

Further evidence for the role of inferential reasoning in forward blocking

STEFAN VANDORPE and JAN DE HOUWER
Ghent University, Ghent, Belgium

and

TOM BECKERS
University of Leuven, Leuven, Belgium

Previous studies on human causal learning (De Houwer & Beckers, 2003; De Houwer, Beckers, & Glautier, 2002) showed that secondary task difficulty (performing an easy vs. a difficult secondary task during the main causal learning task) and ceiling information (outcome occurs with a maximal vs. submaximal intensity) had an influence on forward blocking (i.e., lower causal ratings for cue T when AT+ trials are preceded by A+ trials than when no A+ trials are presented). We extended these studies by also examining self-reports of participants about the reasons behind their causal ratings. Blocking was found only for participants who reported an appropriate blocking inference. Furthermore, secondary task difficulty and ceiling information influenced the number of participants who reported an appropriate blocking inference. These findings point to a major impact of inferential reasoning in human causal learning.

The ability of humans to learn predictive and/or causal relations is essential in daily life because it allows one to anticipate and control events in the environment. There is, however, still much controversy about the processes that underlie human causal learning (HCL; see De Houwer & Beckers, 2002, for a review). One family of accounts of HCL comprises associative models (e.g., Denniston, Savastano, & Miller, 2001; Dickinson & Burke, 1996; Rescorla & Wagner, 1972). These have in common (1) that learning a relation between two events or stimuli is represented by changes in the strength of the association between the representations of those events or stimuli, and (2) that the value of the contingency judgment is at least partly based on this associative strength.

Such associative models are mechanistic, in that little or no role is assigned to controlled higher order reasoning processes. In recent years, however, there is increasing evidence, particularly from research on (forward) blocking, that such processes do play a role in HCL (e.g., Lovibond, 2003; Waldmann, 2000; see De Houwer, Beckers, & Vandorpe, 2005, for a review). When a target cue T is paired with another cue A and followed by the outcome (AT+ trials), ratings of the contingency between cue T and the outcome are lower when the AT+ trials are preceded by A+ trials (cue A followed by the outcome) than when no A+ trials are presented (e.g., Dickinson, Shanks, & Evenden, 1984; Shanks, 1985). Learning about cue T is said

to be *blocked* by cue A. Recent studies have shown that blocking is found more consistently when cues are seen as potential causes of the outcome than when cues are seen as effects of the outcome (Waldmann, 2000; Waldmann & Holyoak, 1992; but see Shanks & Lopez, 1996) or as mere predictors of the outcome (De Houwer, Beckers, & Glautier, 2002). Furthermore, blocking is also influenced by ceiling information. Less blocking occurs when the outcome on A+ and AT+ trials has a maximal intensity (e.g., an impact of 10/10) than when the same submaximal outcome (e.g., an impact of 10/20) occurs on A+ and AT+ trials (De Houwer et al., 2002). Finally, forward blocking also depends on retrospective assumptions about the presence or absence of the blocked cue during the A+ trials (De Houwer, 2002). All these findings raise problems for associative models because they show that specific assumptions about the nature of the cues and outcomes can have an important impact on HCL. Because of their mechanistic, stimulus-driven nature, associative models assign no role to these specific assumptions.

Lovibond (2003) has recently proposed an inferential reasoning account that is capable of explaining the findings just mentioned (see Waldmann, 2000, for a highly related account). He 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” (p. 98). In blocking studies, participants might, for instance, apply the following inferential rule:

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

Correspondence concerning this article should be addressed to S. Vandorpe, Department of Psychology, Ghent University, Henri Dunantlaan 2, B-9000 Ghent, Belgium (e-mail: stefaan.vandorpe@ugent.be).

probability, this implies that cue T is not a cause of the outcome. (De Houwer & Beckers, 2003, p. 346)

It is assumed that participants have conscious (in the sense of reportable) knowledge about this inferential rule and about the conditions under which this rule can lead to valid inferences. They will apply this knowledge in a controlled (in the sense of conscious, intentional, and effortful) manner that allows them to draw conclusions with regard to events that have not been observed directly (e.g., whether T on its own will be followed by the outcome) on the basis of propositional knowledge that they have previously formed as the result of observation, instruction, or deduction. Because inferential reasoning is an effortful process, it will take place only to the extent that participants are motivated and have the opportunity to engage in such reasoning.

The inference that is assumed to underlie blocking is valid only if A and T are potential causes and if the intensity of the outcome on A+ and AT+ trials is not maximal. Suppose that A and T are not potential causes but only potential effects (e.g., Waldmann & Holyoak, 1992) or predictors of an outcome (just as a flickering light and/or a sharp noise at a railway crossing predict but do not cause the coming of a train). In such a situation, people cannot logically infer that T on its own is a poor predictor of the outcome on the basis of the information that A alone predicts the outcome as well as A and T do together. Therefore, on the basis of an inferential account, one would predict that blocking should be less likely to occur when cues are predictors or effects than when they are causes. Likewise, if A and T are indeed potential causes, but the intensity of the outcome is maximal on A+ trials as well as on AT+ trials, people cannot infer with certainty that T is ineffective in causing the outcome above the effect of A because A already causes the outcome to a maximal level. In such a situation, people should be uncertain about the causal status of T. If, however, A as well as A and T together cause the outcome to the same submaximal intensity, then participants can be sure that T has no causal effect if it is assumed that the causes have additive effects. Blocking should thus occur.

