Journal of Experimental Psychology:
Animal Learning and Cognition

2329-8456/19/$12.00

2019, Vol. 45, No. 2, 163–173
http://dx.doi.org/10.1037/xan0000202

Explaining Learned Predictiveness: Roles of Attention and Integration of Associative Structures

Gabriel Rodríguez
University of the Basque Country (UPV/EHU)

Geoffrey Hall
University of York and University of New South Wales

In 3 experiments, participants were trained in an associative learning paradigm in which they learned the relation between consumption of certain foodstuffs and the type of allergic reaction shown by a fictional patient. Experiment 1 demonstrated the learned predictiveness effect, showing that cues that had served as good predictors of outcomes in an initial phase of training were especially effective in a test given after a second phase of training in which learning about the same cues, but with different outcomes, had been required. Experiment 2 showed that this effect could be obtained when the two phases of training occurred in reverse order, so that the critical cues were established as good or bad predictors only after the associations tested in the final test had been acquired. This learned predictiveness effect cannot be explained by an enhancement of the associability of the predictive cues that facilitated learning about them in phase two. This encouraged us to consider 2 alternatives to associability for explaining learned predictiveness: (a) that training a cue as a good predictor increases its effective salience, thus enhancing its power to evoke responding on test and (b) that learned predictiveness is the result of a nonattentional process in which subjects integrate information acquired in the separate phases of training. Support for the latter came from Experiment 3, which showed that a modified test procedure, designed to reduce the tendency to integrate across phases, eliminated the learned predictiveness effect.

Keywords: predictiveness, associability, salience, attention, integration

Learned predictiveness refers to an experimental procedure, devised by Lochmann and Wills (2003) and developed by Le Pelley and McLaren (2003), that has been used for the study of transfer effects in human learning (see Le Pelley, Mitchell, Beesley, George, & Wills, 2016, for an extensive review). In the version of Le Pelley and McLaren, people learn, in the first stage of this procedure, that certain cues (e.g., different fruits: banana, apple, orange, . . .) differ in the accuracy with which they predict the occurrence of possible outcomes (e.g., one or other of two allergic reactions). Thus, they have the opportunity to learn that some cues, the accurate predictors, signal the occurrence of one of the outcomes and the absence of the other; other cues, the inaccurate predictors, sometimes signal the occurrence of one outcome, and sometimes of the other. In a subsequent stage of training, the cues are used to signal a new pair of similar outcomes (e.g., another two possible allergies). During this stage all the cues are reliable predictors of the occurrence of one of the new outcomes and the absence of the other. The learned predictiveness effect is manifested in a final test, when subjects asked to predict the occurrence of these two new outcomes show better performance given the cues pretrained as accurate predictors than in the presence of cues pretrained as inaccurate predictors.

In interpreting their results, Le Pelley and McLaren (2003) advanced an explanation in terms of attention, proposing that the effect was caused by differences between the cues in learning rate during the second stage of training. This interpretation was derived from the attentional account of associative learning proposed by Mackintosh (1975), according to which the *associability* of a cue (i.e., the form of attention that determines the rate of learning about a cue) is directly related to its previous predictive reliability. According to this account, cues trained initially as accurate predictors will have higher associability than the cues pretrained as inaccurate predictors when it comes to the second phase of training. Thus, although all cues are paired equally often with the corresponding outcomes, the associations between previously accurate predictors and the outcomes will be established faster than those involving the previously inaccurate predictors.

It is a problem for this account that there is doubt about whether stimulus associability is determined in the way that is required. According to an alternative analysis, proposed by Pearce and Hall (1980), the principle that governs changes in associability serves to optimize the limited learning resources of the organism, in such a way that a cue will be attended to and learned about only to the extent that there is uncertainty about its consequences. Once an

---

Gabriel Rodríguez, Faculty of Psychology, University of the Basque Country (UPV/EHU); Geoffrey Hall, Department of Psychology, University of York and University of New South Wales.

This research was supported in part by grants from the Spanish Ministerio de Economía y Competitividad (Grant PSI2015-64309-P, MINECO/FEDER) and Gobierno Vasco (Grant IT-694-13).

The trial-level raw data for the present experiments can be requested from Gabriel Rodríguez.

