Journal of Experimental Psychology:
Animal Behavior Processes
2011, Vol. 37, No. 1, 114–120

0097-7403/10/$12.00 DOI: 10.1037/a0019526

BRIEF REPORT

The Influence of Blocking on Overt Attention and Associability in Human Learning

T. Beesley
University College London

M. E. Le Pelley
Cardiff University

Previous studies have demonstrated a retardation in the rate of novel learning about previously blocked cues as compared to appropriate control cues. We report an experiment investigating whether this retardation in novel learning about a blocked cue is accompanied by a reduction in attention to this cue, as anticipated by attentional theories of associative learning. Consistent with these theories, eye gaze measures revealed a reduction in overt attention to the blocked cue both during the compound training phase of the blocking procedure, and also during novel learning with respect to new outcomes. Moreover, the extent of the bias in overt attention away from blocked cues was positively correlated with the subsequent reduction in rate of novel learning about these cues.

Keywords: cue-competition, associability, attention, eye-tracking

Blocking refers to attenuated learning of the contingency between a cue and an outcome, when that cue is trained in compound with a second cue that has previously been established as predictive of the outcome. Thus, if pairings of an AB compound with a particular outcome (AB+) are preceded by pairings of A alone with that same outcome (A+), little is learned about the contingency between cue B and the outcome: learning about B is blocked by prior learning of the A–outcome relationship. Blocking provides an important demonstration that learning about simultaneously presented cues does not proceed independently, but rather, is determined by the competition between them.

Accounts of blocking that are offered by acquisition-based models of associative learning can be broadly divided into two classes. The first sees blocking as resulting from a deficit in processing of the outcome on AB+ trials (e.g., Rescorla & Wagner, 1972). This approach proposes that A+ pretraining renders the outcome unsurprising (and hence unable to support further learning) on AB+ trials, because it is already predicted by A. The second class of models sees blocking as resulting from a deficit in the processing of B on AB+ trials (e.g., Mackintosh, 1975; Pearce & Hall, 1980). Such models are often referred to as "attentional" theories, in that changes in the processing of cues are described as reflecting changes in attention to those cues, and the amount of attention paid to a cue is assumed to modulate how rapidly that cue is learnt about. Mackintosh suggested that the allocation of attention to cues is determined by the relative predictiveness of those cues; that more attention will be devoted to cues that are more accurate predictors of the current outcome, compared to cues that are less accurate predictors (Pearce and Hall's rather different approach will be described later). Applied to a blocking procedure, Mackintosh's account suggests that pretraining A as a good predictor of the outcome ensures that B is a poorer predictor of the outcome on AB+ trials than is A. Consequently, attention to the blocked cue B will be reduced, which will slow learning about the B–outcome relationship on these AB+ trials.

The suggestion that blocking is associated with a decrement in the processing of the blocked cue has received empirical support in both animals (Mackintosh & Turner, 1971), and humans (Le Pelley, Beesley, & Suret, 2007). These studies indicate that the rate at which a blocked cue forms an association with a novel outcome Y is reduced if that cue has previously been blocked with respect to a different outcome X. Such findings add to a growing corpus of empirical evidence demonstrating that the prior predictiveness of a cue can influence the rate of novel learning about that cue (see Le Pelley, 2004, for a review). However, the stronger claim—that these changes in learning rate reflect changes in *attention* to the cues—receives only indirect support, since these experiments do not measure attention directly. Instead, they use rate of novel learning as a proxy for attention, under the assumption that strongly attended cues will be learnt about more rapidly than weakly attended cues. This approach leaves open the possibility that what is influenced by predictiveness, and what in turn influences novel learning, is not attention but rather a learning rate parameter that modulates the readiness with which a cue will enter into associations—commonly referred to as the *associability* of a cue. Indeed, Honey, Close, and Lin (in press) describe a model that

This article was published Online First August 16, 2010.

T. Beesley, Division of Psychology and Language Sciences, University College London; and M. E. Le Pelley, School of Psychology, Cardiff University.

This work was supported by Grant RES000230983 from the Economic and Social Research Council awarded to Michael E. Le Pelley. We thank Amanda Li for her help in running the experiments and John Patrick and the Communication Research Centre in Cardiff for the use of the eye tracking equipment. We also thank Mark Haselgrove and Andy Wills for helpful comments on an earlier draft.