Using a task similar to that task developed by Dickinson et al. (1984), De Houwer and Beckers (2003) tested the inferential reasoning account in a more direct way. On each trial, an army tank moved across the screen from left to right and could explode on a fixed position (outcome). At the bottom of the screen, there were squares that represented weapons (cues). On each trial, one or two weapons fired (represented by a light that appeared in the square) and the tank did or did not explode. After observing all trials, participants were asked to make causal judgments for each cue separately. Within this tank explosion paradigm, De Houwer and Beckers (2003) implemented a simple forward-blocking procedure (with a submaximal outcome) and asked the participants, meanwhile, to perform an easy or a difficult secondary task. If blocking in HCL is determined by inferential reasoning processes, then secondary task difficulty should modulate blocking. This is because

inferential reasoning, in contrast to associative processing, can be considered an effortful, controlled process (Sloman, 1996). In line with their hypothesis, De Houwer and Beckers (2003, Experiment 2) found smaller blocking effects in the difficult secondary task condition than in the easy secondary task condition. Importantly, the ratings of the other cues were not affected by secondary task difficulty, so that the smaller blocking effect could not have been due to inferior functioning of associative processes during performance of a demanding secondary task.

Our aim was to test the inferential reasoning account further. Because inferential reasoning is a controlled process, participants should, at least in principle, be able to report the inferences that they have made. Therefore, self-reports can provide a way of examining whether inferences play a role in HCL. In the present experiment, we asked our participants to report on how they arrived at their causal ratings. According to the inferential reasoning account, blocking occurs because people infer that T cannot be a cause of the outcome if the same submaximal outcome occurs on the A+ and AT+ trials. Therefore, participants who show blocking should be able to report this inference (assuming that they do not forget making it), whereas those who do not show blocking should not report this inference. Furthermore, variables that discourage people from making the blocking inference, such as secondary task difficulty and ceiling information, should also reduce the number of participants who report a blocking inference. We tested these predictions in two experiments.

Our Experiment 1 was similar to Experiment 1 of De Houwer and Beckers (2003) in which secondary task difficulty was manipulated. If secondary task difficulty modulates blocking because of its influence on the number of participants who are able to engage in inferential reasoning, it should influence the number of participants reporting a blocking reasoning. Furthermore, blocking effects should be significantly larger for participants who report a blocking inference.

In Experiment 2, we extended Experiment 1 of De Houwer et al. (2002), who found blocking when the outcome occurred to a submaximal extent but not when the outcome occurred to the maximal extent on A+ and AT+ trials. According to the inferential reasoning account, participants should be unsure about the causal status of T in the maximal outcome condition. If they do make a blocking inference under this condition, they should be unsure about it. Therefore, an inferential reasoning account predicts that a smaller number of participants will report a blocking inference when a ceiling is imposed than when the outcome occurs to a submaximal extent. Furthermore, participants who still report a blocking inference under ceiling conditions should be less sure about it.

In both experiments, we used a standard blocking design: A+ and Z- trials followed by AT+, KL+, and Z- trials. In such a design, K and L are control cues for the blocked cue T. The only difference between K and L on the one hand and T on the other hand was that T was presented in compound with a cue that had already been paired with the outcome in the previous learning stage.

EXPERIMENT 1

Method

Participants. Twenty-six first-year psychology students at Ghent University participated for course credit. They were randomly assigned to the easy or to the difficult secondary task condition.

Stimuli and Materials. The contingency learning task and the secondary task were presented on separate IBM-compatible 486 PCs with 15-in. SVGA screens. Both tasks were implemented using a custom-made Turbo Pascal 5.0 program and were similar to the tasks used by De Houwer and Beckers (2003). The participants faced the computer, on which the contingency learning task was presented. A tank moved in a continuous manner from the left to the right side of the computer screen on a straight line that was situated 10 cm from the top of the screen. It took about 6 sec for a tank to get from the left to the right side of the screen. If a tank exploded, this always occurred 2 sec after the tank appeared. The explosion took about 1 sec. At the same time, the message "IMPACT 10/20" appeared on the screen for 3 sec. If the tank did not explode, the message "IMPACT 0/20" appeared on the screen until the tank, which drove on, had reached the right side of the screen.

Five rectangles, 2.5 cm wide \times 1.7 cm high, were situated at the bottom of the screen. A weapon was said to be firing when a solid white rectangle appeared in the rectangle that represented that particular weapon. The solid square appeared 1,700 msec after the beginning of the trial and was presented for 300 msec. All stimuli were white and were presented on a black background.

Procedure. At the beginning of the experimental session, the participants received written instructions that provided the following information: They would perform two tasks, a learning task and a reaction time (RT) task, and the latter task would be presented on the other computer than the one they faced. The instructions for the RT task were given first. The participants who were assigned to the easy secondary task condition were told that a tone would be presented at certain moments in time. Their task was to press the key "Q" as quickly as possible after hearing a tone. In the difficult secondary task condition, the participants were informed that a high or a low tone would be presented at certain moments in time. They were also asked to react as quickly as possible after hearing a tone, but the appropriate reaction was determined by the pitch of the previous tone. If the previous tone was high, participants had to press the key "Q," and if the previous tone was low they had to press the key "M." All participants were told that they would hear a very low feedback signal (two short tones of 30 msec with a frequency of 100 Hz and an interval of 30 msec) if they pressed too early or too late, or if they pressed an incorrect key. They were asked not to react to the feedback signal. After indicating that they understood these instructions, the participants in the easy secondary task condition completed 30 practice trials. Participants in the difficult secondary task condition completed between one and five blocks of 30 practice trials until they reached a percentage of 70% correct responses. In the easy task, a 750-Hz tone was presented for 60 msec every 900 msec. In the difficult task, the computer also presented tones for

60 msec each. However, whether this tone had a pitch of 500 Hz or 1000 Hz was determined randomly on each trial, with the restriction that the pitch of the tone could not be the same on more than two consecutive trials. The interstimulus interval in the difficult secondary task was fixed at 1,500 msec. If the response was too early or too late, or if the participants pressed an incorrect key, the feedback tone was presented.

