

Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude

PETER F. LOVIBOND, SARA-LEE BEEN, and CHRIS J. MITCHELL
University of New South Wales, Sydney, Australia

and

MARK E. BOUTON and RUSSELL FROHARDT
University of Vermont, Burlington, Vermont

When two causes for a given effect are simultaneously presented, it is natural to expect an effect of greater magnitude. However many laboratory tasks preclude such an additivity rule by imposing a ceiling on effect magnitude—for example, by using a binary outcome. Under these conditions, a compound of two causal cues cannot be distinguished from a compound of one causal cue and one noncausal cue. Two experiments tested the effect of additivity on cue competition. Significant but weak forward blocking and no backward blocking were observed in a conventional “allergy” causal judgment task. Explicit pretraining of magnitude additivity produced strong and significant forward and backward blocking. Additivity pretraining was found to be unnecessary for another cue competition effect, release from overshadowing, which does not logically depend on additivity. The results confirm that blocking is constrained when effect magnitude is constrained and provide support for an inferential account of cue competition.

When humans judge the causal relationship between a cue and an effect, their judgments are influenced systematically by the presence of other potential causal cues. Such *cue competition* effects have played an important role in shaping current theories of causal judgment (see Table 1). One of the most commonly studied examples, associative blocking, was first identified in animal conditioning (Kamin, 1969). In a blocking procedure, one cue—say, A—is consistently paired with the effect (A+). In a second phase, a compound of two cues, A and B, is then paired with the effect (AB+). Judgments of the causal strength of the target cue (B) are discounted or *blocked* by the presence of the established causal cue (A), as shown by comparison with an overshadowing control condition in which A is not separately paired with the effect (Larkin, Aitken, & Dickinson, 1998; Shanks, 1985). Some experiments have reversed the order of the two phases and have demonstrated *backward blocking*, in which the causal status of the target cue (B) is reduced retrospectively by subsequent experience with its partner cue (A) alone (e.g., Shanks, 1985; Wasserman & Berglan, 1998). It has been argued that blocking and other cue competition effects are adaptive in that they allow organisms to ignore redundant

cues and focus on the most valid causes of important effects in their environment (e.g., Dickinson, 1980).

A puzzling feature of human causal judgment research, however, is that blocking effects are often modest in magnitude, both in absolute terms and in comparison with animal conditioning experiments. The weakness of blocking is demonstrated most dramatically when its retrospective form, backward blocking, is contrasted with another retrospective revaluation effect, *release from overshadowing*. In both of these procedures, a compound (AB) is initially paired with the effect. Whereas in backward blocking, A is subsequently paired with the effect, in release from overshadowing, A is presented without the effect (see Table 1). The typical outcome is that participants increase their causal ratings for the target cue (B) by a substantial amount. By contrast, in backward blocking, participants typically decrease their causal ratings for B by only a small amount, and this effect often fails to reach significance (e.g., Larkin et al., 1998). Why should such apparently similar procedures yield such asymmetric results?

One possibility is that people believe that there can be two causes of an effect, but not zero causes. In release from overshadowing, if A is not a cause, then B must be. In backward blocking, however, even if A is a cause, B might or might not also be a cause. We suggest that resolution of this causal ambiguity depends on the assumptions made about effect magnitude when multiple causes are present. If a person believes that effects arising from multiple causes are additive, the fact that the effect is no larger on AB trials than on A trials allows the conclusion

This research was supported by Australian Research Council Grant A10007157 to P.F.L. We thank Patricia Cheng, Barbara A. Spellman, and Fred Westbrook for their helpful comments on an earlier version of this article. Correspondence concerning this article should be addressed to P. F. Lovibond, School of Psychology, University of New South Wales, Sydney NSW 2052, Australia (e-mail: p.lovibond@unswedu.au).

Table 1
Cue Competition Designs and Their Typical Outcomes
in Human Causal Judgment Research

Design	Phase 1	Phase 2	Phase 3 (Test)	Causal Strength for B Relative to Control
Forward blocking	A+	AB+	B	modest decrease
Backward blocking	AB+	A+	B	small or no decrease
Release from overshadowing	AB+	A-	B	large increase
Overshadowing control	AB+		B	

Note—A and B refer to cues; AB refers to the simultaneous presentation of A and B; + refers to presence of the effect; - refers to absence of the effect.

that B is not causal (i.e., blocking). On the other hand, if a ceiling is imposed on effect magnitude, so that magnitude is the same regardless of the number of causes present, additivity is precluded, and blocking is not a rational inference. Rather, participants should be uncertain concerning the causal status of B, because it could be either causal or noncausal.