Correspondence concerning this article should be addressed to Gabriel Rodríguez, Facultad de Psicología, Universidad del País Vasco (UPV/EHU), Avenida de Tolosa 70, 20018 San Sebastián, Spain. E-mail: gabriel.rodriguez@ehu.es

organism has learned that a stimulus is an accurate predictor, there will be no such uncertainty, and no need to allocate learning resources to it (i.e., it will have a low associability). There is evidence from studies of simple conditioning, most of them conducted with nonhuman animals (e.g., Hall & Pearce, 1979; Holland, 1997; Hall & Pearce, 1982), but some with humans as participants (e.g., Griffiths, Johnson, & Mitchell, 2011), supporting this view of associability. How can this be reconciled theoretically with the occurrence of the learned predictiveness effect? We can reject the suggestion that different principles govern changes in associability in human and nonhuman animals given that a version of the learned predictiveness effect (known as the acquired distinctiveness of cues) was first demonstrated in experiments with nonhuman animals (Lawrence, 1949).

Accordingly, we will explore the proposal that the learned predictiveness effect is generated by some process other than change in stimulus associability (Hall & Rodríguez, 2010). Associability change may well occur in the learned predictiveness procedure, but if the process responsible for learned predictiveness is powerful it could obscure any effects based on associability change, even if these are of the sort postulated by Pearce and Hall (1980). It is possible to distinguish two possible approaches to this issue in the existing literature. One involves a closer analysis of the concept of attention as applied to these procedures, with associability being distinguished from other aspects. The other is the suggestion that the effect depends not on attention, as conceived by these theories, but on a learning process in which information acquired during both stages of training is integrated to produce the final response in the test.

Mackintosh (1975) restricted his attentional parameter to a role in determining the speed of new learning, but he allowed the possibility that it might also control other aspects of behavior, including the vigor of performance. Others who have developed his attentional account (e.g., Esber & Haselgrove, 2011; Kruschke, 2001; Le Pelley et al., 2016) have taken this step, allowing the possibility that the learned predictiveness effect is (in part) a consequence of the enhanced ability of a stimulus that is well associated with an outcome to elicit behavior. And even if we suppose that associability declines during training (as suggested by the Pearce-Hall theory Pearce & Hall, 1980) it must be accepted that a stimulus that is firmly associated with an outcome is still well attended to, given that it has the ability to control performance. Hall and Rodríguez (2017, 2019) have proposed a two-factor account of attention (see also George & Pearce, 2012; Le Pelley, 2004; Pearce & Mackintosh, 2010), distinguishing attention-for-learning (i.e., associability) from attention-for-performance, which they identified as an enhancement of the effective salience of a cue that has high associative strength. The learned predictiveness effect was attributed to the latter form of attention, that is, to an attentional mechanism that enhances the ability of the organism to respond to (but not necessarily to learn about) cues that predicts accurately relevant outcomes.

This interpretation is not incompatible with the results of experiments using eye-tracking techniques, that have been taken to support the associability account. These studies (e.g., Beesley, Nguyen, Pearson, & Le Pelley, 2015; see also Mitchell, Griffiths, Seetoo, & Lovibond, 2012) show that people trained with the learned predictiveness procedure, spend more time looking at the good rather than at the bad predictors in both stages of training. These results have been taken to support an interpretation in terms of associability, with attention being controlled by events that need to be learned about. But they equally may be interpreted as showing that eye-gaze, in this procedure, is determined by attention-for-performance—that people are attending more to the reliable predictors to be able respond to them.

We next consider the second alternative to the associability notion that has been offered as an explanation for the learned predictiveness effect. Bonardi, Graham, Hall, and Mitchell (2005) pointed out that the two-stage procedure used in these experiments allows the possibility that the effect might be a consequence of some nonattentional process that operates at the time of test. That is, subjects might be able to look back on information acquired about the cues in both stages and make a response on this basis. We will refer to this as the *integration* account, as it requires the subject to integrate the information acquired in both stages of training. This interpretation has been developed by Mitchell et al. (2012) who argued that the effect is not determined by an automatic, bottom-up, attentional mechanism of the sort described previously but rather that it is the consequence of an "inference-based controlled attentional process." (In terms of the distinction made by Shiffrin & Schneider, 1977, it involves "controlled" rather than automatic processing.). Thus, when faced with new outcomes in the second phase of a learned predictiveness study, subjects infer that the cues that were predictive in the first phase are likely to be predictive in the second. Mitchell et al. noted that this process requires subjects to integrate across the phases of training, and suggest that the tendency to do so may be based on the similarity of the outcomes used in the two phases (see Le Pelley, Oakeshott, Wills, & McLaren, 2005). Finally, they point out that this interpretation, too, is compatible with the results of studies of eye-gaze—that people will orient selectively toward predictive cues does not, they argue, imply the operation of a selective attentional mechanism.

