction" in a red color or the message "no allergic reaction" in a green color was added at the bottom of the screen. This message stayed on the screen for 3,000 msec, after which both the food and the outcome message were erased. The intertrial interval (ITI) was 3,000 msec. After presentation of all trials, the following preference question (translated from Dutch) appeared on the screen:

If you had the possibility to see one additional event, would you like to see what would happen if the patient only ate X or would you like to see what would happen if the patient only ate Y?

Whether food X or Y came first in this question was counterbalanced across participants. They answered this question by pressing key A (preference for cue X) or key K (preference for cue Y) on the keyboard. Then the following usefulness question (translated from Dutch) appeared on the screen:

How useful (on a scale of 1 to 10 where 1 means *not very useful* and 10 *very useful*) would it be for you to see what would happen if the patient only ate . . . ?

A rating scale underneath this question was presented. The participants could make their usefulness rating by a click with the mouse on a digit of the rating scale. After this click, the usefulness question for the second cue appeared on the screen together with the rating scale. Whether X or Y was presented first for this question was counterbalanced across participants, in a manner orthogonal to the counterbalancing of the sequence of X and Y in the preference question.

Results and Discussion

All participants preferred to see the blocked cue X on its own. The mean usefulness score for seeing cue X on its own was 8.63 ($SD = 1.67$); the mean usefulness score for seeing cue Y on its own was only 3.31 ($SD = 2.52$). A paired samples t test showed that this difference was significant [$t(15) = 6.40, p < .001$].

As we predicted on the basis of a higher order reasoning account of human causal learning, all participants preferred to see the blocked cue X on its own over the alternative of seeing the reduced overshadowing cue Y on its own. This preference was not merely the result of a forced choice. Participants found it also significantly more useful to see cue X on its own than to see cue Y on its own, indicating that participants were more sure about the causal status of cue Y than that of cue X.

EXPERIMENT 2

In Experiment 2, we used a simple forward blocking design (A+ trials followed by AX+ and KL+ trials). There were two conditions: In the ceiling condition, participants received during the learning trials only information about whether the outcome occurred. In the no-ceiling condition, participants were also told that if the outcome occurred, it occurred with an intensity of 10/20. After presentation of all trials, participants were asked whether they would prefer to see what would happen when cue X and cue K were presented on their own (option A) or what would happen when cue K and cue L were presented on their own (option B). Because the causal status of the blocked cue X could be inferred in the no-ceiling condition, the higher order reasoning account predicted that people should prefer option B rather than option A, whereas in the ceiling condition, the causal status of cue X was uncertain, and so people should prefer option A over option B.

The reason why both options involved two cues (cue X and cue K vs. cue K and cue L) was that there were three cues (X, K, and L) in the ceiling condition whose causal status was unsure. For instance, if we were to ask only for the preference for X over K, the rational preference in the ceiling condition should be cue K: If one had additional information about cue K, one might have the possibility to make inferences about cue L also.¹ That is, if one had the choice between seeing the effect of cue X or the effect of cue K, one could possibly remove the uncertainty about the causal status of two cues by preferring to see cue K on its own, whereas one would remove the uncertainty about the causal status of only one cue when preferring to see cue X on its own. Therefore, it would always be more rational to choose for cue K, irrespective of whether the

causal status of cue X was uncertain (ceiling condition) or not (no-ceiling condition). If, however, one had to choose between cues X and K (option A) versus cues K and L (option B), the higher order reasoning account would predict that relatively more participants would choose option A in the ceiling condition and relatively more participants would choose option B in the no-ceiling condition. This is because in the ceiling condition, as just argued, there might be a possibility to remove the uncertainty about the three uncertain cues (K, L, and X) by choosing option A. In the no-ceiling condition, however, the causal status of cue X could be inferred with certainty. This would remove the need to see X on its own while the causal status of K and L remained unsure. Therefore, relatively more participants in the no-ceiling condition should prefer option B.

Method

Participants. Thirty-two first-year psychology students at Universiteit Gent participated for course credit. They were randomly assigned to the ceiling and no-ceiling conditions. All were native Dutch speakers.