Correspondence should be addressed to Tom Beesley, Division of Psychology and Language Sciences, University College London, 26 Bedford Way, London WC1H 0AP, United Kingdom. E-mail: t.beesley@ucl.ac.uk

accounts for many findings regarding the relationship between predictiveness and learning rate in terms of stimulus representation, without making recourse to changes in attention.

We can decide between these alternatives by taking a more direct measure of attention during learning. Perhaps the most obvious feature of visual attention is that it tends to coincide with where we are looking. It is, of course, possible to make covert shifts of attention that are not accompanied by eye movements (Posner, 1980). Nevertheless, eye movements and attentional shifts are generally tightly coupled (Deubel & Schneider, 1996), especially when dealing with relatively complex stimuli such as words, which constitute the stimuli in many human associative learning experiments.

Two prior studies have demonstrated a reduction in overt attention (as measured by eye gaze) to the blocked cue on the AB+ trials of an A+, AB+ blocking procedure, relative to an appropriate control cue (Kruschke, Kappenman, & Hetrick, 2005; Wills, Lavric, Croft, & Hodgson, 2007). However, a clear further prediction of attentional theories of learning is that the retardation in novel learning about a blocked cue *following* a blocking procedure will also be accompanied by a reduction in attention to that cue. In contrast, if no decrement in attention to B were observed during novel learning after blocking training, we would have to posit a nonattentional source for the observed reduction in novel learning about B during this phase. This issue provides the focus for the current experiment.

Participants played the role of a physician and were asked to predict the side effects that would occur in patients who had taken pills containing fictitious chemicals. The design of the experiment was based on Le Pelley et al. (2007) and is shown in Table 1. Letters A–T refer to the names of chemicals contained in the pills (cues), while o1-o4 refer to side effects that occurred as a result of taking those pills (outcomes). For example, "AB → o1" indicates that a pill contained chemicals A and B, and that this pill produced side effect o1. Participants' eye gaze was monitored throughout the experiment using an eye tracker.

With regard to stages 1 and 2, the top four rows of Table 1 instantiate four replications of a blocking contingency: in each case, one element of the stage 2 compound was pretrained alone in stage 1 as predictive of the outcome with which the compound was paired. The lower four rows show control contingencies in which neither cue of the stage 2 compound received pretraining, and hence in which blocking would not occur. During stage 2 blocked cues (B, D, F, H) should develop weaker associations to their respective stage 2 outcomes than the equivalent control cues (J, K, M, N, P, Q, S, T). Indeed, we have demonstrated in prior research that exactly this procedure yields the standard blocking effect, with the blocked cues rated as significantly weaker causes of the stage 2 outcomes with which they are paired (Le Pelley et al., 2007).[1]

Table 1
Design of the Experiment

Stage 1	Stage 2	Stage 3	Test
A → o1	AB → o1	BK → o3	BT?
C → o1	CD → o1	DN → o4	DQ?
E → o2	EF → o2	FQ → o3	FN?
G → o2	GH → o2	HT → o4	HK?
I → o1	JK → o1		BN?
L → o1	MN → o1		DK?
O → o2	PQ → o2		FT?
R → o2	ST → o2		HQ?

Note. Letters A to T represent different chemicals contained in a pill; o1 to o4 refer to the type of side effect that a fictitious patient suffered after taking the pill. Blocking contingencies are shown above the dotted line in stages 1 and 2; control contingencies are shown below the line. Stage 3 also involved filler trial types UV → o3, WX → o4, YZ → o3, and αβ → o4, not shown in the table.

On the basis of prior research (Kruschke et al., 2005; Wills et al., 2007), we expected to observe an unequal distribution of eye gaze on the stage 2 trials of these blocking contingencies. Specifically, overt attention to blocked cues should be reduced relative to pretrained cues from the blocking contingencies, or cues from the control contingencies.

