Journal of Experimental Psychology: Animal Learning and Cognition
2014, Vol. 40, No. 1, 116–132

© 2013 American Psychological Association
2329-8456/14/$12.00 DOI: 10.1037/xan0000006

Relative Salience Versus Relative Validity: Cue Salience Influences Blocking in Human Associative Learning

Mike E. Le Pelley, Tom Beesley, and Oren Griffiths
University of New South Wales

Two studies of human contingency learning investigated the influence of stimulus salience on the cue competition effect of blocking. These studies demonstrated that blocking (defined as a difference in responding to blocked and control cues) was greater for target cues that had high "semantic salience" than those of lower salience. Moreover participants showed weaker responding to high salience blocked cues than low salience blocked cues, but a corresponding difference was not observed for control cues. These findings suggest that the influence of relative salience on associative learning depends on the relative validity of the cues in question. Use of eye tracking in Experiment 2 demonstrated that participants' overt attention to cues was also influenced by both relative salience and relative validity. We describe three associative learning models, based on the attentional theory proposed by Mackintosh (1975), that are able to account for our key findings.

Keywords: associative learning, cue competition, blocking, salience, eye tracking

Cue competition phenomena demonstrate that when stimuli (cues) are presented in compound and paired with an outcome, associative learning about cue−outcome relationships does not proceed independently for each cue in isolation. Instead, cues interact and seem to compete for a limited amount of a learning resource that the outcome can support, often termed *associative strength*.[1] Cue competition phenomena have been hugely important not only in the development of theoretical models of associative learning, but also in studies of the neural underpinnings of learning (Waelti, Dickinson, & Schultz, 2001), learning impairments in patients with schizophrenia (e.g., Moran, Al-Uzri, Watson, & Reveley, 2003), and social psychological studies of causal attributions of behavior (Kelley, 1972; Le Pelley et al., 2010). Consequently, it is important to strengthen our understanding of the psychological mechanisms underlying cue competition.

Foremost among demonstrations of cue competition are *overshadowing* and *blocking* effects. Overshadowing denotes the finding that if a cue compound XY is repeatedly paired with an outcome (XY+), then following this training, conditioned responding to stimulus Y will be weaker than if it had been paired with the outcome the same number of times but in isolation (Y+). Hence the presence of X during XY+ training is said to overshadow learning about Y (and, reciprocally, the presence of Y can overshadow learning about X). Overshadowing is well established in both animal (e.g., Mackintosh, 1976; Pavlov, 1927) and human (Waldmann, 2001) associative learning.

Blocking refers to the finding that conditioned responding to an element of a reinforced cue compound is reduced if another element of that compound has previously been established as a reliable predictor of reinforcement. Thus if trials on which A is paired with an outcome are followed by pairings of an AB compound with the same outcome (A+ then AB+), less responding to B is subsequently observed than to a control cue (D) trained in a compound in which neither element received prior conditioning (CD+). Hence prior learning about A is said to block learning about B on AB+ trials. Once again, blocking is well established in animals (e.g., Kamin, 1968) and humans (e.g., Shanks, 1985).

Blocking and overshadowing are of particular interest because they illustrate different determinants of cue competition. Blocking demonstrates that the relative *validity* of a cue influences the extent to which it will gain associative strength. That is, on the first AB+ trial of a blocking procedure, cue A is a better predictor of the outcome than is cue B by virtue of its previous pairings with the outcome. Consequently, the redundant cue B gains little associative strength. In contrast, in the case of overshadowing both cues X and Y begin XY+ training with equal validity, since neither has previously been paired with the outcome. What instead determines cue competition here is the relative *salience* of the cues. Animal studies have demonstrated that if cue X is a more intense or salient stimulus than cue Y, then cue X will develop stronger conditioned responding and cue Y will develop correspondingly less, suggesting that the more salient cue X has captured more of the available associative strength (Mackintosh, 1976).

This article was published Online First October 7, 2013 under the previous title (*Journal of Experimental Psychology: Animal Behavior Processes*).

Mike E. Le Pelley, Tom Beesley, and Oren Griffiths, School of Psychology, University of New South Wales, Sydney, Australia.

This work was supported by Grant FT100100260 from the Australian Research Council. Experiment 1 was conducted while Mike E. Le Pelley and Tom Beesley were based at Cardiff University, Cardiff, United Kingdom.

Correspondence concerning this article should be addressed to Mike E. Le Pelley, School of Psychology, University of New South Wales, Sydney NSW 2052, Australia. E-mail: m.lepelley@unsw.edu.au

[1] Although alternative models exist in which cue competition effects arise because cues compete at the time of performance, rather than during learning (e.g., Miller & Matzel, 1988). We discuss this issue further in the General Discussion; until that point, and for the sake of simplicity and brevity of exposition, we will describe cue competition effects in the widely used language of learning-based theories.

We might ask, then, how these two properties—relative validity and relative salience—interact to determine the amount of cue competition that occurs. Perhaps the simplest way to do this, and the approach taken in the current article, is to observe how the magnitude of blocking changes as the salience of the cues involved is varied. Recall that blocking is assessed by comparing responding to B and D after training on a blocking contingency (A+ then AB+) and a control contingency (CD+). In this example cues B and D are the *target cues*, and cues A and C are the *competing cues*, of blocking and control contingencies, respectively. The question of interest is how the magnitude of blocking (difference in responding to B and D) changes as the salience of the target cues is varied while that of the competing cues is held constant.

Feldman (1975) was the first to address this question in a study with rats, and he found that the magnitude of blocking decreased as the salience of the target cues increased. More specifically, as target cue salience increased, responding to the blocked target cue (B) increased, but responding to the control target cue (D) did not differ, although this may have been the result of a ceiling effect. To the best of our knowledge, only one previous study has systematically examined the impact of salience on blocking in humans. Denton and Kruschke (2006) found a tendency toward a reduction in the difference between blocked and control cues—that is, a reduction in blocking—as target cue salience increased; however, this did not reach statistical significance ($p = .18$). Like Feldman, they also found that responding to the blocked cue increased (significantly) as its salience increased. Denton and Kruschke's data suggest that responding to the control target cue also increased as its salience increased, though no inferential statistics were reported for the appropriate comparison.

Notably, the stimuli used by Denton and Kruschke (2006) differed in what might be termed their *perceptual* salience. Each stimulus constituted a horizontal stripe containing a random pixel pattern of a certain color and pixel density, superimposed on a background density level. Figure 1 shows an example stimulus display, which contains two cues: one salient, corresponding to the densely spotted fourth stripe of the array (that appeared in red), and one nonsalient, corresponding to the relatively sparsely dotted first stripe of the array (that appeared in cyan). All the other stripes in the array represent absent cues; dots in these stripes appear at a very sparse, background density.

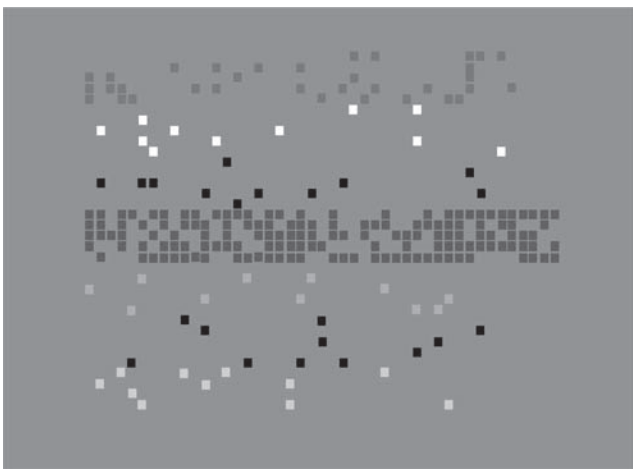


Figure 1. Example stimulus display from the study by Denton and Kruschke (2006). The seven horizontal stripes of dots in this figure were each presented in a different color (represented here as different grayscale shades). This display contains two cues: one salient, corresponding to the densely spotted fourth stripe of the array, and one nonsalient, corresponding to the relatively sparsely dotted first stripe of the array.

Each stimulus display in Denton and Kruschke's (2006) study was presented for only one second. Notably, when looking at Figure 1 it is immediately clear that the fourth stripe differs from the background and hence represents a stimulus that is present on this trial. It is much less clear, however, that the first stripe also represents a presented stimulus. It seems plausible that participants would immediately orient to the salient stimulus and may not even realize that the nonsalient stimulus was present. Furthermore, the salient, high-density stimulus may produce a simultaneous density contrast effect (MacKay, 1973) that will make the nonsalient stimulus appear even less dense than it really is, and so less distinguishable from the background.

The important point here is that if the cues differ greatly in their perceptual salience, to the extent that the high salience cue effectively masks the low salience cue, then we should not be surprised to find a decrease in blocking with increased target cue salience. Suppose we pretrained a quiet tone as a predictor of an outcome, before pairing a compound of that tone and a loud noise with the same outcome. Clearly, blocking would be weak or nonexistent in this case, because the established predictor of the outcome would be perceptually overshadowed (masked) by the added cue. If the added cue were loud enough, participants would not know that the quiet tone had occurred at all, and hence it could not block learning about the loud noise.

The implication is that in Denton and Kruschke's (2006) study, the influence of salience on cue competition may reflect a purely perceptual interaction between the stimuli. Consequently, this means that any theory of associative learning that can explain blocking can also account for the finding of reduced blocking for high salience targets. If there is a perceptual interaction between high and low salience cues presented in compound, such that a low salience competing cue is not perceived as being present, then learning about the high salience target cue will proceed equally in both blocking and control contingencies since in the absence of the competing cue these contingencies are equivalent. As such, this finding cannot discriminate between alternative models of cue competition in associative learning.

We might ask, then, whether a similar influence of salience on cue competition occurs in a situation in which we can be confident that both stimuli have been perceived and processed up to the level of determining their identity (at least). To answer this question, the current experiments assessed blocking in humans using a preparation in which some cues had higher salience than others, but without differing greatly in perceptual characteristics. We used the well-established allergy prediction scenario (e.g., Le Pelley & McLaren, 2001; Mitchell, Lovibond, Minard, & Lavis, 2006; Wasserman, 1990). This is a causal learning task in which participants play an allergist whose task is to judge the likelihood with which foods will cause allergic reactions in a fictitious patient. In this design the foods are the cues, and the allergic reactions are the outcomes. Using the allergy prediction procedure allowed us to manipulate what we term the semantic salience of the food cues; the unusualness, or notableness of the foods. For example, in the

semantic context of foods, a flamingo is more unusual and notable than an apple, but the perceptual characteristics of the word *flamingo* are not markedly more salient than for the word *apple*.

We also manipulated cue salience within subjects, rather than between subjects as in Denton and Kruschke (2006), to ensure that differences in responding are not a consequence of between-subjects differences in task difficulty, motivation and so forth.

Pilot Study

The first step was to establish the relative saliences of various foods, some of which could then be used to investigate the effect of salience on blocking. Nineteen Cardiff University students (aged 19–30) were given a list of 60 foods, and were asked to rate "how salient each food is—that is, how much it 'stands out' or how notable it is as a food" on a scale from 0 (*least salient*) to 10 (*most salient*). The order in which the foods appeared was randomized for each participant. The Appendix shows the full list along with each food's mean rated salience and standard deviation.

Experiment 1

Table 1 shows the design of Experiment 1. Letters in lower case (*a* to *p*) represent low salience foods; upper case bold letters (***Q*** to ***T***) represent high salience foods. O1 and O2 refer to types of allergic reaction (outcomes) suffered by Mr. X: outcome O1 was dizziness, and O2 was sweating. So for example, "*ab* → O1" indicates that low salience foods *a* and *b* were eaten together, and that reaction O1 (dizziness) occurred. On each training trial, participants were told the food(s) eaten by Mr. X and predicted which reaction he would suffer, with immediate feedback. After several blocks of training with the eight trial types of Stage 1, participants moved to a second training phase with the eight trial types of Stage 2. This was followed by a test of responding to each cue individually.

