THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY
2007, 60 (11), 1468–1476

Psychology Press
Taylor & Francis Group

Short article

Blocking of human causal learning involves learned changes in stimulus processing

M. E. Le Pelley, T. Beesley, and M. B. Suret
Cardiff University, Cardiff, UK

Several theories of associative learning propose that blocking reflects changes in the processing devoted to learning about cues. The results of the only direct test of this suggestion in human learning (Kruschke & Blair, 2000) could equally well be explained in terms of, among others, interference in learning or memory. The present study tested this suggestion in a situation in which processing-change and interference accounts predict opposing results. Results support the idea that blocking in human learning can reflect a change in processing of the cues involved.

Blocking refers to the finding that responding to an element of a reinforced stimulus compound is reduced if another element of that compound has previously been established as a predictor of reinforcement. Thus if pairings of A with an outcome are followed by pairings of an AB compound with that same outcome, less responding to B (the blocked cue) is subsequently observed than to a control cue D, trained in a CD compound in which both elements are novel. Blocking is well established in both animal conditioning (e.g., Kamin, 1968) and human learning (e.g., Aitken, Larkin, & Dickinson, 2000).

Several associative learning theories propose that blocking results from changes in the processing afforded to the cues involved in learning, altering their ability to form new associations (their "associability": Kruschke, 2001; Mackintosh, 1975; Pearce & Hall, 1980). These theories propose that, for reasons particular to each model, the associability of B decreases as a consequence of being paired with an outcome in the presence of a established predictor of that outcome, reducing B's ability to develop an association with this outcome. Consistent with this idea, Mackintosh and Turner (1971) demonstrated with rats that a previously blocked cue was slower to enter into novel associations than was a control cue, supporting the idea that its associability had been reduced through blocking treatment. The evidence from human learning is less clear-cut. The only direct evidence comes from a study by Kruschke and Blair (2000), the design of which is shown in Table 1.

Correspondence should be addressed to Mike Le Pelley, School of Psychology, Cardiff University, Tower Building, Park Place, Cardiff, CF10 3AT. UK. E-mail: lepelleyME@cf.ac.uk

This work was supported by Grant RES000230983 from the Economic and Social Research Council to M. E. Le Pelley.

http://www.psypress.com/qjep DOI:10.1080/17470210701515645

Table 1. *Design of Kruschke and Blair's (2000) Experiment 1*

Phase	*Blocking*	*Control to assess attenuation*	*Control to assess blocking*
Training I	A → 1	D → 3	
Training II	AB → 1	D → 3	HI → 6
Test for blocking		BI?	
Training III	A → 1	D → 3	G → 5
	ABC → 2	DEF → 4	GHI → 6
Test for attenuation		BE?	

Note: Letters A to I represent different symptoms; numbers 1 to 6 represent different diseases from which patients could be suffering.

This study used a multiple-outcome disease diagnosis paradigm. On each trial, participants were told the symptom(s) displayed by a patient (e.g., "insomnia") and had to diagnose which of six diseases the patient was suffering from. Letters A–I in Table 1 represent different symptoms that patients might display; numbers 1–6 represent diseases suffered by those patients. After making a diagnosis on each training trial, participants received corrective feedback, allowing learning of the correct diagnosis for each pattern of symptoms. The first column of Table 1 shows a blocking contingency, in which A (during Training I) and AB (during Training II) are followed by the same outcome, Disease 1. This should lead to blocking of Cue B. In the "control for blocking" contingency HI is paired with Disease 6, but as Cue H was not pretrained, Cue I should not be blocked. Thus we might expect the association from B to Disease 1 to be weaker than that from I to Disease 6.

In the test for blocking, sets of symptoms were presented for diagnosis without feedback. One trial presented a combination of B and I. If B were blocked, people should diagnose the disease paired with control symptom I over that paired with blocked symptom B. As predicted, when presented with the BI compound participants made significantly more "Disease 6" than "Disease 1" diagnoses.

