Journal of Experimental Psychology:
Animal Behavior Processes
2009, Vol. 35, No. 3, 312–327

0097-7403/09/$12.00 DOI: 10.1037/a0014315

Learned Predictiveness Effects in Humans: A Function of Learning, Performance, or Both?

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

Many previous studies of animal and human learning indicate a processing advantage for cues previously experienced as good predictors of outcomes over those experienced as poorer predictors. Four studies of human associative learning investigated whether learned predictiveness acts at the level of learning (modulating the rate at which cue–outcome associations form), performance (modulating the strength of behavioral responses), or both. In Experiments 1–3, it was found that retrospectively altering the learned predictiveness of cues influenced responding to those cues, demonstrating that learned predictiveness influences performance. Experiment 4 indicates that learned predictiveness also influences learning by demonstrating that the learned predictiveness of a cue affects the acquisition of an association between a novel cue and the outcome with which it is paired.

Keywords: human associative learning, cue–outcome association, learned predictiveness

It is well established in the field of animal conditioning that the amount of processing power devoted to learning about a given conditioned stimulus (CS) can be influenced by its past history of predictiveness at an associative level. Establishing a CS as a predictor of a reinforcing event seems to alter the readiness with which that stimulus will engage in later learning (see Le Pelley, 2004, for a review). One model of such learned predictiveness effects is that of Mackintosh (1975), which states that the change (Δ) in associative strength of cue P (V_P) on each learning episode is given by

$$\Delta V_P = \theta\, \alpha_P\, (\lambda - V_P) \quad (1)$$

where θ is a constant learning rate parameter, λ is the asymptote of conditioning supportable by the outcome occurring on that trial, and α_P is the associability of cue P. Mackintosh proposed that the associability of a cue varies as a function of that cue's experienced predictive ability. Specifically, α_P increases if P is a better predictor of the outcome occurring on a given trial than are all other presented cues; α_P decreases if P is a poorer predictor of the outcome than are other presented cues. The extent to which the outcome is predicted by P is given by the discrepancy between the current state of the outcome (λ) and the extent to which P predicts that outcome (V_P), that is, the absolute value of the error term $|\lambda - V_P|$. Hence, on each trial, the α value for each presented cue is updated according to the following rules:

$$\mathrm{D}a_P > 0 \text{ if } |\lambda - V_P| < |\lambda - V_Q|$$

$$\Delta\alpha_P < 0 \text{ if } |\lambda - V_P| > |\lambda - V_Q| \quad (2)$$

where V_Q is the associative strength of all stimuli other than P present on that trial.

Several recent studies have indicated that learned predictiveness processes also exert an influence on human associative learning (e.g., Bonardi, Graham, Hall, & Mitchell, 2005; Griffiths & Le Pelley, in press; Le Pelley, Beesley, & Suret, 2007; Le Pelley & McLaren, 2003; Le Pelley, Oakeshott, & McLaren, 2005; Lochmann & Wills, 2003). The results of these experiments are consistent with the idea that people learn more rapidly about cues that have previously been established as reliable predictors of outcomes than those established as poor predictors. As such, these findings agree with the predictions of the Mackintosh (1975) model.[1]

We consider Le Pelley and McLaren's (2003) study in detail here, as it forms the focus of the present experiments. The basic design of this study is shown in Table 1, where letters A–U refer to cues and O1–O4 refer to outcomes that can be paired with those cues. Thus, "AR–O1" indicates that cues A and R were presented together and were paired with outcome O1. The experiment used a multiple outcome allergy prediction paradigm, in which participants played an allergist predicting the type of allergic reaction that a fictitious patient would suffer after eating different foods.

M. E. Le Pelley, M. B. Suret, and T. Beesley, School of Psychology, Cardiff University, Cardiff, Wales.

This work was supported by Grant RES000230983 from the Economic and Social Research Council to M. E. Le Pelley. We thank Mark Haselgrove for his valuable comments on a draft of this article.

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

[1] This theoretical analysis is rather selective. Sutherland and Mackintosh (1971) and Kruschke (2001) have proposed alternative accounts of learned predictiveness effects, but these accounts are fundamentally similar to the Mackintosh (1975) model. The theoretical approach offered by Pearce and Hall (1980), on the other hand, takes a very different view of such effects (see Le Pelley, 2004). However, certain existing experiments (e.g. Bonardi et al, 2005; Griffiths & Le Pelley, in press; Le Pelley & McLaren, 2003), and those described in the present article, provide support for the general account of such effects offered by the Mackintosh model, but are inconsistent with the view taken by Pearce and Hall. Consequently, our discussion focuses solely on the general approach offered by the former.

Table 1
Design of Le Pelley and McLaren (2003)

Stage 1	Stage 2	Test
AR–O1	AT–O3	AC
AS–O1	BU–O4	BD
BR–O2	CR–O3	RT
BS–O2	DS–O4	SU
CT–O2		
CU–O2		
DT–O1		
DU–O1		

Note. Letters A to U represent different food types; O1 to O4 refer to the type of allergy produced (outcome) when the food was eaten by a fictitious patient. On test, ratings of the compounds were obtained with respect to outcomes O3 and O4. Filler trials used in Stage 2 by Le Pelley and McLaren (2003) are omitted for clarity.

Hence, the foods constituted the cues and the different types of allergic reactions were the outcomes.

During the first stage of this experiment, cues A and D were consistently paired with O1, cues B and C were consistently paired with O2, and cues R–U provided no basis for discrimination between the two outcomes, being paired with O1 and O2 an equal number of times. As such, during Stage 1, cues A–D were the best available predictors of the outcome occurring on each trial and hence, according to the Mackintosh model, should have maintained a high α. Meanwhile, the α of cues R–U should have decreased, as these were the poorer predictors of the outcome occurring on each trial.

On each of the Stage 2 trial types shown in Table 1, a predictive cue from Stage 1 (A, B, C, or D) was paired with a nonpredictive cue (R, S, T, or U) with which it had not been presented in Stage 1, and this novel compound was paired with a novel outcome; compounds AT and CR with O3 and compounds BU and DS with O4. After Stage 2, participants were asked to rate how likely each of outcomes O3 and O4 was to follow cue compounds AC, BD, RT, and SU. Following Dickinson, Shanks, and Evenden (1984), these ratings are taken to provide an index of the strength of the cue–outcome associations developed over the course of training.

The Mackintosh model predicts that, at the end of Stage 1, cues A–D (predictive in Stage 1) will have higher associabilities than cues R–U (nonpredictive in Stage 1). Assuming that these associabilities will generalize between the contexts of Stages 1 and 2 (see Le Pelley, Oakeshott, Wills, & McLaren, 2005), this will promote more rapid learning of associations between cues A–D and the Stage 2 outcomes than between cues R–U and Stage 2 outcomes. Therefore, participants should develop strong associations from A and C to O3, strong associations from B and D to O4, weak associations from R and T to O3, and weak associations from S and U to O4. In line with these predictions, participants in Le Pelley and McLaren's (2003) study rated compound AC as a strong predictor of O3 and compound BD as a strong predictor of O4, whereas RT and SU were perceived to be only weak predictors of O3 and O4, respectively.

This finding clearly indicates a difference in the processing of cues A–D and cues R–U and that this difference arises as a result of the difference in the learned predictiveness of these cues during Stage 1. The question then becomes one of exactly where this learned predictiveness exerts its influence. The Mackintosh model as presented earlier (and as presented originally by Mackintosh, 1975) states that α influences *learning*, determining how rapidly a cue undergoes changes in associative strength (as indicated by the term *associability*, i.e., the readiness with which a cue will enter into an association). If two cues with different α values are presented simultaneously and reinforced, the cue with the higher α will develop a stronger association to the outcome than will the cue with the lower α. In this conceptualization of the model, responding to cue P, R_P, is simply a function of that cue's associative strength. In other words,

$$R_P = k V_P \quad (3)$$

where k is a constant.[2]

There exists an alternative view of the locus of learned predictiveness, however. Mackintosh (1975) raised the possibility that learned predictiveness may also influence performance, modulating the response to a cue. That is, responding may also be a function of α. In other words,

$$R_P = k \alpha_P V_P. \quad (4)$$

In the absence of compelling experimental evidence to support the idea that α influences performance as well as learning, however, Mackintosh remained agnostic on this issue. Note that here we are referring to a *direct* influence of α on performance, separate from its effect on learning. Even in the original Mackintosh model, which combines the learning rule of Equation 1 with the performance rule of Equation 3, α has an indirect effect on performance that results from its direct effect on learning. That is, a cue with high α will be learned about rapidly (i.e., it will form strong associations) and hence, according to Equation 3, will support strong responding. This is quite different from the influence of α on performance suggested by Equation 4, in which α has a direct effect in modulating the expression of a learned association.

This raises the issue of how best to interpret the findings of Le Pelley and McLaren (2003). In the discussion earlier, we have appealed to a model that implicates α in learning only. That is, cues that were experienced as predictive in Stage 1 were assumed to engage the learning process more strongly in Stage 2 than those experienced as nonpredictive. Responding to predictive cues would then be greater on test, as these cues would have higher associative strengths. It is, however, also possible to account for these data using a model that implicates α in performance only without influencing learning at all. Suppose that, as before, cues A–D develop higher α values than do cues R–U during Stage 1 (by virtue of the fact that the former are consistently paired with the same outcomes, whereas the latter are not). If α does not influence learning, then, during Stage 2, all cues will form equally strong associations to the outcomes with which they are paired. If α influences responding as in Equation 4, then responding to the predictive cues from Stage 1 on test will be greater than to the nonpredictive cues. For example, for cues A and T on test:

[2] Strictly, this should be expressed as $R_P = f(V_P)$, but as long as the function relating R and V is monotonic, an increase in V will always lead to an increase in R. For the sake of simplicity, we characterize this in the form of a linear relationship as in Equation 3 throughout this article.

$$\text{If } \alpha_A > \alpha_T \text{ and } V_A = V_T,$$

then, by Equation 4, $R_A > R_T$.

Thus, it is theoretically possible to account for Le Pelley and McLaren's findings using a model that makes no recourse to α in the learning mechanism.