The first four rows of Table 1 show two blocking contingencies, and two corresponding control contingencies, with low salience target cues. The amount of blocking for low salience cues is given by comparing test-phase responding to the target cues of Blocking Low (*m*/*n*) and Control Low (*o*/*p*) contingencies. The next four rows of Table 1 show two blocking contingencies, and corresponding control contingencies, with high salience target cues. Hence the amount of blocking for high salience cues is given by comparing test-phase responding to the target cues of Blocking High (***Q***/***R***) and Control High (***S***/***T***) contingencies. The trial types at the bottom of Table 1 were filler items to make Stage 1 more challenging, and are not discussed further.

Experiment 1 was run in two replications (Experiment 1A and Experiment 1B), which used the same experimental design but with minor procedural differences as described below.

Method

Participants, apparatus, and stimuli. Forty-one Cardiff University students participated in Experiment 1A, and 32 in Experiment 1B, for course credit. Participants were tested individually, and stimulus presentation was controlled by a Visual Basic program.

In Experiment 1A, 12 low salience foods (bread, rice, potato, apples, cucumber, onion, chicken, grapes, pasta, eggs, banana, and mushrooms) were randomly assigned as cues *a* to *l* for each participant. The four foods used as low salience target cues were tomato, lettuce, carrots, and ham; these were randomly assigned as cues *m* to *p*. The four high salience foods were brains, caterpillars, ducks' tongues, and flamingo, randomly assigned as cues ***Q*** to ***T***. Outcomes O1 and O2 were dizziness and sweating, respectively.

In Experiment 1B cue assignment was similar, but the specific foods used were slightly different. Foods used as cues *a* to *l* were rice, potato, apples, cucumber, onion, chicken, grapes, pasta, eggs, mushrooms, yoghurt and fish; foods used as cues *m* to *p* were

Table 1
Experiment Design

Contingency	Stage 1	Stage 2	Test Experiment 1		Test Experiment 2
Blocking low	*a* → O1	*am* → O1	*a*? *m*?	*mp*?	(blocked low vs. control low)
	b → O2	*bn* → O2	*b*? *n*?	*no*?	
Control low		*co* → O1	*c*? *o*?	***QT***?	(blocked high vs. control high)
		dp → O2	*d*? *p*?	***RS***?	
Blocking high	*e* → O1	*e**Q*** → O1	*e*? ***Q***?	*m**R***?	(blocked low vs. blocked high)
	f → O2	*f**R*** → O2	*f*? ***R***?	*n**Q***?	
Control high		*g**S*** → O1	*g*? ***S***?	*o**T***?	(control low vs. control high)
		*h**T*** → O2	*h*? ***T***?	*p**S***?	
Fillers (Experiment 1 only)	*i* → O1				
	j → O1				
	k → O2				
	l → O2				

Note. Letters in lower case (*a* to *p*) represent low salience foods; upper case, bold letters (***Q*** to ***T***) represent high salience foods. O1 and O2 refer to the type of allergic reaction suffered by Mr. X after eating these foods. Filler trial types in the lower section of the table were included only in Experiment 1. The column labeled "Test Experiment 1" shows the individual cues that were presented during the test phase of Experiment 1; the column labeled "Test Experiment 2" shows the cue compounds that were presented during the test phase of Experiment 2.

tomato, lettuce, carrots, and ham; and foods used as cues ***Q*** to ***T*** were caterpillars, blood, fried spiders, and sheep's eyeballs.[2]

Procedure. Participants were told they were to play an allergist whose task was to discover which foods caused different allergic reactions in Mr. X by observing meals that he ate and the reactions he suffered, and that later they would be tested on this knowledge. Each Stage 1 trial displayed the name of the food eaten by Mr. X on that trial. In Experiment 1A only, this was accompanied by a picture of the food. Participants predicted which reaction would occur by clicking one of two radio buttons, labeled *Dizziness* and *Sweating*. For each trial type, in a random half of the blocks the Dizziness option appeared above Sweating; for the other half of blocks this was reversed. Immediate feedback was provided. If participants' prediction was correct, the word *Correct* appeared; if it was incorrect, *Wrong* appeared and the computer beeped. Stage 1 comprised eight blocks in Experiment 1A and 12 blocks in Experiment 1B, with each of the eight trial types occurring once per block in random order.[3]

Stage 2 followed immediately from Stage 1 with no break. Trials were identical in form to those of Stage 1, except that now two foods were presented on each trial (arranged horizontally in Experiment 1A and vertically in Experiment 1B). For each trial type, presentation order of the two foods was counterbalanced across blocks. Stage 2 comprised four blocks in Experiment 1A and eight in Experiment 1B, with each of the eight trial types appearing once per block in random order.

In the final test, participants rated the strength of food−allergy associations. Each test trial presented a single food. Participants rated how likely it was that Mr. X would suffer from one of the two types of allergic reaction (say, dizziness) after eating this food, on a scale from 0 (*Food very unlikely to cause this reaction*) to 10 (*Food very likely to cause this reaction*). On the succeeding trial, participants rated the ability of the same food to cause the other type of allergic reaction (sweating in this example). Whether participants rated dizziness before sweating or vice versa was determined randomly for each participant but was consistent across all test trials. Participants rated all target cues (*m* to ***T*** in Table 1) and competing cues (*a* to *h*) in random order.

Results

Training accuracy. A blocking effect could be expected only if participants managed to learn the cue−outcome associations during Stage 1. Following Le Pelley and McLaren (2003), a criterion was imposed of 60% correct responses across all Stage 1 trials (chance performance = 50% correct). One participant in Experiment 1B failed to reach this criterion (in fact, this participant scored below 50% in Stages 1 and 2), and so this participant's data were excluded from further analysis. Figure 2 shows mean accuracy of remaining participants' predictions during Stages 1 and 2. Data have been averaged over trial types belonging to the same contingency in each stage (e.g., data labeled *Blocking Low* represent mean accuracy on $a \rightarrow$ O1 and $b \rightarrow$ O2 trials in Stage 1, and mean accuracy on $am \rightarrow$ O1 and $bn \rightarrow$ O2 trials in Stage 2). Learning is evident in both stages of Experiment 1A and 1B, with performance well above chance by the end of each training stage.

A concern noted earlier was that the introduction of high salience cues could cause participants to overlook the presence of low salience cues (just as a loud noise will mask a quiet tone). The training data rule out this interpretation of the current results. If such an account applied, we would expect poorer performance in the first block of Stage 2 for Blocking High contingencies (where the added high salience cues would mask the pretrained low salience cues) than for Blocking Low contingencies. This was not the case: mean performance in block 1 of Stage 2 was numerically higher on Blocking High than Blocking Low trials (i.e., opposite in direction to what is proposed above), though the difference was not significant in Experiment 1A, $t(40) = 1.36$, $p = .18$, 95% CI

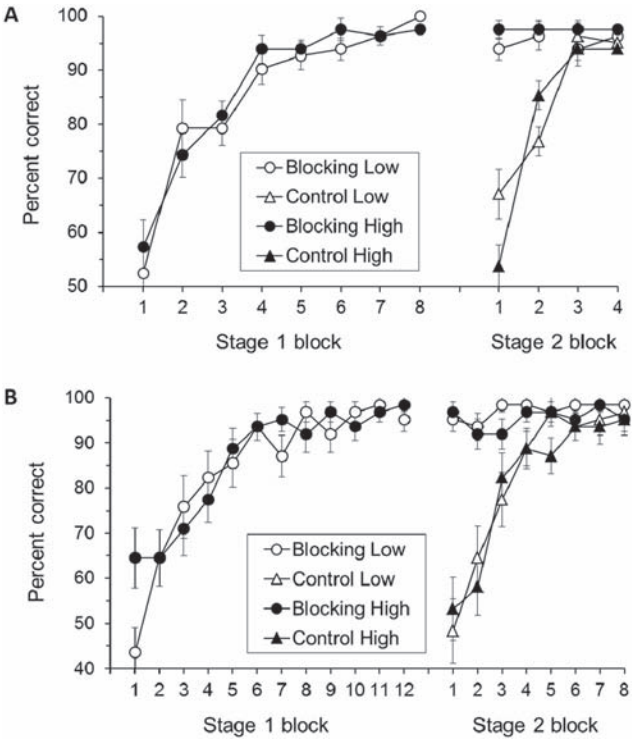


Figure 2. Percent correct responses across training in Experiment 1A (A) and Experiment 1B (B). Data have been averaged over trial types belonging to the same contingency in each stage. Chance responding corresponds to 50% correct. Error bars show standard error of the mean.

[2] The sets of foods used in Experiments 1A, 1B, and 2 were slightly different. This was largely a consequence of the change from presenting stimuli as words only (in Experiment 1B, which was chronologically the first study to be run), to including pictures of the foods (in Experiments 1A and 2). For example, Experiment 1B had "sheep's eyeballs" as stimuli, but we could not find a suitable picture for use in the other experiments so this was changed for a different high salience food. [The move from words only to words and pictures was part of a general change in the way that we programmed experiments around the time that Experiment 1A was run, largely motivated by our increasing use of eye tracking which is well suited to larger, more complex stimuli such as pictures.]

[3] The difference in training duration between Experiments 1A and 1B was a consequence of the change from using words only as stimuli in Experiment 1B to words and pictures in Experiment 1A. Our experience has been that participants typically learn faster—and hence require less training—when pictures are provided, a notion that is borne out by the training data shown in Figure 2.

[−1.80, 9.11], or Experiment 1B, $t(31) = .44$, $p = .66$, 95% CI [−5.85, 9.08].[4]

Test phase rating data. Of primary interest are participants' judgments of the strengths of the trained cue−outcome relationships. That is, the critical data relate not to how strongly each food was perceived to predict allergic reactions in general, but instead to how strongly that food was perceived to predict the specific reaction with which it had been paired. Hence, following Le Pelley and McLaren (2003), participants' test phase ratings were used to calculate difference scores for each cue. We took the rating for each cue with respect to the correct outcome (i.e., the outcome O1 or O2 with which it had been paired in Stage 2), and subtracted from that the rating with respect to the incorrect outcome (the outcome with which it had not been paired). For example, the difference score for cue *a* was calculated by taking the rating for *a* with respect to O1 (dizziness) minus the rating for *a* with respect to O2 (sweating), because *a* was paired with O1 in Stage 2 and so O1 was the correct outcome. Likewise, the difference score for cue *b* was calculated by taking the rating for *b* with respect to O2 minus the rating for *b* with respect to O1, because *b* was paired with O2 in Stage 2. These difference scores index the extent to which a cue was perceived as predicting the outcome with which it was paired in Stage 2 more than an outcome with which it was not paired. High difference scores (maximum = 10) indicate strong, selective responding, while a score of 0 indicates no selective responding. Recall that there were two versions of each contingency in the experimental design (see Table 1). Difference scores were averaged across equivalent cues from these two versions (e.g., the score for Blocking Low target cues is the mean of cues *m* and *n*; the score for Blocking Low competing cues is the mean of *a* and *b*; etc).