In fact, the design of most laboratory causal judgment tasks imposes just such a ceiling on effect magnitude, because they present the effect as binary—in other words, as either present or absent (see White, 2001). For example, researchers have employed hypothetical causal scenarios in which the effect is the presence or absence of a medical disease (Gluck & Bower, 1988), the presence or absence of a shift in the value of the stock market (Chapman & Robbins, 1990), or the presence or absence of an allergic reaction (Van Hamme & Wasserman, 1994). Although each of these effects could vary along a magnitude dimension in the real world (e.g., a person could have a mild, moderate, or severe form of a disease), in the laboratory the effect has a single indeterminate magnitude whenever it is present. Under these conditions, participants cannot distinguish between a compound of two causal cues and a compound of one causal cue and one noncausal cue, and it may be for this reason that they are reluctant to conclude that the target cue (B) is noncausal in a blocking procedure.¹

A parallel argument to the one advanced here has recently been put forward by Cheng (1997). She points out that when the probability of an effect is 1 in both the presence and the absence of a putative causal cue, as is the case for the added cue B in a typical causal judgment blocking experiment, the causal status of that cue is indeterminate. Only when the probability of the effect is greater on AB trials than on A trials will a cue be judged to have causal strength with respect to the effect. However Cheng's (1997) argument and her mathematical model of causal power are intended to apply to binary variables, which vary only in terms of probability. As was argued earlier, lifting the ceiling effect on magnitude as well as probability can potentially disambiguate the blocking task. In other words, even when the probability of an effect is 1 on both A and AB trials, if the magnitude of the effect is greater on AB trials, it is reasonable to infer that B has causal strength. Note that strict arithmetic additivity is not required to resolve the causal status of the blocked cue; all that is needed is for the magnitude of the effect on AB trials to be larger than the effect on A trials.

Accordingly, we set out to investigate whether establishing an additive rule for effect magnitude would enhance blocking and whether it would do so to the point at which blocking was comparable in strength to release from overshadowing, a cue competition effect that does not logically depend on magnitude additivity. All the participants were exposed to a blocking treatment in a conventional causal judgment task modeled on Larkin et al. (1998), in which they were asked to play the role of an allergist dealing with a hypothetical patient, Mr. X. Each trial consisted of a description of a meal in which the patient ate either one or two foods. After predicting the outcome of the trial, the participants were informed whether the patient experienced an allergic reaction or not.

In a pilot study, we had attempted to establish an additive rule for effect magnitude via instruction. Some participants received standard instructions indicating that the outcome (allergic reaction) was binary and that a compound of two causal cues (two foods Mr. X was allergic to) would produce the same allergic reaction as one causal cue (nonadditive condition). Other participants received instructions indicating that a compound of two causal cues would produce a reaction larger than that produced by a single causal cue (additive condition). In the standard (nonadditive) condition, forward blocking was weak but significant, and backward blocking was nonsignificant. However, the same pattern was observed in the additive group; the instructions designed to enhance blocking had no impact at all.

In the present research, therefore, a different strategy was adopted to try to alter participants' additivity assumptions and a manipulation check was included to test whether this procedure had succeeded. The strategy employed was to provide a direct demonstration of either magnitude additivity or nonadditivity in a pretraining phase, using a separate set of cues, prior to the blocking procedure. Half of the participants received the blocking trials in a forward order, and half in a backward order. Thus, there were four groups in total in a 2×2 factorial design: pretraining (additive or nonadditive) \times order (forward or backward).

EXPERIMENT 1

Method

Participants. The participants were 80 undergraduate students from the University of New South Wales, 55 females and 25 males, who volunteered for the experiment in return for course credit.

Apparatus. Testing was conducted in a laboratory containing 26 personal computers running Filemaker Pro 5 software. The Filemaker program was used to present all instructions and experimental trials and to record predictive and causal ratings via the computer mouse.

Procedure. The general procedure followed that of Larkin et al. (1998). On arrival at the computer laboratory, the participants were told that the experiment was designed to investigate how people judge relationships between events. They were seated at individual computers and were asked to follow the instructions on the screen. The first five screens consisted of instructions about the allergy task. The initial screens included a demonstration of the possible outcomes of each meal: the words “no reaction” in a single box, all in green, or the words “an allergic reaction” in a single box, all in red. Unlike Larkin et al., the outcomes were not accompanied by any auditory stimulus. The participants in the additive condition were additionally exposed to an outcome consisting of the words “a STRONG allergic reaction” in a double box, all in red. The two conditions were further differentiated by the nature of the prediction task. The non-additive participants were asked to predict the outcome of each trial, using the same binary format (no reaction or allergic reaction) as in Larkin et al. The additive participants were instead asked to predict the outcome using a three-choice format: no reaction, allergic reaction, or strong allergic reaction. This procedure was adopted in order to remind the additive participants throughout the experimental trials that three outcomes were possible on each trial.

Table 2 shows the trial sequence for the experiment. The design for the blocking component (element and compound phases) involved eight cues, A–H. These phases were preceded by two new pretraining phases involving four additional cues, I–L. All the phases were presented sequentially without interruption, and the order of trials within phases was random. In the pretraining element phase, the four cues were each presented four times. Two of these cues (I and J) were followed by the allergic reaction, and two (K and L) were followed by no reaction. The foods employed were cheese, ham, pears, and bread, and the assignment of foods to the pretraining cues was counterbalanced across participants. In the pretraining compound phase, three compounds of the pretraining cues were presented: IJ, JK, and KL. For all the participants, JK was followed by the allergic reaction, and KL was followed by no reaction. The additivity manipulation in the experiment involved the IJ compound. For the additive condition, IJ was followed by a strong allergic reaction, whereas in the nonadditive condition, IJ was followed by a regular allergic reaction. Thus, the additive participants were given an explicit demonstration that a compound of two foods, each of which is capable of producing an allergic reaction, will cause a larger allergic reaction. By contrast, the nonadditive participants were given an explicit demonstration that such a compound produced only the same outcome as either food by itself. Thus, in the nonadditive condition, there was no difference between the outcomes of a compound

of two causal cues (IJ) and a compound on one causal and one non-causal cue (JK).