After completing the secondary task practice trials, participants read the instructions for the contingency learning task. Instructions said that drawings of army tanks would ride across the computer screen. There would also be five weapons represented by five squares at the bottom of the screen. Participants were told that the firing of a weapon would be indicated by a white light appearing in the square representing that weapon. They were asked to determine the effectiveness of each weapon in destroying tanks (see De Houwer & Beckers, 2003, for more details). Finally, it was stressed that the learning task and the RT task were equally important.

When the participant indicated that he or she had fully understood the instructions, the experimenter started the secondary task and, after a few sec, the contingency learning task. The contingency learning task consisted of the following events. First, 10 A+ and 10 Z- trials were presented. Then, 10 AT+, 10 KL+, and 10 Z- trials were presented. There was no break between the two phases. The order of the trials within each phase was determined randomly for each participant. Which square functioned as which cue was counterbalanced with the restriction that cues that were presented in compound never appeared next to each other (see De Houwer & Beckers, 2003).

When all 50 contingency learning trials had been presented, the instructions for the rating phase appeared on the screen. At that time, the experimenter stopped the secondary task. In the instructions, participants were asked to indicate for each weapon separately how effective it was in destroying tanks. They could do so by entering a score between 0 (*very ineffective*) and 100 (*very effective*). After entering an effectiveness rating, participants expressed how certain they were that their effectiveness rating was accurate. They did so by entering a score between 0 (*very unsure*) and 100 (*very sure*). All participants first rated the cue that was represented by square 1 (the square on the far left side), then the cue represented by square 2, and so on (see De Houwer & Beckers, 2003, for more details). After all cues had been rated, the experimenter asked the following question: "You have now made the ratings for the different cues. Could you tell me how you arrived at those ratings?" The experimenter classified a participant as making a blocking inference if the participant mentioned that T was less likely to be a cause of the outcome because it always went together with A and A already produced the outcome with an impact of 10 out of 20.

Results

Overall ratings. The mean effectiveness and confidence ratings for each cue are given in Table 1. The table also contains mean blocking scores and mean confidence

Table 1
Mean Effectiveness (Effect.) and Confidence (Conf.) Ratings for the
Different Cues as a Function of Secondary Task Difficulty in Experiment 1

Task	A		T		K		L		Z		Blocking		Conf. Blocking	
	M	SE	M	SE	M	SE	M	SE	M	SE	M	SE	M	SE
Effect. easy	78	7	17	6	47	4	48	4	0	0	31	7		
Effect. difficult	71	6	41	7	55	6	58	8	4	3	15	13		
Conf. easy	84	6	77	6	62	4	56	5	99	1			18	7
Conf. difficult	66	8	42	10	30	6	43	9	80	8			6	10

blocking scores. The blocking score was calculated by subtracting the effectiveness rating for T from the mean of the effectiveness ratings for K and L. A high blocking score thus indicates that T was given a lower effectiveness rating than were K and L, which shows that blocking had occurred. The confidence blocking score corresponds to the difference between the confidence rating for T and the mean of the confidence ratings for K and L. A high confidence blocking score thus indicates that participants had more confidence in their rating for T than in their ratings for K and L.

One-sample *t* tests revealed that the blocking score was different from zero in the easy secondary task condition [$t(12) = 4.14, p = .001$] but not in the difficult secondary task condition [$t(12) = 1.2, p > .25$]. Similar results were obtained for the confidence blocking scores [$t(12) = 2.73, p < .05$, and $t < 1$] for the easy and difficult secondary task conditions, respectively. Independent-samples *t* tests showed that neither the blocking scores [$t(24) = 1.03, p > .30$] nor the confidence blocking scores differed significantly between secondary task conditions [$t(24) = 1.01, p > .30$]. Further tests showed that the effectiveness rating of T was affected by secondary task difficulty [$t(24) = 2.50, p = .02$], whereas the effectiveness ratings of the other cues were not (all t s < 1.45). Confidence ratings were higher in the easy secondary task condition, with significant differences [for cue T, $t(24) = 2.93, p < .01$; cue K, $t(24) = 4.11, p < .001$; cue Z, $t(24) = 2.47, p < .05$; to a lesser extent for cue A, $t(24) = 1.77, p < .10$].

Self-reports. Eight out of 13 participants in the easy task condition spontaneously mentioned that they rated weapon T zero because T fired together with weapon A and A had already produced a tank explosion with an impact of 10 out of 20. In the difficult task condition, there was only 1 participant who reported this blocking inference. This difference was significant [$\chi^2(1, N = 26) = 8.33, p < .01$]. On the basis of these reports, we divided our participants into a group who reported a blocking inference (reasoning group) and a group who did not (no-reasoning group). One-sample *t* tests showed that the mean blocking score was significantly different from zero for the reasoning group ($M = 50.84, SE = 5.56, n = 9$) [$t(8) = 9.14, p < .001$] but not for the no-reasoning group ($M = 8.24, SE = 9.14, n = 17$) [$t(16) < 1$]. The difference in blocking scores between both groups was significant [$t(24) = 3.20, p < .01$]. Also, the confidence blocking score was significantly different from zero for the reasoning group ($M = 37.22, SE = 9.25, n = 9$) [$t(8) = 4.03, p < .01$] but not for the no-reasoning group ($M = -1.6, SE = 6.01, n = 17$) [$t(16) < 1$]. The difference in confidence blocking scores between the two groups was significant [$t(24) = 3.65, p = .001$]. Finally, all 9 participants in the reasoning group gave an effectiveness rating of zero for cue T, whereas only 1 of the 17 participants in the no-reasoning group did so.¹

Secondary task results. Analyses of secondary task performance showed that participants in the easy sec-

ondary task responded more quickly [$t(24) = 4.57, p < .001$] and more accurately [$t(24) = 9.62, p < .001$] to the tones.