Figure 3A shows mean difference scores for target cues. A blocking effect was present for both high and low salience target cues in that difference scores were lower for blocking target cues than for control target cues. Notably, this effect was larger for the high salience cues. These data were analyzed using 2 × 2 × 2 analysis of variance (ANOVA), with factors of experiment (1A vs. 1B), salience (low vs. high), and contingency (blocking vs. control). The experiment factor did not exert a main effect or interact with any other factor, largest $F(1, 70) = 2.19$, $p = .14$. There was a main effect of contingency, $F(1, 70) = 28.8$, $MSE = 12.3$, $p < .001$, $\eta_p^2 = .29$, but no main effect of salience, $F(1, 70) = .95$, $MSE = 9.72$, $p = .33$. Crucially, the contingency × salience interaction was significant, $F(1, 70) = 13.7$, $MSE = 8.76$, $p < .001$, $\eta_p^2 = .16$, confirming a significantly stronger blocking effect for high salience target cues than for low salience targets.

Collapsing across experiments, analysis of simple effects found that the difference between blocking and control contingencies for low salience targets approached significance, $F(1, 71) = 3.42$, $MSE = 10.4$, $p = .069$, $\eta_p^2 = .05$, 95% CI [-.078, 2.06]; the difference for high salience targets was highly significant, $F(1, 71) = 40.4$, $MSE = 10.6$, $p < .001$, $\eta_p^2 = .36$, 95% CI [2.37, 4.53]. Moreover, high salience blocked target cues received significantly lower difference scores than low salience blocked cues, $F(1, 71) = 8.13$, $MSE = 11.7$, $p = .006$, $\eta_p^2 = .10$, 95% CI [.49, 2.76]. There was a tendency toward high salience control target cues receiving higher scores than low salience control targets; this approached significance, $F(1, 71) = 3.64$, $MSE = 6.86$, $p = .060$, $\eta_p^2 = .05$, 95% CI [−.037, 1.70].

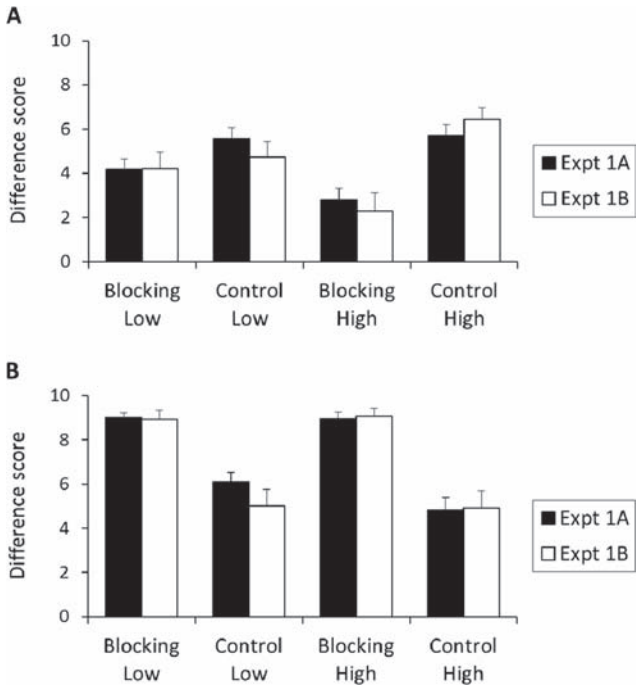


Figure 3. Mean difference scores for the test cues of Experiment 1. Panel A shows difference scores for target cues; Panel B shows difference scores for competing cues. Difference scores for each cue were calculated by taking the rating for the outcome with which that cue was paired in Stage 2 and subtracting from it the rating for the outcome with which it was not paired. Hence high difference scores (maximum = 10) indicate strong, selective learning. Scores were averaged across equivalent cues from the copies of each contingency in the experimental design. Error bars show standard error of the mean.

We explored these latter findings further by analyzing participants' ratings before they were subtracted to yield difference scores, collapsing across experiments. Due to a computer error, these raw ratings were lost for one participant in Experiment 1A. For remaining participants, ratings with regard to the correct outcome for high salience blocked target cues ($M = 5.40$, SEM = .30) were significantly lower than for low salience blocked targets ($M = 6.29$, SEM = .28); $t(70) = 2.71$, $p = .008$, $\eta_p^2 = .10$, 95% CI [−1.55, −.24]. Ratings with regard to the incorrect outcome for high salience blocked targets ($M = 2.82$, SEM = .27) were significantly higher than for low salience blocked targets ($M = 2.09$, SEM = .23); $t(70) = 2.32$, $p = .023$, $\eta_p^2 = .07$, 95% CI [.10, 1.35]. For control target cues, ratings with regard to the correct outcome for high salience cues ($M = 7.58$, SEM = .23) were higher than for low salience cues ($M = 7.14$, SEM = .27), and this difference approached significance, $t(70) = 1.68$, $p = .097$, $\eta_p^2 = .04$, 95% CI [−.08, .95]. Ratings with regard to the incorrect outcome did not differ significantly for high salience ($M = 1.54$, SEM = .20) and low salience ($M = 1.94$, SEM = .23) control targets; $t(70) = 1.62$, $p = .11$, 95% CI [−.89, .09].

[4] Here and throughout, "95% CI" refers to the 95% confidence interval on the difference between the means of the two conditions being compared, or (for one-sample tests) the 95% confidence interval for the mean of the condition in question.

Returning to difference scores, Figure 3B shows data for the competing cues (note that all competing cues were of low salience; *Low* and *High* in Figure 3B refer to the salience of the target cues with which competing cues were paired). Competing cues from blocking contingencies (pretrained in Stage 1) received high difference scores near the maximum of 10 regardless of the salience of the target cues with which they were paired, and these scores were higher than the scores for competing cues from control contingencies. ANOVA conducted in the same manner as for target cues revealed a main effect of contingency, $F(1, 70) = 144$, $MSE = 7.02$, $p < .001$, $\eta_p^2 = .67$. No other effects were significant, largest $F(1, 70) = 1.37$, $p = .25$.

A final analysis assessed whether the salience of the target cue influenced the extent to which responding was unequally distributed between this target cue and the competing cue with which it was paired in the Control contingencies. The analysis of target cues presented above provided some evidence that responding to the target cue of a Control contingency was greater if that cue was of high salience. If high salience cues were better able to overshadow low salience cues with which they were paired, then there should be a concomitant reduction in responding to the low salience competing cues on Control High trials as compared to Control Low trials. In other words, we would expect a greater difference in responding to target and competing cues for Control High contingencies than that for Control Low contingencies. This pattern was confirmed by an ANOVA carried out using the data from the Control contingencies with factors of experiment (1A vs. 1B), salience (low vs. high), and cue type (target vs. competing). This found a significant interaction between salience and cue type, $F(1, 70) = 8.12$, $MSE = 5.83$, $p = .006$, $\eta_p^2 = .10$. No other effects were significant, largest $F(1, 70) = 2.78$, $p = .10$.

Discussion

In Experiment 1, the magnitude of blocking depended on the salience of the target cues involved: blocking was significantly stronger for high salience than low salience targets. This pattern is inconsistent with previous findings in humans (Denton & Kruschke, 2006) and rats (Feldman, 1975), demonstrating reduced blocking with high salience target cues. We return to this discrepancy in the General Discussion.

An intriguing finding of Experiment 1 was that responding to high salience blocked target cues was significantly weaker than to low salience blocked targets; participants seem to have learned less about a high salience cue than a low salience cue, which runs counter to intuition. In support of this view, high salience blocked cues received significantly lower ratings with regard to the correct outcome than did low salience blocked cues. That is, participants did view the high salience cues as less likely to cause the outcome with which they had been paired. We also found that high salience blocked cues received higher ratings than low salience blocked cues with regard to the incorrect outcome. This may reflect participants reasoning that all cues caused either O1 or O2, because these were the only response options. Since they had given the high salience blocked cues a low rating with regard to the correct outcome, they would therefore give these cues a higher rating with regard to the incorrect outcome.

The finding of reduced responding to high salience cues was specific to target cues of Blocking contingencies. For target cues of Control contingencies, there was instead a tendency (approaching significance) toward stronger responses to high salience cues than low salience cues. This tendency was supported by the finding that, for Control contingencies, increasing the salience of the target cue increased responding to the target cue relative to its competing cue partner. This is consistent with previous findings from nonhuman animals suggesting that high salience target cues produce more overshadowing of the low salience competing cues with which they are paired (Mackintosh, 1976).

Experiment 2

Experiment 2 had two main aims. The first was to replicate the two key findings: (1) The size of the blocking effect (given by the difference in responding to target cues of Blocking and Control contingencies) is greater when the target cues have high salience than when they have low salience; and (2) Responding to the target cue of a Blocking contingency is reduced when this cue has high salience compared to when it has low salience. The second aim was to assess whether these differences were accompanied by differences in overt attention, measured via eye tracking. While blocking has previously been shown using eye gaze (Beesley & Le Pelley, 2011; Kruschke, Kappenman, & Hetrick, 2005; Wills, Lavric, Croft, & Hodgson, 2007), the role of cue salience in this relationship has not been explored.

The trial types experienced during Stage I and 2 training phases were as for Experiment 1, except that Experiment 2 omitted the "Filler" trials shown in Table 1 in order to simplify the training procedure. Experiment 2 used a different type of test phase that allowed for a direct comparison of the strength of responding to different cues, by pitting them against one another. On each test trial a pair of foods was presented and participants chose which outcome, O1 or O2, they thought was more likely to occur.

The final column of Table 1 shows the cue compounds presented during the test phase of Experiment 2. Each compound contained one cue that had been paired with outcome O1 in Stage 2, and one that had been paired with O2. Consequently, choices on these test trials indicate which of the two cues was perceived as a stronger cause of its respective outcome. Consider compound *mp*. Cue *m* was paired with O1 in Stage 2, and *p* was paired with O2. If participants had acquired a stronger association between *m* and O1 in Stage 2 than between *p* and O2, then they should select O1 as the outcome more likely to be produced by *mp*, and vice versa.

The first two test compounds shown in Table 1 (*mp* and *no*) each have a target cue from a Blocking Low contingency as the first element (*m* and *n*), and a target cue from the Control Low contingency as the second element (*p* and *o*). To the extent that blocking occurs for low salience targets, we would expect *m* and *n* to have weaker associations to their respective outcomes than *p* and *o*. Consequently participants should choose the outcome associated with the second element of each of these compounds (O2 for *mp*; O1 for *no*). That is, responding to these compounds should be dominated by the second element.

The second pair of test compounds in Table 1 (***QT*** and ***RS***) each have a Blocking High target cue as first element, and a Control High target cue as second element. For the same reasons as in the previous paragraph, to the extent that blocking occurs for high salience targets we would expect responding to these compounds to be dominated by the second element. And if—as suggested by

Experiment 1—blocking is stronger for high salience targets than for low salience targets, the extent of this dominance by the second element should be greater for high salience compounds ***QT*** and ***RS*** than for low salience compounds *mp* and *no*.

The third pair of test compounds (*m**R*** and *n***Q**) have a Blocking Low target cue as first element, and a Blocking High target cue as second element. If—as suggested by Experiment 1—participants learn more about low salience blocked cues than high salience blocked cues, then responding to these compounds should be dominated by the first element.

The final pair of test compounds (*ot* and *ps*) have a Control Low target cue as first element, and a Control High target cue as second element. Experiment 1 found a tendency toward weaker responding to low salience control target cues than high salience cues; hence, in this case we might expect responding to these compounds to be dominated by the second element.

Method

Participants, apparatus, and stimuli. Forty-two students of the University of New South Wales (29 female, ages 18–31) completed this and an unrelated experiment for $20 AUD. They were tested either individually or in pairs, using two Tobii TX300 monitor-mounted eye trackers (Tobii Technology, Sweden) that recorded gaze location throughout the experiment at a sampling rate of 300 Hz. Participants rested their head on a chinrest approximately 55 cm from the 58.4-cm widescreen monitor.