Although Mitchell et al. (2012) presented integration as a being result of top-down processing, other interpretations are possible. For example, integration could be the consequence of an automatic tendency to exploit an existing associative structure in learning about the relationships among events occurring in a novel situation, when the latter is to some extent similar to the situation in which the old structure was acquired. From this perspective, what Mitchell et al. regarded as an act of inference (in which participants arrive at a logical conclusion with regard to the second stage from a series of premises based on their learning in the first stage) is seen as an involuntary tendency to organize the new information acquired in the second stage (the relationship among the cues and the new outcomes) according to the associative structures acquired in the first stage (among the cues that are good and bad predictors of outcomes). Although there are relevant differences between these two versions of the integration account, they concur in predicting that the tendency to integrate will depend on the similarity between the two phases; and both agree that it should be possible to obtain the learned predictiveness effect in procedures in which attentional processes, whether these involve changes in associability or in effective salience, cannot be responsible.

The experiments to be reported here address these issues as follows. Experiment 1 provides a demonstration of the basic learned predictiveness effect. Experiment 2 provides a direct test of the proposal that the effect depends on the enhancement of the associability of predictive cues (the account derived from Mackintosh's, 1975, theory). It shows, to anticipate, that a powerful effect can be obtained even when the procedure is modified so as to preclude a role for changes in stimulus associability. Experiment 3, although it does not rule out the possible contribution of other forms of attention, provides a demonstration of the role of a nonattentional process in producing the effect, one that is consistent with the notion of integration.

Experiment 1

This study was intended to demonstrate that we could obtain the basic learned predictiveness effect with our subject population and procedures. It was modeled, as closely as we could manage, on the work of Le Pelley and McLaren (2003). In this procedure, participants are required to learn to predict what outcomes will occur after presentation of certain cues. Specifically, the participant was asked to play the role of an allergist, and to try to determine what allergic reactions might be suffered by fictitious patients after eating various foods. The experimental design is summarized in Table 1. The letters represent cues (foods) and the numbers represent the outcomes (negative reactions) that occur after the presentation of the cues. For example, AW->1 indicates a trial in which the cues A and W are simultaneously presented, and followed by the occurrence of the Outcome 1. The cues were presented in pairs throughout the experiment.

In the first stage of training one of two outcomes (1 or 2) could occur on each trial. One cue of each pair was consistently paired with a given outcome (A and D with 1, and B and C with 2), the other cue of the pair (V and W, X and Y) was equally paired with 1 and 2. We refer to this stage as *differential training* as participants are exposed to conditions in which they can learn that there are differences between the cues in their predictive accuracy. The second stage of training involved a similar task, but with the cues presented in different pairs and the introduction of two new outcomes, 3 and 4. Critically, in this stage, both cues in each pair were consistently followed by a given outcome; that is, there was now no difference in predictive accuracy between cues A–D and cues W–Z (hence, this sort of training is referred to as *nondifferential* in Table 1).

In a final test the same cues were presented in four different novel pairs, in such a way that two pairs were formed by cues paired with Outcome 3 (AC and VX), and the other two pairs by cues paired with Outcome 4 (BD and WY). Two of these pairs (AC and BD) were formed by cues that had been accurate predictors

Table 1
Experimental Designs

Experiment 1

Differential training	Test: O1 or O2?	Nondifferential training	Test: O3 or O4?
Av→1	Av	Ax→3	AC
Bv→2	Bv	By→4	BD
Aw→1	Aw	Cv→3	vx
Bw→2	Bw	Dw→4	wy
Cx→2	Cx		
Dx→1	Dx		
Cy→2	Cy		
Dy→1	Dy		

Experiment 2

Nondifferential training	Test: O3 or O4?	Differential training	Retest: O3 or O4?	Test: O1 or O2?
Ax→3	AC	Av→1	AC	Av
By→4	BD	Bv→2	BD	Bv
Cv→3	vx	Aw→1	vx	Aw
Dw→4	wy	Bw→2	wy	Bw
		Cx→2		Cx
		Dx→1		Dx
		Cy→2		Cy
		Dy→1		Dy

Experiment 3

Differential training	Test: O1 or O2?	Nondifferential training	Test: O1? O2? O3? O4?
Av→1	Av	Ax→3	AC
Bv→2	Bv	By→4	BD
Aw→1	Aw	Cv→3	vx
Bw→2	Bw	Dw→4	wy
Cx→2	Cx		
Dx→1	Dx		
Cy→2	Cy		
Dy→1	Dy		

Note. Each letter represents a cue (a foodstuff) and numbers represent outcomes (Os: allergic reactions).

during the initial differential training; the other two pairs (VX and WY) were formed by cues that had been inaccurate predictors in that stage. The learned predictiveness effect is obtained when, in this final test, participants show evidence of expecting the appropriate outcome more confidently in the presence of the pairs formed by the cues that were accurate predictors during the initial differential training (e.g., AC strongly evokes the response O3) than after the presentation of previously inaccurate predictors (e.g., VX is less effective at evoking O3).