Discussion

The results of Experiment 1 support an inferential reasoning account of blocking. The mean blocking and confidence blocking scores of participants in the reasoning group were substantial and significantly different from zero, whereas the mean blocking and confidence blocking scores of the no-reasoning group were very close to and not significantly different from zero (see note 1). Moreover, both blocking scores differed significantly between groups. This is in line with the hypothesis that blocking depends on participants' making a blocking inference. Furthermore, secondary task difficulty modulated the number of participants who reported a blocking inference, as was expected. Note that although we did not find a significant effect of secondary task difficulty on blocking, additional analyses showed that secondary task difficulty did modulate blocking when we took into account only the data of the participants for whom the manipulation had the intended effect (i.e., participants who reported reasoning in the easy secondary task condition and participants who reported no reasoning in the difficult secondary task condition).

One could argue that participants reported a blocking inference *ex post facto* in order to justify why they had given cue T a low causal rating. The ratings of the confidence blocking scores, however, indicate that this was not the case. Participants in the reasoning group had high confidence blocking scores at the moment of rating, whereas participants in the no-reasoning group did not have more confidence in their rating for T than they had in their rating for the control cues K and L. This indicates that the self-reports reflected the actual inference (or absence of inference) that participants made during the learning or test stage.

EXPERIMENT 2

De Houwer et al. (2002) found significant blocking when the outcome had a submaximal intensity but not when the outcome had a maximal intensity on the A+ and AT+ trials. In Experiment 2, we extended Experiment 1 of De Houwer et al. (2002) in three ways. The first modification was related to our main research question, whereas the other two aimed at broadening the findings about the effect of ceiling information on blocking.

First, we wanted to investigate the influence of ceiling information on the self-reports of the participants. We therefore interviewed the participants after the experimental session about the reasons for their causal ratings. Moreover, the interviewing and coding procedures were conducted more systematically and formally than in Experiment 1. If ceiling information influences reasoning processes, this should manifest itself in the self-reports of the participants. Second, we used a food allergy paradigm

instead of the tank explosion paradigm. This allowed us to further test the generality of the findings of De Houwer et al. (2002). Finally, we wanted to investigate the influence of ceiling information on forward blocking when participants have no information at all about the intensity of the outcome. Although the ceiling effects that were found by De Houwer et al. (2002) support an inferential reasoning account, one could argue that the message “impact 10/10” artificially introduced concerns about ceiling information, but that this is not necessarily the case when the outcome just occurs and no information about its intensity is given. We therefore compared a condition where participants received the message “allergic reaction: 10/20” (information condition) with a condition where participants just received the message “allergic reaction” (no-information condition). Whereas in the former condition, participants could easily infer that the to-be-blocked cue T had no effect, the causal status of T was unsure in the latter condition. In this situation, an inferential reasoning account would predict that fewer participants will report a blocking inference and that participants who still report a blocking inference will be unsure about it.

Method

Participants. Thirty-two first-year psychology students at Ghent University participated for course credit. They were randomly assigned to one of the two conditions.

Stimuli and Materials. The task was presented on a 486 PC and was implemented using a custom-made Turbo Pascal 5.0 program. As in Experiment 1, there were 10 A+ and 10 Z- trials in Stage 1 and 10 AT+ trials, 10 KL+ trials, and 10 Z- trials in Stage 2. Five names of foods (the Dutch words for beans, carrots, leek, broccoli, and lettuce) were randomly assigned to the different cues for each participant. The food names were presented in white against a black background. The feedback messages “allergic reaction: 10/20” and “allergic reaction: 0/20” in the information condition and “allergic reaction” in the no-information condition also appeared in a white color. The feedback message “no allergic reaction” in the no-information condition was presented in red. Each letter was 5 mm wide and 7 mm high.

Procedure. In the no-information condition, the following instructions appeared on the screen at the beginning of the experiment (translated from Dutch):

In this experiment, you will receive information about a patient who is allergic to some but not all foods. On each trial, you will see which foods the patient has eaten and after that you will be told whether the patient showed an allergic reaction. The patient always eats the same portion of each type of food. Sometimes the patient eats two types of food within the same meal and thus two portions (one portion of each type of food). If this happens and the patient has an allergic reaction, then you do not

know which of the two foods was responsible for the allergic reaction. You nevertheless have to determine for every single food to what extent it causes an allergic reaction in your patient. First, you will see 50 trials. Every trial stands for a separate moment on which the patient ate a meal and the allergic reaction was measured. At the end of the experiment, you will have to judge for each food separately how likely it is to cause an allergic reaction in the patient.

For the information condition, the following lines were added after “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. Keep in mind that the maximal intensity that can be measured corresponds to an intensity of 20.”