During Training III, participants were given information on new symptoms and diseases, including ABC → 2 and DEF → 4. It was argued that, if blocking of B were due to a reduction in its associability, then it should be slower to enter into association with a novel outcome (Disease 2 on ABC → 2 trials) than a control cue that had not been blocked (Cue E from DEF → 4 trials). In a final test, participants were presented with, among others, compound BE. As predicted, participants made significantly more diagnoses of the novel disease paired with the control cue, E, than that paired with the blocked cue, B.

While this finding is consistent with associability-based theories of blocking, alternative interpretations are possible. In particular, the design used by Kruschke and Blair (2000) admits a confound in the different training histories of the blocked (B) and control (E) cues. During Training II, B is paired with Disease 1 on AB → 1 trials. To the extent that B develops any association with Disease 1, we might expect this prior learning to interfere with learning of new information about B during Training III, and/or retrieval of this information during the final test (Underwood, 1957). The control cue, E, being novel at the outset of Training III, would not suffer from proactive interference during learning or retrieval. Hence an appeal to well-established principles of interference in learning or retrieval can account for the BE compound eliciting more responding appropriate to E than to B, without making recourse to associability.

Kruschke and Blair's (2000) Experiment 2 attempted to address problems raised by the different exposure histories of blocked and control cues. One of the "control for attenuation" cues used in Experiment 2 was presented individually during Training II, without corrective feedback. This controlled for one difference between Cues B and E in Experiment 1, namely that E was novel at the outset of Training III while B was not. This additional measure is insufficient, however, as (a) the exposure received by B during Training II (presented as part of a stimulus compound) remained different to that received by the control cue (presented individually), and (b) the absence of corrective feedback for the control cue meant that again this cue, unlike the blocked cue,

Table 2. *Design of Experiment 1*

Contingency	*Stage 1*	*Stage 2*	*Test 1*	*Stage 3*	*Test 2*
Blocking	A → 1	AB → 1	BD?	BK → 3	BF?
	C → 1	CD → 1	FH?	DN → 4	DH?
	E → 2	EF → 2	KN?	FQ → 3	KQ?
	G → 2	GH → 2	QT?	HT → 4	NT?
Control	I → 1	JK → 1		UV → 3	
	L → 1	MN → 1		WX → 4	
	O → 2	PQ → 2		YZ → 3	
	R → 2	ST → 2		$\alpha\beta$ → 4	

Note: Letters A to Z and α to β represent different chemicals; numbers 1 to 4 refer to the type of mutant that was created when these chemicals were mixed with the goo.

would not suffer proactive interference during Training III or on test.[1]

Thus it remains unclear whether the attenuation in learning about a blocked cue observed by Kruschke and Blair (2000) reflects a reduction in processing of that cue, or if it is an artefact of their experimental design (an additional confound is discussed below). Given the importance of blocking as a test-bed for models of cue competition in human learning (e.g., De Houwer & Beckers, 2003) it is important to establish its underlying mechanisms unequivocally. Moreover, the suggestion that common learning mechanisms guide animal and human associative learning would be bolstered by a clear demonstration that blocking in humans, like that in animals (Mackintosh & Turner, 1971), involves a change in processing of the blocked cue.

EXPERIMENT 1

Experiment 1 used a novel "mad scientist" paradigm, in which participants took the role of a scientist specializing in creating mutants. They were told that mutants were created by combining chemicals with a special "goo", and that different chemicals could create different types of mutants. During training, participants received information on the chemicals used on each trial and had to predict what sort of mutant would be created. The design of Experiment 1, which was within subjects, is shown in Table 2. Letters A–Z and α–β represent different chemicals; numbers 1–4 represent different types of mutants that could be created.