Each of the stage 3 compounds shown in Table 1 contains one cue that was blocked during stage 2 (e.g., B) and one "nonblocked" control cue from stage 2 (e.g., K). Note that all stage 3 compounds were paired with new outcomes, and in such a way that there was no consistent mapping between the cue—outcome pairings in stage 2 and those in stage 3. For instance, while cues B and K were paired with o1 in stage 2 and o3 in stage 3, cues D and N, which were also paired with o1 in stage 2, were paired with o4 in stage 3. This statistical independence of the outcomes in stages 2 and 3 ensures that any general rule applying contingency information from stage 2 to stage 3 (e.g., "cues that predicted outcome 1 in stage 2 now predict outcome 3 in stage 3") would lead to chance performance in stage 3.

For the sake of brevity, we will continue to reference the prior predictive status of cues when distinguishing between elements of stage 3 compounds. For example, we will describe the stage 3 compounds BK as containing a "blocked" cue B and a "control" cue K, but it should be remembered that the terms "blocked" and "control" are used with respect to the *different* conditions under which these cues were trained in stage 2; these cues were trained under *equivalent* conditions in stage 3.

If blocking leads to a reduction in the associability of blocked cues (B, D, F, H) relative to control cues (K, N, Q, T), then novel learning about control cues should proceed more rapidly than learning about blocked cues during stage 3. Each test compound in the current study (see Table 1) comprised one blocked cue (e.g., B) and one control cue (e.g., T). Importantly, for each test compound these cues had been paired with different outcomes during stage 3; for example, B was paired with o3 in stage 3, while T was paired with o4. On each test trial, participants selected which of o3 or o4 they believed would be produced by the compound in a forced-choice test, and then provided a rating of their confidence in that decision. With regard to compound BT, if participants acquired a stronger association between T and o4 than between B and o3

[1] We did not test for blocking of participants' causal judgments in the current experiments, for two main reasons, (1) to maintain a smooth progression through the training phases, without having to introduce a new test task between stages 2 and 3; and (2) to avoid participants having to consciously and explicitly evaluate and rate the causal status of blocked and control cues, which might influence the attention that they subsequently paid to those cues during stage 3 (e.g., the simple act of explicitly giving a cue a low causal rating might encourage participants to ignore that cue in future).

during stage 3, then they should select o4 as the outcome most likely to be produced by BT. More generally, if participants learnt less in stage 3 about blocked cues than control cues, then on each test trial we would expect them to select the outcome with which the control cue had been paired in stage 3 (so they should select o4 for test compounds BT, FN, BN, and FT, and o3 for test compounds DQ, HK, DK, and HQ). Crucially, attentional theories of learning predict that any learning advantage for control cues over blocked cues during stage 3 will be accompanied by an attentional bias toward control cues, with participants looking for longer at control cues than blocked cues.

Stage 3 training also involved four "filler compounds" comprising novel cues (trial types UV → o3, WX → o4, YZ → o3, and αβ → o4). These were included to ensure that stage 3 involved the same number of different trial types as, and was of comparable difficulty to, stages 1 and 2, and will not be discussed further.

Method

Participants, Apparatus, and Stimuli

Thirty-one Cardiff University undergraduates participated in exchange for payment or course credit. Eye gaze was measured using a Tobii 1750 Eye Tracker (Tobii Technology); a 43.2 cm monitor with a monitor-mounted eye tracker, recording gaze location every 20 ms. This nonintrusive system can compensate for small head movements. Participants sat ~60 cm from the screen. Stimulus presentation was controlled by a Visual Basic program, with timing determined by Windows API functions QueryPerformanceCounter and QueryPerformanceFrequency for millisecond resolution. Error signals were given over speakers.

Names of 28 fictitious chemicals (e.g., Addexium, Rezaline) were randomly assigned to letters A-β in the experimental design for each participant. Two different drawings of men were used to represent Mr. X and Mr. Y, and were edited to depict the four side effect outcomes; dizziness, itchiness, headache, and nausea. These were randomly assigned to outcomes o1-o4 for each participant.

Procedure

Participants received onscreen instructions telling them that their task was to predict which of two side effects a fictitious patient, Mr. X, would suffer after taking a vitamin pill that contained certain chemicals; that they would start out guessing, but that with the aid of feedback their predictions should become more accurate. The eye tracker was calibrated using a 9-point procedure at the start of the experiment, with additional calibrations conducted before the ninth block of stage 1 and the first block of stage 3.