In the element and compound phases, the additive and nonadditive participants received identical trials. There were eight cues (A–H; yogurt, mushrooms, carrots, grapes, walnuts, noodles, oranges, and chicken, counterbalanced). In the element phase, Cue A was followed by the allergic reaction. In the compound phase, Cue A was combined with Cue B and was followed by the allergic reaction. These trials formed the basic blocking procedure, in which A served as the pretrained cue and B as the target (added) cue. CD trials were included in the compound phase as a within-subjects control. They were comparable to the AB trials, except that neither C nor D had been pretrained. Thus, both Cues C and D served as comparison cues for evaluating blocking to B. The remaining trials provided additional comparison cues (E–H) to prevent the participants from learning inappropriate rules, such as “all compound trials are followed by an allergic reaction,” and to maintain continuity over the element and compound phases. In these filler trials, E was always followed by the allergic reaction, whereas F and GH were always followed by no reaction. The participants in the forward condition received the element phase before the compound phase, and the participants in the backward condition received the compound phase before the element phase.

At the end of all 58 trials, the participants were presented with a set of rating tasks. In the first task, the participants were asked to rate the “likelihood that each food would cause an allergic reaction in Mr. X.” In the additive condition, the phrase “of any magnitude” was inserted in parentheses after “allergic reaction.” The eight foods from the blocking procedure (A–H) were listed on one screen in random order, with nine buttons next to each food labeled *definitely not* at the left extreme (0), *possibly* in the middle position (4), and *definitely* at the right extreme (8). The four pretraining cues were rated in the same way on a subsequent screen. In the second task, the participants were asked to nominate which of a set of 11 food pairs were actually fed to Mr. X. Six of the food pairs corresponded to compounds they had been exposed to (IJ, JK, KL, AB, CD, and GH), whereas the other 5 were distractor pairs (e.g., AD and EF). The third task listed the 6 food pairs that were actually presented, and the participants were asked to select the outcome that had followed each compound from two choices (no reaction and allergic reaction) in the case of the nonadditive group or from three choices (no reaction, allergic reaction, and strong allergic reaction) in the case of the additive group.

Finally, an additional task was included as a manipulation check. The participants were informed about four hypothetical foods—referred to only as “A,” “B,” “C,” and “D,”—and were told that A and B caused an allergic reaction in Mr. X, whereas C and D did not. They were then asked to rate the probability of an allergic reaction (in the case of the additive participants, “of any size”) following each

Table 2
Design of Experiment 1

Pretraining Element Phase	Pretraining Compound Phase	Element Phase	Compound Phase
I+ (4)	IJ++ (2)	A+ (6)	AB+ (6)
J+ (4)	JK+ (2)		CD+ (6)
K- (4)	KL- (2)	E+ (3)	E+ (3)
L- (4)	F- (3)		F- (3)
		GH- (3)	GH- (3)

Note—Letters A to L refer to foods; + refers to an allergic reaction, ++ to a strong allergic reaction, and - to no reaction; numbers in parentheses give the number of trials of each type; within each phase, trial types were intermixed. For the forward groups, the element and compound phases were presented in the order shown; for the backward groups, the two phases were reversed. For the additive groups, the pretraining compound trials were presented as shown. For the nonadditive groups, the IJ trials shown in boldface were followed by an allergic reaction (IJ+), rather than by a strong allergic reaction.

of three compounds: AB, AC, and CD. A 9-point rating scale was given, with 0 labeled as *definitely not*, 4 labeled as *possibly*, and 8 labeled as *definitely would*. The participants in the additive condition only were then asked to nominate the strength of the allergic reaction they would expect Mr. X to show on each of the same three compounds, selecting from three choices: no reaction, allergic reaction, or strong allergic reaction.

Scoring and analysis. To evaluate learning of the food–outcome relationships over trials, the percentage of participants in each group who predicted an allergic reaction was calculated for each trial. The primary measure of causal learning was the postexperimental causal rating for each cue. These ratings were analyzed by a set of planned orthogonal contrasts, using a multivariate repeated measures model (O'Brien & Kaiser, 1985).

Results

Trial-by-trial predictions. The participants rapidly learned the relationships between the pretraining stimuli (I, J, K, and L) and their outcomes. On first presentation of the compound of one causal and one noncausal cue

(JK), 95% of the participants in the additive condition and 95% of the participants in the nonadditive condition predicted an allergic reaction. On first presentation of the compound of two causal cues (IJ), 100% of the participants in the nonadditive condition predicted an allergic reaction. In the additive condition, 22.5% of the participants predicted an allergic reaction, and 77.5% predicted a strong allergic reaction.

Figure 1 shows the percentage of participants who predicted an allergic reaction across the two blocking phases. In the case of the additive condition, predictions of allergic reaction and strong allergic reaction were collapsed. It can be seen that the participants rapidly learned and remembered the experimental contingencies. Approximately 80%–90% of the participants predicted the correct outcome for each cue or cue combination by the second trial, and there was little further improvement over trials. The top panel of Figure 1 shows that in the forward blocking




Figure 1. Percentage of participants who predicted an allergic reaction on each trial in the target phases of Experiment 1, collapsed across the pretraining factor. The forward blocking groups are shown in the top panel, and the backward blocking groups in the bottom panel. For the additive groups, the size of the allergic reaction predicted (allergic reaction or strong allergic reaction) was ignored.

conditions, almost all the participants predicted an allergic reaction on the first AB+ trial, demonstrating generalization from the prior A+ training. By contrast, in the backward blocking conditions (bottom panel), there was a slight drop from the final AB+ trial to the first A+ trial, demonstrating some generalization decrement from the compound to the element.