The suggestion that learned predictiveness influences performance as well as learning has been explicitly formalized by Kruschke (1996, 2001) in his ADIT and EXIT models. Despite such formal statements regarding the locus at which α operates, there currently exists no conclusive empirical evidence bearing on this issue. Moreover, this ambiguity of interpretation applies not only to Le Pelley and McLaren's (2003) study but (to the best of our knowledge) to all previous demonstrations of learned predictiveness effects in both humans and animals that have been taken in support of the Mackintosh model: As far as we are aware, all such effects currently described in the literature could be explained by a model implicating α in learning (the formation of associations) only, in performance (the expression of associations as behavior) only, or in both learning and performance.

The experiments described in this article represent a first attempt to decide between these alternative views of the locus of α in human associative learning. Experiments 1–3 investigate whether α does indeed have a direct influence on performance that is independent of its effect on learning. Experiment 4 then attempts to establish evidence for an influence of α on learning.

Experiment 1

The design of Experiment 1 is shown in Table 2. This experiment used a "mad scientist" paradigm, with each participant playing a scientist who creates mutants by combining certain chemicals with a special "goo" substance. Thus, the letters A–U in Table 2 were represented by different chemicals and outcomes O1-O6 were represented by different types of mutants that could be created. The context for each separate stage of the experiment was provided by a particular color of goo, with different colors giving rise to different types of mutant.

Table 2
Design of Experiment 1

		Stage 3		
Stage 1	Stage 2	Consistent	Inconsistent	Test
AR–O1	AT–O3	AR–O5	AR–O5	AC
AS–O1	BU–O4	AS–O5	AS–O6	BD
BR–O2	CR–O3	BR–O6	BR–O5	RT
BS–O2	DS–O4	BS–O6	BS–O6	SU
CT–O2		CT–O6	CT–O6	
CU–O2		CU–O6	CU–O5	
DT–O1		DT–O5	DT–O6	
DU–O1		DU–O5	DU–O5	

Note. Consistent and inconsistent refer to the different types of Stage 3 training received by the two participant groups of Experiment 1. Letters A to U represent different chemicals; O1 to O4 refer to the type of mutant that was created (outcome) when these chemicals were mixed with the goo. On test, ratings of the compounds were obtained with respect to outcomes O3 and O4.

Stages 1 and 2 of Experiment 1 were similar to those used by Le Pelley and McLaren (2003) and follow the logic described in the introduction. After Stage 1, participants completed a third stage of training, in which cue compounds were paired with novel outcomes (O5 and O6). For participants in the consistent group, during Stage 3, cues A–D were once again predictive of outcomes (A and D were consistently paired with O5; B and C were consistently paired with O6), whereas cues R–U were nonpredictive (being paired with O5 and O6 an equal number of times). For these participants, predictiveness in Stage 3 was consistent with what was learned in Stage 1. For participants in the inconsistent group, the reverse was true; cues A–D were nonpredictive, whereas cues R–U were predictive in Stage 3. For this group, predictiveness in Stage 3 was inconsistent with what was learned in Stage 1. After Stage 3, participants were asked to rate how likely each of the Stage 2 outcomes (O3 and O4) was to follow compounds AC, BD, RT, and SU. This was the same test as used by Le Pelley and McLaren (2003). The question of interest is what effect Stage 3 training had on the pattern of responding to these compounds.

For both groups, at the outset of Stage 2, the α values for cues A–D should be higher than those for R–U as a result of the higher predictiveness of the former cues during Stage 1. If α influences learning, then this will promote more rapid learning of associations between cues A–D and the Stage 2 outcomes than between cues R–U and the same outcomes. The α values of the cues should diverge in the two groups during Stage 3. In the consistent group, cues A–D will maintain a high α during Stage 3 and cues R–U will maintain a low α. In the inconsistent group, on the other hand, we may expect the α of cues R–U to rise over Stage 3 and the α of A–D to fall (reflecting the reversed predictiveness of these cues).

If α affects learning only, then there is no way for these subsequent changes in α to influence participants' reports of the relationship between test cues and Stage 2 outcomes. Therefore, we would expect similar results in the consistent and inconsistent groups, with both groups providing higher ratings for compounds AC and BD than for compounds RT and SU. Note that training in Stage 3 cannot directly influence the cue–outcome associations formed in Stage 2, as the outcomes used in the two stages are (a) different and (b) statistically independent; that is, cues that were paired with outcome O3 during Stage 2 were equally likely to be paired with outcome O5 or O6 in Stage 3 (and the same was true for cues paired with O4 in Stage 2). Consequently, simply learning that a particular chemical predicts O5 in the Stage 3 context itself tells a participant nothing about the effect of that chemical in the Stage 2 context. This rules out the possibility that learning during Stage 3 may have any direct influence on the Stage 2 cue–outcome associations through a process akin to retrospective revaluation of causal judgments (e.g., Melchers, Lachnit, & Shanks, 2004). This statistical independence of outcomes, coupled with our use of difference scores to analyze the rating data on test (as explained later in the *Data analysis* section), also rules out the possibility that learning about specific cue–outcome associations during Stage 3 may have an indirect influence on participants' judgments of the Stage 2 cue–outcome relationships (which involve the same cues but different outcomes) through a process of generalization. To summarize, if response is purely a function of associative strength, then there is no obvious way that changes in α after Stage 2 can influence response on test.

Suppose instead that α affects performance only. In that case, all cues will develop equally strong associations to whichever of outcomes O3 and O4 they are paired with during Stage 2. Changes in α during Stage 3 will influence the extent to which these associations are expressed in ratings provided on test, according to Equation 4. In the consistent group, we would expect stronger responding to compounds made up of cues A–D (as these cues maintain high α during Stage 3) than to compounds made up of cues R–U (which maintain low α during Stage 3). In the inconsistent group, the α values of the cues at the time of test will be quite different. To the extent that the α of cues R–U ends Stage 3 higher than that of cues A–D, we would expect participants to give higher ratings for compounds RT and SU than for AC and BD. That is, the reversal of the α values of the component cues during Stage 3 should interfere with the effect observed by Le Pelley and McLaren (2003), reducing (and possibly reversing) the advantage for AC/BD over RT/SU.

In summary, if the locus of α is at the level of learning only, then we would expect the advantage for AC/BD over RT/SU to be equivalent in the consistent and inconsistent groups. If instead α acts only at the performance level, modulating the expression of learned associations, then we would expect a greater advantage for AC/BD over RT/SU in the consistent group than in the inconsistent group (and possibly even a reversal of this advantage in the latter group). A third possibility, as noted earlier, is that α influences both learning and performance. To the extent that α has *any* influence on performance, we would still expect inconsistent training to reduce the advantage for AC/BD over RT/SU, compared with consistent training, given that changes in α during Stage 3 would still be expected to modulate the expression of any associations that were formed during Stage 3. Consequently, this selective influence of Stage 3 training on the pattern of responding to the different compounds on test provides an assay of whether α has any direct effect on performance. That said, if α also influences learning, we may expect the interfering effect of Stage 3 training in the inconsistent group to be less extreme. This is because any influence of α on learning ensures that cues A–D develop stronger associations during Stage 2 than do cues R–U, which will tend to produce an advantage for AC/BD over RT/SU in the inconsistent group (albeit an advantage that will subsequently be offset by the influence of changes in α during Stage 3 on performance). Consequently, the extent to which the advantage for AC/BD over RT/SU is weakened or reversed in the inconsistent group depends on the relative contribution of α to learning and performance.

Method

Participants, apparatus, and materials. Thirty-eight Cardiff University undergraduates participated in exchange for course credit. Participants were split equally between, and randomly assigned to, the consistent and inconsistent groups and were tested individually with a standard PC. The eight "chemical" names were Bizancrine, Daktyre, Halorite, Kluphane, Nelomine, Ontone, Quezalin, and Yestimox. These were randomly and independently assigned to the letters A–U in the experimental design shown in Table 2 for each participant. The six mutant names were Draguts, Goygle, Jominoid, Necromon, Rargon, and Snarlig, which were again randomly assigned to outcomes O1 to O6 for each participant. Pictures of six different mutant creatures were used, with pictures being randomly assigned to mutant names for each participant. The goo colors used for Stages 1, 2, and 3 were blue, red, and yellow, respectively, created by recoloring the same picture of an amorphous goo substance using image editing software.

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

Stage 1 comprised 14 blocks, with each of the eight trial types shown in Table 2 occurring once per block. Trial order within a block was randomized, with the constraint that there could be no immediate repetitions across blocks. For each trial type, the order of presentation of the chemicals (left/right) was counterbalanced across blocks. The two mutants presented on each Stage 1 trial were always O1 and O2; the left/right order of presentation of mutants was again counterbalanced across blocks for each trial type.

After Stage 1, participants were told that, in the next phase of their research, they would be using a different goo, which created new types of mutants. The form of each Stage 2 trial was the same as that for Stage 1 except that (a) the goo pictured on each trial was red and (b) the two mutants pictured on each trial represented O3 and O4. There were six blocks in Stage 2, with each of the four trial types appearing once per block; counterbalancing and randomization were the same as for Stage 1. After Stage 2, participants were told that the next phase of their research would use a new goo, which again created new types of mutants. The form of each Stage 3 trial was the same as for the previous stages, except that (a) the goo pictured on each trial was yellow, and (b) the two mutants on each trial represented O5 and O6. Stage 3 comprised 10 blocks, with counterbalancing and randomization as in previous stages.

After Stage 3, instructions told participants that, as a test of their understanding, they would be asked to make decisions for new chemical combinations. For each combination, they were to rate how likely different types of mutants were to be created on a scale ranging from 0 (*chemicals very unlikely to create that type of mutant*) to 10 (*chemicals very likely to create that type of mutant*).

Each of the four test compounds (AC, BD, RT, and SU) was presented for rating, in random order. Each test trial gave the names of two chemicals, pictured being poured onto the red goo that had been used during Stage 2. Below that appeared the message "How likely is it that the following mutant will be created?" along with a picture and name of one of the Stage 2 mutants (O3 or O4). Participants entered their rating by clicking one of 11 radio buttons labeled from 0 to 10, with the leftmost

being 0 (*"Chemicals very unlikely to create this mutant"*) and the rightmost being 10 (*"Chemicals very likely to create this mutant"*). Participants rated the ability of a pair of chemicals to create one type of mutant (e.g., O3) and, on the immediately succeeding trial, rated the ability of that same compound to create the other type of Stage 2 mutant (O4, in this case). Whether participants rated mutant O3 or O4 first was consistent across all test compounds, and was determined randomly for each participant.