Each training trial displayed the name of one (in stage 1) or two (in stages 2 and 3) chemicals, each encased in a white rectangle measuring 9.5 × 6.75 cm (visual angle ≈ 9° × 6.4°). During stage 1, the single chemical was presented in the horizontal center of the screen; during stages 2 and 3, the two chemicals were presented side-by-side, separated by 10.2 cm (≈ 9.7°). Participants made their prediction as to which of two side effects was likely to occur by clicking on one of two side effect pictures (which were shown side-by-side at the bottom of the screen). If participants made a correct prediction, the word "Correct" appeared; if they made an incorrect prediction, the word "Wrong" appeared and the computer beeped. A blue box also framed the correct answer.

Stage 2 followed immediately from stage 1 with no break. During stages 1 and 2, outcome o1 was always presented on the left of the screen and o2 was always presented on the right. After stage II, participants received instructions concerning a new patient, Mr. Y, who suffered from different side effects to Mr. X. Pictures on each stage 3 trial showed Mr. Y suffering from outcomes o3 (on the left of the screen) and o4 (on the right).

Stages 1, 2, and 3 consisted of 10, 8, and 6 blocks, respectively, with each of the 8 trial types of each stage occurring once per block. Trial order within a block was randomized, with the constraint that there could be no immediate repetitions across blocks. For each trial type during stages 2 and 3, the left/right order of presentation of the chemicals was counterbalanced across blocks. For example, for trial type AB → o1 in stage 2, there would be four presentations with cue A to the left of cue B, and four presentations with B to the left of A (the order of these presentations was randomized).

After stage 3, participants were told that their knowledge regarding Mr. Y would be tested. On each test trial, the names of two chemicals were presented. Participants first selected which outcome (o3 or o4) they thought would occur by clicking on the appropriate picture. A confidence scale then appeared and participants provided a rating of their confidence in that judgment on a scale from 1 (labeled "Not confident at all") to 10 ("Very confident"). Each of the eight test trial types shown in Table 1 was shown once, in random order. No feedback was provided during the test phase.

Results

The eye tracker monitors the gaze location of each eye independently. However, it occasionally fails to register gaze location for one or both eyes (e.g., as a result of blinks or head movements), resulting in missing gaze data. For each participant, the proportion of missing data for each eye was calculated across stages 2 and 3, and gaze data from the single eye with less missing data were used for all further analyses.

Despite several attempts, the eye tracker could not be calibrated for two participants. For three of the remaining 29 participants, the proportion of missing data for the "better" eye was greater than 20% and so these participants were excluded from further analysis. The mean percentage of missing data for the better eye for the final set of 26 participants was 8.25% ($SD = 4.72$).

Figure 1 shows mean percent correct of participants' outcome predictions during each of the three training stages. Data have been averaged over equivalent trial types in each stage. Learning is evident in each stage, in that the percentage of correct predictions rises steadily.

Of more interest are the data relating to participants' choices during the test phase. To provide a continuous dependent variable, for each test trial the outcome choice and confidence rating were combined into a single score. Recall that each test compound contained a blocked cue and a control cue that had been paired with different outcomes during stage 3. If a participant chose the outcome with which the blocked cue had been paired in stage 3, the confidence rating was scored as a negative value; if they chose the outcome with which the control cue had been paired, the

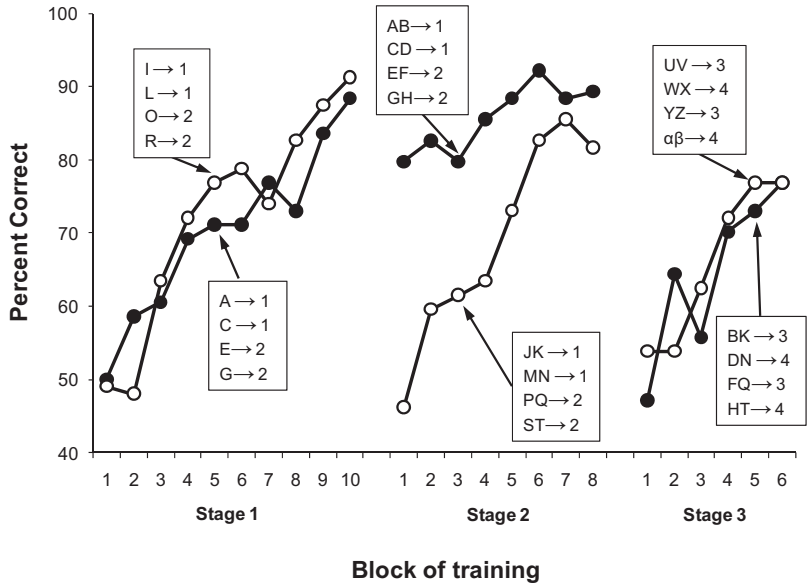

Figure 1. Percentage of correct responses for the various trial types over the 10 blocks of stage 1, 8 blocks of stage 2, and 6 blocks of stage 3. Data are averaged over equivalent trial types in each stage. Chance responding corresponds to 50% correct.