Foods used as competing cues *a* to *h* were bread, apples, cucumber, onion, pasta, eggs, carrots, and banana; foods used as low salience target cues M to P were rice, potato, chicken, and lettuce; and foods used as high salience target cues ***Q*** to ***T*** were pig's ears, caterpillars, brains, and flamingo. Outcomes O1 and O2 were dizziness and itchiness respectively.

Procedure. The eye tracker was initially calibrated using a 9-point procedure. After every 16 trials, a red dot appeared in the center of the screen; participants fixed their eyes on this dot and clicked the mouse to continue. Gaze data from these dot calibration presentations were used to correct for drift from the initial calibration (see Eye tracking data analysis section below).

The form of Stage 1 and Stage 2 training trials was as for Experiment 1. Figure 4 shows an example screenshot from a Stage 2 trial. Stage 1 comprised eight blocks of trials, and Stage 2 comprised six blocks. On each test trial a pair of foods appeared and participants chose which of outcomes O1 or O2 was more likely to occur. Each of two blocks in the test phase contained each of the eight test compounds in Table 1, in random order. The left/right presentation order of the elements of each compound was counterbalanced across these two test blocks.

Eye tracking data analysis. For each dot calibration presentation, mean gaze location was calculated over a 1,000-ms period beginning 500 ms before participants clicked the mouse. This mean gaze location was used to update the calibration for the screen center for all trials until the next dot calibration, unless the mean location was more than 150 pixels different from the previous dot calibration (which might indicate that participants had failed to fixate the dot correctly), in which case the previous dot calibration was used instead.

The key analyses relate to eye gaze on Stage 2 training trials in the period between cue presentation and response (selecting an outcome). Gaze was defined as falling on a cue if it fell within the square border surrounding the cue's picture and name (see Figure 4). Following procedures that we have used previously (Beesley & Le Pelley, 2011; Le Pelley, Beesley, & Griffiths, 2011; Le Pelley, Mitchell, & Johnson, 2013), gaze data were analyzed in terms of dwell time: the summed duration for which gaze fell on a cue. For each Stage 2 trial we calculated proportion of dwell time on the target cue, denoted *Prop_target*. For example, on an *am* → O1 trial *Prop_target* on the target cue (*m*) was given by dwell time on cue *m* divided by summed dwell time on cues *a* and *m*. Consequently a *Prop_target* above .5 indicates that participants spent longer looking at the target cue on that trial, and a value below .5 indicates they spent longer looking at the competing cue. Data for trials with response latencies over 10 s were excluded (0.68% of Stage 2 trials), as were those for which dwell time summed across both cues was less than 200ms (14.6% of Stage 2 trials); these missing trials were assigned a value of .5 for *Prop_target*.[5]

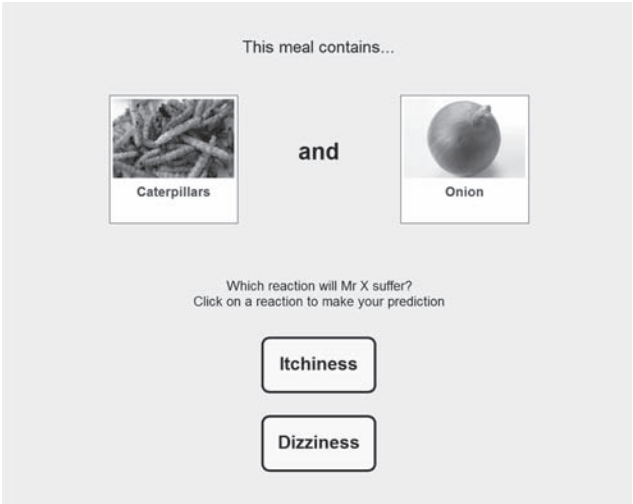


Figure 4. Example screenshot from a Stage 2 trial of Experiment 2. Eye gaze was defined as falling on a cue if it fell within the square border surrounding the cue's picture and name.

Results

Training accuracy. Two participants failed to meet the criterion of 60% correct responses during Stage 1, and so their data were excluded from further analysis. Figure 5 shows mean accuracy of remaining participants' predictions across training, averaged as for Experiment 1. Learning is evident in both stages. As in Experiment 1, there is no suggestion that the Stage 2 introduction of high salience cues in the Blocking High contingencies caused participants to overlook the pretrained, low salience competing

[5] Since many critical analyses of the eye gaze data are within-subjects, we could not simply drop these missing trials from analyses. If we had, then if by chance a particular participant had both trials of a particular type (e.g., Blocking High) excluded in just one block of the experiment, we would have to drop all of that participant's data from these critical analyses. In order to avoid this problem, we instead gave all these trials *Prop_target* = .5. This is the most conservative solution, since it can only reduce our sensitivity to detecting differences between conditions.

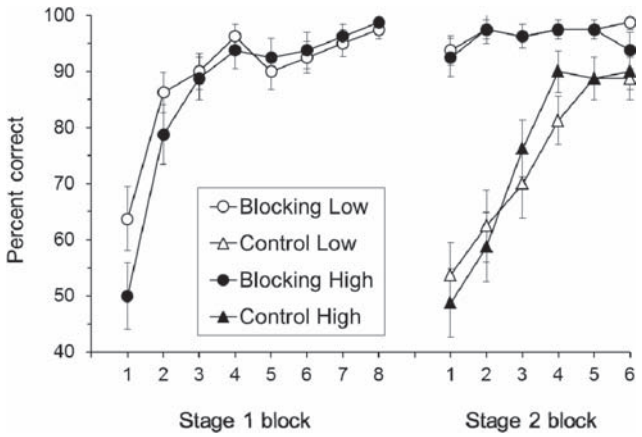


Figure 5. Percent correct responses across training in Experiment 2. Data have been averaged over trial types belonging to the same contingency in each stage. Chance responding corresponds to 50% correct. Error bars show standard error of the mean.

cues. Accuracy on early Blocking High trials of Stage 2 was very similar to that for Blocking Low trials; $t(39) = .30$, $p = .77$ for the first block.

Test phase choice data. Consider again the list of test compounds shown in Table 1. For each compound, if a participant chose the outcome with which the first cue of that compound (as listed in the table) had been paired, this choice was assigned a score of −1. If they chose the outcome with which the second cue of the compound had been paired, the choice was assigned a score of +1. For example, for test compound *mp* choice of outcome O1 would be scored as −1 since O1 was the outcome paired with the first element (*m*) of this compound. Choice of O2 for this compound would be scored as +1 since O2 was the outcome paired with second element (*p*) of the *mp* compound.

These choice scores were averaged over the two presentations of each compound in the test phase, and over the two equivalent compounds of each pair shown in Table 1 (*mp*/*no*; ***QT***/***RS***; *m****R***/*n****Q***; *o****T***/*p****S***). The mean choice scores provide an index of the extent to which responding to each type of compound was dominated by each element, ranging from −1 (total domination by the first element) through 0 (both elements weighted equally) to +1 (total domination by the second element). Figure 6 shows the resulting mean choice scores for each pair. The score for *mp*/*no* was significantly greater than zero, $t(39) = 2.05$, $p = .047$, $\eta_p^2 = .10$, 95% CI [.002, .40], indicating that blocking had occurred for the low salience cues. The score for ***QT***/***RS*** was also significantly above zero, $t(39) = 5.38$, $p < .001$, $\eta_p^2 = .43$, 95% CI [.28, .62], demonstrating a blocking effect for high salience cues. The choice score for ***QT***/***RS*** was significantly greater than for *mp*/*no*, paired $t(39) = 2.24$, $p = .031$, $\eta_p^2 = .11$, 95% CI [.02, .48], indicating that the blocking effect was greater for high salience target cues than low salience target cues.

The mean choice score for *m****R***/*n****Q*** was significantly below zero, $t(39) = 2.45$, $p = .019$, $\eta_p^2 = .13$, 95% CI [-.37, −.03], indicating stronger responding to low salience blocked cues than high salience blocked cues. The score for *o****T***/*p****S*** did not differ significantly from zero, $t(39) = .42$, $p = .67$, 95% CI [−.14, .22]; that is, there was not a significant difference in responding to high salience and low salience control cues. Finally, the difference in mean choice score for *m****R***/*n****Q*** and *o****T***/*p****S*** was significant using a one-tailed test (which is justified because this direction of effect is explicitly anticipated on the basis of the results of Experiment 1), $t(39) = 1.92$, $\eta_p^2 = .09$, $p = .031$, 95% CI [−.01, .49]. Hence increasing salience had a different effect on blocked versus control cues.

Eye tracking. Figure 7A shows *Prop_target* across Stage 2 training. Data were averaged across the two equivalent exemplars of each trial type in Stage 2 (e.g., Blocking Low data reflect the average of *am* → O1 and *bn* → O2 trials). Recall that a value above .5 indicates greater dwell time on the target cue, and a value below .5 indicates greater dwell time on the competing cue.

Initial analysis was by ANOVA with factors of contingency (blocking vs. control), target salience (low vs. high), and block. This revealed a main effect of contingency, $F(1, 39) = 45.4$, $p < .001$, $\eta_p^2 = .54$, with lower *Prop_target* on blocking trials than control trials. There was also a main effect of salience, $F(1, 39) = 7.18$, $p = .011$, $\eta_p^2 = .16$, with higher *Prop_target* for trials involving high salience targets. A main effect of block, $F(5, 195) = 3.30$, $p = .007$, $\eta_p^2 = .08$, reflected a reduction in mean *Prop_target* across blocks. A significant block × salience interaction, $F(5, 195) = 2.97$, $p = .013$, $\eta_p^2 = .07$, reflected the finding that *Prop_target* decreased across blocks for trials with high salience targets, but did not change (on average) for trials with low salience targets. There was also a significant block × contingency interaction, $F(5, 195) = 2.39$, $p = .039$, $\eta_p^2 = .06$, but the salience × contingency interaction, $F(1, 39) = .53$, and the three-way interaction, $F(5, 195) = 1.03$, were not significant, $ps > .40$. Further analysis revealed a significant downward linear trend in *Prop_target* across blocks for Blocking High and Control High contingencies, $F(1, 39) = 11.0$, $p = .002$, $\eta_p^2 = .22$ and $F(1, 39) = 4.82$, $p = .034$, $\eta_p^2 = .11$, respectively, but not for Blocking Low and Control Low contingencies, $Fs < 1$. Collapsing across blocks

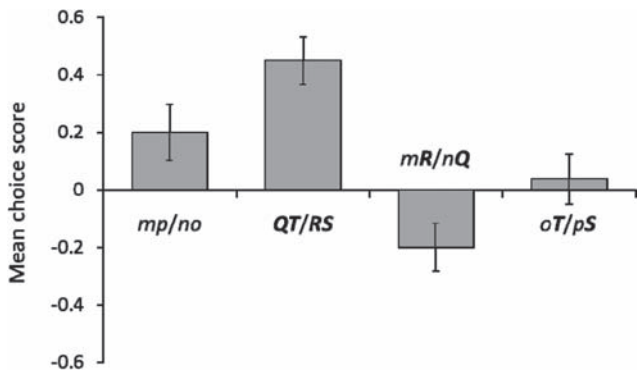


Figure 6. Mean choice scores for the test compounds of Experiment 2. For each compound, if a participant chose the outcome with which the first cue of that compound (as listed in Table 1) had been paired, this choice was assigned a score of −1. If they chose the outcome with which the second cue of the compound had been paired, the choice was assigned a score of +1. Data were averaged over the two presentations of each compound in the test phase, and over the two equivalent compounds of each pair shown in Table 1. Mean choice scores index the extent to which responding to each type of compound was dominated by each element, ranging from −1 (total domination by the first element) through to +1 (total domination by the second element). Error bars show standard error of the mean.