Data analysis. Following Le Pelley and McLaren's (2003) method, we used these ratings to calculate difference scores for each compound. This was done by taking the rating for each compound, with respect to the outcome (O3 or O4) with which its constituent cues were paired in Stage 2, and subtracting from that the rating for the same compound with respect to the outcome with which its cues were not paired in Stage 2. For example, the difference score for AC is given by the rating for AC with respect to O3 minus the rating for AC with respect to O4, because A and C were paired with outcome O3 during Stage 2. Likewise, the difference score for BD is given by BD's rating for O4 minus its rating for O3, because B and D were paired with O4 during Stage 2. These difference scores index the differential predictiveness of compounds with respect to Stage 2 outcomes—the extent to which a compound predicts the outcome with which it was paired in Stage 2 more than it predicts an outcome with which it was not paired. High difference scores (maximum = 10) indicate strong, selective performance, whereas a difference score of zero indicates no selective performance.

These difference scores are free from influences of generalization that would render any analysis based on raw rating data uninterpretable. Consider Stage 1 of the present experiment. During this stage, participants will learn that cue A is predictive of O1 and that cue T is not predictive of either O1 or O2. Suppose that, immediately after Stage 1, participants were asked how likely cues A and T are to cause outcome O3 when mixed with the red goo. As O3 is a novel outcome at this point, neither A nor T will have a direct association with it. However, O3 has a degree of similarity to O1 (both are types of mutants). Thus, it seems likely that participants would generalize from A's ability to cause O1 and so also view it as a potential cause of O3. Similarly, participants may generalize from the knowledge that T has no effect in creating mutants O1 and O2 to think that it will also have no effect in creating mutant O3. Hence, on the basis of generalization alone and in the absence of further training, we may expect participants to perceive A as a more likely cause of O3 than is T. More generally, using raw ratings could generate a spurious difference in response to previously predictive and previously nonpredictive cues that has nothing to do with differences in the α values of those cues; therefore, we cannot trust any such effect seen in the raw rating data as conclusive evidence to support the operation of α mechanisms.

Using difference scores, on the other hand, allows us to disentangle responding based on direct learning from that based on generalization. Suppose that, in our hypothetical single-stage experiment, we were also to ask people how strongly A predicted outcome O4. Random assignment of names and pictures to outcomes O1 to O6 in Table 2 ensures that, on average, O1 will be as similar to O3 as it is to O4, so that generalization from O1 to O3 is the same as that from O1 to O4. Therefore, if perception of A as a predictor of O3 were purely a consequence of generalization from its association with O1, we would expect an equally high rating for the A–O4 relationship (yielding a difference score of zero). Suppose, in contrast, that participants rated A as a better predictor of O3 than of O4 (yielding a nonzero difference score). This would indicate that the A–O3 rating is not simply based on generalization but that there is also some direct association between A and O3. The stronger this direct association, the greater the magnitude of this difference score.

Results and Discussion

Figure 1 shows mean percent correct of participants' predictions during each block of the three training stages (chance = 50% correct). Performance clearly adapts to the prevailing contingencies in each stage. Collapsing across blocks, independent t tests revealed that the consistent and inconsistent groups did not differ significantly during Stages 1 and 2, $t < 1$, but that the consistent group performed significantly better than the inconsistent group did during Stage 3, $t(36) = 2.27$; significance in this and all subsequent analyses was assessed against a Type I error rate of $\alpha = 0.05$. The difference between the two groups decreased as Stage 3 training continued, however, with no significant difference on the final block, $t < 1$.

This difference in performance of the two groups during the early part of Stage 3 is predicted by both views of the locus of α in learning or performance. According to both models, cues that were predictive during Stage 1 (A–D) will begin Stage 3 with higher α values than those that were earlier nonpredictive (R–U). In the learning-based view of α, this tends to favor learning about cues A–D over learning about cues R–U during Stage 3, producing faster learning in the consistent group (for which cues A–D predict the correct answer and so must be learned about for performance to improve) than in the inconsistent group (for which cues A–D are irrelevant, and instead R–U must be learned about). In the performance-based view of α, response to, for example, the A component of compound AR will initially be amplified relative to response for cue R. This will enhance performance in the consistent group, as A is the cue that must eventually come to control responding. In contrast, performance will be relatively impaired in the inconsistent group, as cues R–U (which will be only weakly responded to by virtue of their low α) are those that must ultimately come to control responding. As such, the observation of a performance difference during Stage 3 training cannot lead to deciding between learning- and performance-based views of α. Nevertheless, given that our previous demonstrations of learned predictiveness effects (e.g., Le Pelley & McLaren, 2003) have been manifest in postacquisition causal judgment ratings, it is encouraging to observe similar effects in the acquisition data itself (see also Bonardi et al., 2005). The fact that the difference between groups is also present in the first training block of Stage 3, $t(36) = 2.15$, presumably reflects within-block learning, given that each cue element appears twice in each training block (e.g., cue A appears in AR and AS).

The results of main interest from this study concern the difference scores derived from ratings provided during the test phase, shown in Figure 2. These scores have been averaged for compounds AC and BD (which both comprise cues that were predictive during Stage 1) and for compounds RT and SU (which both comprise cues that were nonpredictive during Stage 1). For the

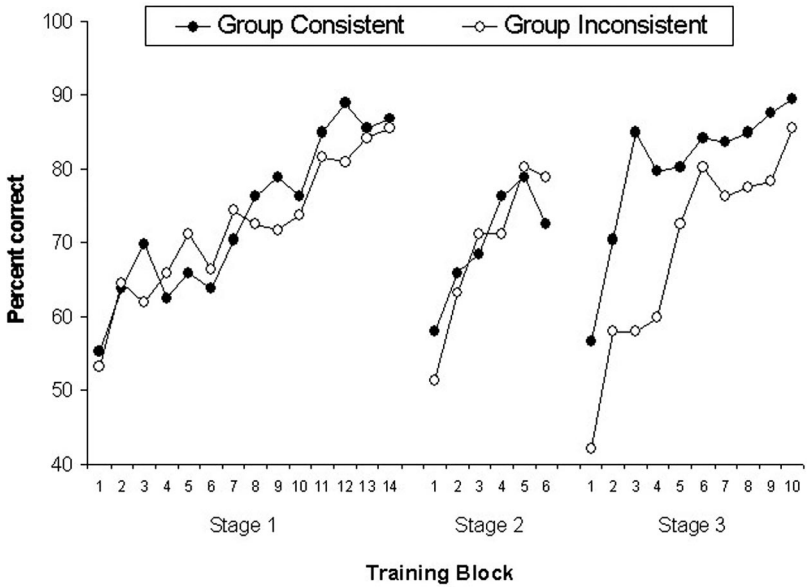

Figure 1. Mean percentage of correct responses for the various trial types over the training blocks of Stages 1–3 of Experiment 1. Data are averaged over all trial types, separately for the consistent group and the inconsistent group.

consistent group, responding to compounds AC/BD is considerably better than that for compounds RT/SU. The response pattern in the inconsistent group is quite different—response to AC/BD is, if anything, poorer than that to RT/SU. These data were analyzed using an analysis of variance (ANOVA), with group (consistent vs. inconsistent) and compounds (AC/BD vs. RT/SU) as factors. This revealed that there was no main effect of Group, $F < 1$; or Compounds, $F(1, 36) = 1.33$, $MSE = 6.69$, $p = .26$. Crucially, the interaction of these two factors was significant, $F(1, 36) = 8.06$, $MSE = 6.69$, indicating a difference in the pattern of performance in the two groups. We used preplanned paired *t* tests to analyze this interaction further. This revealed that, in the consistent group, compounds AC/BD received difference scores that were significantly higher than those of RT/SU, $t(18) = 3.20$; whereas in the inconsistent group, the apparent reversal in performance to these compounds (with mean difference score for RT/SU higher than that for AC/BD) failed to reach significance, $t(18) = 1.08$, $p = .30$.

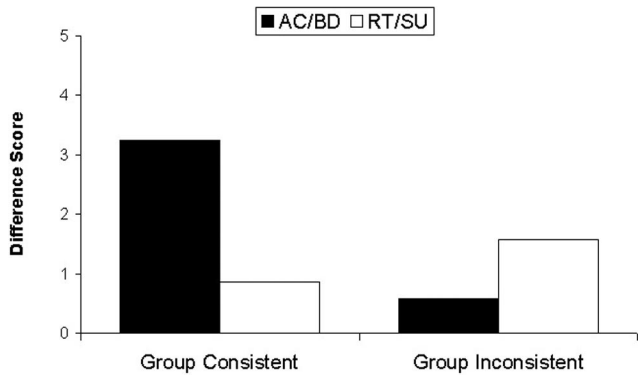

Figure 2. Mean difference scores for the test compounds of Experiment 1. Scores are shown separately for compounds made up of predictive cues from Stage 1 (AC/BD) and those made up of nonpredictive cues from the same stage (RT/SU).

The consistent group showed a clear learned predictiveness effect in line with that observed by Le Pelley and McLaren (2003), with better response to compounds made up of cues that were predictive during Stages 1 and 3. This confirms the influence of α in this study, but it does not tell us whether this influence was on learning or performance. The Stage 3 training received by the inconsistent group—during which cues that had previously been experienced as predictive (A–D) were now found to be nonpredictive, and vice versa—exerted a selective influence on the pattern of results. In this latter group, there was no longer any advantage for compounds AC/BD over RT/SU on test. Thus, it appears that changes in α of cues after the critical learning phase during Stage 2 are sufficient to alter responding to those cues. These results therefore seem to lie beyond a model in which α affects only learning. Instead, they suggest that α is able to influence performance by modulating responding to cues.