Design, Stimuli, Materials, and Procedure. Experiment 2 was identical to Experiment 1 except as follows. First, the learning stage consisted of A+ and Z- trials followed by AX+ and KL+ trials. The same four foods as in Experiment 1 were used and were again cycled among A, X, K, and L. Eggs were always allocated to cue Z. Cue Z was held constant because it was just a filler cue. Second, in the instructions for the no-ceiling condition, a paragraph was added (see the Appendix). Third, in the no-ceiling condition, the messages "allergic reaction: 0/20" or "allergic reaction: 10/20" were presented in a black color. Fourth, the preference question was presented as in Figure 1. The participants answered this question by pressing either the letter A (for option A) or the letter B (for option B) on the keyboard. The sequence of both options was counterbalanced across participants for each no-ceiling condition separately. The sequence of which food came first in both options was also counterbalanced. Finally, there were now three usefulness questions for X, K, and L, respectively. Whether the usefulness question for X was presented first or last was also counterbalanced across participants for each no-ceiling condition separately, orthogonally to the counterbalancing of the preference question.

Results

The results are summarized in Table 1. The preference of participants is almost equally divided between option A and B, but there is a clear difference between conditions. More participants preferred option B in the no-ceiling condition and vice versa in the ceiling condition. This interaction was significant [$\chi^2(1, N = 32) = 10.17, p = .001$].

Because there were no differences between usefulness ratings for the control cues K and L for each condition both ($t_s < 1$), we computed the mean usefulness ratings for cue K and L and will report this measure simply as the usefulness of cue K/L. A MANOVA of usefulness scores with cue (X vs. K/L) as a within-subjects variable and condition (ceiling vs. no ceiling) as a between-subjects variable showed a main effect of cue [$F(1,30) = 23.41, p < .001$], indicating lower usefulness ratings for cue X than for cue K/L. The interaction of this effect with condition, however, approached significance [$F(1,30) = 4.00, p < .06$]. Independent-samples t tests showed that usefulness ratings for cue X were significantly higher in the ceiling

If you had the possibility to see two additional events,
would you like to see what happens if

Option A: 1) the patient only eats X; and 2) the patient only eats K

OR

Option B: 1) the patient only eats K; and 2) the patient only eats L?

(Press A if you prefer option A, press B if you prefer option B.)

Figure 1. Preference question for Experiment 2 (translated from Dutch).

condition than in the no-ceiling condition [$t(30) = 2.62$, $p < .05$]. This was not the case for the usefulness ratings of cue K/L [$t(30) < 1$]. Finally, paired-samples t tests revealed that the difference in usefulness of cue X and cue K/L was highly significant in the no-ceiling condition [$t(15) = 4.79$, $p < .001$], but only marginally significant in the ceiling condition [$t(15) = 2.03$, $p < .07$].

Discussion

As predicted on the basis of higher order reasoning accounts of human causal learning, more participants preferred to see the blocked cue X and the control cue K on their own rather than both control cues K and L on their own when the causal status of cue X was unsure (ceiling condition). This pattern was completely reversed when participants could infer with certainty that cue X was not a cause of the outcome (no-ceiling condition). Furthermore, the usefulness of seeing cue X on its own was significantly higher in the ceiling condition than in the no-ceiling condition. This was not the case for the control cues K and L. The usefulness for cue X, however, still tended to be lower than the usefulness for cue K/L in the ceiling condition, which is at first sight somewhat surprising because the causal status of both cue X and cue K/L should be unsure. However, the participants in the ceiling condition may have found it relatively more useful to see cue K/L on its own than to see cue X, because having causal information about one control cue (K or L) could possibly also result in having causal information about the other control cue.