After the participant indicated that the task was clear, the experimenter started the learning stage. Each trial started with the presentation of the food(s) the patient had eaten. Single foods were presented in the center of the screen during 2,000 msec. On compound trials, one food was presented immediately above and one immediately below the center of the screen, also during 2,000 msec. For each compound stimulus, each food appeared equally often above and beneath the center of the screen. After the presentation of the food(s), information about the allergic reaction appeared during 3,000 msec. The intertrial interval was 3,000 msec. After the 50 trials of the learning stage, participants were asked to judge for each food how likely it was to cause an allergic reaction. They were also asked to indicate how sure they were about that judgment. The ratings were given on a scale from 0 to 100. At the end of the experiment, a semistructured interview was administered that focused on the following questions: (1) “How did you make your ratings for the different foods?” (2) “How much have you given K/L/T and why?” and (3) “Why have you given T as much/less/more than K/L?” The interviews were recorded.²

Results

Overall ratings. The results are summarized in Table 2. One-sample *t* tests showed that the blocking score in the no-information condition was only marginally significant [$t(15) = 1.86, p = .08$] and that participants in this condition did not have more confidence in their ratings for T than in their ratings for K and L ($t < 1$). In the information condition, one-sample *t* tests revealed that both blocking and confidence in blocking were highly significant [$t(15) = 10.81, p < .001$, and $t(15) = 4.98, p < .001$, respectively]. Independent-samples *t* tests revealed that both blocking and confidence in blocking differed significantly between conditions [$t(30) = 3.26, p < .01$, and $t(30) = 3.38, p < .01$, respectively].

A closer look at the individual blocking scores revealed some additional interesting data. The marginally significant blocking score in the no-information condition was due to the fact that only 5 out of 16 participants had a positive blocking score. The other participants had a zero

Table 2
Mean Effectiveness (Effect.) and Confidence (Conf.) Ratings for the Different Cues as a Function of Information Condition in Experiment 2

Condition	A		T		K		L		Z		Blocking		Conf. Blocking	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Effect. no info	100	0	42	6	55	4	55	4	0	0	13	7		
Effect. info	80	6	2	1	43	3	41	3	0	0	40	4		
Conf. no info	99	7	63	6	68	6	66	6	100	0			-3	27
Conf. info	86	5	90	4	62	6	64	6	97	3			27	5

(9 out of 11) or even negative (2 out of 11) blocking score. This was in sharp contrast to the information condition, in which all participants had a positive blocking score. The difference between information conditions in the number of participants who had a positive blocking score was significant [$\chi^2(1, N = 32) = 16.76, p < .001$].

Further tests showed that the causal ratings of all cues except cue Z were lower in the information condition than in the no-information condition [$t(30) = 3.30, p < .01$; $t(30) = 6.77, p < .001$; $t(30) = 2.50, p < .05$; $t(30) = 2.90, p < .01$; and $t < 1$, respectively, for cues A, T, K, L, and Z]. Finally, the confidence rating of cue A was lower in the information condition than in the no-information condition [$t(30) = 2.57, p < .05$], whereas the reverse was true for cue T [$t(30) = 3.58, p = .001$]. There was no difference for the other cues ($ts < 1$).

Self-reports. The interviews were judged by two independent raters who were unaware of the research questions of the experiment and of which condition each participant was in. They were asked to judge (1) whether participants reported a blocking inference (we use the term “blocking inference judgment” to refer to this judgment) and, if the answer was “yes,” (2) whether participants mentioned something about the fact that they could not be sure about their blocking inference (for which we will use the term “uncertainty judgment”). The raters were told that we understood under a blocking inference “that the participant realized or even merely considered that the A+ trials could demonstrate or imply that T was a lesser cause of an allergic reaction.” We also pointed out that we spoke of blocking inference even if the participant made this consideration but finally did not bring this into account when making his or her causal ratings. The agreement between the two raters was 91% for the blocking inference judgment. When we excluded the three cases on which there was no agreement for the blocking inference judgment, the agreement for the uncertainty judgment was 97%.

The self-report data are summarized in Table 3. A participant was categorized as “reporting a blocking inference” and “being sure about his or her blocking inference” if one rater judged that the participant reported a blocking inference and judged that the participant was sure about it. An alternative criterion for classification would have been to classify a participant as reporting a blocking inference and being sure about his or her blocking inference if both raters agreed on this. However, we chose the first criterion because this was more conservative with respect to the predictions of the inferential reasoning account. Because there were more disagreements between raters in the no-information condition, the first criterion led to the fact that more participants in the no-information condition were judged to report a blocking inference and to be sure about it because raters did not have to agree. This counteracts the hypotheses of an inferential reasoning account (i.e., more participants in the information condition will report a blocking inference and will be sure about it than will participants in the no-information condition).

Table 3
Judgments of Self-Reports in Function of Information Condition (Computed Over Interviewers) in Experiment 2

	Without Information		With Information	
	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>
Blocking judgment*	.50	.13	1.00	.00
Uncertainty judgment**	.50	.19	.94	.06

*Proportions of participants. **The higher the proportion, the more participants judged themselves to be sure about their reasoning.

Both raters judged that all participants in the information condition reported a blocking inference. This is in sharp contrast with the no-information condition, where only half of the participants reported a blocking inference according to one or both of the raters. This difference in blocking inference between conditions is significant [$\chi^2(1, N = 32) = 10.67, p < .001$]. Furthermore, uncertainty judgments also differed. The raters judged that a greater proportion of participants were certain about their blocking inference in the information condition than in the no-information condition [$\chi^2(1, N = 24) = 6.19, p < .05$].