Before we can be completely confident in this conclusion, however, we must rule out two alternatives. Figure 1 indicates that participants in the inconsistent group found Stage 3 somewhat harder than did those in the consistent group. This may have caused these participants to give up at the task, with performance on test falling to floor levels so that no advantage for AC/BD over RT/SU could be observed. Two aspects of the data contradict this suggestion. The first is the failure to find a significant main effect of group in the ANOVA, indicating that overall level of performance of the two groups on test was comparable. The second is that performance to RT/SU is, if anything, better in the inconsistent group than in the consistent group. One-sample *t* tests of the RT/SU difference score against a hypothesized mean of zero (indicating no learning) reveal a significant effect for the inconsistent group, $t(18) = 2.13$; but not for the consistent group, $t < 1$. (Note that one-tailed tests are appropriate here, as there are no

circumstances under which we could expect a difference score that is significantly below zero for any cue; hence, no meaning could be attached to such a finding.) Thus, we have evidence for appropriate responding to RT/SU in the inconsistent group but not in the consistent group (although direct comparison of the two groups on responding to RT/SU fails to reach significance, $t < 1$). If the failure to observe a difference in responding to AC/BD and RT/SU in the inconsistent group were the result of a generally lower level of responding than in the consistent group, then we would expect performance on RT/SU to be nonsignificant, given that it is nonsignificant in the consistent group. That this is not observed is a sign that the difference between these two groups is more selective, in that responding to AC/BD is impaired in the inconsistent group, compared with the consistent group, but responding to RT/SU is, if anything, enhanced in the former group.

A second alternative is that the difference between groups reflects associative interference arising from Stage 3 training. Up to this point, we have assumed that the only influence of Stage 3 training on response to the Stage 2 relationships is in terms of changes in α exerting an influence on the extent to which Stage 2 associations are expressed. It is possible, however, that Stage 3 training exerts a more direct effect in terms of retroactive interference: Associations developed during Stage 3 could interfere with retrieval of information learned in Stage 2 at the time of test. The predicted impact of retroactive interference on the results is unclear, however, as two contradictory interpretations are possible.

The first interpretation argues that, during Stage 3, cues that are predictive will develop stronger associations to their respective outcomes than will cues that are nonpredictive. Consequently, Stage 2 cue–outcome associations involving cues that become predictive during Stage 3 will be subject to more retroactive interference than will cues that become nonpredictive. In the inconsistent group, this would produce greater interference for cues R–U than for cues A–D, with the prediction that (compared with the consistent group), we should see a greater impairment in responding to RT/SU than to AC/BD. This is, of course, the opposite of the results observed.

The second interpretation (see Lochmann & Wills, 2003) argues instead that the crucial factor is the number of sources of interference. On this view, Stage 2 cue–outcome associations involving cues that become predictive in Stage 3 will be subject to less retroactive interference than will associations involving cues that become nonpredictive. This is because predictive cues in Stage 3 are paired with only one outcome, giving one source of interference, whereas nonpredictive cues are paired with two different outcomes, giving two sources of interference. Consider the design for the inconsistent group as shown in Table 2. Memory of the T → O3 association from Stage 2 will receive interference from T → O5 (because T is consistently paired with O5 in Stage 3), whereas memory of the A → O3 association will receive interference from A → O5 and A → O6 (because A is paired with both O5 and O6 in Stage 3). Consequently, this interpretation anticipates greater interference for cues A–D than for cues R–U in the inconsistent group, with the attendant prediction that (compared with the consistent group), there will be a greater impairment in responding to AC/BD than to RT/SU. This is the result observed empirically.

These interpretations differ in terms of whether the number of sources of interference (which is greater for nonpredictive cues than for predictive cues) or the strength of each source of interference (greater for predictive cues than for nonpredictive cues) is considered as the dominating factor. To the extent that the number of sources has a stronger influence than the strength of each source, an interference account can explain the results of Experiment 1. In Experiment 2, we aimed to decide between α-based and interference-based accounts of these learned predictiveness effects by using conditional discriminations during Stages 1 and 3 to equate the number of sources of interference for predictive and nonpredictive cues.

Experiment 2

Our basic aim in Experiment 2 was the same as that in Experiment 1—namely to test whether changes in the predictiveness of cues after a critical learning phase could exert a selective influence on response to those cues and thus implicate α in modulating performance. The design of Experiment 2 for the consistent group is shown in Table 3. In this table, letters A, B, R, S, J, K, L, and M denote cues, O1–O6 denote outcomes, and X_{1-5} represents the color of the goo context present on a given trial. Unlike in Experiment 1, this context could vary from trial to trial within a given stage of the experiment. In Stage 1C, for example, AR:X_1–O1 indicates that chemicals A and R, mixed with goo color X_1, created mutant type O1, whereas AR:X_2–O2 indicates that the same chemicals mixed with goo color X_2 instead created mutant type O2. This experiment therefore required participants to learn

Table 3
Design of Experiment 2, Consistent Group

Stage 1A	Stage 1B	Stage 1C	Stage 2	Stage 3A	Stage 3B	Stage 3C	Test
AR:X_1–O1	AR:X_2–O2	AR:X_1–O1	AR:X_3–O3	AR:X_4–O5	AR:X_5–O6	AR:X_4–O5	A:X_3
AS:X_1–O1	AS:X_2–O2	AS:X_1–O1	BS:X_3–O4	AS:X_4–O5	AS:X_5–O6	AS:X_4–O5	B:X_3
BR:X_1–O2	BR:X_2–O1	BR:X_1–O2	*JK:X_3–O3*	BR:X_4–O6	BR:X_5–O5	BR:X_4–O6	R:X_3
BS:X_1–O2	BS:X_2–O1	BS:X_1–O2	*LM:X_3–O4*	BS:X_4–O6	BS:X_5–O5	BS:X_4–O6	S:X_3
		AR:X_2–O2				AR:X_5–O6	*J:X_3*
		AS:X_2–O2				AS:X_5–O6	*K:X_3*
		BR:X_2–O1				BR:X_5–O5	*L:X_3*
		BS:X_2–O1				BS:X_5–O5	*M:X_3*

Note. A, B, R, S, J, K, L, and M represent different chemicals; O1 to O6 refer to the type of mutant that was created (outcome) when these chemicals were mixed with the goo; X_1 to X_5 represent the color of the goo context present on a trial. On test, ratings of the cues were obtained with respect to outcomes O3 and O4. Filler trials are shown in italics.

conditional discriminations, as the effect of the chemicals was conditional on the presented context. The design for the inconsistent group differed only in Stages 3A–3C, where R was consistently paired with outcome O5 in context X_4 and with O6 in context X_5, S was paired with O6 in context X_4 and with O5 in context X_5, and cues A and B were paired equally with both outcomes in both contexts.

Pilot work indicated that participants had difficulty in learning the conditional discriminations of Stage 1C (and 3C) from scratch, with many failing to perform significantly above chance. Clearly, if participants do not learn the discriminations, we cannot expect to see any influence of α (which will only change through learning). To improve performance, we therefore introduced these discriminations in stages before combining them. Thus, in Stage 1A, participants began with only those trial types involving context X_1. Subsequently, Stage 1B involved only those trial types using context X_2. Finally, Stage 1C combined both of the previous learning phases to give a true conditional discrimination. A similar approach was used in Stage 3.

Following the rationale of Experiment 1, Stages 1A–C were designed to establish certain cues as predictive and others as nonpredictive; during Stage 2, all cues were equally predictive of novel outcomes; and predictiveness during Stages 3A–C, again with respect to novel outcomes, could either be consistent with Stages 1A–C (consistent group) or inconsistent with these stages (inconsistent group). The question of interest was whether this Stage 3 training would have a selective influence on participants' ratings of the Stage 2 cue–outcome associations as measured in a final test.

The reasoning behind this design can best be understood by considering the combination of Stages 1C, 2, and 3C. During Stage 1C, cue A was consistently paired with O1 in context X_1 and with O2 in context X_2; cue B was paired with O2 in X_1 and with O1 in X_2. Thus, A and B were, in a sense, consistent predictors of the outcomes with which they were paired, conditional on the context. In other words, participants could feasibly predict the correct answer on a given trial from knowledge of the context and whether A or B was present. Thus, A and B were relevant to the solution of the discrimination; hence, we may expect α_A and α_B to remain high. In contrast, cues R and S were paired equally often with outcomes O1 and O2 in both contexts X_1 and X_2. Consequently, these cues provided no useful information as to the correct answer on each trial: They were irrelevant to the solution of the discrimination; hence, we may expect α_R and α_S to decline.

All Stage 2 trials were conducted in context X_3 and involved novel outcomes O3 and O4. During this phase, all cues were equally predictive of the outcomes with which they were paired: A and R were paired with O3, and B and S were paired with O4. Filler trials involving novel cues (JK:X_3–O3 and LM:X_3–O4) were introduced in Stage 2 purely to make the difficulty of Stage 2 more on a par with that of previous and subsequent learning phases and will not be discussed further here.

Stage 3C training used novel outcomes O5 and O6. For participants in the consistent group, relevance during Stage 3C was consistent with that during Stage 1C: Cues A and B were relevant to the solution of the Stage 3 discrimination, whereas R and S were irrelevant. For participants in the inconsistent group, relevance in Stage 3C was instead inconsistent with that in Stage 1C: for this group, cues A and B were irrelevant to the Stage 3 discrimination, whereas cues R and S were relevant. After Stage 3C, participants rated how likely each of the Stage 2 outcomes (O3 and O4) was to follow individual cues A, B, R, and S, all presented in context X_3, the context previously used in Stage 2.

We first consider the predictions made by α-based theories with respect to this final test, ignoring for the moment any potential impact of interference. A similar argument to that applied to Experiment 1 anticipates that, regardless of whether α influences learning or performance, participants in the consistent group should perceive A and B as better predictors of the Stage 2 outcomes than R and S (because the higher α of cues A and B promotes more rapid learning of associations involving these cues, promotes stronger expression of those associations, or both). Also as for Experiment 1, an approach in which α influences only learning predicts a similar advantage for A and B over R and S in the consistent and inconsistent groups, as this account has no way for subsequent changes in α to influence an association once that association has been formed. In contrast, an account in which α influences performance predicts a greater advantage for A and B over R and S in the consistent group than in the inconsistent group. In the inconsistent group, the reversed relevance of the cues in Stage 3C should cause α_A and α_B to fall relative to α_R and α_S. To the extent that α influences the expression of associations formed in Stage 2, these changes in α during Stage 3 should lead to a decline in ratings for A and B relative to R and S. Hence, we would expect the change in the α values of the component cues during Stage 3 to reduce (and possibly reverse) the advantage for A and B over R and S.