Causal ratings. Figure 2 shows mean causal ratings for the four target cues for the participants in the forward blocking conditions (top panel) and the backward blocking conditions (bottom panel). Within each panel, the data are separated by the pretraining factor (additive vs. non-additive). Averaged across all the participants, there was a significant blocking effect. That is, causal ratings to B were lower than the average of those to C and D [$F(1,76) = 32.2, p < .05$]. However, as is shown in Figure 2, this comparison interacted with both of the group factors. The size of the blocking effect was significantly larger in the forward conditions than in the backward conditions [$F(1,76) = 5.2, p < .05$] and in the additive than in the nonadditive conditions [$F(1,76) = 7.0, p < .05$]. Although the impact of the additivity pretraining appeared to be greater in the backward than in the forward conditions, the triple interaction did not reach significance ($F < 1, p > .05$). Separate analyses indicated that the blocking effect was significant in all the individual groups except the backward nonadditive group [$F(1,19) < 1, p > .05$].

Overall, the participants rated Cue A higher than the other three target cues [$F(1,76) = 655.9, p < .05$], and the size of this difference was slightly greater in the additive conditions [$F(1,76) = 10.3, p < .05$]. Finally, there was a main effect in which the participants in the additive conditions gave lower causal ratings, averaged across all four target cues, by comparison with the participants in the nonadditive conditions [$F(1,76) = 15.1, p < .05$], largely owing to the lower ratings to the blocked cue B. Causal ratings for the filler stimuli E, F, G, and H closely tracked their pairings with allergic reaction or no reaction, and there were no significant interactions involving groups. Causal ratings for the pretraining cues I, J, K, and L also closely followed their experimental contingencies, with ratings for I and J near ceiling and ratings for K and L near floor [$F(1,76) = 1,396.8, p < .05$]. There were no significant interactions involving groups.

Memory for compounds. Averaged across all four groups, 60.0% of the participants correctly identified all three of the pretraining phase compounds from the list of 11, 31.2% identified 2 compounds, 7.5% identified 1 compound, and 1.3% identified none of the compounds. When presented with Compounds JK and KL, 2.5% of the participants nominated the wrong outcome for JK, and 5% for KL. In the case of the critical compound IJ, 92.5% of the participants in the additive condition correctly nominated the strong allergic reaction, and the remaining 7.5% nominated the allergic reaction. For the nonadditive condition, 97.5% correctly nominated the allergic reaction.

Averaged across all four groups, 78.7% of the participants correctly identified all 3 of the compound phase

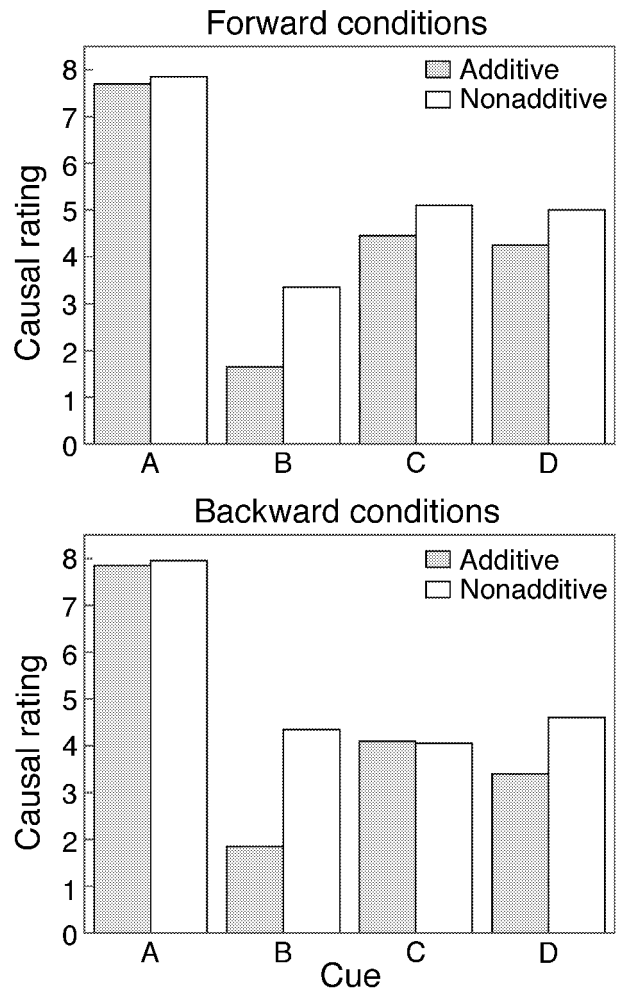


Figure 2. Mean causal ratings for the four target foods in Experiment 1. The forward blocking groups are shown in the top panel, and the backward blocking groups in the bottom panel.

compounds from the list of 11, 20.0% identified 2 compounds, and 1.3% identified only 1 compound. When presented with the three compounds AB, CD, and GH, 5% of the participants nominated the wrong outcome for AB, 11.25% for CD, and 5% for GH. There were no differences in identification or outcome selection across the pretraining or order factors.