Looking at Stages 1 and 2, the first four rows in Table 2 show four blocking contingencies, in which one element of the Stage 2 compound is pretrained as a predictor of the outcome that follows that compound. The lower four rows show four control contingencies, in which neither element of the Stage 2 compound receives pretraining. We might expect that the blocked cues (B, D, F, H) would develop weaker associations to the outcomes with which they were paired in Stage 2 than would the controls (K, N, Q, T). Following Stage 2, participants had to rate the predictive strength of certain compounds for each of Outcomes 1 and 2; these test

[1] Some authors have also raised issues with the medical diagnosis paradigm in studies of cue competition (e.g., Waldmann, 2000). This has the peculiarity that participants diagnose the disease that a patient suffers from on the basis of their symptoms, whereas the causal relationship is reversed: The disease causes the symptoms. Waldmann argues that cue competition, including blocking, is fundamentally different in diagnostic learning (where cues are effects, and outcomes are causes) and predictive learning (where cues are causes, and outcomes are effects). Our view is that associative learning should be expected to be robust under different experimental paradigms and that one should be wary of ascribing unexpected results to such differences. Nevertheless we acknowledge the concerns expressed by others in the field, and hence the present experiment uses a predictive learning paradigm.

compounds are shown in Table 2. Compound BD comprised two blocked cues that were paired with Outcome 1 during Stage 2, while compound KN comprised two control cues that were paired with Outcome 1. If blocking occurred, then BD would be perceived as a weaker predictor of Outcome 1 than was KN. Likewise, compound FH comprised two blocked cues paired with Outcome 2 during Stage 2, while compound QT comprised two control cues paired with Outcome 2. Blocking would be evidenced if FH were seen as a weaker predictor of Outcome 2 than was QT. In general, blocking would be reflected by a greater perceived predictiveness of compounds KN and QT during Test 1 than of BD and FH.

In each of the first four trial types shown for Stage 3, a blocked cue from Stage 2 was compounded with a control cue, and this compound was paired with a novel outcome (Outcome 3 or 4). If the blocked cues had lower associability than the controls, then they would be slower to form associations with these novel outcomes. Test 2 probed the perceived predictiveness of various cue compounds. Compound BF comprised two cues that were blocked in Stage 2 and were paired with Outcome 3 in Stage 3, while compound KQ comprised two controls from Stage 2 that were paired with Outcome 3. If blocking led to a reduction in associability, we would expect BF to be perceived as a weaker predictor of Outcome 3 than was KQ. Likewise, DH comprised two cues that were blocked in Stage 2 and were paired with Outcome 4 in Stage 3, while NT comprised two controls from Stage 2 that were paired with Outcome 4. As above, we might expect that DH would be perceived as a weaker predictor of Outcome 4 than was NT. In general, a role of associability in blocking would be demonstrated by a greater perceived predictiveness of compounds KQ and NT than of BF and DH during Test 2.

Four concurrent blocking contingencies and four control contingencies were used during Stages 1 and 2 to ensure that (a) no outcome was experienced more frequently than any other, and (b) all Stage 2 and Stage 3 outcomes had the same relationship to one another in terms of the cues with which they were paired. Consider the first blocking contingency (A $\rightarrow$ 1, AB $\rightarrow$ 1) and the first control contingency (JK $\rightarrow$ 1). Both B and K were paired with Outcome 1 in Stage 2, and both were paired with Outcome 3 in Stage 3. This treatment might cause the representation of Outcome 1 to become associated with that of Outcome 3 (Hall, Mitchell, Graham, & Lavis, 2003). Given that blocking would render B less able than K to activate the representation of Outcome 1, on this basis alone we might expect B to be perceived as a poorer predictor of Outcome 3 than was K on test. However, use of the other three blocking and control contingencies rules this out. Cues D and N, which were also paired with Outcome 1 during Stage 2, were paired with Outcome 4 during Stage 3. Therefore during Stage 3 the associatively activated representation of Outcome 1 occurred equally often in the presence of Outcomes 3 and 4. Likewise F and Q (paired with Outcome 2 in Stage 2) were paired with Outcome 3 in Stage 3, while H and T (also paired with Outcome 2) were paired with Outcome 4. Hence the associatively activated representation of Outcome 2 also occurred equally often in the presence of Outcomes 3 and 4. Consequently there is no way for these associatively activated representations to influence the development of discrimination between compounds in Stage 3. Kruschke and Blair (2000) did not balance the relationships between outcomes in the different stages of their experiments, such that "associative mediation" could have influenced their results.