We saw earlier that an account based on retroactive interference is potentially able to explain the findings of Experiment 1 if it is assumed that the number of sources of interference has a stronger interference than the strength of each source. In Experiment 2, however, the number of sources of interference was the same for all cues—in both groups, all of cues A, B, R, and S were paired with two different outcomes in Stage 3C. Thus, if the number of sources of interference were the sole factor generating the results of Experiment 1, then we would expect no difference between the consistent and inconsistent groups in Experiment 2. In fact, if anything, we would expect the influence of retroactive interference to be greater for those cues that were relevant in Stage 3 than those that were irrelevant. Looking at the inconsistent group, we may expect R to be implicated in strong associations to outcomes O5 and O6 (as it is predictive of both of them, conditional on the context), whereas A would be only weakly associated with these outcomes (as it is not predictive of either). To the extent that the strength of each source of interference also exerts an influence, then, we would expect greater interference for cues R and S than for cues A and B in the inconsistent group, with the attendant prediction that (compared with the consistent group), there will be a greater impairment in responding to R and S than to A and B. This is the opposite of the prediction made by accounts implicating α in performance as described earlier.

Method

Participants, apparatus, and materials. Forty-two Cardiff University students participated for course credit. Participants were split equally between the consistent and inconsistent groups. Other details were as for Experiment 1, with the exception that this

study used five different goo colors: blue, red, yellow, green, and brown. These were randomly assigned to contexts X_1 to X_5 for each participant.

Procedure. Instructions to participants at the outset of the experiment were similar to those of Experiment 1, as was the form of each training trial. The only substantive difference was that the goo context was made more salient by labeling its picture, with the label being written in the color of the goo; for example, the label "Brown goo" would appear in brown below the picture of the brown goo.

Stages 1A and 1B comprised eight blocks, Stage 1C had six blocks, Stage 2 had eight blocks, Stages 3A and 3B had six blocks, and Stage 3C had four blocks, with trial types in each block as shown in Table 2. Instructions preceding each stage informed participants of the goo colors that they would be dealing with in that stage. All other details of the training phases were as for Experiment 1.

Test trials were as for Experiment 1, although in Experiment 2 these trials involved individual chemicals. Each of the eight cues that had appeared in Stage 2 was presented in random order for rating; the context on each test trial was provided by goo color X_3. Difference scores were calculated from these ratings as described for Experiment 1.

Data exclusion. To be sure that the number of sources of retroactive interference had indeed been equated for all cues, we needed to be sure that participants had learned both conditional discriminations in Stages 3A and 3B and, more particularly, that they could maintain both simultaneously during Stage 3C. The alternative is that participants simply learn one of these discriminations and ignore the other, so that the number of sources of retroactive interference is not matched (as for Experiment 1).

If participants had learned one of the Stage 3 discriminations perfectly and performed randomly on the other, they would be expected to achieve a mean accuracy of 75% correct during training in Stage 3C. Consequently, the data for participants scoring 75% correct or less in Stage 3C (3 participants in both the consistent and inconsistent groups) were excluded from all further analyses.

Results and Discussion

Figure 3 shows mean percent correct of participants' predictions during each training block. Accurate performance is observed by the end of each stage. Collapsing across blocks, independent *t* tests revealed that the consistent group performed significantly better than the inconsistent group during Stage 3A, $t(34) = 3.09$; and that the advantage for the consistent group in Stage 3B approached significance, $t(34) = 1.92$, $p = .063$; but that performance in the two groups did not differ significantly in any other stage, ($t = 1.65$, $p = .11$ for Stage 3C; $t < 1$ for all other stages). The difference in training performance during Stages 3A and, to a lesser extent, 3B mirrors the difference between groups observed during the early blocks of Stage 3 in Experiment 1, and presumably stems from the same source: Cue processing in both groups begin this stage focused on cues A–D at the expense of cues R–U. For the consistent group, this will aid learning and/or responding during Stage 3, whereas for the inconsistent group (given the reversal in the predictiveness of the cues), it will be a hindrance. Participants' performance in Stage 3C is clearly far above 75% correct for both groups; $t = 28.7$ for the consistent group and 18.9 for the inconsistent group. Hence, we can be confident that participants are able to maintain the conditional discriminations of Stage 3A and 3B simultaneously.

Difference scores derived from the ratings given to cues during the test phase are shown in Figure 4. These scores have been averaged for cues A and B (which were both predictive during Stage 1) and for cues R and S (which were both nonpredictive

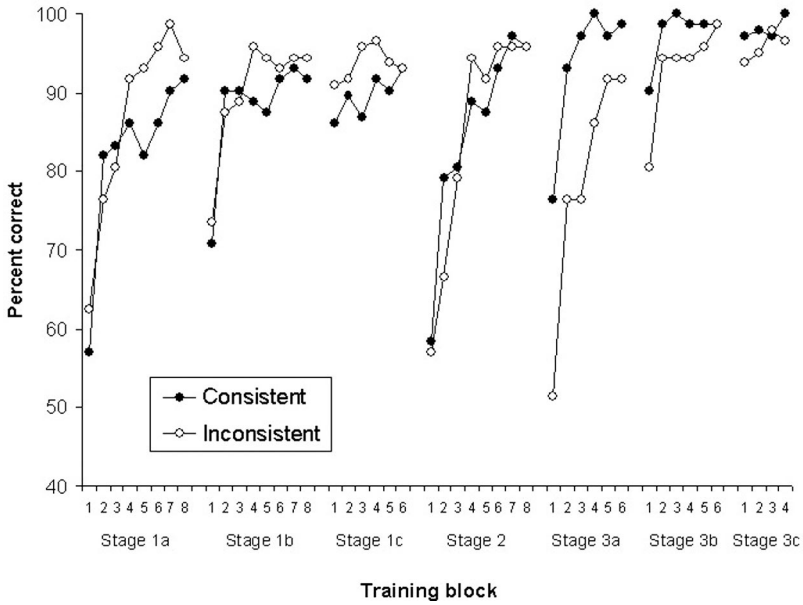

Figure 3. Mean percentage of correct responses for the various trial types over the training blocks of Stages 1a–3c of Experiment 2. Data are averaged over all trial types, separately for the consistent group and the inconsistent group.

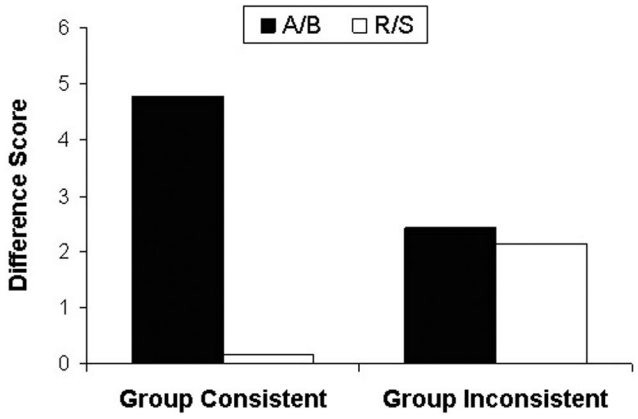

Figure 4. Mean difference scores for the test cues of Experiment 2. Scores are shown for predictive cues from Stage 1 (A/B) and nonpredictive cues from Stage 1 (R/S), separately for the consistent group and the inconsistent group.

during Stage 1). For the consistent group, responding to predictive cues from Stage 1 (A/B) is considerably better than that for nonpredictive cues (R/S). In contrast, in the inconsistent group, difference scores for A/B and R/S are very similar. We analyzed these data using an ANOVA with group (consistent vs. inconsistent) and cue (A/B vs. R/S) as factors. This revealed that there was no main effect of group, $F < 1$; but there was a main effect of cue, $F(1, 34) = 7.24$, $MSE = 14.9$; with A/B receiving higher difference scores than R/S overall. Crucially, the interaction between these two factors was significant, $F(1, 34) = 5.69$, $MSE = 14.9$, indicating a difference in the pattern of performance in the two groups. Preplanned paired *t* tests were used to analyze the results of each group separately. This revealed that, in the consistent group, cues A and B received difference scores that were significantly higher than those of cues R and S, $t(17) = 3.77$; whereas in the inconsistent group, there was no significant difference in the mean difference scores for A–B and R–S, $t < 1$.

As for Experiment 1, it seems unlikely that the Group × Cue interaction stems from participants in the inconsistent group responding at chance levels on test, with difference scores thus falling to floor levels. First, the lack of a main effect of group indicates that the overall performance level in the two groups was comparable. Second, response to R–S is, if anything, better in the inconsistent group than in the consistent group. Although a direct comparison of the R–S score in both groups failed to reach significance, $t(34) = 1.39$, $p = .17$; one-sample tests revealed that the R–S score was significantly greater than chance (zero) in the inconsistent group, $t(17) = 1.96$ (one-tailed; see Experiment 1's Results and Discussion), but not in the consistent group, $t < 1$.

In summary, as for Experiment 1, the results of Experiment 2 indicate that the different Stage 3 training received by the two groups exerted a selective influence on participants' ratings of the strength of earlier-learned cue–outcome associations. Specifically, Stage 3 training in the inconsistent group led to a decline in the perceived causal strength of cues experienced as being irrelevant, relative to cues experienced as being relevant. Most important is that this finding is, if anything, in the opposite direction to that expected by an account based on differences in retroactive interference suffered by the different cues. Our results therefore seem to demand that α is able to influence performance as suggested in Equation 4.

A possible concern, however, is that Experiments 1 and 2 both rely on a between-groups comparison to establish this influence of α on performance. In both experiments, difference scores for the inconsistent group are, in general, numerically (but not significantly) lower than those for the consistent group, raising the possibility that a floor effect in the inconsistent group may contribute to the Group × Cue interaction in each case. Although we have demonstrated (by comparing difference scores for R–S in the two groups) that such an influence is unlikely, Experiment 3 was designed to rule out this possibility using a wholly within-subjects design.