The analyses of causal ratings described above were repeated for the participants who were able to identify both AB and CD and who also nominated the correct outcome for each of these compounds ($n = 43$). The pattern of significant results was the same as that in the analysis based on all the participants. However, the overall blocking effect was somewhat stronger in these participants [$F(1,39) = 82.0, p < .05$], and the interaction that assessed the impact of the additivity pretraining on the extent of blocking was also stronger [$F(1,39) = 16.0, p < .05$].

Manipulation check. In this test, the participants were informed that two hypothetical cues (A and B) caused an al-

lergic reaction and that two further cues (C and D) did not cause an allergic reaction in Mr. X. Averaged across the four groups, the mean probability ratings for an allergic reaction (in the case of the additivity participants, "of any size") were 7.9 for the AB compound, 7.1 for AC, and 0.1 for CD. Interestingly, the difference between AB and AC was significant [$F(1,76) = 26.1, p < .05$]. This difference did not interact with the pretraining factor [$F(1,76) = 1.6, p > .05$]; 18% of the additive participants and 28% of the nonadditive participants gave a rating for AD of less than 7. When the additive participants were asked to nominate the strength of the outcome for each compound, 97.5% nominated a strong allergic reaction for AB (1 participant nominated an allergic reaction), 100% nominated an allergic reaction for AC, and 100% nominated no reaction for CD.

Discussion

The results for the nonadditive participants in this experiment were highly similar to those obtained in Larkin et al. (1998). Blocking, as measured by causal ratings for Cue B relative to the control cues C and D, was absent in the backward condition and was present but small in magnitude in the forward condition. This replication is important because it was obtained under conditions of explicit nonadditivity pretraining. That is, nonadditive participants in the present experiment received a direct demonstration that the outcome of a compound comprising one causal and one noncausal cue was identical to the outcome of a compound comprising two causal cues. The similarity of the results to those of Larkin et al. provides support for the claim that in a conventional causal judgment study with a binary outcome, participants assume a lack of magnitude additivity.

In the present experiment, the attempt to manipulate the participants' assumptions about magnitude additivity was successful, suggesting that the strategy of direct demonstration was more powerful than the instructional strategy employed in our pilot study. The postexperimental memory tests and the manipulation check confirmed that the majority of the participants in the additive conditions had learned the additive outcome rule for the pretraining cues and that they could generalize this knowledge to a new set of hypothetical cues. Almost all of the participants in the additive condition correctly predicted that a strong allergic reaction would occur when two causal cues were combined in the manipulation test.

Given that the additivity manipulation appears to have been effective, the critical question is whether this manipulation affected blocking. As can be seen in Figure 2, the additivity manipulation had the predicted effect of markedly increasing the strength of blocking. This effect was most dramatic in the backward conditions, where the nonadditive participants showed no evidence of blocking, whereas the additive participants showed blocking that was as strong as that in the forward additive group. The mean causal rating for the blocked cue B in the additive conditions was 1.75 on the 0–8 scale, and when the analysis was restricted to those participants who could identify AB and

CD and their outcomes, the mean rating for B was 1.11. By comparison, mean ratings for the control cues C and D were greater than 4, indicating that under conditions of additivity, causal ratings to the blocked cue were closer to the floor of the scale than to the control cues.

The additivity manipulation also produced a small decrease in the causal ratings for the control cues C and D, which was reflected in the interaction between pretraining and the comparison between A and the other three cues. This pattern is understandable in terms of the basic additivity manipulation: In the additive conditions only one of C and D could have been causal, whereas in the nonadditive conditions one or both could have been causal. The effect of the additivity training on Cue B was, however, much greater than the effect on C and D. Thus, additivity enhanced the overall blocking effect (difference between B and C/D) despite the reduction in ratings to the control cues.

EXPERIMENT 2

Experiment 1 demonstrated that explicit pretraining of magnitude additivity greatly enhances the strength of the blocking effect. This outcome is consistent with the argument that the causal status of a blocked cue is ambiguous whenever magnitude additivity is precluded—for example, by the use of a binary outcome. If this argument is correct, other cue competition effects that do not logically require magnitude additivity should be observed regardless of whether additivity is permitted or not. One such effect is release from overshadowing, which is identical to backward blocking, except that the partner cue is presented without the outcome rather than with the outcome (AB+, A–). In this design, information that A is noncausal allows an unambiguous inference that B is causal. Experiment 2 was thus designed to test the proposition that release from overshadowing should be unaffected by the same manipulation of participants' assumptions about additivity that was successful in modulating blocking in Experiment 1. The design was identical to the backward conditions of Experiment 1, except that the partner cue A was followed by no reaction, rather than by an allergic reaction. It was expected that this design would lead to an increase in causal ratings to the partner cue B, by comparison with the control cues C and D (i.e., release from overshadowing). There were two groups: additive and nonadditive.

Method

Participants. The participants were 40 undergraduate students from the University of New South Wales, 23 females and 17 males, who volunteered for the experiment in return for course credit.

Apparatus. The experimental software and testing conditions were the same as those in Experiment 1.

Procedure. The design was identical to the backward conditions of Experiment 1, except that the six A+ trials in the element phase were replaced by A– trials.