The major conceptual difference between this design and that of Kruschke and Blair (2000) is that we used the same control cues in assessing blocking and subsequent attenuation of learning, setting interference at odds with associability. At the outset of Stage 3, the control cues should have stronger associations to the Stage 2 outcomes than do the blocked cues (assessed by the test of blocking in Test 1). Consequently, the control cues would suffer more interference during Stage 3 than would the blocked cues. Interference therefore predicts greater responding to compounds

composed of blocked cues than those composed of controls during Test 2, the opposite of the pattern predicted by associability processes.

The lower four Stage 3 trial types were fillers to ensure that Stage 3 involved the same number of different trial types as, and was of comparable difficulty to, Stages 1 and 2.

Method

Participants, apparatus, and stimuli

A total of 8 Cardiff University undergraduates participated in the experiment in exchange for course credit. Participants were tested individually, using a standard desktop PC.

The 28 chemicals were Ulginate, Renphane, Nelomine, Kluphane, Bizancrine, Alzaze, Quezalin, Xentine, Frestix, Trizopane, Lobinz, Zapotyne, Fazakane, Jintsone, Eframide, Sistax, Kikaran, Ontone, Pukintz, Gratix, Ventox, Cucose, Yestimox, Halorite, Prental, Goladine, Ilomine, and Daktyre. These were randomly and independently assigned to letters A–Z and α–β in the experimental design for each participant. The four mutant names were Jominoid, Draguts, Goygle, and Necromon, which were randomly assigned to Outcomes 1–4 for each participant. Cartoon pictures of four mutant creatures were obtained from the web; these were randomly assigned to mutant names for each participant.

Procedure

Participants initially received on-screen instructions describing the task: that they had been given a newly discovered set of chemicals to experiment with and were to predict which mutant would be created when the chemicals used on each trial were mixed with a blue goo, and that they would have to start out guessing, but that with the aid of feedback their predictions should become more accurate. Figure 1 shows a screenshot of a Stage 1 trial. On each trial participants were given the name of a chemical and were asked what sort of mutant would be created. Below this were pictures of two mutants, with their names. Participants entered predictions by clicking on one of these pictures and then clicking an OK button. Immediate feedback was provided: A blue box highlighted the correct answer. If participants made a correct prediction, the word "Correct" appeared; if they had made an incorrect prediction, the word "Wrong" appeared, and the computer beeped.

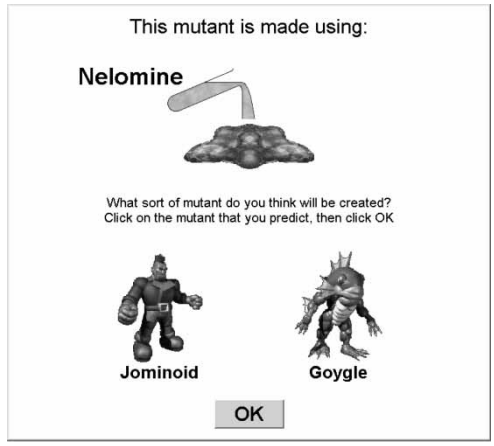

Figure 1. *Screenshot of a typical trial during Stage 1 of Experiment 1. Participants would click on the picture of the mutant that they thought would be created on that trial. The message "Correct" or "Wrong" would then appear in place of the "What sort of mutant do you think will be created?" text, and a blue box would frame the correct mutant picture.*

Stage 1 comprised 14 blocks, with each of the eight trial types occurring once per block. Trial order within a block was randomized, with the constraint that there could be no immediate repetitions across blocks. The two mutants presented on each trial were always Types 1 and 2. For each trial type, the order of presentation of mutants (left/right) was counterbalanced across blocks. So for trial type A → 1, there would be seven presentations with Mutant 1 to the left of Mutant 2 and seven presentations with Mutant 2 to the left of Mutant 1 (the order of these presentations was randomized).