Experiments 1 and 2 both looked at the ability of retrospective changes in α to modulate the magnitude of an existing learned predictiveness bias. In Experiment 1, for example, learning about differential predictiveness during Stage 1 would create a bias toward previously predictive cues in Stage 2 (as demonstrated by Le Pelley & McLaren, 2003), and it was the ability of subsequent changes in predictiveness during Stage 3 to modulate this bias that was of central interest. In contrast, Experiment 3 examined whether retrospective changes in predictiveness could *create* a bias; that is, whether such changes could result in unequal performance to cues that were otherwise equivalent.

Experiment 3

Table 4 shows the design of Experiment 3, which is essentially generated by removing Stages 1A–1C from the consistent group of Experiment 2. During Stage 1 of Experiment 3, all cues were equally predictive of their respective outcomes (O1 and O2). Subsequently in Stage 2, cues A and B were trained as predictive cues with respect to outcomes O3 and O4, whereas R and S were nonpredictive. Stage 2 training once again used conditional discriminations depending on the context (X_2 or X_3) to equate the number of sources of retroactive interference for all cues. After Stage 2, participants rated how likely each of the Stage 1 outcomes (O1 and O2) was to follow individual cues A, B, R, and S, all presented in context X_1, the context previously used in Stage 1.

Table 4
Design of Experiment 3

Stage 1	Stage 2A	Stage 2B	Stage 2C	Test
AR:X_1–O1	AR:X_2–O3	AR:X_3–O4	AR:X_2–O3	A:X_1
BS:X_1–O2	AS:X_2–O3	AS:X_3–O4	AS:X_2–O3	B:X_1
JK:X_1–O1	BR:X_2–O4	BR:X_3–O3	BR:X_2–O4	R:X_1
LM:X_1–O2	BS:X_2–O4	BS:X_3–O3	BS:X_2–O4	S:X_1
			AR:X_3–O4	*J:X_1*
			AS:X_3–O4	*K:X_1*
			BR:X_3–O3	*L:X_1*
			BS:X_3–O3	*M:X_1*

Note. A, B, R, S, J, K, L, and M represent different chemicals; O1 to O4 refer to the type of mutant that was created (outcome) when these chemicals were mixed with the goo; X_1 to X_3 represent the color of the goo context present on a trial. On test, ratings of the cues were obtained with respect to outcomes O1 and O2. Filler trials are shown in italics.

Given that all cues were equally predictive during Stage 1, we can be sure that *learning* about each, with respect to the relevant Stage 1 outcome, will be equal. If α influences performance, then any subsequent change in α during Stage 2 will create a bias in participants' perceptions of the different cues. Specifically, we would expect participants to rate cues A and B (experienced as predictive during Stage 2) as more predictive of their respective Stage 1 outcomes than cues R and S (experienced as nonpredictive during Stage 2).

Method

Participants, apparatus and materials. Nineteen Cardiff University students participated in exchange for £5. Other details were as for Experiment 1. The three goo colors used were blue, red, and yellow; these were randomly assigned to contexts X_1–X_3 for each participant.

Procedure. Stage 1 comprised 10 blocks of training, Stages 2A and 2B had eight blocks, and Stage 2C had six blocks, with trial types in each block as shown in Table 4. All other details of the training phases were as for Experiment 2. Test trials were as for Experiment 2, with the exception that the two mutants for which ratings were provided represented O1 and O2, and the context on each trial was given by goo color X_1. Difference scores were calculated from these ratings as described for Experiment 1.

Data exclusion. As for Experiment 2, data were excluded for participants who failed to average above 75% correct over the training trials of Stage 2C, resulting in removal of 3 participants.

Results and Discussion

Training data revealed that participants' responses rapidly adapted to the prevailing cue–outcome contingencies in each stage of Experiment 3. Collapsing across blocks, and across all trial types in each stage (as all trial types within each stage were equivalent), mean percent correct was 82.3% in Stage 1, 81.6% in Stage 2A, 92.6% in Stage 2B, and 96.1% in Stage 2C. Crucially, accuracy during Stage 2c was far above 75% correct, $t(15) = 19.1$; hence, we can be confident that participants were able to maintain the conditional discriminations of Stages 2A and 2B simultaneously.

The mean difference score for cues A and B was 4.38, while that for R and S was 1.88. This difference was significant, $t(15) = 2.31$, indicating that experience of the differential predictiveness of cues during Stage 2 did indeed influence participants' perceptions of the causal strength of these cues with respect to the Stage 1 outcomes. Once again this finding is consistent with the suggestion that α influences performance, and as for Experiment 2, these results cannot be accounted for in terms of retroactive interference (which would predict, if anything, a higher score for R/S than for A/B).

The demonstration, in Experiments 1–3, that α exerts an influence on the responding to cues, rather than (or perhaps in addition to) the learning about those cues, raises the possibility that many, if not all, previously described learned predictiveness effects in both humans and animals also reflect the operation of α at the response level, rather than at the learning level as has often been assumed.

Consider the learned predictiveness effect observed by Le Pelley and McLaren (2003), wherein previously predictive cues support stronger responding on test than do previously nonpredictive cues (this is analogous to the significant advantage for predictive cues over nonpredictive cues observed in the consistent group of Experiments 1 and 2). The results of Experiments 1 and 2 show that α has an influence on performance in these studies. The question, therefore, becomes one of whether this effect is best explained wholly in terms of the influence of α on performance; that is, by a model in which α has no effect on learning, so that predictive and nonpredictive cues form equally strong associations during Stage 2 (but in which associations for predictive cues are expressed more strongly than those for nonpredictive cues). The alternative is that α influences both learning and performance so that predictive cues form stronger associations in Stage 2, and the expression of these cues is also magnified relative to nonpredictive cues. The fact that the level of response to AC/BD and RT/SU in Experiment 1, and to A/B and R/S in Experiment 2, did not reverse significantly in the inconsistent group may be seen as indicating that there is a persistent advantage for AC/BD (or A/B) in terms of higher associative strengths developed as a result of the higher α values of these cues during Stage 2. The influence of α on responding may then be insufficient to completely reverse this advantage.

That said, it is also possible to reconcile this finding with an approach incorporating α at the response level only, with all cues developing equally strong associations during Stage 2, by making additional assumptions about the way in which α changes. Such an account could explain the lack of a significant reversal in the inconsistent group by suggesting that changes in α during Stage 3 are slow, and hence, α values do not change sufficiently to produce a reversal in responding. As it stands then, on the basis of Experiments 1–3, a model incorporating α at the response level only appears able to incorporate all of our present data and those of all other studies of learned predictiveness effects that have been taken as support of the general approach offered by the Mackintosh (1975) model in both human and animal learning.

Experiment 4

Experiment 4 attempted to determine whether there is also a role for learned predictiveness in modulating learning. The design of Experiment 4 is shown in Table 5. Stages 1 and 2A are identical to Stages 1 and 2 of Experiment 1. Thus, during Stage 1, we again

Table 5
Design of Experiment 4

Stage 1	Stage 2A	Stage 2B	Test
AR–O1	AT–O3	Aa–O3	Aa
AS–O1	BU–O4	Bb–O4	Bb
BR–O2	CR–O3	Cc–O3	Cc
BS–O2	DS–O4	Dd–O4	Dd
CT–O2		Rr–O3	Rr
CU–O2		Ss–O4	Ss
DT–O1		Tt–O3	Tt
DU–O1		Uu–O4	Uu

Note. Uppercase letters represent trained cues used as pretrained blocking cues in Stage 2B; lowercase letters represent novel blocked cues in Stage 2B. On test, ratings of the cues were obtained with respect to outcomes O3 and O4.

expected the predictive cues (A–D) to maintain high α while the α of the nonpredictive cues (R–U) decreases. On each Stage 2A trial, a compound of a previously predictive cue and a previously nonpredictive cue is paired with a novel outcome, O3 or O4. If, as proposed by the Mackintosh model, learned predictiveness modulates the learning of cue–outcome associations (as opposed to merely influencing responding to cues), then during Stage 2A cues A–D would be expected to form stronger associations to their respective outcomes than would cues R–U. We tested this idea by including an additional stage of training, Stage 2B, during which we looked at the ability of cues A–D and R–U to block learning about novel cues.

Blocking is a well-established phenomenon of both animal and human learning (e.g., Dickinson et al., 1984; Kamin, 1969) and refers to the finding that responding to an element of a reinforced stimulus compound is reduced if another element of that compound has previously been established as a predictor of reinforcement. Thus, if pairings of K (the blocking cue) with an outcome are followed by pairings of a KL compound with that same outcome, less responding to L (the target cue) is subsequently observed than to a control cue N, trained in an MN compound in which both elements are novel.

Clearly, the occurrence of blocking depends on there being a preexisting association between the blocking cue and the outcome at the outset of compound training, this being the difference between blocking and control conditions, as outlined earlier. As an extension of this idea, Dickinson et al. (1984) demonstrated in a study of human learning that the amount of blocking obtained to a target cue was directly related to the predictiveness of the blocking cue: The stronger the preexisting association between the blocking cue and the outcome at the outset of compound training, the less was learned about the target cue. This finding forms the basis of the crucial empirical manipulation in Experiment 3. On each trial of Stage 2B, one of the stimuli trained in Stage 2A was compounded with a novel cue (labeled in lowercase in Table 5), and this compound was paired with the same outcome as that with which the familiar cue was paired in Stage 2A. For example, the trained cues A and T were combined with novel cues a and t, to form two new compounds, Aa and Tt, and each of these compounds was paired with outcome O3 (as each of the trained elements, A and T, had been paired with O3 during Stage 2A).

If differences in α of cues A–D and R–U resulting from Stage 1 training lead to cues A–D developing stronger associations to their respective outcomes in Stage 2A than do cues R–U, we would expect cues A–D to be better able to block learning about the novel cues with which they are paired in Stage 2B. That is, novel cues a–d would be subject to stronger blocking than would novel cues r–u: This would be evidenced by lower difference scores for cues a–d than for cues r–u.