Results

Trial-by-trial predictions. As in Experiment 1, the participants rapidly learned the relationships between the cues and their outcomes (see Figure 3). In this experiment,




Figure 3. Percentage of participants who predicted an allergic reaction on each trial in the target phases of Experiment 2, collapsed across the pretraining factor. For the additive group, the size of the allergic reaction predicted (allergic reaction or strong allergic reaction) was ignored.

67.5% of the participants predicted an allergic reaction on the first A- trial. This percentage dropped and remained below 10% for the remainder of the A- trials. There were no group differences in predictive ratings across either the pretraining phases or the experimental phases, with the exception of the critical IJ trials in the pretraining compound phase. On the first IJ trial, 100% of the participants in the nonadditive group predicted an allergic reaction. In the additive group, 60% of the participants predicted an allergic reaction, 35% predicted a strong allergic reaction, and 5% (1 participant) predicted no reaction.

Causal ratings. Figure 4 shows mean causal ratings for the four target cues separated by the pretraining factor (additive vs. nonadditive). Reflecting the change in design, causal ratings for Cue A in this experiment were lower than the average of those for the other three cues [$F(1,38) = 121.7, p < .05$], and this difference did not interact with groups ($F < 1, p > .05$). There was a clear release from overshadowing effect, demonstrated by stronger causal ratings to the target cue B than to the control cues C and D, whose partner cue had not been presented alone [$F(1,38) = 45.6, p < .05$]. Furthermore, the strength of the release-from-overshadowing effect did not vary across the two groups ($F < 1, p > .05$). Separate analyses indicated that the release-from-overshadowing effect was significant in both individual groups [additive, $F(1,19) = 20.5, p < .05$; nonadditive, $F(1,19) = 25.1, p < .05$]. Causal ratings for the pretraining cues (I, J, K, and L) and for the filler cues (E, F, G, and H) were highly similar to those obtained in Experiment 1 and did not vary across groups.

Memory for compounds. Identification and memory for outcomes for the pretraining cues I, J, K, and L were similar to those obtained in Experiment 1, with 75% of the participants in the additive group correctly recalling a strong allergic reaction on the critical IJ trials. Averaged

across groups, 60% of the participants correctly identified all 3 of the compound phase compounds from the list of 11, 30% identified 2 compounds, 7.5% identified 1 compound, and 1.5% identified none of the compounds. When presented with the three compounds AB, CD, and GH, 25% of the participants nominated the wrong outcome for AB, 27.5% for CD, and 7.5% for GH. There were no differences in identification or memory for outcomes across the two groups.

The analyses of causal ratings described earlier were repeated for the participants who were able to identify both AB and CD and who also nominated the correct outcome for each of these compounds ($n = 14$). The overall pattern was similar to that shown in Figure 4, and the results of the statistical tests were the same.

Manipulation check. The same pattern of results was found on this test as in Experiment 1. Averaged across groups, the mean probability ratings for an allergic reaction (in the case of the additivity participants, “of any size”) were 8.0 for the AB compound, 6.4 for AC, and 0.1 for CD. The difference between AB and AC was again significant [$F(1,76) = 33.8, p < .05$]. This difference did not interact with groups [$F(1,76) < 1, p > .05$]; 40% of the additive participants and 35% of the nonadditive participants gave ratings for AD of less than 7. When the additive participants were asked to nominate the strength of the outcome for each compound, 95% nominated a strong allergic reaction for AB (1 participant nominated an allergic reaction), 100% nominated an allergic reaction for AC, and 100% nominated no reaction for CD.

Between-experiment analysis. Since both experiments employed the same within-subjects control stimuli (C and D), a direct comparison was possible between the two forms of retrospective reevaluation studied. The additive and nonadditive backward blocking groups from Ex-

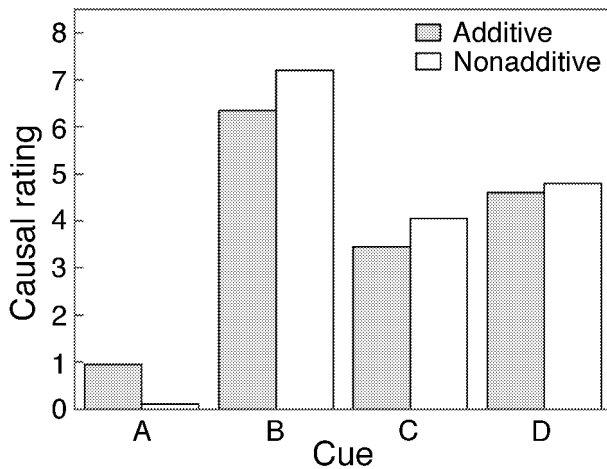


Figure 4. Mean causal ratings for the four target foods in Experiment 2.

periment 1 were analyzed together with the additive and nonadditive release-from-overshadowing groups from Experiment 2. An initial analysis compared ratings for the control cues C and D for the backward blocking groups ($M = 4.04$) and the release-from-overshadowing groups ($M = 4.23$). These ratings did not differ [$F(1,76) < 1, p < .05$], confirming that they provided a stable reference point for the between-experiments comparison. A *revaluation score* was then calculated for all the participants. For the backward blocking participants, this score was calculated by subtracting the rating for B from the mean rating for C and D. For the release-from-overshadowing participants, this score was calculated by subtracting the mean rating for C and D from the rating for B. In both cases, therefore, a positive score indicated revaluation in the expected direction.