Stage 2 followed immediately from Stage 1 with no break. Trials were as described for Stage 1, the only exception being that two chemicals were shown on each trial, with the name of one on the left of the screen and the other on the right. Ten blocks of Stage 2 trials were presented. For each

trial type the order of presentation of the chemicals (left/right) was counterbalanced across blocks.

Following Stage 2, instructions told participants that as a test of their understanding of the mutant experiments, they would be asked to make decisions for new chemical combinations. For each combination, they were to rate how likely different types of mutants were to be created, on a scale from 0 (chemicals very unlikely to create that type of mutant) to 10 (chemicals very likely to create that type of mutant).

Each of the four Test 1 compounds shown in Table 2 was presented in random order for rating. On each test trial, participants were given the names of two chemicals, pictured being poured onto the blue goo. Below that came the message "How likely is it that the following mutant will be created?", along with a picture and name of one of the Stage 2 mutants (Type 1 or 2). Participants entered their rating by clicking one of 11 radio buttons labelled from 0 to 10, the leftmost being 0 (labelled "Chemicals very unlikely to create this mutant"), and the rightmost 10 ("Chemicals very likely to create this mutant"). Participants rated the ability of a pair of chemicals to create one type of mutant (e.g., Type 1) and on the immediately succeeding trial rated the ability of that same compound to create the other type of Stage 2 mutant (Type 2 in this case). Half of the participants provided ratings for Mutant 1 before ratings for Mutant 2 during Test 1, and the other half provided ratings for Mutant 2 before Mutant 1.

After Test 1 the participants were told that in the next phase of their research they would be using a different, red goo, which created new types of mutants. The form of each Stage 3 trial was the same as that for Stage 2, except that (a) the goo pictured on each trial was red, rather than blue, and (b) the two mutants pictured on each trial were of Types 3 and 4. There were six blocks in Stage 3. Other details were as for Stage 2.

After Stage 3 participants were told that again they would be tested by being asked to make decisions for new chemical combinations and were given a reprise of the instructions relating to the rating scale as used in Test 1. The form of Test 2 was the same as that of Test 1, except that (a) the goo pictured was red, and (b) the two mutants used were of Types 3 and 4. As for Test 1, the order in which ratings were provided for the different types of mutants was counterbalanced across participants.

Results

Figure 2 shows mean percentage correct of participants' predictions during each block of the three training stages. Data have been averaged over equivalent trial types in each stage. Learning is evident throughout.

More important are the perceived predictive strengths of the compounds in the two test phases. Looking first at Test 1, the question of interest was the extent to which participants had learnt the mappings between chemicals and the specific mutants with which they were paired during Stage 2: that is, not how strongly each cue was perceived to predict the creation of mutants in general, but instead how much it was perceived to predict the creation of the mutant with which it had been paired more than the mutant with which it was not paired—that is, the *selective* learning about associations between cues and specific outcomes. The technique used to assess this selective learning was that of Le Pelley and McLaren (2003; see also Le Pelley, Oakeshott, & McLaren, 2005). For each test compound, we took the predictiveness rating for that compound with respect to the outcome with which its constituent elements were paired in Stage 2 and subtracted the rating for the same compound with respect to the outcome with which its elements were not paired in Stage 2, to yield a difference score. Thus the difference score for compound BD was calculated by taking the rating for compound BD with respect to Outcome 1 (as B and D were paired with Outcome 1 in Stage 2) and subtracting from that the rating for BD with respect to Outcome 2 (as B and D were not paired with Outcome 2 in Stage 2). High difference scores (maximum = 10) indicate strong, selective learning, while a score of zero indicates no selective learning. The