GENERAL DISCUSSION

In this article, we have tested in a direct way the explanation of higher order reasoning accounts for ceiling effects in forward blocking—that is, the reduction in blocking when the outcome is merely present on A+ and AX+ trials as opposed to when the outcome has the same submaximal intensity on these trials. According to this explanation, ceiling effects are due to the fact that the causal status of a blocked cue X is unsure when the outcome is always fully present on A+ and AX+ trials, whereas the causal status of X can be inferred with a high degree of certainty when the outcome is submaximal on A+ and AX+ trials. Contrary to the uncertainty regarding the causal status of cue X when the outcome is merely present on A+ and AX+ trials, the causal status of a reduced overshadowing cue Y (B–, BY+) can always be inferred with certainty, even when the outcome on BY+ trials is fully present. In line with these hypotheses, we found that (1) after an A+, B–, AX+, BY+ learning stage, all participants preferred to see the effect of the blocked cue X on its own as opposed to seeing the effect of the reduced overshadowing cue Y (Experiment 1) and (2) after an A+, Z–, AX+, KL+ learning stage, relatively more participants preferred to see the blocked cue X on its own when the outcome was fully present than when the outcome was submaximal on A+ and AX+ trials (Experiment 2).

Although these results fit well within accounts that assign a role to controlled reasoning processes in human

Table 1
Number of Participants Who Preferred Option A or B
and Mean Usefulness Ratings and Standard Deviations
As a Function of Ceiling Condition in Experiment 2

Condition	Option A	Option B	Usefulness X		Usefulness K		Usefulness L	
			<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
No ceiling	4	12	3.69	3.16	8.00	2.45	8.56	1.71
Ceiling	13	3	6.63	3.18	8.44	1.67	8.63	1.50

Note—Option A is a preference for cues X and K, and option B is a preference for cues K and L.

causal learning, most associative models are silent about these predictions, certainly a priori. However, associative models that incorporate a role for attention (e.g., Kruschke, 2001, Kruschke & Blair, 2000; Mackintosh, 1975) may have an explanation for the results of Experiment 1. According to these models, participants pay little attention to the blocked cue X or learn not to pay attention to it. When participants are asked after the learning stage which cue they would prefer to see on its own, they may prefer the blocked cue over the reduced overshadowing cue because of this (learned) inattention to cue X during the learning stage. However, this explanation does not hold for the results of Experiment 2. According to attention models, there should be a preference for the blocked cue X over the control cues, independently of whether the outcome always occurs with a maximal (ceiling condition) or a submaximal (no-ceiling condition) intensity.

The power PC theory (Cheng, 1997), an extension of the probabilistic contrast model (e.g., Cheng & Holyoak, 1995; Cheng & Novick, 1990), also incorporates uncertainty about the causal status of a blocked cue X when the probability of the outcome is 1 on A+ and AX+ trials. In this way, the power PC theory can account for the results of Experiment 1 and the results of the ceiling condition of Experiment 2. The power PC model, however, can deal only with variations in the probability of the outcome but not with variations in the intensity of the outcome. Therefore, the power PC model cannot account for the data from the no-ceiling condition of Experiment 2. Also note that the power PC model is a normative model. This means that causal judgments are assumed to reflect the outcome of probabilistic contrasts, but no assumptions are made about the processes that are responsible for the ratings (e.g., it is not assumed that people actually come to causal judgments by computing probabilistic contrasts in a controlled manner).

Another class of models of causal inference is based on causal maps and Bayes nets (see, e.g., Glymour, 2003; Gopnik et al., 2004; Sobel, Tenenbaum, & Gopnik, 2004; Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003). A variant of these models (Tenenbaum & Griffiths, 2003) have built-in assumptions whereby the knowledge that A is a cause of the outcome makes the causal status of the blocked cue X uninformative. Bayes nets can, just like the power PC model, incorporate uncertainty about the blocked cue X when the outcome is merely present on A+ and AX+ trials. How Bayesian models deal with variations in intensity of the outcome, however, remains unclear. Furthermore, it is also not very clear whether these models assume that the underlying processes of causal judgments are based on controlled reasoning.