Method

Subjects. Twenty-four students (17 female; $M_{\text{age}} = 19.71$ years, range: 18–33) from the University of the Basque Country agreed to participate after being informed that they would take part in an experiment involving cognitive tasks. All of them had normal or corrected-to-normal vision. The Research Ethics Committee of the University of the Basque Country (CEISH) approved the experimental protocol.

Apparatus and stimuli. The participants were tested individually, sitting at approximately 50 cm from the 17-in screen of a standard PC. The eight cues (A to D and W to Z) were images of fruits (apple, orange, melon, grapes, banana, strawberry, pear, and cherry) presented on a white background. Assignment of these images to specific cue roles in the design shown in Table 1 was randomized for each participant. Outcomes were pictures of negative or aversive reactions, presented on the same white background, and accompanied by a text box with the name of the reaction (in Spanish) underneath. The outcomes in the stage of differential training were "stomach ache" and "rash," and the outcomes in the nondifferential training stage were "headache" and "conjunctivitis." Assignment of these negative reactions as Outcomes 1 or 2, and as Outcomes 3 or 4, was also randomized.

Procedure. The participants were first informed that they were to play the role of an allergist who had to learn to predict what reactions would be suffered a patient, Mr. X, after eating pairs of fruits.

Differential training. This stage comprised 14 blocks of trials, with each of the eight trial types shown in Table 1 occurring once per block. Each trial began with the simultaneous presentation of the pictures of two cues (arranged horizontally at the center of the screen, 3 cm apart) and a question (arranged horizontally at the center of the screen, 3 cm underneath the pictures of the cues): *Stomachache or rash?* (for participants for whom stomachache was Outcome 1), or *Rash or stomachache?* (for participants for whom rash was Outcome 1). After 6 s, a picture illustrating the appropriate reaction, along with the name of that reaction, was presented for 3 s. In each block, the trial order was randomized apart from the restriction that the same pair of cues could not occur on consecutive trials (i.e., it was not permitted for the same pair of fruits to be presented on the last trial of a block and the first trial of the next block). For each trial type the position (left/right) of the fruits on the screen was counterbalanced across blocks, and the order in which these positions varied was randomized across the experiment.

Test of differential training. Following the initial stage of training, participants were told that they would be tested on what they had learned. The test consisted of eight trials, one for each type of trial of Stage 1 (see Table 1). Each trial consisted of the 20-s presentation of a pair of fruits. The left/right position of each fruit and the order of presentation of each type of trial were randomized across participants. Participants were informed that Mr. X would again eat pairs of fruits and that they should rate, for each pair, the likelihood of occurrence of Outcome 1 and similarly, the likelihood of occurrence of Outcome 2. Responses were made on paper sheets showing two scales, (one for each outcome) ranging from 0 (*cues very unlikely to predict the outcome*) to 10 (*cues very likely to predict the outcome*). On these, and subsequent test trials, subjects were allowed 20 s to make their response. No feedback was provided in this test.

Nondifferential training. Immediately after differential training test, participants were told that they were now to deal with a different patient, Mr. Y, who suffered different reactions to the same cues. There were four blocks in this stage in each of which the four nondifferential trial types shown in Table 1 appeared once per block in random order. The structure of trials in this stage of training was identical to that of trials in the previous stage.

Test of nondifferential training. Following the second stage of training, participants received instructions for a new test. This consisted of four trials with new combinations of cues, AC and BD being composed of cues that had been predictive in the first stage of training, VX and WY composed of cues that had been poor predictors in Stage 1 (see Table 1). The response sheet required ratings for Outcomes 3 and 4 for each test compound. In other respects, the structure of these trials was the same as that of the trials of the test following the first stage of training.

Data treatment and analysis. For the first test a difference score was computed for each subject for each compound by taking the rating for the correct outcome with which that compound had been paired during training and subtracting from it the rating for the incorrect outcome, that with which that compound had not been paired. For example, the score for compound AV was calculated as the rating for Outcome 1 minus the rating for Outcome 2. Higher scores (with a maximum score = 10) indicate better learning. Participants that over the total of eight test trials exhibited three or more negative scores (i.e., rated the occurrence of the incorrect outcome more likely than the occurrence of the correct outcome) were excluded from the analysis, on the grounds that they had failed to learn the contingencies in effect during differential training.