Just as in Experiment 1, we divided our participants into a reasoning group and a no-reasoning group. The mean blocking and confidence blocking scores for the reasoning group ($M = 35.73, SE = 4.92$, and $M = 16.88, SE = 6.36$, respectively, $n = 24$) were significantly different from zero [$t(23) = 7.26, p < .001$, and $t(23) = 2.65, p < .05$]. This was not the case for the no-reasoning group ($M = -0.63, SE = .63$, and $M = -3.75, SE = 5.57$, respectively, $n = 8$), both $ts \leq 1$. The difference in blocking scores between both groups was significant [$t(30) = 4.21, p < .001$], and the difference in confidence blocking scores was marginally significant [$t(30) = 1.78, p < .09$]. For participants in the reasoning group, the mean confidence blocking score was $-3.13 (SE = 13.72, n = 8)$ in the no-information condition and $26.88 (SE = 5.40, n = 16)$ in the information condition. This difference was significant [$t(22) = 2.45, p < .05$]. There was no difference at all between mean confidence blocking scores of participants in the no-information condition who were in the reasoning group and those who were not ($t < 1$). Finally, 16 participants in the reasoning group gave a causal rating of zero for cue T, whereas none did so in the no-reasoning group. Of the 6 participants in the reasoning group who did not give a causal rating of zero, 4 were in the no-information condition.

Discussion

The self-report data were clearly in line with an inferential reasoning account. Fewer participants spontaneously reported a blocking inference in the no-information condition than in the information condition, and those who did reported more often that they could not be sure about it. Furthermore, blocking scores

significantly differed between the reasoning and no-reasoning groups. Actually, the mean blocking score for the no-reasoning group was (virtually) zero.

The mean confidence blocking score of the reasoning group was higher than the mean confidence blocking score of the no-reasoning group. Furthermore, for the participants in the reasoning group, the confidence blocking scores were significantly higher in the information condition than in the no-information condition. Finally, in the no-information condition, there was no difference in confidence blocking scores between the reasoning and no-reasoning groups. In fact, both confidence blocking scores were close to zero. This is completely in line with the hypothesis of an inferential reasoning account that participants who report a blocking inference in the no-information condition cannot be sure about it because the causal status of the target cue T cannot be inferred with certainty. The fact that confidence differed between the two information conditions for the participants who reported a blocking inference makes an explanation of the self-reports as an *ex post facto* justification of the causal ratings very implausible. All these confidence results indicate that the self-report data reflect inferential reasoning during the experimental session and not inferential reasoning *ex post facto* that participants might have engaged in in order to justify their causal ratings.

The present results also replicate and extend findings about ceiling effects in forward blocking (De Houwer et al., 2002). Blocking and confidence blocking scores were much higher in the information condition than in the no-information condition. Our results extend previous results because participants received only the message “allergic reaction” in the no-information condition, whereas in the study by De Houwer et al. (2002), participants were explicitly pointed to the fact that the outcome occurred to a maximal extent on A+ and AT+ trials (“allergic reaction: 10/10”).

At first sight, these results seem to be at variance with the studies in the literature where “blocking” was found even when the outcome was either present or absent (see, e.g., Aitken, Larkin, & Dickinson, 2001; Chapman & Robbins, 1990; Dickinson & Burke, 1996). However, these studies used an A+, B− (first stage), AT+, BY+ (second stage) design and measured “blocking” by comparing the rating for T with the rating for Y. This comparison actually reflects forward cue competition rather than blocking. The question remains in such studies whether forward cue competition is due to blocking (a decrease in ratings for T), to reduced overshadowing (an increase in ratings for Y), or to both. According to an inferential account, when the outcome is always present or absent, forward cue competition should, for the most part, be due to reduced overshadowing rather than blocking. This is because one can infer that cue Y is the cause of the outcome on the basis of the information B−, BY+, but one will be unsure about the causal status of cue T on the basis of A+, AT+ information (as is demonstrated by our results).

In a recent study at our lab (Vandorpe & De Houwer, in press-a), we directly compared forward blocking with reduced overshadowing by means of an A+, B− (first stage), AT+, BY+, KL+ (second stage) design and an outcome that was all or nothing. In line with an inferential reasoning account, reduced overshadowing was more than seven times larger than blocking and the confidence in reduced overshadowing was more than four times larger than the confidence in blocking. Both differences were very significant and blocking and confidence in blocking were not significantly different from zero.

We are aware of only two studies (Dickinson, Shanks, & Evenden, 1984; Shanks, 1985)³ in which an appropriate control for measuring forward blocking was included in the design and in which blocking was found even though the outcome was always either fully present or fully absent. These studies, however, used probabilistic cue–outcome relations (i.e., the contingency between the cue and outcome was always lower than one) rather than deterministic cue–outcome relations (i.e., the contingency between each cue and the outcome is either one or zero). In probabilistic studies, ceiling considerations do not apply because A does not always cause the outcome. Participants can therefore (at least in principle) verify that T does not change the probability of the outcome, and thus infer that T is not a cause of the outcome (see Cheng, 1997).

The only results that seem somewhat awkward are the lower ratings for A, K, and L in the information condition than in the no-information condition. It is, however, likely that some participants “translated” the intensity information of the outcome (10/20) into a score of 50 instead of the usual 100 for the A cue and a score of approximately 25 for the K and L cues instead of the usual 50. The reason for this translation might have been that some participants thought that the message “allergic reaction: 10/20” meant that the outcome was caused in only half of the situations instead of that the outcome was always caused with an intensity of 10 on 20. They might also have thought that a rating of 100 should be reserved for a cause with an impact of 20/20. A closer look at the data revealed that this was the case. Five participants in the information condition did not give ratings of more than 50, whereas in the no-information condition, all participants gave a rating of 100 for cue A.