confidence rating was scored as a positive value. Consequently, if participants had learnt equally about blocked cues and control cues during stage 3, we would expect the mean confidence score across all compounds to equal zero, indicating no systematic tendency to be guided by the blocked or control cues. In fact, the mean confidence score across all test compounds was +2.03, and a one-sample test revealed that this score was significantly different from zero, $t(25) = 2.74$, $p < .05$, indicating a reliable tendency toward selection of the outcome that was paired with the control cue, rather than the blocked cue. Consistent with previous research (Le Pelley et al., 2007), this indicates that participants had learnt less about blocked cues than about control cues; the implication is that learning about the novel cue—outcome relationships in effect during stage 3 proceeded more slowly for blocked cues than for control cues.

Dwell time on a cue was defined as a recording of gaze within the white rectangle surrounding the cue name, and was summed across the period from cue onset to response, before being natural log transformed. Trials were excluded from the analysis if response latencies were longer than 10 s (0.5% of trials in stage 2; 0.2% of trials in stage 3). Figure 2A shows mean dwell times on cues across stage 2 training, averaged across equivalent pretrained cues (A, C, E, G), blocked cues (B, D, F, H), and control cues (J, K, M, N, P, Q, S, T). Overall, dwell time decreased over the course of training. Presumably, this reflects participants' increased familiarity with the task in general, and learning of the specific cue—outcome relationships, allowing them to make accurate predictions more and more rapidly (and thus decreasing the window over which dwell time is summed).

Figure 2A indicates that dwell times on pretrained cues were greater than on blocked cues across stage 2 training. Dwell times on control cues were initially similar to those for pretrained cues, but declined to an intermediate level across stage 2. These data were subjected to repeated measures analysis of variance (ANOVA) with factors of cue type (pretrained vs. blocked vs. control) and training block. This revealed a significant effect of cue type, $F(2, 50) = 39.68$, $p < .001$, a significant effect of training block, $F(7, 175) = 16.07$, $p < .001$, and a significant interaction, $F(14, 350) = 2.42$, $p < .01$. Analysis of simple effects revealed that dwell time on blocked cues was significantly less than on pretrained cues, $F(1, 25) = 58.78$, $p < .001$, and on control cues, $F(1, 25) = 27.14$, $p < .001$.

However, mean response latencies on stage 2 trials of the blocking contingencies (1456 ms) were significantly shorter than for the control contingencies (1590 ms), $t(25) = 2.52$, $p < .05$, and this response latency determines the window over which dwell time is summed. A comparison of the absolute dwell times for blocked and control cues is therefore not strictly valid, as the baseline differs in each case. Consequently, following Wills et al. (2007) we also analyzed these data in terms of the proportion of the total dwell time (e.g., proportional dwell time on cue A = dwell time on cue A / summed dwell time on cues A and B). Averaging over stage 2 blocks for compound trials of the blocking contingencies, mean proportional dwell time on the pretrained cue was .53. A one-sample *t* test revealed that this was significantly greater than .5, $t(25) = 6.19$, $p < .001$ (and hence the mean proportional dwell time on blocked cues was significantly below .5).[2]

Figure 2B shows mean dwell time during stage 3 for blocked cues (B, D, F, H) and control cues (K, N, Q, T). Repeated measures ANOVA with factors of cue type (blocked vs. control) and training block revealed a main effect of cue type, $F(1, 25) = 6.52$, $p < .05$, which indicates that dwell times on blocked cues were signifi-

[2] Note that the use of log transformed dwell times leads to smaller proportional differences than in the nontransformed data, for which proportional dwell time on pretrained cues was .62.