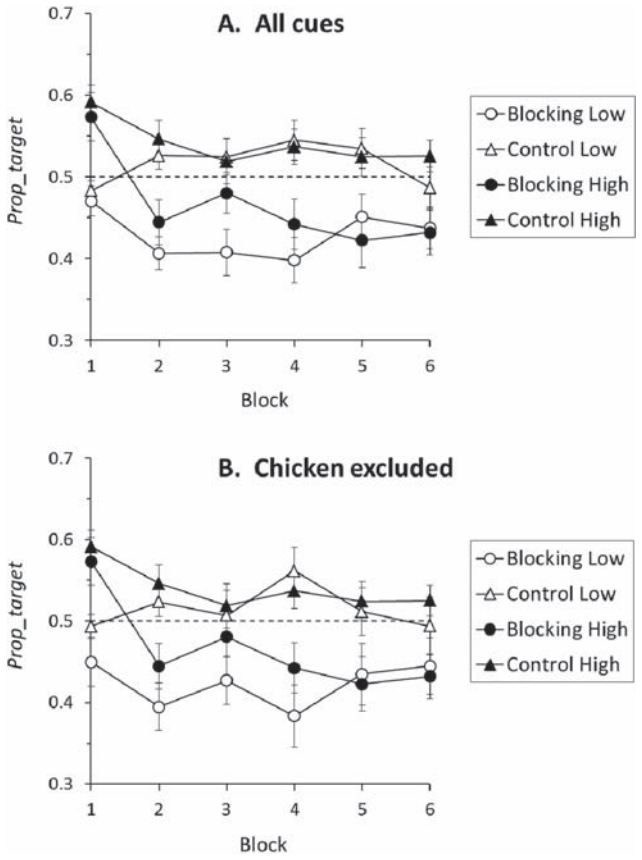


Figure 7. Eye gaze data across Stage 2 compound training blocks of Experiment 2. The dependent variable is the proportion of dwell time on the target cue (*Prop_target*) calculated as described in the text. *Prop_target* above .5 indicates that participants spent longer looking at the target cue on a trial, and a value below .5 indicates that they spent longer looking at the competing cue. A: Data from all trial types, averaging across the two equivalent exemplars of each trial type in each block. B: Data from all trial types apart from those featuring chicken as a low salience target cue. Error bars show standard error of the mean.

1–6, *Prop_target* for Blocking Low and Blocking High contingencies was significantly less than .5, indicating greater attention to the pretrained competing cue than the target cue; $t(39) = 5.77$, $p < .001$, $\eta_p^2 = .46$, 95% CI [.40, .45], and $t(39) = 2.15$, $p = .038$, $\eta_p^2 = .11$, 95% CI [.43, .50] respectively. *Prop_target* for Control Low and Control High contingencies was significantly greater than .5, indicating greater attention to the target cue; $t(39) = 2.07$, $p = .045$, $\eta_p^2 = .10$, 95% CI [.50, .53], and $t(39) = 3.21$, $p = .003$, $\eta_p^2 = .21$, 95% CI [.52, .57] respectively.

Given that the pattern of data in block 1 appeared quite different to that in blocks 2–6, follow-up analyses split the data along these lines. ANOVA using the block 1 data revealed a highly significant effect of salience, $F(1, 39) = 29.1$, $p < .001$, $\eta_p^2 = .43$, but no effect of contingency or salience × contingency interaction, Fs < 1. *Prop_target* in block 1 for Blocking High and Control High contingencies was significantly greater than .5; $t(39) = 2.50$, $p = .017$, $\eta_p^2 = .14$, 95% CI [.51, .63], and $t(39) = 4.45$, $p < .001$, $\eta_p^2 = .34$, 95% CI [.55, .63], respectively. *Prop_target* in block 1 for Blocking Low and Control Low contingencies did not differ significantly from .5; $t(39) = 2.50$, $p = .12$, 95% CI [.43, .51] and $t(39) = 1.30$, $p = .20$, 95% CI [.46, .51], respectively.

Collapsing across blocks 2–6, ANOVA revealed a highly significant effect of contingency, $F(1, 39) = 54.7$, $p < .001$, $\eta_p^2 = .58$, but no main effect of salience or salience × contingency interaction, $F(1, 39) = 1.46$ and .80, respectively, ps > .2. *Prop_target* in blocks 2–6 for Blocking High and Blocking Low contingencies was significantly less than .5; $t(39) = 3.70$, $p < .001$, $\eta_p^2 = .26$, 95% CI [.41, .47], and $t(39) = 5.59$, $p < .001$, $\eta_p^2 = .44$, 95% CI [.39, .45] respectively. *Prop_target* in blocks 2–6 for Control High and Control Low contingencies was significantly greater than .5; $t(39) = 2.29$, $p = .028$, $\eta_p^2 = .12$, 95% CI [.50, .56], and $t(39) = 2.26$, $p = .030$, $\eta_p^2 = .12$, 95% CI [.50, .54], respectively.

In sum, gaze behavior in block 1 of Stage 2 was driven by the salience of target cues. Across subsequent training, contingency (blocking or control) exerted more influence on eye gaze such that, by the end of Stage 2, gaze was largely determined by contingency and not by salience.

Discussion

Participants' test phase responses replicated the key findings of Experiment 1. The blocking effect was greater for high salience than low salience target cues. Moreover, a direct comparison between high and low salience blocked cues revealed that the latter exerted more influence on responding, suggesting that cue–outcome associations were weaker for high salience than low salience target cues. And once again this effect was specific to blocked cues—the same pattern was *not* observed for control cues, where there was no evidence for a difference in the influence of low and high salience target cues on responding. This is a slight departure from Experiment 1, where there was some evidence (albeit rather weak) for greater responding to high salience than low salience control target cues.

Previous work has shown that, on Stage 2 trials of a blocking contingency, people pay more overt attention to the pretrained competing cue than the novel target cue (Beesley & Le Pelley, 2011; Kruschke et al., 2005; Wills et al., 2007). The eye gaze data of Experiment 2 replicated this pattern, with greater proportional dwell time on the competing cue than the target cue in both Blocking Low and Blocking High contingencies when averaged across Stage 2. However, this finding hid a more interesting pattern of changes in attention over blocks. On Blocking High trials, high salience target cues initially captured attention in the first block of Stage 2 before a rapid reorienting of attention, such that in later blocks participants paid more attention to pretrained competing cues. No such change in attention across trials occurred for Blocking Low contingencies.

Turning to the Control High contingency, there was an initial gaze bias toward the high salience target cue, which was reduced in later blocks of Stage 2, as for Blocking High trials. However, the reduction was not as great in the case of Control High trials, such that over the latter blocks overt attention still remained greater to the target cue than the competing cue.

More generally, analysis demonstrated that in the first block of Stage 2 eye gaze was determined by the salience of the target cue (with greater gaze on the target cue on trials with high salience targets) but was not greatly influenced by contingency (blocking vs. control). In contrast, in blocks 2–6 gaze was determined by

contingency (greater gaze on the competing cue on blocking trials than control trials), but was not greatly influenced by target cue salience.

The findings of the Control Low contingency were unexpected. For this contingency both cues in the compound were novel at the outset of Stage 2, and both had low salience. Consequently we would expect equal gaze to target and competing cues throughout Stage 2, since the cues were effectively equivalent. Figure 7A shows that, while *Prop_target* for Control High trials began and ended Stage 2 near .5, during intervening blocks it rose above .5 such that averaged over blocks there was a significant bias toward the target cue. It is unclear why this occurred, but we believe the most likely explanation lies in a poor choice made when selecting cues. As noted in the Method section, foods used as low salience targets came from a different pool from foods used as competing cues. It is possible that by chance the foods used as low salience targets actually had higher salience (on average) than the competing cues they were paired with. Notably, and unlike Experiments 1A and 1B, chicken was used as a low salience target in Experiment 2. While the rated salience of chicken is relatively low (see Appendix A), this rating relates to the salience of chicken as a food in general, and not in the context of how likely it is to cause adverse reactions. In this latter sense chicken is a more notable food, since it is a frequent cause of food poisoning in the real world. Figure 7B shows the pattern of eye gaze data when trials involving chicken were excluded for each participant. Analysis of these data replicated all of the findings described in the Results section, with the following exceptions: (i) critically, *Prop_target* on Control Low trials did not differ from .5 either over all blocks, or just blocks 2–6, $t(39) = 1.49$ and 1.62, respectively, $ps > .1$; (ii) in the omnibus ANOVA with factors of contingency, salience and block, the block $\times$ contingency interaction was nonsignificant, $F(5, 195) = 1.75$, $p = .13$, and the block $\times$ salience interaction approached significance, $F(5, 195) = 1.95$, $p = .087$, $\eta_p^2 = .05$.

General Discussion

Two experiments examined the influence of the salience of target cues on blocking. Previous research in humans (Denton & Kruschke, 2006) and rats (Feldman, 1975) found a reduced blocking effect (i.e., smaller difference in responding to target cues of blocking and control contingencies) for high salience target cues than low salience targets. Contrary to this prior research, our experiments found a larger blocking effect for high salience target cues than low salience target cues.

We believe the most likely reason for this discrepancy lies in the way that salience was manipulated in these studies. Previous studies (Denton & Kruschke, 2006; Feldman, 1975) manipulated cues' perceptual salience. High salience cues were more perceptually intense than low salience cues and hence they may have interacted at a perceptual level: the more intense high salience cues may have caused participants to overlook the presence of low salience cues, in which case blocking would not occur. Indeed, evidence from Feldman's study supports this interpretation. In this study, introducing high salience target cues on Stage 2 trials of the blocking contingency led to a significant drop in responding, relative to blocking trials involving low salience target cues. That is, introduction of high salience cues, but not low salience cues, caused rats to overlook the presence of the pretrained competing cue that would otherwise have generated a high level of responding. Denton and Kruschke (2006) do not report the block-by-block training data that would allow us to assess this account; however, given the stimuli they used (see Figure 1) it remains plausible.

In contrast the current studies manipulated the so-called semantic salience of the stimuli. High salience foods were unusual members of the "food" category. However, their perceptual salience was similar, as shown by the finding that introduction of high salience targets (relative to low salience targets) did not cause participants to overlook the presence of pretrained competing cues in the blocking contingencies of either experiment. This implies that participants were aware of the presence of the target and competing cue from the outset of Stage 2, and processed them up to the level of determining their identity (at least). And yet participants' pattern of responding during the test phase was still influenced by the salience of these cues.

Perhaps the most interesting and counterintuitive finding of Experiments 1 and 2 is that test phase responding was lower for a high salience blocked cue than a low salience blocked cue. To the best of our knowledge, all of the most influential models of associative learning anticipate that participants will, if anything, learn more about a high salience cue than a low salience cue, regardless of whether it is blocked or not (e.g., Mackintosh, 1975; Pearce, 1987; Pearce & Hall, 1980; Rescorla & Wagner, 1972; Wagner, 1981), and hence these models cannot account for this finding.