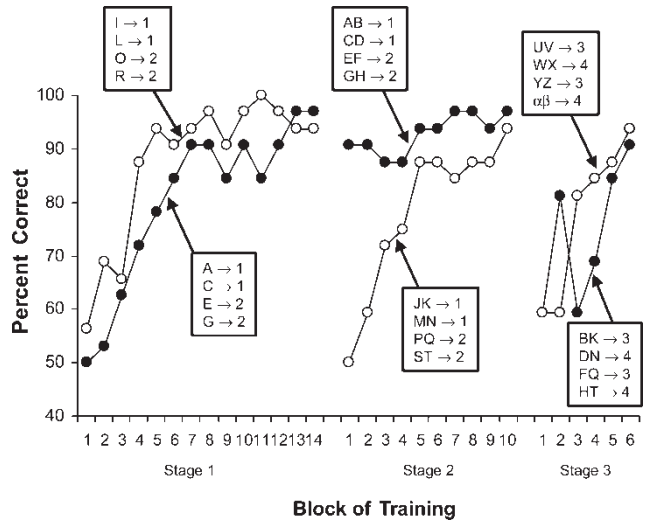

Figure 2. *Percentage of correct responses for the various trial types over the 14 blocks of Stage 1, 10 blocks of Stage 2, and 6 blocks of Stage 3 of Experiment 1. Data are averaged over equivalent trial types in each stage. Chance responding corresponds to 50% correct.*

advantage of using difference scores over raw rating data is that the former are free from influences of generalization that render the latter uninterpretable (see Le Pelley et al., 2005).

Figure 3A shows mean difference scores for the Test 1 compounds. These data were analysed using repeated measures analysis of variance (ANOVA), with factors of type (compound composed of blocked or control cues) and outcome (cues paired with Outcome 1 or Outcome 2 in Stage 2). This revealed a significant effect of type, $F(1, 7) = 17.50$, $p < .01$, with compounds composed of blocked cues yielding lower difference scores than those composed of control cues. The main effect of outcome, and the interaction, was nonsignificant, Fs < 1.

Difference scores were calculated for the Test 2 compounds by taking the predictiveness rating for each compound with respect to the outcome with which its constituent elements were paired in Stage 3 and subtracting from that the rating with respect to the outcome with which its elements were not paired. Figure 3B shows these mean difference scores, which were analysed as above (the outcome factor now discriminates between Outcomes 3 and 4). This revealed a significant effect of type, $F(1, 7) = 7.00$, $p < .05$, with compounds composed of blocked cues yielding lower differences scores than those composed of control cues. The main effect of outcome, and the interaction, was nonsignificant, $F < 1$, and $F(1, 7) = 2.03$, $p = .20$, respectively.

Discussion

The results of our experiment provide unequivocal support for the conclusion drawn by Kruschke and Blair (2000), that blocking involves a reduction in associability of the blocked cue: Blocked cues were significantly slower to form novel associations than were otherwise equivalent control cues, in a situation in which interference processes would, if anything, predict the opposite. This is not to say that a reduction in cue processing is the only source of blocking in human learning: It is possible that observed blocking reflects the operation of several independent mechanisms (e.g., Le Pelley, 2004). Our results simply show that at least one of these involves a change in cue processing.

Our results also address nonassociative, reasoning-based models of learning (e.g., De Houwer & Beckers, 2003), which posit that