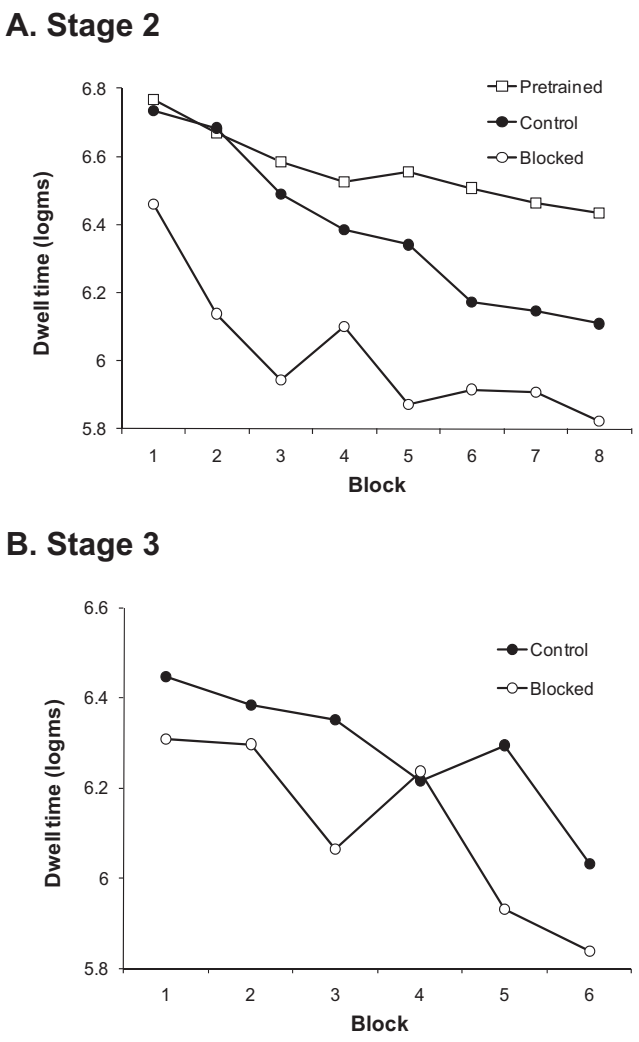

Figure 2. Dwell times (in logms) for pretrained, blocked, and control cues across stage 2 (A) and stage 3 (B) training. Dwell time on a cue is calculated as the log transform of the summed duration on each trial for which eye gaze falls on that cue.

cantly shorter than on control cues (there is no issue with differences in response latency here, since blocked and control cues are presented on the same trial and hence the baseline is identical for each). The main effect of block was also significant, $F(5, 125) = 4.99$, $p < .001$, indicating that dwell time on cues decreased across stage 3. The cue type × block interaction was not significant, $F(5, 125) = 1.40$, $p = .23$.

Recall that the primary aim of this experiment was to examine whether the influence of prior predictiveness on novel cue—outcome learning reflects a change in attention to cues, as suggested by attentional theories of learning. It seems plausible that there will be individual differences in the extent to which participants distribute attention unequally between the elements of a stimulus compound during stage 3 (e.g., see Kruschke et al., 2005; Wills et al., 2007). If this is the case, then attentional theories anticipate that there should be corresponding differences in the extent to which participants learn about these cues at different rates. Specifically, those participants showing the greatest advantage in overt attention for control cues over blocked cues should also show the greatest advantage in rate of learning for control cues over blocked cues.

For each participant, we calculated: (1) the predictiveness effect observed on our attentional measure (given by mean dwell time on control cues minus mean dwell time on blocked cues, collapsed across stage 3 blocks); and (2) the predictiveness effect observed on our learning measure (given by mean confidence score across all test compounds). On each of these measures, a higher positive score indicates a greater effect of predictiveness, with less attention to, or learning about, blocked cues than control cues. There was a significant positive correlation between these two measures, Spearman's $r_s(26) = .58$, $p < .01$, indicating that those participants showing the greatest bias in attention toward control cues during stage 3, also demonstrated the strongest preference for outcomes associated with these control cues in their confidence ratings at test.