Up to this point, we have considered cue competition as reflecting a competition for *learning*, wherein blocking reflects a failure to learn about the blocked cue. An alternative approach in which blocking is instead ascribed to competition between cues at the point of *responding* could—at least in theory—account for lower responding to a high salience blocked cue. According to comparator theory (Miller & Matzel, 1988; Stout & Miller, 2007), on Stage 2 trials of an [A+, AB+] blocking contingency associations form between B and the outcome, and between B and A. On test, presentation of B retrieves the representation of the outcome via two routes; a direct route based on the B–outcome association, and an indirect route wherein B retrieves A via the B–A within-compound association, and A retrieves the outcome via an A–outcome association formed during Stage 1 A+ training. The stronger the retrieval of the outcome via the direct route, the greater the responding to B; the stronger the retrieval via the indirect route, the weaker the responding to B. Thus, overall responding to B relies on a comparison of the strength of the direct and indirect routes. In Stout and Miller's formalization of this model, a cue's salience modulates the rate at which it forms associations with the outcome, and with other cues. Consequently increasing the salience of the blocked cue should equally strengthen the direct and indirect routes and hence have no overall impact on responding to that cue. However, if the model could be altered such that the influence of salience on the indirect route was greater than on the direct route, then increasing the salience of the blocked cue would result in a disproportionate strengthening of the indirect route and a concomitant reduction in responding to this cue on test; the result observed in the current experiments.

The problem with comparator theory's account of our data is that it is constrained to predict that salience will have the same influence on blocking and control contingencies. Consider a CD+ control contingency where neither cue has been pretrained. Ac-

cording to comparator theory, target cue D forms associations with the outcome and with C, and competing cue C forms associations with the outcome and with D. Responding to D on test will depend on a comparison of the direct (D → outcome) and indirect (D → C → outcome) routes. If the model has been parameterized such that increasing salience strengthens the indirect route relative to the direct route, then it must predict that increasing the salience of the control target cue (D) will reduce responding to this cue. But this was not observed empirically. In Experiment 2, the salience of control cues had no significant influence on test phase responding to these cues; in Experiment 1, evidence suggested that there had been greater learning about high salience than low salience control cues. Contrary to this modification of comparator theory, the influence of salience depended on the predictive status of the cues.

The fact that salience had a different influence in blocking and control contingencies also rules out various other explanations of the finding of reduced responding to high salience blocked cues. For example, it might be argued that participants viewed what we had labeled as high salience foods (caterpillars, flamingo, etc.) as less remarkable than the supposedly low salience foods (lettuce, tomato, etc.), and consequently they learned less about what we had termed the *high salience* blocked foods. Another possibility would be that the names and pictures of the high salience foods were, on average, more unpleasant than for the low salience foods. If participants avoided looking at these unpleasant foods, then this could explain a reduction in learning about high salience than low salience cues. A third possibility relates to preconceptions about the foods. Participants may have viewed the unusual high salience foods as being likely to cause adverse reactions in general, rather than any one specific allergic reaction more than others. Given that the test phase of our experiments assessed selective responding to each cue (i.e., the extent to which it was perceived as causing the outcome [O1 or O2] with which it had been paired more strongly than the outcome with which it had not been paired), this again could explain evidence for weaker selective responding to high salience cues than low salience cues.[6] The important point, however, is that each of these explanations applies equally to blocked and control cues. Consequently, each account is ruled out by the finding that high salience control target cues do not show the same reduction in responding when compared to low salience control cues.

Blocking and Attentional Theories of Associative Learning

So the question remains, how are we to explain the current findings? The account that we favor is based on an approach wherein blocking is (at least in part) a consequence of a change in attention to the cues involved (Kruschke, 2003; Le Pelley, 2004; Mackintosh, 1975). This idea is supported by previous demonstrations in both humans and rats that blocked cues are slower than control cues to form new associations (often termed "blocking of unblocking": see Griffiths & Le Pelley, 2009; Kruschke & Blair, 2000; Le Pelley, Beesley, & Suret, 2007; Mackintosh, 1978; Mackintosh & Turner, 1971). Such findings imply that the blocked cue has undergone a reduction in attention, on the assumption that attention influences the rate of learning about a cue. More direct support for an attentional account of blocking comes from studies that use online markers such as eye tracking (Beesley & Le Pelley, 2011; Kruschke et al., 2005; Wills et al., 2007: and the current Experiment 2) and event-related potentials (Wills et al., 2007) to demonstrate a reduction in attention to blocked cues over the course of Stage II training.

Such findings follow naturally from the theory proposed by Mackintosh (1975), which states that the strength of the association between a cue X and an outcome (denoted V_X) is updated on each trial according to the expression:

$$\Delta V_X = \alpha_X \beta (\lambda - V_X) \quad (1)$$

where ΔV_X represents the change in V_X on the current trial and β is a learning rate parameter relating to the salience of the outcome. The error term $(\lambda - V_X)$ represents the discrepancy between the observed magnitude of the outcome (λ) and the magnitude of the outcome expected on the basis of the presence of cue X (V_X); that is, the extent to which X predicts the outcome. α_X represents the associability of cue X; the readiness with which X forms associations. α_X is a variable, and depends on how well X predicts the outcome occurring on that trial. Specifically, α_X increases if X is a better predictor of the outcome than are all other presented cues, that is, if $|\lambda - V_X| < |\lambda - V_Q|$, where V_Q is the associative strength of all cues other than X present on that trial. α_X decreases if X is a poorer predictor of the outcome than are all other presented cues Q, that is, if $|\lambda - V_X| > |\lambda - V_Q|$. Mackintosh suggested that the size of the change in α_X should be proportional to the magnitude of the relevant inequality, but gave no specific algorithm. Le Pelley (2004) proposed the following expression to determine changes in α in a manner consistent with the principles suggested by Mackintosh:

$$\Delta\alpha_X = \theta\left(\left|\lambda - V_Q\right| - \left|\lambda - V_X\right|\right) \quad (2)$$

where θ is a constant learning-rate parameter, and α is constrained to lie between .05 and 1.

Consider how Mackintosh's model applies to blocking and control contingencies in which target and competing cues have equal salience. In an A+, AB + blocking contingency, pretraining with A+ will establish A as a good predictor of the outcome. Thus on subsequent AB+ trials the target cue B will be a poorer predictor of the outcome than is A, and so α_B will fall, slowing the rate at which B develops associative strength according to Equation 1. In contrast, in a CD+ control contingency neither cue is preestablished as a better predictor of the outcome and hence both will maintain a relatively high α, such that both can develop associative strength relatively rapidly.

Mackintosh proposed that the salience of a cue determined its starting value of α, with high salience cues having a higher starting α than low salience cues. With regard to the current experiments, this constrains the model to predict greater learning about high salience than low salience blocked cues. If a high salience blocked cue starts with higher α, more will be learned about this cue on the first compound trial than about an otherwise-equivalent low salience cue (by Equation 1). This in turns means that the decline in

[6] Further evidence against this interpretation comes from the finding from Experiment 1 that participants did not perceive the high salience blocked cues as more likely to cause allergic reactions in general. These cues were perceived as significantly less likely than low salience blocked cues to cause the specific outcome with which they had been paired.

α for the high salience blocked cue will be slower (by Equation 2): since more has been learned about the high salience cue, the difference in validity (that is, predictiveness) between this cue and the pretrained competing cue will be smaller. Therefore the high salience blocked cue will maintain a higher α than a low salience blocked cue across compound trials (until α falls to its floor of .05), such that the former will always develop a stronger associative strength. For a similar reason, the model predicts that high salience control target cues will gain more associative strength than low salience control target cues.

However, we can amend this model so that it is more successful in accounting for the current data, by allowing a cue's salience to influence the rate at which α changes. This means that a high salience blocked cue will undergo a more rapid decline in α than a low salience blocked cue, limiting the amount that is subsequently learned about the high salience cue and potentially resulting in—across all compound trials—less learning about the high salience than the low salience blocked cue. This modification of the model is clearly post hoc, but it does have a degree of plausibility. In Mackintosh's (1975) original formulation, a cue's salience influences the rate at which associative strength (*V*) develops (via its influence on starting α values), and hence the rate at which conditioned responses develop. It has recently been argued that allocating attention to a stimulus is a response that is learned like any other (Le Pelley et al., 2013), and in that sense it is consistent to have salience modulating learning of both attentional and behavioral responses.

We must retain from the original Mackintosh model the idea that cue salience also has a more direct, positive influence on the development of conditioned responses. This is so the resulting model can account for the well-established finding that in an overshadowing contingency, a high salience cue will tend to develop a stronger conditioned response than a low salience cue with which it is paired (Mackintosh, 1976); a pattern replicated in the current Experiment 1.

Below we describe three models that combine these two principles—namely (1) salience influences the rate of change of α, and (2) salience also exerts a positive influence on the rate at which conditioned responding develops. The models differ in how they implement the second of these principles. Model A does so by having salience influence the starting value of α, consistent with Mackintosh's (1975) original formulation. Model B instead has salience directly influence responding, following the approach suggested by Le Pelley, Suret, and Beesley (2009). Model C has salience influence the rate of change of associative strength (*V*), as suggested by Rescorla and Wagner (1972; see also Pearce & Hall, 1980). Each of these approaches will, on its own, cause high salience cues to develop conditioned responses more rapidly than low salience cues. As such they implement a positive influence of salience on conditioned responding.

These models are not intended to be definitive accounts of our data, but provide "existence proofs" of the ability of attentional theories of associative learning to account for our key findings.

Model A: Salience Influences Starting α and Rate of Change of α

This model uses the equation for changes in associative strength suggested by Mackintosh (1975):

$$\Delta V_X = \alpha_X \beta(\lambda - V_X) \tag{1}$$

The expression for changes in α is amended to:

$$\Delta \alpha_X = \theta \cdot S_X \cdot \left(\left|\lambda - V_Q\right| - \left|\lambda - V_X\right|\right) \tag{2A}$$

where S_X represents the semantic salience of cue X. Model A follows Mackintosh (1975) in that high salience cues have higher starting α values than low salience cues; this latter assumption is the means by which the model implements the positive influence of salience on responding.

In Model A, the response to a cue X ($Resp_X$) when it is presented is given simply by the associative strength of that cue, i.e.:

$$Resp_X = V_X \tag{3}$$

Strong empirical evidence suggests that eye gaze is a function of both the semantic salience of a cue (Daffner, Scinto, Weintraub, Guinessey, & Mesulam, 1994; Loftus & Mackworth, 1978) and the α value of that cue (Beesley & Le Pelley, 2011; Le Pelley et al., 2011; Rehder & Hoffman, 2005; Wills et al., 2007). Hence we model eye gaze on cue X (E_X) as:

$$E_X = S_X \cdot \alpha_X \tag{4}$$

For each trial, we then calculated proportional gaze strength for the target cue (*Prop_target*) X within compound XY as:

$$Prop_target = \frac{E_X + \varepsilon}{(E_X + \varepsilon) + (E_Y + \varepsilon)} \tag{5}$$

This value is an analogue of the *Prop_target* data calculated from empirical dwell times (see Figure 7); a value greater than .5 indicates greater dwell time on the target cue, and a value less than .5 indicates greater gaze on the competing cue. The constant ε added to each term acts to compress the simulated data; in effect it simulates participants looking at each cue for a fixed minimum period on each trial in addition to the gaze represented by *E*.

Simulations using the model described above (data available on request) provided a good match to participants' behavioral response data, and most of the eye gaze data. However, unlike in the empirical data (Figure 7A), for the Control High condition simulated *Prop_target* did not decrease toward .5 as training progressed; it began high and increased gradually. In order to better fit the empirical data, we allowed for habituation to cues in the model. Over repeated presentations of a salient stimulus, the unconditioned response to that stimulus declines (Phan, Liberzon, Welsh, Britton, & Taylor, 2003; Sokolov, 1963; Thompson & Spencer, 1966); in effect, the salience of the stimulus decreases. In the context of our experiments, while a picture of a flamingo as food might be noteworthy the first time it is seen, it will not be so noteworthy on its sixth presentation. To allow for habituation in our simulations, after each presentation the salience of each cue X was updated according to:

$$\Delta S_X = -H(S_X - S_{\min}), \tag{6}$$

where *H* is a constant controlling habituation rate, between 0 (no habituation) and 1 (total habituation in a single trial). The constant S_{min} reflects the minimum salience toward which cues converge via habituation.