Experiments 1–3 allow for the alternative possibility that α solely modulates responding, as in Equation 4, and plays no part in the algorithm governing learning itself. If this is indeed the case, then the associations developed by cues A–D during Stage 2A will be of identical strength to those developed by cues R–U; hence, novel cues a–d and r–u will be subject to identical amounts of blocking. On this basis, we should see no difference in response to these novel cues on test. In more general terms, if α has no effect on learning, then there is no way for the difference in α values of cues A–D and R–U to differentially influence the acquisition of associative strength by target cues a–d and r–u. We return to this issue in the General Discussion.

Method

Participants, apparatus, and materials. Twenty-four Cardiff University students participated in exchange for £6. The 16 chemical names were Alzaze, Bizancrine, Daktyre, Frestix, Gratix, Jintsone, Kluphane, Lobinz, Ontone, Pukintz, Quezalin, Renphane, Sistax, Trizopane, Ventox and Xentine. The mutant names were Draguts, Goygle, Jominoid and Necromon. Other details were as for Experiment 1.

Procedure. Stages 1 and 2A were as for Experiment 1, although Stage 2A was lengthened to 10 blocks. Stage 2B followed without interruption from Stage 2A, so it appeared to participants to be one continuous set of trials. Stage 2B comprised four blocks. During the subsequent test phase (which followed the procedural details of Experiment 2), participants were required to rate each of the 16 chemicals individually. Difference scores were calculated from these ratings as for Experiment 1.

Results and Discussion

Figure 5 shows mean percent correct of participants' predictions during each training block. Accurate performance is observed by the end of each stage. Collapsing across blocks, a paired *t* test revealed that accuracy during Stage 3 did not differ between compounds containing previously predictive blocking cues (Aa, Bb, Cc, and Dd) and those containing previously nonpredictive blocking cues (Rr, Ss, Tt, and Uu), $t < 1$.

The results of the test phase are shown in Figure 6. This figure shows mean difference scores collapsed across equivalent cues from the three training phases. Hence, results were averaged across cues A, B, C, and D as previously predictive blocking cues (labeled A–D in Figure 6); R, S, T, and U as previously nonpredictive blocking cues (labeled R–U); a, b, c, and d as target cues paired with previously predictive cues (labeled a–d); and r, s, t, and u as target cues paired with previously nonpredictive cues (labeled r–u).

Preplanned paired *t* tests were used to analyze this data. For the blocking cues, the difference between predictive (A–D) and nonpredictive (R–U) cues failed to reach significance, $t < 1$. Although, on the basis of Experiments 1 and 2 (and Le Pelley & McLaren, 2003), we might have expected cues A–D to support stronger responding than cues R–U at the end of Stage 2A, the additional training received by these cues during Stage 2B would likely weaken this difference. By virtue of their previous training, both A–D and R–U would be better predictors of the outcomes occurring on Stage 2B trials than would the novel cues with which they were paired. Hence, we may expect the α values of both classes of blocking cue to rise over the course of Stage 2B. Consequently, response to both A–D and R–U on test will be relatively strong, and this, coupled with possible ceiling effects generated by our finite rating scales, may well mask any difference between them.

The crucial comparison in this study relates to responding to the target cues. Target cues paired with predictive cues (i.e., cues a–d) received mean difference scores that were significantly lower than those paired with nonpredictive cues (i.e., cues r–u), $t(23) = 3.77$.

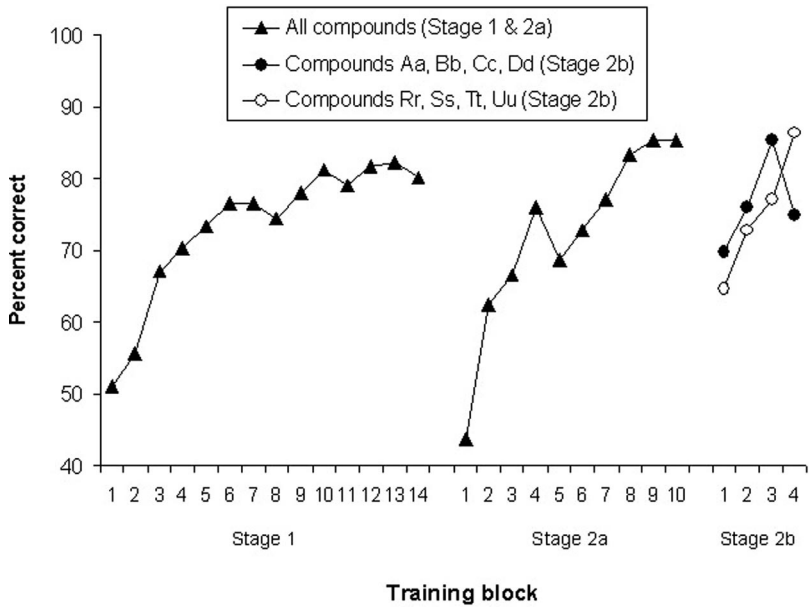

Figure 5. Mean percentage of correct responses for the various trial types over the training blocks of Stages 1–2b of Experiment 2. Data are averaged over all trial types in Stages 1 and 2a. For Stage 2b, data are shown separately for compounds containing previously predictive blocking cues (Aa, Bb, Cc, and Dd) and those containing previously nonpredictive blocking cues (Rr, Ss, Tt, and Uu).

This finding contradicts a model in which α exerts its effects only at the level of performance by modulating the strength of the response made to a cue, as it demonstrates that the α of one cue (the blocking cue) can influence the acquisition of associative strength by another (the target cue); hence, α must exert some effect at the level of learning.

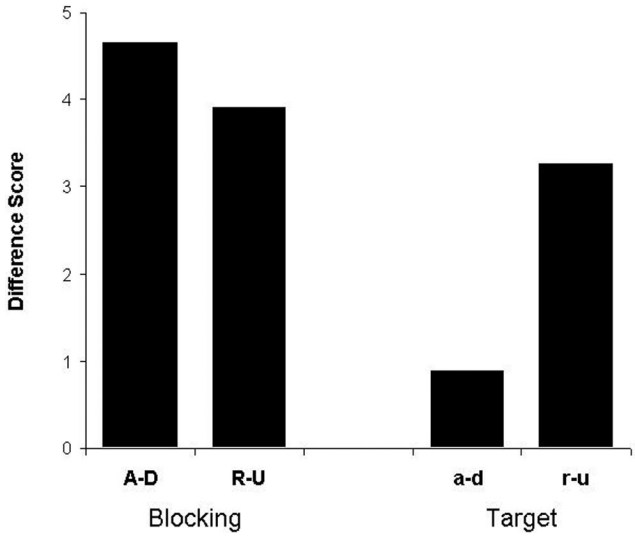

Figure 6. Mean difference scores for the individual cues during the test stage of Experiment 4. Data are averaged over the four equivalent trial types in the experiment: previously predictive blocking cues (Cues A–D); previously nonpredictive blocking cues (Cues R–U); target cues paired with previously predictive cues (Cues a–d); and target cues paired with previously nonpredictive cues (Cues r–u).

General Discussion

Four experiments investigated the locus of learned predictiveness effects in human learning. Experiment 1 demonstrated that changes in the learned predictiveness of a stimulus after acquisition of a critical association can influence the extent to which that association is expressed on test. This experiment indicates that α exerts an effect at the level of performance, modulating the strength of the response to a cue. Thus, Experiment 1 indicates that the behavioral response to cue P, R_P, is given by:

$$R_P = k \alpha_P V_P \quad (5)$$

Experiment 2 replicated the basic finding of Experiment 1 under circumstances in which any impact of retroactive interference from Stage 3 on memory of the Stage 2 cue–outcome associations would, if anything, tend to produce the opposite pattern of results, thus supporting an analysis of the effect of Stage 3 training in terms of changes in α. We note also that the observation of a significant advantage for A/B over R/S in Experiment 2's consistent group runs contrary to accounts of more standard learned predictiveness effects in terms of proactive interference. Lochmann and Wills (2003) have argued that the learned predictiveness effect observed by Le Pelley and McLaren (2003) could reflect differences in the proactive interference suffered by cues: Memory of Stage 2 information involving previously predictive cues is subject to only one source of proactive interference (as these cues were paired with only a single outcome in Stage 1), whereas memory of Stage 2 information involving previously nonpredic-

tive cues is subject to two sources of proactive interference (as these cues were paired with two different outcomes in Stage 1). The demonstration of an analogous effect in the consistent group of Experiment 2, in which proactive interference would, if anything, be greater for predictive cues than for nonpredictive cues, lies beyond such an account. Our present findings therefore strengthen the case for a role of changes in cue processing, as suggested by α-based theories, in these learning preparations.

Experiments 1 and 2 both examined the ability of retrospective changes in α to modulate the magnitude of an existing learned predictiveness bias. In contrast, Experiment 3 demonstrated that retrospective changes in α can also create a bias, producing unequal performance to cues that were otherwise equivalent. Moreover, Experiment 3 used a wholly within-subject design and, hence, ruled out explanations of the findings of Experiments 1 and 2 in terms of between-groups motivational differences.

Although the results of Experiments 2 and 3 are consistent with the spirit of α changes as suggested by the Mackintosh (1975) model, they also highlight a limitation of this model. As an explicitly elemental model, which treats cues as separable elements in the learning process, this model is unable to account for participants' learning of the conditional discriminations used in Stages 1C and 3C of Experiment 2, and Stage 2C of Experiment 3. This is because none of the elements (cues or contexts) are individually predictive of the outcome on each trial of this training—each cue, and each context, is paired equally often with both types of outcome. Instead, only certain configurations of cues and context are predictive; for example, in Stage 1C, the configuration of cue A and context X_1 is predictive of outcome O1. The fact that such training leads to a change in processing of the separable cue elements involved may be taken as indicating that α acts at the level of cue configurations and that the α value developed by a particular configuration will generalize to influence learning about similar configurations (an idea suggested by George & Pearce, 1999). For example, a high associability developed by the configuration of A and context X_1 might generalize to influence processing of the configuration of A and context X_2. Other interpretations are also possible, however. One is that the context acts as an "occasion setter" (see Holland, 1992), a stimulus that does not elicit a response itself but instead sets the occasion for the response to occur, by telling the participant which of multiple possible responses should be made to a cue. On this account, it is the separable cue element that "owns" the α value, but learning about the cue's predictiveness is determined in relation to the occasion setter. Another possibility is that it is not correlation with a specific reward that is crucial in determining α, but rather the relevance of a stimulus to a discrimination (George & Pearce, 1999). Thus, although cue A considered alone is not correlated with either outcome during Stage 1C, it is relevant to the solution of the discrimination in that stage. The present experiments do not allow us to decide between these alternative views of stimulus representation with respect to α.