Discussion

Eye tracking was used to examine the impact of blocking on overt attention during a human contingency learning task. Consistent with prior findings reported by Kruschke et al. (2005) and Wills et al. (2007), we found that participants spent less time looking at blocked cues than either pretrained cues or control cues during the compound training phase of a blocking procedure (stage 2). Consistent with Le Pelley et al. (2007), we also found a retardation in the rate of novel learning about blocked cues relative to control cues during stage 3. The crucial finding of the current study is that this retardation in rate of novel learning was accompanied by a reduction in overt attention to blocked cues as compared to control cues. Moreover, those participants showing a larger attentional bias during stage 3 tended to show a larger bias in learning. These findings are clearly consistent with the central premise of attentional theories of learning (e.g., Mackintosh, 1975; Pearce & Hall, 1980), that attention to a cue is modulated by the perceived predictiveness of that cue, and that differences in attention can give rise to differences in rate of learning.

Notably, in the prior studies by Kruschke et al. (2005) and Wills et al. (2007), the blocked cue was paired with the same outcome on every trial, and eye gaze to this cue was assessed while it was paired with that outcome. This raises the possibility that any resulting attentional bias is outcome-specific, that is, conditional on the presence of that particular outcome. In the current experiment, the blocking treatment during stage 2 had an influence on the distribution of attention when the critical cues were paired with novel outcomes, during stage 3. Hence, our study demonstrates that the learnt attentional distribution cannot be entirely outcome-specific (see also Le Pelley, Oakeshott, Wills, & McLaren, 2005). Moreover, these prior studies demonstrated attentional biases in the case in which the objective cue—outcome contingency differed for blocked and control cues. For example, Cheng and Novick (1990) noted that the statistical contingency between the blocked cue B and the outcome (computed over the "focal set" of trials on which the competing cue A is present) is reduced by prior training of A. In contrast, the current study demonstrates a bias in overt attention during stage 3, when the objective statistical cue—outcome contingency for blocked and control cues is identical.

It was noted above that the current results are consistent with the predictions of attentional theories of learning. Up to this point we have largely considered such theories as a single generic class of model, but in fact different attentional theories take rather different views of the relationship between attention and predictiveness.

One general class of attentional theory is that exemplified by the model of Mackintosh (1975), wherein the attention to a cue decreases if it is a poorer predictor of the current outcome than are other presented cues. Consider the A → o1, AB → o1 contingency of the current design. On AB → o1 trials, the pretrained cue A is an established predictor of o1, and hence the added, but redundant, cue B will undergo a reduction in attention. In contrast, the stage 2 control compound JK → o1 contains two novel cues, neither of which is more predictive than the other; consequently neither J nor K will undergo the rapid decrease in attention that applies to B. As a result, following stage 2, attention to K will be greater than that to B. If this unequal attentional distribution persists to stage 3, then attention to B (on trial BK → o3) will remain lower than that to K (as indicated by our eye gaze data), and learning about the B–o3 relationship in stage 3 will proceed more slowly than for the K–o3 relationship (as indicated by our causal judgment data).

An alternative attentional account was proposed by Pearce and Hall (1980), who suggested that more attention will be devoted to stimuli that are followed by surprising outcomes, than to stimuli that are followed by well-predicted and hence unsurprising outcomes. During initial A → o1 training, outcome o1 should come to be well predicted by the presence of A, and hence attention to A should decline. The outcome occurring on subsequent AB → o1 trials is also well predicted by the presence of A, and hence the model predicts a rapid reduction in attention to B on these trials.[3] The outcome occurring on JK → o1 trials, however, is not predicted at the outset of stage 2 since J and K are novel; while attention to J and K will decline as stage 2 learning proceeds, this decline will not be as great as for the blocked cue B. As a result, the Pearce–Hall model correctly predicts that attention to the control cue K will be greater than to the blocked cue B at the start of stage 3 training, and hence that learning about K will be more rapid than learning about B during stage 3.

In fact Pearce and Hall's (1980) original model predicts equal attention to blocked and control cues in all blocks of stage 3 apart from the first. Figure 2B indicates that the difference between blocked and control cues persists beyond the first block of stage 3; taking the data for blocks 2–6 only, and collapsing across these blocks, still yields an attentional advantage for control cues, $t(25) = 8.34$, $p < .05$. This rules out Pearce and Hall's original formulation. Pearce, Kaye, and Hall's (1982) refinement of the model predicts that the attentional advantage for control cues will persist across stage 3 training, although the difference between the cues should diminish as training proceeds. It is somewhat more difficult to assess whether the size of the attentional advantage for control cues diminishes as stage 3 training proceeds, since the decrease in overall dwell time across training means that we would have to rely on comparisons at different points on the performance scale. Consequently, Pearce et al.'s refinement remains a viable account of our findings.