Simulations of blocking contingencies had eight trials with the competing cue, followed by six compound trials; control contingencies had six compound trials only. Parameters were $\lambda = 1$, $\beta = .2$, $\theta = .6$, $\varepsilon = .3$, $H = .5$. Low salience cues had $S = .3$ and starting $\alpha = .8$; high salience cues had $S = .9$ and starting $\alpha = .9$. S_{min} was set to .3, the salience value of low salience cues.

Figure 8 (left-hand column) shows results of a simulation using this model; changes in response strength (*Resp*; Panel A), α (B), and *Prop_target* (C) for target cues across compound training. For the Control Low contingency, both target and competing cues begin with equal salience and equal associative strength ($V = 0$ for both, since they are novel). Hence these cues are equivalent and are learned about at the same rate. Since neither is more or less predictive than the other at any stage, their α remains unchanged throughout compound training. For the Control High contingency, the higher starting α of the high salience target cue means initial learning about this cue is slightly more rapid than for the low salience competing cue. Therefore on subsequent trials this target cue is a (slightly) better predictor of the outcome than is the competing cue, and so α for the target cue rises gradually across training. The upshot is that at the end of training the response to the Control High target cue will be slightly greater than to the Control Low target cue.

Turning to the Blocking Low contingency, the target cue begins compound training as a much poorer predictor of the outcome than is the pretrained competing cue. Consequently the target cue's α falls over training, limiting its gain of associative strength. In comparison, the Blocking High target cue has an advantage on the first compound trial due to its higher starting α, but its higher salience means that α for this cue declines more rapidly (by Equation 2A), greatly slowing learning on later compound trials. The overall result is that, at the end of training, the response to the Blocking High target cue is considerably lower than that to the Blocking Low target cue.

With regard to responses, this model mirrors several aspects of our empirical data, namely: (i) Blocking (difference between control and blocked target cue) is greater for high salience targets than for low salience targets (Experiments 1 and 2); (ii) Responding to high salience blocked target cues is lower than to low salience blocked targets (Experiments 1 and 2); (iii) Responding to high salience control target cues is greater than to low salience control targets (approached significance in Experiment 1). Moreover, the difference

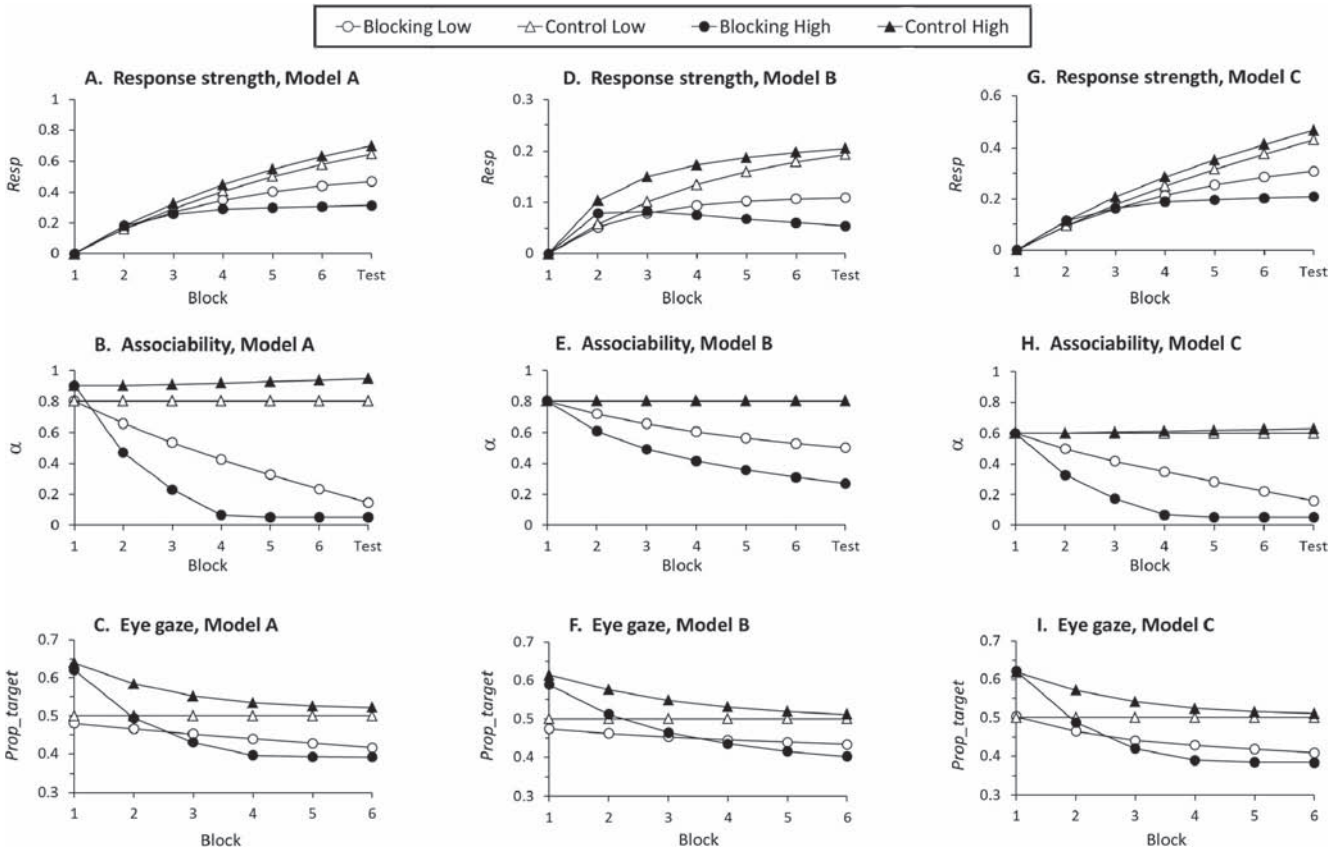


Figure 8. Results of simulations using Model A (Panels A-C), Model B (Panels D-F), and Model C (Panels G-I). Panels A, D and G show response strength (*Resp*) for target cues of each contingency in each block of compound training, and on test (i.e. following the final block of compound training). Panels B, E, and H show associability (α) of target cues of each contingency in each block of compound training, and on test. Panels C, F, and I show proportion of eye gaze on the target cue (*Prop_target*) calculated according to Equation 5, for each block of compound training.

in responding to low and high salience control targets is simulated to be smaller than the difference for low and high salience blocked targets, which could account for the fact that empirical evidence for this difference in control targets was relatively weak in Experiment 1, and no significant difference was found in Experiment 2.

Figure 8C shows that *Prop_target* is high in block 1 for trials involving high salience targets because these targets have greater S and starting α than the competing cues with which they are paired. *Prop_target* then falls rapidly on Blocking High trials due to the rapid reduction in α for high salience blocked target cues (see Figure 8B). *Prop_target* for Blocking Low trials begins near .5 and then falls more gradually as a consequence of the slower reduction in α for low salience blocked cues. *Prop_target* for Control Low trials remains at .5 throughout, since cues on these trials do not change in α during training. Finally, *Prop_target* falls across Control High trials due to decreases in salience S for these cues (by Equation 6), since S influences eye gaze by Equation 4. These patterns are similar to those observed empirically, especially for the data with chicken as a cue excluded (Figure 7B); namely: (i) Gaze in block 1 is primarily determined by salience and not by contingency; (ii) Gaze in later blocks is determined more strongly by contingency than by salience; (iii) *Prop_target* falls across trials for Blocking High and Control High contingencies, but this decline is greater for Blocking High trials; (iv) *Prop_target* in block 1 is greater than .5 for trials with high salience targets, and does not differ from .5 for trials with low salience targets; (v) *Prop_target* across blocks 2–6 is less than .5 for blocking contingencies, greater than .5 for the Control High contingency, and does not differ from .5 for the Control Low contingency. The simulation is not perfect; notably, it predicts that from block 3 onward *Prop_target* will be lower for Blocking High trials than for Blocking Low trials, whereas the empirical data show a trend in the opposite direction in blocks 3 and 4. However, the difference between these trial types was not significant in any of blocks 2–6, even using p values uncorrected for multiple comparisons (largest $t(39) = 1.96$ for data in Figures 7A and 7B; all $ps > .05$), and the simulated difference between them is also relatively small. Hence we do not see this as a fatal weakness of the model.

Model B: Salience Influences Responding and Rate of Change of α

We now present an alternative attentional model that can also account for our empirical findings. As for Model A, Equations 1 and 2A determine changes in V and α respectively. In Model B, all cues begin with equal α. The positive influence of salience on responding is instead implemented by having salience modulate response strength, *Resp*. We also assume *Resp* is modulated by α:

$$Resp_X = V_X \cdot S_X \cdot \alpha_X \tag{7}$$

The idea that *Resp* is influenced by α is supported empirically (Le Pelley et al., 2009). This increases the extent to which α modulates behavior, since it influences the rate of association-formation (Equation 1), and expression of these associations (Equation 7). Eye gaze E is determined as for Model A (Equations 4 and 5), as is habituation to cues (Equation 6).

Figure 8D-F shows simulation results from Model B, with parameters $\lambda = 1$, $\beta = .3$, $\theta = .3$, $\varepsilon = .3$, $H = .4$, $S_{min} = .3$. Low salience cues began with $S = .3$ and high salience cues with $S = .7$. All cues began with $\alpha = .8$. Changes in *Resp* and *Prop_target* across training are similar to those for Model A, and Model B has similar success in explaining the key aspects of our empirical data.

Model C: Salience Influences Learning and Rate of Change of α

In Model C the positive influence of salience on responding is implemented by having it modulate changes in associative strength, V (cf. Rescorla & Wagner, 1972); Equation 1 becomes:

$$\Delta V_X = (S_X)^k \alpha_X \beta(\lambda - V_X) \tag{8}$$

The simulation reported here used $k = 0.2$. A parameterization with $0 < k < 1$ is necessary to ensure that the effect of salience on changes in α for blocked cues is not outweighed by its effect on changes in V according to Equation 8. All other details of Model C are as for Model A.

Figure 8G-I shows simulation results for Model C with $\lambda = 1$, $\beta = .2$, $\theta = .6$, $\varepsilon = .3$, $H = .5$, $S_{min} = .2$. Low salience cues began with $S = .3$; high salience cues with $S = .8$. All cues began with $\alpha = .6$. Changes in *Resp* and *Prop_target* across training are similar to those for Models A and B, and Model C accounts for the same aspects of the empirical data.

Modeling Summary

Models A−C all incorporate the idea that salience influences changes in α (Equation 2A), allowing each model to explain our finding of weaker responding to high salience than low salience blocked cues. The models differ in how they implement a more direct, positive influence of salience on responding. Regardless of how this positive influence is implemented, with regard to response strength all models successfully accounted for the key findings of our experiments.

It is worth noting that Models A−C tend to behave similarly to the standard Mackintosh (1975) model in situations where cue salience is not systematically varied, which is the case in most studies of associative learning. This is because in such cases S is equal for all cues and so will not differentially affect the model's predictions for different cues. Hence Models A−C inherit the Mackintosh model's previously established explanatory abilities (see Hall, 1991; Le Pelley, 2004; Mitchell & Le Pelley, 2010, for reviews).

It is likely that other models implementing the general principles described above could also account for our data. For example, the model of Pearce and Hall (1980) could be adapted in a similar fashion to allow for more rapid changes in attention to high salience cues. It is even possible that retrieval-based models of cue competition (e.g., Miller & Matzel, 1988) could be adapted along analogous lines, although we have not explored this possibility. We have focused on Mackintosh's (1975) theory here because it provides a relatively simple model of attentional processes, and because previous findings from human contingency learning have tended to support this account over Pearce and Hall's theory (see Le Pelley, 2010, for a review).