GENERAL DISCUSSION

The present experiments were designed to further investigate the role of inferential reasoning in forward blocking. After the rating stage, we asked participants to explain how they arrived at their causal ratings. The rationale for this was that inferential reasoning is a controlled process that can, in principle, be reported on verbally by the participants. Our approach was successful in that the results clearly supported predictions of the inferential reasoning account of forward blocking. First, blocking effects were found only in participants who reported a

blocking inference. Second, the manipulations that we assumed would have an influence on inferential reasoning in forward blocking (difficulty of secondary task and information about the intensity of the outcome) also had an influence on the number of participants who reported a blocking inference after the experimental session. Finally, participants who reported a blocking inference and who could be sure about it were much more confident in their causal rating for T (relative to the causal ratings for the control cues K and L) than participants who did not report a blocking inference or could not be sure about the validity of the inference.

One might be surprised by the fact that the verbal self-reports were in line with the actual observed behavior (i.e., ratings). Nisbett and Wilson (1977), for instance, argued that people can most often not give accurate verbal reports about the processes that underlie their behavior. Even in cases in which verbal reports do seem to be accurate (i.e., in line with observed behavior), it might be that this is due to “the incidentally correct employment of a priori causal theories” (Nisbett & Wilson, 1977, p. 233) rather than a true insight into the underlying processes. Smith and Miller (1978), however, correctly pointed out that the position of Nisbett and Wilson becomes unfalsifiable when even evidence for accurate self-reports is dismissed. More importantly, Smith and Miller summarized a number of studies that show that self-reports about underlying processes can provide accurate and important information about these processes. They therefore put forward the more constructive position that the question is not *whether* people can have access to mental processes but *when* people have access to such processes. Smith and Miller argued that the validity of self-reports depends on

the degree to which the subject is asked to report on tasks that are novel and of interest. Tasks that are novel and engaging for subjects, such as choosing a college or solving challenging problems, often seem to evoke accurate introspective awareness of process. (p. 361)

The tasks in our experiments certainly were novel for the participants. Moreover, they posed a clear problem (i.e., which foods cause an allergic reaction?) that participants were asked to solve. Our contingency learning tasks therefore created the conditions under which verbal reports are likely to be accurate.

One should also note that we did not ask our participants to verbalize the processes that underlie inferential reasoning itself. The question of whether the processes that underlie inferences are, for example, similar to those proposed by rule theory (e.g., Rips, 1983), mental model theory (e.g., Johnson-Laird, 1983), or any other theory of human inferential reasoning is independent from the question of whether or not inferential reasoning as a controlled (i.e., conscious, intentional, effortful) process plays a role in HCL.

In the present experiments, self-reports were given after the rating stage. One advantage of this is that the reports cannot bias the causal and confidence ratings. One

could, however, also ask participants for verbal protocols before or during the rating stage. For example, Brown (1995) followed this approach to study the processes that underlie judgments of event frequency. In a recent study on HCL that was conducted at our lab (Vandorpe & De Houwer, in press–b), we used a similar approach by asking participants which additional information they wanted to see *before* the causal rating stage. The general idea was that if the causal status of a blocked cue T was unsure, participants would prefer to see what would happen if this cue was presented on its own. In line with this idea, we found that participants preferred to see the blocked cue T on its own above a reduced overshadowing cue Y (B– trials in the first stage, BY+ trials in the second stage; Experiment 1) and that more participants preferred to see the blocked cue T on its own when the outcome on A+ and AT+ trials was merely present than when the outcome occurred with an intensity of 10/20 (Experiment 2).

The self-report data that we have presented in this article cannot be explained on the basis of associative models of HCL. These models focus on the role of automatic (i.e., unconscious, unintentional, effortless, or uncontrollable) associative processes (see De Houwer, Vandorpe, & Beckers, 2004). As such, there is nothing in these models that would allow one to make predictions about the self-reports that participants will generate when asked to justify their ratings, nor about how these self-reports will relate to the actual contingency ratings or to variables such as secondary task difficulty or intensity of the outcome. One could, however, again revert to the idea that the self-report data were merely *ex post facto* justifications and argue that the crucial ratings themselves were due to associative processes. Above, we have argued that the confidence data support the assertion that the self-reports did provide an accurate reflection of the processes that led to the ratings. In addition, it needs to be pointed out that associative models do not provide an acceptable explanation for the ratings themselves. Because associative models incorporate the assumption that contingency learning is based on automatic (and thus effortless) processes, they cannot explain the fact that blocking is (at least in some cases) influenced by the difficulty of the secondary task (see De Houwer & Beckers, 2003, for a more detailed discussion of why the secondary task findings contradict associative models). As was pointed out by De Houwer et al. (2002), existing associative models are also incompatible with the fact that blocking (i.e., lower ratings for T than for K and L) is stronger when outcomes occur at a submaximal level than when the outcomes are always fully present.