Experiment 4 investigated the idea that response modulation may be the only way in which learned predictiveness exerts its effects. We found that cues differing in learned predictiveness and trained in compound subsequently differed in their ability to block learning about novel cues. This finding demonstrates that α cannot solely be involved in response modulation. If this were the case, then there would be no way for the difference in predictiveness of these cues to influence the acquisition of responding (which must be based, at some level, on learning) by the novel target cues.

We believe that our data leave open two classes of interpretation for the locus of learned predictiveness effects. The first proposes that α contributes directly and independently to the rate of learning about a cue and to the level of responding to that cue. Thus, the results of Experiment 1 could be captured by a model in which learning is determined by the standard Mackintosh (1975) model

$$\Delta V_P = \theta\, \alpha_P\, (\lambda - V_P) \tag{6}$$

in combination with the modified response rule of Equation 5. It is easy to see how Equation 5 allows the resulting model to account for the finding that changes in α after learning of a critical association can modulate the expression of that association, as demonstrated in Experiment 1; how this model accounts for the results of Experiment 4 deserves more explanation. The differences in α developed by cues A–D and cues R–U during Stage 1 ensure that, according to Equation 6, the former develop stronger associations to their respective outcomes in Stage 2A. On the initial trials of Stage 2B, the difference in predictiveness between cues A–D and the novel target cues paired with them (a–d) are relatively large, causing the α of these novel target cues to fall rapidly and preventing them from developing (Equation 6) and expressing (Equation 5) strong associations to the outcome. The difference in predictiveness between cues R–U and the novel target cues paired with them (r–u) will be considerably smaller, as cues R–U formed only weak associations during Stage 2A. Therefore, the α values of the novel cues r–u will fall more slowly, allowing them to develop and express significant associations with the outcome.

Although this view of α having direct and independent effects on learning and performance is appealing, we must acknowledge an alternative view that is also consistent with our data. On this latter account, explained later, α does not have such a direct influence in modulating learning.

In Mackintosh's (1975) model, as shown in Equation 6, the error term $(\lambda - V_P)$ represents the extent to which the outcome occurring on a given trial is predicted by stimulus P alone. If two cues, P and Q, are presented on the same learning episode, the error term for P will be $(\lambda - V_P)$ and that for Q will be $(\lambda - V_Q)$: This model has a separate error term for each presented stimulus. As such, the learning undergone by each cue will be independent of the current associative strength of the other. In contrast, Rescorla and Wagner (1972) suggested that the error governing learning for a given cue should be based on the combined associative strength of all cues present on that trial. Hence, their model took the general form:

$$\Delta V_P = \theta\, (\lambda - \Sigma V), \tag{7}$$

where ΣV is the summed associative strength of all currently presented cues. Although such a model in combination with the modified response rule of Equation 5 would be able to account for the results of Experiment 1, it is unable to explain the results of Experiment 4. At the outset of Stage 2A, cues A–D and R–U will differ in their α levels. However, given that α does not appear in Equation 7, there is no way for this difference to be reflected in learning about the cues; hence, both will form equal associations to their respective outcomes during Stage 2A. Consider the Aa and Tt trials of Stage 2B. If cues A and T begin Stage 2B with equal

associative strengths, and there is no way for differences in their α values to influence learning of new associations to the novel cues a and t, then the associative strengths of these target cues will remain equal. Furthermore, the calculation determining the α values of these target cues (Equation 2) makes reference only to associative strengths. Given that $V_A = V_T$ and $V_a = V_t$, this means that the α values for a and t will also remain equal to one another. Essentially, because α is a parameter based on the learned predictiveness of cues, if there is no difference in the learning about the blocking cues, they cannot give rise to differences in the α values of novel target cues with which they are paired. Consequently, the associative strengths and α values of a and t will remain identical throughout Stage 2B; hence, this model is constrained to incorrectly predict no difference between these cues on test.

However, a slight modification allows this model to fare better. The error term in Equation 7 represents the discrepancy between the current state of the outcome and the summed associative strength of all presented cues. If this is changed so that the error term instead represents the discrepancy between the current state of the outcome and the summed response strength of all presented cues, with response strength modulated by α as in Equation 5, then this model becomes able to account for the results of Experiment 4. Thus:

$$\Delta V_P = \theta\,(\lambda - \Sigma R) \tag{8}$$

where ΣR is the summed response strength of all presented cues, with the response strength for each cue P given by $R_P = k\,\alpha_P\,V_P$.

Given that a common error term applies to all cues presented on a given trial, these cues will always undergo the same change in associative strength on that trial. For example, applying Equation 8 to each AT trial during Stage 2A gives the following:

$$\Delta V_A = \theta\,(\lambda - [k\,\alpha_A\,V_A + k\,\alpha_T\,V_T])$$

$$\Delta V_T = \theta\,(\lambda - [k\,\alpha_A\,V_A + k\,\alpha_T\,V_T])$$

$$\rightarrow \Delta V_A = \Delta V_T.$$

Consequently, cues A–D and R–U enter Stage 2B with equal associative strengths. However, using response strengths rather than association strengths in the error term means that any differences in the α values of cues A–D and R–U can now influence learning about their accompanying target cues. In effect, this account suggests that the ability of one cue to block learning about another depends on the α of the blocking cue. Consider the initial Aa and Tt trials of Stage 2B. Applying Equation 8 gives the following:

$$\Delta V_a = \theta\,(\lambda - [k\,\alpha_A\,V_A + k\,\alpha_a\,V_a])$$

$$\Delta V_t = \theta\,(\lambda - [k\,\alpha_T\,V_T + k\,\alpha_t\,V_t]).$$

Because a and t are novel on these initial trials, $V_a = V_t = 0$. Hence:

$$\Delta V_a = \theta\,(\lambda - k\,\alpha_A\,V_A)$$

$$\Delta V_t = \theta\,(\lambda - k\,\alpha_T\,V_T).$$

In addition, from earlier we have $V_A = V_T$ on these initial trials. Therefore, as long as the higher α value of cues A–D (compared with that of cues R–U) is assumed to persist from Stage 1 so that $\alpha_A > \alpha_T$ during Stage 2B, the model predicts that $\Delta V_a < \Delta V_t$. Hence, it correctly predicts weaker response to cues a–d than to cues r–u on test. More generally, allowing learning to depend on response strength, and response strength to depend on α, in turn allows α to influence learning about cues. Specifically, the higher response strength of A–D cues after Stage 2A ensures that there is little room for cues a–d to develop additional response strength on Stage 2B trials, compared with cues r–u.

Whether, in such a model, learned predictiveness can be truly said to influence learning in addition to performance is a moot point. An α term does certainly feature in the learning equation of such a model (Equation 8)—the results of Experiment 4 show that this must be the case for any successful model. However, as noted earlier, the use of a common error term governing learning ensures that two cues presented on the same trial will always undergo the same change in associative strength, even if those cues have very different α values; for example, A and T developed equally strong associations during Stage 2A, even though $\alpha_A > \alpha_T$. This use of a combined error term means that the learning undergone by a specific, individual stimulus need not be directly related to the α of that particular stimulus. This can be contrasted with a model based on Equation 6, in which the α of each individual stimulus has an independent and direct influence on learning about that cue. The closely entangled relationship between learning about a stimulus and making a response to that stimulus makes it difficult, however, to see a way in which these two potential accounts could be teased apart. Given the widespread assumption among attentional theories of learning that α has a selective and direct influence on learning (e.g., Kruschke, 1996, 2001; Mackintosh, 1975; Sutherland & Mackintosh, 1971), it is perhaps remarkable that providing experimental evidence to support this idea unequivocally is not at all straightforward.

Regardless of which view is correct, however, the present data provide the first conclusive evidence that α must influence the expression of the response to a cue and also must play some role in the learning equation, however indirect. Moreover, although they conflict with its detail, our results provide support for the spirit of the Mackintosh (1975) model, whereby reliable predictors maintain higher α values than do poorer predictors. The relationship between learning and performance is clearly a complex one, but the present results bring us closer to a more complete understanding of how learned predictiveness may influence these processes in humans and, potentially, animals as well.

References

Bonardi, C., Graham, S., Hall, G., & Mitchell, C. (2005). Acquired distinctiveness and equivalence in human discrimination learning: Evidence for an attentional process. *Psychonomic Bulletin & Review, 12,* 88–92.

Dickinson, A., Shanks, D. R., & Evenden, J. L. (1984). Judgement of act-outcome contingency: The role of selective attribution. *Quarterly Journal of Experimental Psychology, 36*A, 29–50.

George, D. N., & Pearce, J. M. (1999). Acquired distinctiveness is controlled by stimulus relevance not correlation with reward. *Journal of Experimental Psychology: Animal Behavior Processes, 25,* 363–373.

Griffiths, O., & Le Pelley, M. E. (in press). Attentional changes in blocking are not a consequence of lateral inhibition. *Learning & Behavior*.

Holland, P. C. (1992). Occasion setting in Pavlovian conditioning. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 28, pp. 69–125). San Diego, CA: Academic Press.

Kamin, L. J. (1969). Predictability, surprise, attention and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.

Kruschke, J. K. (1996). Dimensional relevance shifts in category learning. *Connection Science, 8,* 225–247.

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

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

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

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

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

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

Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In F. Schmalhofer, R. M. Young, & G. Katz (Eds.), *Proceedings of EuroCogSci 03: The European Conference of the Cognitive Science Society* (pp. 217–222). Mahwah, NJ: Erlbaum.

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

Melchers, K. G., Lachnit, H., & Shanks, D. R. (2004). Within-compound associations in retrospective revaluation and in direct learning: A challenge for comparator theory. *Quarterly Journal of Experimental Psychology, 57*B, 25–53.

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.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Sutherland, N. S., & Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.

Received April 18, 2008
Revision received September 17, 2008
Accepted September 23, 2008 ■