Higher order reasoning accounts, on the other hand, postulate that participants will actually try to make a controlled inference about the causal status of cues. When the outcome is always fully present on A+ and AX+ trials, a valid inference regarding the causal status of X is not possible, and hence participants should be uncertain about the status of X. If, however, the outcome occurs to the same

submaximal extent on A+ and AX+ trials and if one assumes that causes have additive effects, one can infer with certainty that X is not an effective cause of the outcome. Previous (Vandorpe, De Houwer, & Beckers, 2005) and present evidence suggest that this is indeed the standard assumption of humans. If, however, one has good reason to believe that the effect of cues is not additive, one will not make the inferences just described.

The results of our experiments are in line with other recent evidence in favor of higher order reasoning accounts (e.g., De Houwer, Beckers, & Vandorpe, 2005, for a review). Work of Waldmann and colleagues (Waldmann, 2000, 2001; Waldmann & Holyoak, 1992), for instance, has revealed that blocking is modulated by whether cues are seen as causes or as effects of the outcome. Similarly, De Houwer et al. (2002) found significant blocking when cues were causes but not when cues were predictors. Cue competition also seems to depend on the assumption that causal cues have additive effects on the outcome (Beckers et al., 2005; Lovibond et al., 2003) and on retrospective assumptions about the presence or absence of cues (De Houwer, 2002). Whether the outcome on A+ and AX+ trials occurs with an intensity equal to or smaller than the maximal intensity experienced also modulates blocking (Beckers et al., 2005). Furthermore, De Houwer and Beckers (2003) have found that blocking is modulated by the amount of available working memory resources, whereas the causal ratings of the other cues are not affected by less available working memory resources. Finally, Vandorpe et al. (2005) have shown that blocking is found only in participants who report an appropriate blocking inference. All this evidence shows that higher order reasoning processes play a major role in cue competition in human causal learning.

Associative models, on the other hand, leave little or no room for reasoning processes and are therefore unable to account for all the evidence just mentioned. Higher order reasoning accounts of human causal learning, however, although better capable of explaining this evidence, have one big disadvantage in comparison with associative models. They lack a formalized precision and are therefore very unconstrained. Nevertheless, by making specific a priori predictions, higher order reasoning accounts can improve our knowledge of the underlying processes of human causal learning, and it is therefore worthwhile to investigate them further.

REFERENCES

- BECKERS, T., DE HOUWER, J., PINEÑO, O., & MILLER, R. R. (2005). Outcome additivity and outcome maximality influence cue competition in human causal learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **31**, 238-249.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, **104**, 367-405.
- CHENG, P. W., & HOLYOAK, K. J. (1995). Complex adaptive systems as intuitive statisticians: Causality, contingency, and prediction. In H. L. Roitblat & J.-A. Meyer (Eds.), *Comparative approaches to cognitive science* (pp. 271-302). Cambridge, MA: MIT Press.
- CHENG, P. W., & NOVICK, L. R. (1990). A probabilistic contrast model