For the final test, a difference score was computed for each of the four test compounds by taking the rating for the correct outcome with which the elements of that compound were paired during the nondifferential training and subtracting from it the rating for the incorrect outcome, that with which the elements of that compound were not paired during training. For example, the score for compound AC was calculated as the rating for Outcome 3 (as A and C were paired with this outcome during the nondifferential training phase) minus the rating for Outcome 4 (A and C were not paired with this outcome during training). Difference scores were averaged for compounds AC and BD (comprising cues that were accurate predictors during the first stage of training), and for VX and WY (comprising cues that were inaccurate predictors during initial training).

Data were analyzed by the analysis of variance (ANOVA) and a criterion of statistical significance of p less than .05 was adopted. Effect sizes for ANOVAs are reported as partial eta squared and those for pairwise comparisons are reported using Cohen's d. The

95% confidence intervals (CIs) around the effect sizes are also reported in parentheses following the effect size.

Results and Discussion

Three participants were excluded because of poor performance exhibited in the test that followed differential training. The remaining participants had a mean difference score of 9.6 on this test (just short of the possible maximum of 10) indicating that the cue-outcome relationships had been well learned.

Figure 1 shows performance in the test that followed nondifferential training. It is evident that compounds consisting of cues that had been accurate predictors during the differential-training stage evoked more accurate ratings than compounds consisting of inaccurate cues. Statistical analysis confirmed this conclusion, $t(20) = 3.41$, $p = .005$, $d = 1.05$ (0.36–1.73).

Although the design is identical to that used in the experiment by Le Pelley and McLaren (2003), the procedure used in this experiment differs in a number of details. The events used as cues and outcomes were different, and, in our experiment, the subjects were simply exposed to them during the two main stages of training, with no response being required (in the original version of the experiment, the subjects were asked to guess which outcome would follow a cue on each training trial). Nonetheless, the learned predictiveness effect was clearly demonstrated, and we use the same basic procedure in subsequent experiments intended to analyze the nature of the effect.

Experiment 2

The aim of this experiment was to provide a demonstration of the learned predictiveness effect using a modified procedure that would preclude explanation in terms of the account offered by Le Pelley and McLaren (2003) based on the notion of associability derived from Mackintosh (1975). Even if we accept the (disputed) proposal that training a cue as a good predictor enhances its associability, there are other reasons to question the adequacy of this notion as the (sole) source of the effect. Experiments by Le Pelley, Suret, and Beesley (2009) have shown that training given to the cues after the usual two stages of training in the learned predictiveness design could modify the size of the learned predictiveness effect. Specifically, training the cues in a third stage as either good or bad predictors of a new set of outcomes modulated the response given when the subjects were asked again about the relationship between these cues and the outcomes used in the second stage. The implication is that the ability of these cues to control performance on the final test was modified by training in which they were made good or bad predictors of (other) events. This does not prove, however, that the standard learned predictiveness effect is a product of performance factors; it is quite possible that differences in associability, and hence in the rate of learning during the second stage of training, were also operating in these experiments and were responsible for the basic effect that was modified by later training. The aim of the present experiment, therefore, was to modify the standard experimental design in a way that would preclude any possibility that differences in associability could cause the observed learned predictiveness effect.

The design of this experiment is presented in the middle panel of Table 1. It is essentially the same as that of Experiment 1 except that the order of the training phases was reversed, with differential training following nondifferential training. With this arrangement, any difference in associability between good and bad predictors produced by differential training with Outcomes 1 and 2 would be irrelevant given that learning about the relationship between the cues and Outcomes 3 and 4 had already occurred. The critical test was that given after differential training ("retest" in Table 1). Would the learned predictiveness effect (superior performance with AC and BD than with VX and WY) be obtained in these conditions? A test with these cues was also given after the first stage of nondifferential training to confirm that the contingencies had been learned about. A final test, asking about Outcomes 1 and 2, was included to assess the effectiveness of the differential training phase.