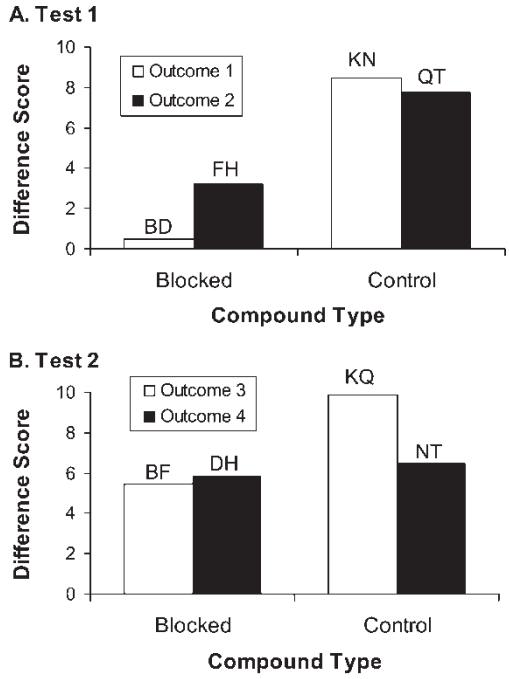

Figure 3. *Mean difference scores for the test compounds of Experiment 1. Panel A shows difference scores for compounds presented in Test 1. Scores are shown separately for compounds made up of blocked cues (BD and FH) and those made up of control cues (KN and QT); scores are also shown separately for compounds made up of cues paired with Outcome 1 in Stage 2 (BD and KN) and those paired with Outcome 2 in Stage 2 (FH and QT). Panel B shows difference scores for compounds presented in Test 2. Scores are shown separately for compounds made up of blocked cues (BF and DH) and those made up of control cues (KQ and NT); scores are also shown separately for compounds made up of cues paired with Outcome 3 in Stage 3 (BF and KQ) and those paired with Outcome 4 in Stage 3 (DH and NT). Columns are labelled with the name of the compound that they represent. Difference scores for each compound were calculated by taking the rating for the outcome with which the component cues of that compound were paired and subtracting from it the rating for the outcome from the same stage of the experiment with which the cues were not paired.*

human causal learning involves the operation of controlled processes of rational inference, with blocking resulting from participants drawing the following inference: "Cues A and B together cause the outcome to occur with the same intensity and probability as does A alone; therefore B is not a cause of the outcome." This relies on participants having access to memories of the different cues and the outcomes with which they were paired in order to draw the appropriate inferences.

Note, however, that during Stage 3 of the current experiment all cues were objectively equally predictive of the outcomes with which they were paired—for example, the relationship between B (a previously blocked cue) and Outcome 3 was exactly the same as that between K (a control cue) and Outcome 3. If participants based their ratings on memories of these cues and the outcomes with which they were paired during Stage 3, they should rate both cues as equally predictive of Outcome 3. The fact that participants did not treat the different cues equally seems to argue against this kind of "rational", inference-based approach.

We acknowledge, however, that an inference-based model could be elaborated to capture our data. For example, the inference process could be influenced by the attention paid to cues, with attention in turn influenced by the cues' predictive ability in much the same way as suggested by associability based associative models. This cognitive account of our data still relies on the operation of cue-processing mechanisms in blocking and so does not undermine the conclusions drawn above. In some sense, it is simply a redescription of the processes at work in the associative models of these effects discussed earlier.

Original manuscript received 13 September 2006
Accepted revision received 16 April 2007
First published online 20 August 2007

REFERENCES

Aitken, M. R. F., Larkin, M. J. W., & Dickinson, A. (2000). Super-learning of causal judgements. *Quarterly Journal of Experimental Psychology, 53B*, 59–81.

De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *Quarterly Journal of Experimental Psychology, 56B*, 345–357.

Hall, G., Mitchell, C. J., Graham, S., & Lavis, Y. (2003). Acquired equivalence and distinctiveness in human discrimination learning: Evidence for

associative mediation. *Journal of Experimental Psychology: General, 132*, 266–276.

Kamin, L. J. (1968). Attention-like processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the Prediction of Behavior: Aversive Stimulation* (pp. 9–32). Coral Gables, FL: University of Miami Press.

Kruschke, J. K. (2001). Towards a unified model of attention in associative learning. *Journal of Mathematical Psychology, 45*, 812–863.

Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review, 7*, 636–645.

Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology, 57B*, 193–243.

Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *Quarterly Journal of Experimental Psychology, 56B*, 68–79.

Le Pelley, M. E., Oakeshott, S. M., & McLaren, I. P. L. (2005). Blocking and unblocking in human causal learning. *Journal of Experimental Psychology: Animal Behavior Processes, 31*, 56–70.

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82*, 276–298.

Mackintosh, N. J., & Turner, C. (1971). Blocking as a function of novelty of CS and predictability of UCS. *Quarterly Journal of Experimental Psychology, 23*, 359–366.

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian conditioning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*, 532–552.

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review, 64*, 49–60.

Waldmann, M. R. (2000). Competition among causes but not effects in predictive and diagnostic learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 53–76.