An analysis of the revaluation scores for the four groups showed a main effect for experiment [$F(1,76) = 8.70, p < .05$]. This effect indicates that the overall magnitude of release from overshadowing (Experiment 2) was greater than the magnitude of blocking (Experiment 1). The main effect for additivity was not significant [$F(1,76) = 1.8, p > .05$]. Importantly, however, the interaction between experiment and additivity was significant [$F(1,76) = 4.72, p < .05$], indicating that the relative magnitude of the two revaluation effects differed as a function of additivity training. A direct comparison of the two additive groups was nonsignificant ($F < 1, p > .05$), confirming that backward blocking was comparable in magnitude to release from overshadowing under additivity conditions. However, the two nonadditive groups differed significantly [$F(1,38) = 13.1, p < .05$], indicating that backward blocking was weaker than release from overshadowing under nonadditivity conditions.

Discussion

The results of Experiment 2 were very clear: The participants in both groups showed a strong upward revalua-

tion of causal ratings for the target cue B after its partner cue A had been presented without being followed by any allergic reaction. The additivity pretraining manipulation again successfully altered the participants' assumptions about the effects of combining two causal cues. Thus, release from overshadowing does not appear to require magnitude additivity: It is observed with the same strength regardless of whether additivity is permitted or not. This finding suggests that the causal status of the target cue in a release-from-overshadowing procedure is unambiguous, even when magnitude additivity is violated. The between-experiments analysis directly demonstrated that under conditions of nonadditivity, release from overshadowing was stronger in absolute magnitude than was backward blocking, whereas under conditions of additivity the two effects were equivalent in magnitude.

GENERAL DISCUSSION

The present results provide support for the argument put forward in the introduction that the causal status of a blocked cue is ambiguous when magnitude additivity is precluded. Standard laboratory causal judgment procedures prevent additivity by employing a binary outcome, thus placing a ceiling on magnitude. When nonadditivity was explicitly demonstrated in the nonadditive conditions of Experiment 1, forward blocking was weak, and backward blocking was nonexistent. By contrast, when additivity was demonstrated in the additive conditions of Experiment 1, both forward blocking and backward blocking were strong. Under additivity, backward blocking was equivalent in magnitude to release from overshadowing (Experiment 2), resolving the paradox that exists in the literature in which these two highly similar procedures produce very different degrees of cue competition.

Since the present research was conducted, we have become aware of independent research along similar lines by J. De Houwer and colleagues (De Houwer, Beckers, & Gauthier, 2002). These researchers employed a dimensional outcome and found strong forward and backward blocking of causal judgments when the outcomes presented were submaximal (thus allowing additivity), but not when they were at ceiling (thus preventing additivity). Their results therefore parallel those of the present experiments. Other studies that have shown blocking using dimensional outcomes (e.g., Busmeyer, Myung, & McDaniel, 1993) can also be interpreted in terms of allowing magnitude additivity.

How well can available models of cue competition account for the importance of magnitude additivity? Traditional associative models, such as the Rescorla-Wagner model (Rescorla & Wagner, 1972), can account for forward blocking. Variations of this model, such as those proposed by Dickinson and Burke (1996) and Van Hamme and Wasserman (1994), as well as other associative models, such as that of Miller and Matzel (1988), can additionally account for backward blocking and release from overshadowing. The Rescorla-Wagner model and, pre-

sumably, therefore also its variations, has as one of its basic assumptions the notion of additivity of effects. The model assumes that the associative strength of a compound equals the algebraic sum of associative strengths of its elements (Rescorla & Wagner, 1972, p. 76). It could therefore be argued that when a ceiling is imposed on effect magnitude, a basic premise of the model is violated, and cue competition is prevented. However, there are two difficulties that arise in this argument. First, it is not clear how instructions that establish a ceiling on effect magnitude (e.g., concerning the binary nature of the effect) or additivity pretraining of the sort given in Experiment 1 that removes this ceiling make contact with the basic associative mechanism (see, also, Lovibond, 2003). Second, the argument suggests that all cue competition effects would be compromised by violation of magnitude additivity, but as Experiment 2 demonstrated, release from overshadowing is unaffected by this violation.

Statistical models based on comparisons of conditional probabilities have recently been extended to deal with the question of causal ambiguity under ceiling effects (Cheng, 1997). However, as was described earlier, Cheng's (1997) model assumes that effects are binary and considers only the possibility of resolving ambiguity by establishing a probability contrast between the compound (AB) and the element (A) trials. Nonetheless, the present results are logically consistent with Cheng's (1997) argument and can be seen as extending this argument to apply to magnitude as well as to probability.

A final potential approach to cue competition is to view it as a product of a more general inferential reasoning process. According to such an inferential approach, causal judgments may be derived from the same sorts of inferential or deductive processes that participants employ in other complex reasoning tasks (e.g., Johnson-Laird & Byrne, 1991). These deductive processes may be thought to operate on information obtained from the environment through inductive processes (e.g., associative or statistical processes) or the causal models derived from such processes (e.g., Cheng, 1993; Waldmann, 2000). Because this approach assumes that cue competition arises from higher order propositional processes, it can incorporate the impact of instructions and prior additivity training. Language and prior experience can both be thought of as generating propositional knowledge or beliefs that can then serve as inputs to the inferential system (see also Lovibond, 2003; Lovibond & Shanks, 2002). In the case of blocking, an inferential approach would allow for the importance of additivity rules, because when additivity is precluded, participants cannot distinguish between two hypothetical causal scenarios: Either A is causally effective and B is not, or both A and B are causally effective. Conversely, this approach can account for the lack of impact of additivity on release from overshadowing, since there is only one causal account consistent with the evidence regardless of additivity.