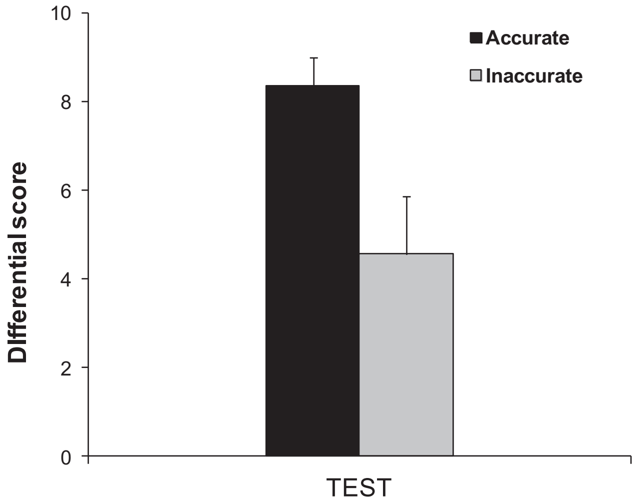

Figure 1. Averaged difference scores for compounds AC and BD (comprising cues that were Accurate predictors during the initial nondifferential training), and for VX and WY (comprising cues that were Inaccurate predictors during that initial training) during the final test (O3 or O4?) of Experiment 1. Difference scores for each of the four test compounds were first computed by taking the rating for the correct outcome with which the elements of that compound were paired during the nondifferential training and subtracting from it the rating for the incorrect outcome, that with which the elements of that compound were not paired during that training. Vertical bars represent standard errors of the mean.

Method

The experiment was carried out in two identical replications. Thirty-five students (21 female; mean age = 23.48 years, range: 18–43) in the first replication, and 60 students (39 female; mean age = 22.27 years, range: 18–36) in the second replication, all of them from the University of the Basque Country, participated after being informed that they would take part in an experiment involving cognitive tasks. All of the participants had normal or corrected-to-normal vision. The apparatus, the cues, and the outcomes used in this experiment were identical to those used in Experiment 1.

The training procedure was identical to that used in Experiment 1, except for reversal of the order in which the two types of training were given. In this experiment, participants first received the four blocks of nondifferential training, and then the 14 blocks of differential training. (Although they were experienced first in this experiment, the outcomes used in nondifferential training are still referred to as 3 and 4; those in differential training as 1 and 2). As in Experiment 1, participants were first told that they were treating a person called Mr. X, and then, at the start of differential training that they were now dealing with a new patient, Mr. Y.

The first test, introduced after nondifferential training, asked about Outcomes 3 and 4. This test was repeated after the phase of differential training. It was followed immediately by a test asking about Outcomes 1 and 2, the procedure being that described for Experiment 1.

Results and Discussion

As in Experiment 1, performance on the test asking about Outcomes 1 and 2 was used to exclude subjects who had failed to learn during the phase of differential training. Five participants in the first replication, and nine in the second replication, were excluded according to the criterion described for Experiment 1. The remaining participants exhibited a good level of performance on this test, with an average difference score (rating of correct outcome minus rating of incorrect outcome) of 8.7 in the first replication, and 8.4 in the second replication.

Figure 2 shows performance (difference scores computed as described in Experiment 1) on the tests asking about Outcomes 3 and 4. Separate means are presented for cues that were trained as accurate or as inaccurate predictors in the differential training phase that followed. Unsurprisingly, on the first test, given prior to any differential training, there is no difference between the two classes of cue. Both showed high difference scores indicating that the previous stage of training had been successful. Ratings on the retest decreased for both sets of cues, suggesting that the time interval and/or the differential training interposed between the test and the retest had a deleterious effect on retrieval of the information learned during the initial nondifferential training. More critical for our present purposes is that, in the retest, the compounds consisting of cues established as accurate predictors during differential training evoked more accurate ratings than compounds consisting of cues established as inaccurate predictors.

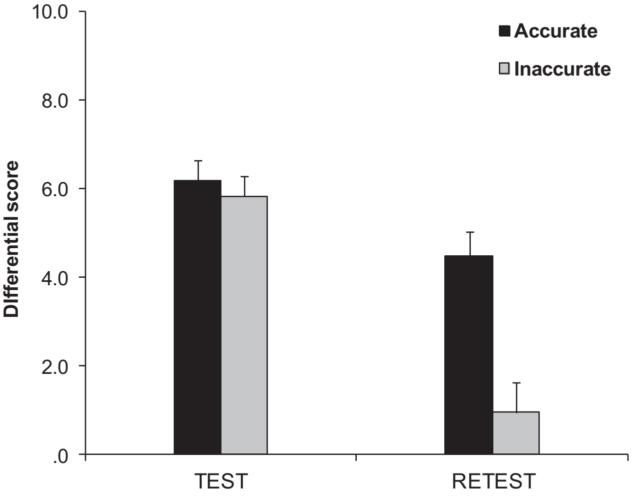

Figure 2. Averaged difference scores for compounds AC and BD (comprising cues that were Accurate predictors during the nondifferential training), and for VX and WY (comprising cues that were Inaccurate predictors during that training) during the test and retest (O3 or O4?) of Experiment 2. Difference scores for each of the four test compounds were first computed by taking the rating for the correct outcome with which the elements of that compound were paired during the nondifferential training and subtracting from it the rating for the incorrect outcome, that with which the elements of that compound were not paired during that training. Vertical bars represent standard errors of the mean.