So both the Mackintosh (1975) and Pearce et al. (1982) models can account for the critical stage 3 findings of the current study, despite taking very different views of the relationship between predictiveness and attention. It is worth noting that a previous study by Hogarth et al. (2006), which also monitored eye gaze during a human associative learning task (but not a blocking task) found results that seemed more consistent with Pearce et al.'s approach than with Mackintosh's model. That said, the results of certain prior studies of associability effects in humans (e.g., Le Pelley & McLaren, 2003) conflict with Pearce et al.'s model, but follow naturally from Mackintosh's theory. This is not the place to try to resolve these discrepancies; it is sufficient to note that this is currently an unsettled question, which future research will aim to address.

More generally, while our findings are consistent with certain attentional theories of learning, a degree of caution is required. The stage 3 data assess a correlation between overt attention and learning, and hence allow us to state that the observed bias in learning is *associated with* a bias in attention, but they do not allow us to assess the stronger claim of attentional theories, that biases in learning can be *caused by* biases in attention. To illustrate this point, consider an alternative account in which the associability of a cue is determined by its relative predictiveness, but in which associability is not determined by attention. The higher associability of control cues will ensure faster learning about these cues than about blocked cues during stage 3. We need only assume then that participants pay more attention to those cues that they have learned more about, to explain the overall bias in eye gaze toward control cues during stage 3 training.

This problem of establishing causality is inherent in any attempt to test the influence of attention on learning using a correlative technique. To be sure of a truly causal connection, an experimental approach is required. Suppose that there exists a manipulation that will reduce the extent to which a person is able to use selective attention. If this manipulation also reduces the influence of blocking on novel learning, this would support the suggestion that the influence of blocking on learning is via attention. Identification and testing of such manipulations remains a task for future research (see Le Pelley, in press).

[3] Pearce and Hall's (1980) original model predicts equal attention to A and B on all stage 2 trials but the first. Pearce, Kaye, and Hall's (1982) later refinement of the model predicts greater attention to B than to A throughout stage 2. Neither version of the model is able to predict that attention to B will be less than to A, and yet this is the result observed empirically in stage 2 (see Figure 2B). However, the difference in predictiveness between A and B is confounded with a difference in familiarity, since A has been experienced more times than has B at any point during stage 2. Hence, it would be possible to appeal to differences in the familiarity of cues as an additional parameter in the model to account for the stage 2 deficit in attention to B relative to A.

References

Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction. *Journal of Personality and Social Psychology, 58,* 545–567.

Deubel, H., & Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision Research, 36,* 1827–1837.

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: The role of uncertainty. *Quarterly Journal of Experimental Psychology, 61,* 1658–1668.

Honey, R. C., Close, J., & Lin, E. (in press). Acquired distinctiveness and equivalence: A synthesis. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and Associative Learning: From Brain to Behaviour.* Oxford: Oxford University Press.

Kruschke, J. K., Kappenman, E. S., & Hetrick, W. P. (2005). Eye gaze and individual differences consistent with learned attention in associative blocking and highlighting. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 830–845.

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. (in press). Attention and human associative learning. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and learning.* Oxford: Oxford University Press.

Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *Quarterly Journal of Experimental Psychology, 60,* 1468–1476.

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., Wills, A. J., & McLaren, I. P. L. (2005). The outcome-specificity of learned predictiveness effects: Parallels between human causal learning and animal conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 31,* 226–236.

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.

Pearce, J. M., Kaye, H., & Hall, G. (1982). Predictive accuracy and stimulus associability. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Acquisition* (Vol. 3). Cambridge, MA: Ballinger.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

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.

Wills, A. J., Lavric, A., Croft, G. S., & Hodgson, T. L. (2007). Predictive learning, prediction errors, and attention: Evidence from event-related potentials and eye tracking. *Journal of Cognitive Neuroscience, 19,* 843–854.

Received November 26, 2009
Revision received March 3, 2010
Accepted March 3, 2010 ■