One difficulty raised by the present experiments for an inferential approach concerns the forward blocking data.

According to a strict application of the above argument, there should have been no forward blocking at all in the nonadditive condition of Experiment 1. Instead, these participants should have been uncertain about the status of Cue B, because it could have been either causal or non-causal, and thus they should have rated it like C and D. There would seem to be two possible ways to account for the forward blocking data. First, it may be that there are two processes that can lead to blocking: an inferential process that operates in both forward and backward blocking and a second process that operates only in forward blocking—for example, one based on inattention or a failure of associative processing (Kamin, 1969; Rescorla & Wagner, 1972). Alternatively, it may be possible to account for the forward blocking data within an inferential model by appealing to some form of asymmetry resulting from the order of presentation of the information. For example, in forward blocking, participants may discount the possibility that B is causal by virtue of the strongly primed knowledge that A is causal; by contrast, in backward blocking, participants are exposed to the AB compound first and must consider the possibility that B is causal. Separation of these two possible accounts would require a clearer specification both of how associative models give rise to causal judgments and of how an inferential model is to be applied to causal judgment.

In conclusion, the present results confirm that blocking is weak or nonexistent under conditions in which additivity of effect magnitude is precluded. Establishing additivity leads to strong forward and backward blocking. We have recently replicated the impact of additivity on blocking in human autonomic conditioning (Mitchell & Lovibond, 2002), another domain in which blocking has been difficult to obtain (Lovibond, Siddle, & Bond, 1988). Both sets of results, along with those of De Houwer et al. (2002), provide broad support for an inferential approach to cue competition and associative learning.

REFERENCES

- BUSEMEYER, J. R., MYUNG, I. J., & McDANIEL, M. A. (1993). Cue competition effects: Empirical tests of adaptive network learning models. *Psychological Science*, *4*, 190-195.
- CHAPMAN, G. B., & ROBBINS, S. J. (1990). Cue interaction in human contingency judgment. *Memory & Cognition*, *18*, 537-545.
- CHENG, P. W. (1993). Separating causal laws from casual facts: Pressing the limits of statistical relevance. In D. L. Medin (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 30, pp. 215-264). San Diego: Academic Press.
- CHENG, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- DE HOUWER, J., BECKERS, T., & GLAUTIER, S. (2002). Outcome and cue properties modulate blocking. *Quarterly Journal of Experimental Psychology*, *55A*, 965-985.
- DICKINSON, A. (1980). *Contemporary animal learning theory*. Cambridge: Cambridge University Press.
- DICKINSON, A., & BURKE, J. (1996). Within-compound associations mediate the retrospective reevaluation of causality judgments. *Quarterly Journal of Experimental Psychology*, *49B*, 60-80.
- GLUCK, M. A., & BOWER, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, *117*, 227-247.

- JOHNSON-LAIRD, P. N., & BYRNE, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.
- KAMIN, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York: Appleton-Century-Crofts.
- LARKIN, M. J. W., AITKEN, M. R. F., & DICKINSON, A. (1998). Retrospective reevaluation of causal judgments under positive and negative contingencies. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **24**, 1331-1352.
- 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.
- LOVIBOND, P. F., & SHANKS, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, **28**, 3-31.
- LOVIBOND, P. F., SIDDLER, D. A. T., & BOND, N. (1988). Insensitivity to stimulus validity in human Pavlovian conditioning. *Quarterly Journal of Experimental Psychology*, **40B**, 377-410.
- MILLER, R. R., & MATZEL, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 22, pp. 51-92). San Diego: Academic Press.
- MITCHELL, C. J., & LOVIBOND, P. F. (2002). Backward and forward blocking in human autonomic conditioning requires an assumption of outcome additivity. *Quarterly Journal of Experimental Psychology*, **55B**, 311-329.
- O'BRIEN, R. G., & KAISER, M. K. (1985). MANOVA method for analyzing repeated measures designs: An extensive primer. *Psychological Bulletin*, **97**, 316-333.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- SHANKS, D. R. (1985). Forward and backward blocking in human contingency judgments. *Quarterly Journal of Experimental Psychology*, **37B**, 1-21.
- 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.
- 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.
- WASSERMAN, E. A., & BERGLAN, L. R. (1998). Backward blocking and recovery from overshadowing in human causal judgment: The role of within-compound associations. *Quarterly Journal of Experimental Psychology*, **51B**, 121-138.
- WHITE, P. A. (2001). Causal judgments about relations between multi-level variables. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 499-513.

NOTE

1. Consistent with the present argument, animal conditioning studies have shown that increasing the magnitude of the effect (unconditioned stimulus) on the compound AB+ trials has the effect of reducing blocking, referred to in that literature as *unblocking* (Kamin, 1969).

(Manuscript received July 23, 2001;
revision accepted for publication September 4, 2002.)