The scores summarized in Figure 2 were subjected to an ANOVA with stimulus type (accurate vs. inaccurate), phase (test vs. retest), and replication (first vs. second) as the variables. The main effect of replication was significant, $F(1, 79) = 13.36$, $p < .01$, $\eta_p^2 = 0.14$ (0.03–0.29), indicating that overall performance was superior in the second replication. None of the interactions involving the factor replication was significant, $Fs(1, 79) < 2.37$, $p > .12$. The main effect of stimulus was significant, $F(1, 79) = 19.98$, $p < .001$, $\eta_p^2 = 0.20$ (0.07–0.35). Both the effect of test, $F(1, 79) = 33.68$, $p < .001$, $\eta_p^2 = 0.30$ (0.14–0.44), and the interaction, Stimulus × Test, $F(1, 79) = 14.44$, $p < .001$, $\eta_p^2 = 0.16$ (0.04–0.30), were also significant. Further analysis in order to reveal the source of this interaction showed that the two types of stimulus differed in the retest, $t(80) = 5.57$, $p < .001$, $d = 0.88$ (0.54–1.20), but not in the initial test, $t(80) = 1.04$, $p = .303$. There were also significant effects of test versus retest, both for the compounds consisting of accurate cues, $t(80) = 2.39$, $p = .019$, $d = 0.44$ (0.12–0.75), and for the compounds consisting of inaccurate cues, $t(80) = 6.78$, $p < .001$, $d = 1.07$ (0.71–1.41). The decisions on null hypothesis rejection in these latter four comparisons remained the same when controlling the false discovery rate ($q < .05$; Benjamini & Hochberg, 1995).

This pattern of results shows that the learned predictiveness effect survives a manipulation in which the order of the differential and nondifferential training phases is reversed. It was found that training cues as accurate or inaccurate predictors could modify the effectiveness of these cues in controlling responses about information already acquired in a previous stage of training. We conclude that a learned predictiveness effect can be obtained when effects depending on changes in associability (of the sort envisaged by the theory of Mackintosh, 1975) cannot be responsible. This example of the effect is, however, consistent with both of the alternative accounts discussed in the introduction. The notion of attention-for-performance assumes that this aspect of attention will be enhanced during differential training for cues that are consistently followed by a given consequence. Such training will make them particularly effective in eliciting responding appropriate to events with which they have become associated. That the differential training that establishes differences among cues in their ability to control performance comes after the training that established the associations being tested will be irrelevant. The integration account is similarly neutral about the order in which the stages of the procedure are presented. The central notion is that the learned predictiveness effect is the result of a bias on the part of the participants to integrate information acquired about the cues in both phases of training, to generate a single response on the test. Thus they learn, (a) that some cues are related to Outcome 3 and other cues are related to Outcome 4 (information acquired during nondifferential training); and (b) that among these cues there are some more strongly associated with outcomes (something learned during the differential training with respect to Outcomes 1 and 2). Integrating these two items of information (no matter in what order they have been acquired) will lead to better performance on test to cues that have been accurate predictors during differential training.

Experiment 3

How can the integration account be distinguished from an account in terms of attention-for-performance? A relevant observation comes from an experiment by Le Pelley et al. (2005) who demonstrated that the learned predictiveness effect was attenuated when the nature of Outcomes 1 and 2 differed from that of Outcomes 3 and 4. The explanation favored by Le Pelley et al. (2005)—that this result reflects the outcome-specificity of the associability parameter—cannot apply if we accept the implication of the results of Experiment 2, that the basic effect does not depend on associability. The result is equally problematic for an account in terms attention-for-performance, as there are no grounds for thinking that enhanced attention of this type is specific to the nature of the outcome that follows the cue. This result accords readily, however, with the integration notion, given the reasonable assumption that people will be more likely to integrate, when there is a similarity between Outcomes 1 and 2 and Outcomes 3 and 4 that encourages them to treat one set of outcomes as equivalent to, or functional substitutes for, the other set.

The experiment to be reported next was designed to provide a test of one interpretation of the integration hypothesis by exploring the effects of another variable that might be expected to influence the likelihood of integration occurring. In the standard experimental design, subjects learn in the second stage about the relationship between the cues and the new outcomes (3 and 4) and they are asked about these outcomes on test. Outcomes 1 and 2, used in the first stage are no longer presented, and the relationship between the cues and these outcomes is no longer an issue. The fact that Outcomes 1 and 2 are absent after the end of the first stage could be critical in leading the participants to perceive the new outcomes as substitutes for the old, so that the cue that reliably predicted Outcome 1 is readily taken to be a reliable predictor of Outcome 3 (and so on). The bias to integrate will be maintained in the final test in which participants are allowed to express their knowledge only about Outcomes 3 and 4. But what would happen if, during the test, subjects were permitted to respond separately about all four of the outcomes experienced across the whole experiment? The basis on which integration occurred—the notion that Outcomes 3 and 4 must be treated as substitutes for 1 and 2—would be seriously cast in doubt. If integration bias is responsible for learned predictiveness, then the effect could be abolished in these conditions. By contrast, the proposal that the learned predictiveness effect is the consequence of a learned change in some aspect of attention has no grounds for supposing that this change would be precluded by providing additional tests about Outcomes 1 and 2.