On a related note, while the models we have presented here are framed in terms of the formation of associations between representations of stimuli and outcomes (what Mitchell, De Houwer, & Lovibond, 2009, refer to as "link-formation" models), it would be possible to implement these ideas in terms of inferential reasoning. Mitchell et al. argued that blocking arises because participants draw an inference that the blocked cue is redundant. In order to account for reduced

responding to high salience blocked target cues, one would have to assume that participants are more likely or faster to draw such inferences regarding high salience cues. The suggestion that salience influences the rate at which people deduce the predictive status of cues is clearly similar to the central idea of Models A−C that salience influences the rate of change of attention to cues, although it lacks the formal specification and detailed description of mechanism offered by the link-based account.

Conclusions

Previous studies of cue competition in humans and nonhuman animals have established a critical role for both relative salience and relative validity in associative learning. The current experiments demonstrate that these properties interact—that the influence of relative salience depends on the relative validity of the cues in question. If cue X is less valid than cue Y (as in the case of blocking), then an increase in the salience of X leads to weaker responding to this cue (consistent with a reduction in learning about X); if X and Y are equal in terms of validity, then an increase in the salience of X leads to (if anything) stronger responding to this cue. We have presented three models based on Mackintosh's (1975) attentional theory of associative learning that can account for this finding. Essentially these models propose that salience influences the rate at which participants determine a cue's predictive (or nonpredictive) status such that if a cue is redundant—as in a blocking contingency—participants will come to this conclusion more rapidly if the cue is of high salience than if it is of low salience. Other accounts of these data may be possible, but we believe that any successful account will ultimately rely on an interaction between salience and validity.

References

Beesley, T., & Le Pelley, M. E. (2011). The influence of blocking on overt attention and associability in human learning. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 114–120. doi:10.1037/a0019526

Daffner, K. R., Scinto, L. F. M., Weintraub, S., Guinessey, J., & Mesulam, M. M. (1994). The impact of aging on curiosity as measured by exploratory eye-movements. *Archives of Neurology, 51,* 368–376. doi:10.1001/archneur.1994.00540160062009

Denton, S. E., & Kruschke, J. K. (2006). Attention and salience in associative blocking. *Learning & Behavior, 34,* 285–304. doi:10.3758/BF03192884

Feldman, J. M. (1975). Blocking as a function of added cue intensity. *Animal Learning & Behavior, 3,* 98–102. doi:10.3758/BF03209108

Griffiths, O., & Le Pelley, M. E. (2009). Attentional changes in blocking are not a consequence of lateral inhibition. *Learning & Behavior, 37,* 27–41. doi:10.3758/LB.37.1.27

Hall, G. (1991). *Perceptual and associative learning*. Oxford, UK: Oxford University Press. doi:10.1093/acprof:oso/9780198521822.001.0001

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.

Kelley, H. H. (1972). Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 1–26). Morristown, NJ: General Learning Press.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science, 12,* 171–175. doi:10.1111/1467-8721.01254

Kruschke, J. K., & Blair, N. J. (2000). Blocking and backward blocking involve learned inattention. *Psychonomic Bulletin & Review, 7,* 636–645. doi:10.3758/BF03213001

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. doi:10.1037/0278-7393.31.5.830

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

Le Pelley, M. E. (2010). Attention and human associative learning. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 187–215). Oxford, UK: Oxford University Press.

Le Pelley, M. E., Beesley, T., & Griffiths, O. (2011). Overt attention and predictiveness in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes, 37,* 220–229. doi:10.1037/a0021384

Le Pelley, M. E., Beesley, T., & Suret, M. B. (2007). Blocking of human causal learning involves learned changes in stimulus processing. *The Quarterly Journal of Experimental Psychology, 60,* 1468–1476. doi:10.1080/17470210701515645

Le Pelley, M. E., & McLaren, I. P. L. (2001). Retrospective revaluation in humans: Learning or memory? *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 54,* 311–352.

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

Le Pelley, M. E., Mitchell, C. J., & Johnson, A. M. (2013). Outcome value influences attentional biases in human associative learning: Dissociable effects of training and of instruction. *Journal of Experimental Psychology: Animal Behavior Processes, 39,* 39–55. doi:10.1037/a0031230

Le Pelley, M. E., Reimers, S. J., Calvini, G., Spears, R., Beesley, T., & Murphy, R. A. (2010). Stereotype formation: Biased by association. *Journal of Experimental Psychology: General, 139,* 138–161. doi:10.1037/a0018210

Le Pelley, M. E., Suret, M. B., & Beesley, T. (2009). Learned predictiveness effects in humans: A function of learning, performance, or both? *Journal of Experimental Psychology: Animal Behavior Processes, 35,* 312–327. doi:10.1037/a0014315

Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 565–572. doi:10.1037/0096-1523.4.4.565

MacKay, D. M. (1973). Lateral interaction between neural channels sensitive to texture density. *Nature, 245,* 159–161. doi:10.1038/245159a0

Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82,* 276–298. doi:10.1037/h0076778

Mackintosh, N. J. (1976). Overshadowing and stimulus intensity. *Animal Learning & Behavior, 4,* 186–192. doi:10.3758/BF03214033

Mackintosh, N. J. (1978). Cognitive or associative theories of conditioning: Implications of an analysis of blocking. In H. Fowler, W. K. Honig, & S. H. Pulse (Eds.), *Cognitive processes in animal behavior* (pp. 155–175). Hillsdale, NJ: Erlbaum.

Mackintosh, N. J., & Turner, C. (1971). Blocking as a function of novelty of CS and predictability of UCS. *The Quarterly Journal of Experimental Psychology, 23,* 359–366. doi:10.1080/14640747108400245

Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. *The Psychology of*

Learning and Motivation, 22, 51–92. doi:10.1016/S0079-7421(08)60038-9

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences, 32,* 183–198. doi:10.1017/S0140525X09000855

Mitchell, C. J., & Le Pelley, M. E. (Eds.). (2010). *Attention and associative learning: From brain to behaviour*. Oxford, UK: Oxford University Press.

Mitchell, C. J., Lovibond, P. F., Minard, E., & Lavis, Y. (2006). Forward blocking in human learning sometimes reflects the failure to encode a cue-outcome relationship. *The Quarterly Journal of Experimental Psychology, 59,* 830–844. doi:10.1080/17470210500242847

Moran, P. M., Al-Uzri, M. M., Watson, J., & Reveley, M. A. (2003). Reduced Kamin blocking in non paranoid schizophrenia: Associations with schizotypy. *Journal of Psychiatric Research, 37,* 155–163. doi:10.1016/S0022-3956(02)00099-7

Pavlov, I. P. (1927). *Conditioned reflexes*. London, UK: Oxford University Press.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review, 94,* 61–73. doi:10.1037/0033-295X.94.1.61

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. doi:10.1037/0033-295X.87.6.532

Phan, K. L., Liberzon, I., Welsh, R. C., Britton, J. C., & Taylor, S. F. (2003). Habituation of rostral anterior cingulate cortex to repeated emotionally salient pictures. *Neuropsychopharmacology, 28,* 1344–1350. doi:10.1038/sj.npp.1300186

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology, 51,* 1–41. doi:10.1016/j.cogpsych.2004.11.001

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, NY: Appleton-Century-Crofts.

Shanks, D. R. (1985). Forward and backward blocking in human contingency judgement. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 37,* 1–21.

Sokolov, E. N. (1963). *Perception and the conditioned reflex*. Oxford, UK: Pergamon Press.

Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval (SOCR): A formalization of the comparator hypothesis. *Psychological Review, 114,* 759–783. doi:10.1037/0033-295X.114.3.759

Thompson, R. F., & Spencer, W. A. (1966). Habituation: A model phenomenon for study of neuronal substrates of behavior. *Psychological Review, 73,* 16–43. doi:10.1037/h0022681

Waelti, P., Dickinson, A., & Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. *Nature, 412,* 43–48. doi:10.1038/35083500

Wagner, A. R. (1981). SOP: A model of automatic memory processing in animal behaviour. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Erlbaum, Inc.

Waldmann, M. R. (2001). Predictive versus diagnostic causal learning: Evidence from an overshadowing paradigm. *Psychonomic Bulletin & Review, 8,* 600–608. doi:10.3758/BF03196196

Wasserman, E. A. (1990). Attribution of causality to common and distinctive elements of compound stimuli. *Psychological Science, 1,* 298–302. doi:10.1111/j.1467-9280.1990.tb00221.x

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. doi:10.1162/jocn.2007.19.5.843

(*Appendix follows*)

Appendix

Mean Salience and Standard Deviation of Each of the Foods Tested in Pilot Work

Superscripts indicate the foods used as cues in Experiments 1A, 1B and 2. An underlined subscript indicates that the food was used as a target cue in that experiment.

Food	Mean salience	Standard deviation	Food	Mean salience	Standard deviation
Bread[1A, 2]	1.53	1.87	Rhubarb	5.00	2.40
Rice[1A, 1B, 2]	1.58	1.64	Passion fruit	5.21	2.49
Potato[1A, 1B, 2]	1.89	2.13	Fennel	5.53	2.61
Apples[1A, 1B, 2]	2.16	1.68	Star Fruit	5.68	2.71
Cucumber[1A, 1B, 2]	2.16	1.57	Lobster	5.74	2.94
Tomato [1A, 1B]	2.16	1.64	Black Pudding	6.11	2.62
Onion[1A, 1B, 2]	2.21	1.58	Prawn Heads	6.58	2.19
Chicken[1A, 1B, 2]	2.26	2.10	Monkey Nuts	6.74	3.21
Grapes[1A, 1B]	2.26	1.76	Whitebait	6.79	2.46
Pasta[1A, 1B, 2]	2.32	2.11	Haggis	6.84	2.71
Eggs[1A, 1B, 2]	2.42	1.80	Ugli Fruit	7.26	2.28
Lettuce [1A, 1B, 2]	2.47	2.01	Stinging Nettles	7.63	2.67
Carrots [1A, 1B, 2]	2.53	2.09	Dandelions	7.68	2.14
Banana[1A, 2]	2.68	2.24	Goose Liver	7.68	2.00
Garlic	2.74	2.02	Snails	7.95	1.58
Mushrooms[1A, 1B]	2.74	2.62	Frog's Legs	8.00	1.73
Yoghurt[1B]	2.79	1.84	Locusts	8.16	1.77
Peaches	3.05	2.15	Ants	8.37	2.56
Noodles	3.11	2.05	Worms	8.42	1.68
Fish[1B]	3.17	2.60	Brains [1A, 2]	8.68	1.73
Ham [1A, 1B]	3.21	2.53	Caterpillars [1A, 1B, 2]	8.79	1.44
Melon	3.37	2.36	Pig's Ears [2]	8.84	1.30
Broccoli	3.42	2.65	Snake	8.89	1.63
Sardines	4.11	2.49	Ducks' Tongues [1A]	8.95	1.35
Steak	4.11	3.21	Blood [1B]	9.00	1.53
Honey	4.21	2.46	Fried Spiders [1B]	9.00	1.97
Lentils	4.32	2.67	Rhinoceros	9.00	1.49
Olives	4.47	2.65	Bull's Testicles	9.05	1.31
Asparagus	4.53	2.76	Flamingo [1A, 2]	9.11	1.52
Dates	5.00	2.56	Sheep's Eyeballs [1B]	9.11	1.10

Received January 23, 2013
Revision received August 27, 2013
Accepted August 27, 2013 ■