One important class of models that we have not yet discussed comprises probabilistic contrast models (e.g., Cheng, 1997; Cheng & Holyoak, 1995; Cheng & Novick, 1990; Waldmann & Holyoak, 1992). According to these models, contingency judgments will be a function of the difference between the probability of the outcome in the presence of a target cue $[P(O/T)]$ and the probability of the outcome in the absence of that cue $[P(O/\sim T)]$. When a target cue T is always paired with an alternative cue A,

contingency judgments for T will reflect the difference between the probability that the outcome is present when both A and T are present [P(O/AT)] and the probability of the outcome when A but not T is present [P(O/A~T)]. Such probabilistic models are, however, normative models that are, as such, silent about the processes that underlie contingency judgments (e.g., it is not assumed that participants actually calculate and compare probabilities). Therefore, they do not generate predictions about what participants will report when they are asked to justify their ratings. One notable exception is the causal model theory of Waldmann (2000). His model is similar to the probabilistic contrast model in that contingency judgments are assumed to reflect the outcome of probabilistic contrasts. The causal model theory focuses, however, on the fact that people have conscious knowledge about the conditions under which the outcome of probabilistic contrasts will provide a good basis for contingency judgments (e.g., whether the cues are causes or effects of the outcome). In more recent papers (e.g., Waldmann, 2000; Waldmann & Walker, 2005), Waldmann has specified that participants use causal models to make inferences about cue–outcome relations. These more recent formulations of the causal model theory are in many ways similar to the inferential reasoning account that was proposed by Lovibond (2003).

Also note that without such additional assumptions about the role of inferential reasoning, probabilistic models are not able to fully account for the rating data. Cheng (1997), for example, proposed a probabilistic model in which uncertainty about the blocked cue T was built in when the outcome is always fully present on A+ and AT+ trials. This model can therefore account for lower (or no) blocking in the no-information condition of Experiment 2. Probabilistic models, however, cannot account for higher blocking effects in the information condition of Experiment 2 because these models are developed to deal with variations in probabilities but not with variations in intensity of the outcome. Furthermore, although probabilistic models may account for the interference of secondary task difficulty on causal judgments, if one makes the additional assumption that these probabilities are actually calculated and compared, it is difficult to see how this interference affects the causal judgment of the blocked cue T but not the causal judgments of the other cues (especially the control cues K and L; see also De Houwer & Beckers, 2003).

There are two additional arguments for the causal ratings being best explained by an inferential reasoning account. First, participants exhibited either strong blocking or no blocking at all. This result can be best explained by assuming that some participants made a blocking inference and others did not. Second, the finding that a large number of participants gave a causal rating of exactly zero for T is much more in line with an inferential reasoning account than it is with an associative account of blocking.

Finally, it should be noted that we do not exclude the possibility that associative processes do play a role in

HCL. The main aim of our experiments was to provide further evidence for the hypothesis that inferential reasoning is (also) an important source of (blocking in) HCL. The fact that both the self-report data and the rating data were in line with the predictions of an inferential reasoning account and the fact that other existing models failed to provide an accurate account of both the self-report and rating data strongly suggest that inferential reasoning is indeed an important component of HCL.

REFERENCES

- AITKEN, M. R. F., LARKIN, M. J. W., & DICKINSON, A. (2001). Re-examination of the role of within-compound associations in the retrospective reevaluation of causal judgements. *Quarterly Journal of Experimental Psychology*, **54B**, 27-51.
- BROWN, N. R. (1995). Estimation strategies and the judgment of event frequency. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **21**, 1539-1553.
- CHAPMAN, G. B. (1991). Trial order affects cue interaction in contingency judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **17**, 837-854.
- CHAPMAN, G. B., & ROBBINS, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, **18**, 537-545.
- 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. (2002). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology*, **55B**, 289-310.
- 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. (2004). 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 reevaluation 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.
- JOHNSON-LAIRD, P. N. (1983). *Mental models*. Cambridge: Cambridge University Press.
- LOVIBOND, P. F. (2003). Causal beliefs and conditioned responses: Retrospective reevaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 97-106.
- NISBETT, R. E., & WILSON, T. D. (1977). Telling more than we can know:

- Verbal reports on mental processes. *Psychological Review*, **84**, 231-259.
- 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.
- RIPS, L. J. (1983). Cognitive processes in propositional reasoning. *Psychological Review*, **90**, 38-71.
- SHANKS, D. R. (1985). Forward and backward blocking in human contingency judgement. *Quarterly Journal of Experimental Psychology*, **37B**, 1-21.
- SHANKS, D. R., & LOPEZ, F. J. (1996). Causal order does not affect cue selection in human associative learning. *Memory & Cognition*, **24**, 511-522.
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- SMITH, E. R., & MILLER, F. D. (1978). Limits on perception of cognitive processes: A reply to Nisbett and Wilson. *Psychological Review*, **85**, 355-362.
- VANDORPE, S., & DE HOUWER, J. (in press-a). A comparison of forward blocking and reduced overshadowing in human causal learning. *Psychonomic Bulletin & Review*.
- VANDORPE, S., & DE HOUWER, J. (in press-b). People want to see information that will help them make valid inferences in human contingency learning. *Memory & Cognition*.
- 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., & 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 human causal learning. *Learning & Behavior*, **33**, 211-229.

NOTES

1. This participant thought that weapon T had once fired on its own and that that event was not followed by an explosion of the tank. When his data were excluded from the analyses, the mean blocking score for the nonreasoning group dropped further, from 8.23 to 2.5.

2. Transcripts of self-reports are available on request.

3. At first sight, the study of Chapman (1991) also seems to provide evidence for forward blocking. However, in this study, participants were not allowed to give two different cues the same causal rating. As a consequence, it is plausible that participants gave a lower rating for the blocked cue T than for the control cues K and L because they could have taken into account the possibility that T was noncausal. The results of the no-information condition of Experiment 2 confirmed this plausibility. Nine participants gave exactly the same causal rating for the blocked cue T as for the control cues K and L.

(Manuscript received December 8, 2003;
revision accepted for publication October 18, 2004.)