- of causal induction. *Journal of Personality & Social Psychology*, **58**, 545-567.
- 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. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology*, **56B**, 345-357.
- DE HOUWER, J., BECKERS, T., & GLAUTIER, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, **55A**, 965-985.
- DE HOUWER, J., BECKERS, T., & VANDORPE, S. (2005). Evidence for the role of higher order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior*, **33**, 239-249.
- DE HOUWER, J., VANDORPE, S., & BECKERS, T. (2005). On the role of controlled cognitive processes in human associative learning. In A. Wills (Ed.), *New directions in human associative learning* (pp. 41-63). Mahwah, NJ: Erlbaum.
- DENNISTON, J. C., SAVASTANO, H. I., & MILLER, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 65-117). Mahwah, NJ: Erlbaum.
- DICKINSON, A., & BURKE, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology*, **49B**, 60-80.
- DICKINSON, A., SHANKS, D., & EVENDEN, J. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology*, **36A**, 29-50.
- GLYMOUR, C. (2003). Learning, prediction and causal Bayes nets. *Trends in Cognitive Sciences*, **7**, 43-48.
- GOPNIK, A., GLYMOUR, C., SOBEL, D. M., SCHULZ, L. E., KUSHNIR, T., & DANKS, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, **111**, 3-32.
- KRUSCHKE, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, **45**, 812-863.
- KRUSCHKE, J. K., & BLAIR, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review*, **7**, 636-645.
- LOVIBOND, P. F. (2003). Causal beliefs and conditioned responses: Retrospective revaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, & 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 judgment 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., & 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: Advances in research and theory* (Vol. 22, pp. 51-92). San Diego: Academic Press.
- 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.
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- SOBEL, D. M., TENENBAUM, J. B., & GOPNIK, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, **28**, 303-333.
- STEYVERS, M., TENENBAUM, J. B., WAGENMAKERS, E.-J., & BLUM, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, **27**, 453-489.
- TENENBAUM, J. B., & GRIFFITHS, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in neural information processing systems* (Vol. 15, pp. 35-42). Cambridge, MA: MIT Press.
- VANDORPE, S., & DE HOUWER, J. (2005). A comparison of forward blocking and reduced overshadowing in human causal learning. *Psychonomic Bulletin & Review*, **12**, 945-949.
- VANDORPE, S., DE HOUWER, J., & BECKERS, T. (2005). Further evidence for the role of inferential reasoning in forward blocking. *Memory & Cognition*, **33**, 1047-1056.
- VAN HAMME, L. J., & WASSERMAN, E. A. (1994). Cue competition in causality judgments: The role of nonpresentation of compound stimulus elements. *Learning & 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. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **26**, 53-76.
- WALDMANN, M. R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin & Review*, **8**, 600-608.
- WALDMANN, M. R., & HOLYOAK, K. J. (1992). Predictive and diagnostic learning within causal models: Asymmetries in cue competition. *Journal of Experimental Psychology: General*, **121**, 222-236.
- WALDMANN, M. R., & WALKER, J. M. (2005). Competence and performance in causal learning. *Learning & Behavior*, **33**, 211-229.
- WU, M., & CHENG, P. W. (1999). Why causation need not follow from statistical association: Boundary conditions for the evaluation of generative and preventative causal powers. *Psychological Science*, **10**, 92-97.

NOTE

1. Note that this is exactly what is going on in a reduced overshadowing design, at least when one assumes that controlled reasoning processes are involved in cue competition. If one knows that B is not a cause of the outcome, one can infer during BY+ trials that Y is the cause of the outcome.

APPENDIX

Learning Instructions

This experiment is about how people learn relations between different events. Try to imagine that you are a doctor. One of your patients suffers from allergic reactions after eating certain foods. To discover which foods lead to an allergic reaction, the patient ate specific foods on different days and this was followed by a test on whether an allergic reaction occurred. In a moment, you will see the results of these daily allergy tests one by one on the screen. On each trial, you will first see what the patient ate that day. On some days, he ate only one food; on other days, he ate two different foods. Look carefully each time at what the patient ate that day. You will also receive information about whether the patient showed an allergic reaction or not. Use this information to determine for each food separately whether it led to an allergic reaction in your patient. Note that if the patient ate two different foods and there was an allergic reaction, you do not know which of the two foods was responsible for the allergic reaction. *“But you will always get information about the total intensity of the allergic reaction, as caused by all consumed foods. If the intensity is zero, this means that there is no allergic reaction; if the intensity is greater than zero, this means that there is an allergic reaction. Note that the maximal intensity that can be measured corresponds to an intensity of 20.”* You nevertheless have to determine for each food separately to what extent it causes an allergic reaction in the patient.

First you will see information about 24 allergy tests. After that, you will have to judge for each food the extent to which you think it is a cause of an allergic reaction in the patient. Notice that only the presented information can help you. The task is to determine to what extent the foods cause an allergic reaction in this specific patient. Your personal experiences with the foods or occasional knowledge about the properties of the foods are not relevant and cannot help you. Only the presented information matters.

Note—These instructions are translated from Dutch. The words in italics were added only in the no-ceiling condition of Experiment 2.

(Manuscript received November 19, 2004;
revision accepted for publication June 10, 2005.)