The design of the experiment is summarized in Table 1 (bottom panel). It differs from that of Experiment 1 only in the details of the test phase, which allowed a test of all four outcomes. This was done in different ways in two subexperiments. In Experiment 3a, there were minimal constraints on the test, in that the subjects were simply asked to say what outcome (or outcomes) might follow a given test compound. In Experiment 3b, the subjects in the critical experimental condition (Group 4) were given rating scales for all four outcomes and asked to respond on any of them. This arrangement allowed us to include a control condition (Group 2) in which, as in Experiment 1, the subjects were given scales only for Outcomes 3 and 4. This latter condition should allow replication of the basic learned predictiveness effect, an effect that might, according to the integration account, be absent in Group 4.

Method

Students from the University of the Basque Country agreed to participate in the experiment after being informed that they would take part in an experiment involving cognitive tasks. There were 48 students (34 female; M_{age} = 24.2 years, range: 18–30) in Experiment 3a, and 72 students (43 female; M_{age} = 24.4 years, range: 18–43) in Experiment 3b. In Experiment 3b they were assigned at random to one of two, equal-sized, experimental groups: Group 2 (tested with just two outcomes, 3 and 4), and Group 4 (tested with all four outcomes). All subjects had normal or corrected-to-normal vision.

The apparatus, the cues, and the outcomes used in these experiments were identical to those used in Experiment 1. The training procedure was as described for that experiment except that the image of the cues presented on training trials was not accompanied by text asking what outcome would occur (e.g., *Rash or stomachache?*). Presenting a question with only two options might encourage the participants to think that only two possible outcomes should be considered; by omitting this question we hoped to provide for a less directed procedure with participants having more open choice when it came to the test. The results of Group 2 of Experiment 3b will allow a demonstration that the learned predictiveness effects can be obtained in these circumstances.

The test procedure for Experiment 3a was as follows. During each trial of each of the test phases (that after differential training and that after nondifferential training), participants were presented with a test sheet and were required to write for each trial the name of the outcome, or outcomes, that they expected to occur, and to rate, with a number from 0 to 10, the strength of their expectancy. The instructions explicitly informed the participants that they could rate any of the outcomes experienced up to that point in the task (and more than one outcome if they wished). As in previous experiments, participants had 20 s to respond on each trial.

The test procedure for Group 2 of Experiment 3b was identical to that described for Experiment 1. Group 4 differed only in that, in the final test, subjects were asked to rate all four of the outcomes. During each trial of this test, participants were given a sheet having four scales, labeled with the name of the outcome to rate, one scale for each of the four outcomes. Participants in Group 2 had only two scales, labeled with the names corresponding to Outcomes 3 and 4.

Results and Discussion

Five participants in Experiment 3a, and six participants in Experiment 3b (two in Group 2 and four in Group 4), were excluded because of their poor performance in the test that followed differential training. The remaining participants exhibited good performance on this test, with an average difference scores (rating of correct outcome minus rating of incorrect outcome) of 8.6 in Experiment 3a, and 8.8 in Experiment 3b.

Experiment 3a. The top panel of Figure 3 shows the results from the final test of Experiment 3a. The scores given for Outcomes 1 and 2 are the absolute mean ratings for those stimuli. For the most part, the subjects were willing to give a rating for these outcomes (three cases in which participants did not offer a rating for these outcomes were computed as zero). As before, test stimuli are labeled as accurate or inaccurate predictors according to their role during differential training. As each element in the compounds made up of inaccurate predictors had been equally associated with Outcomes 1 and 2, there is no reason to expect a difference according to outcome here. The same is true for the accurate predictors, as each contained an element that had been reliably followed by Outcome 1 and another that had been reliably followed by Outcome 2. It is apparent, however, that mean ratings were higher (for both outcomes) for the compounds made up of accurate predictors than the compounds made up of inaccurate predictors. An ANOVA with stimulus type (accurate or inaccurate) and outcome (1 or 2) as the variables revealed a main effect of stimulus, $F(1, 42) = 13.92$, $p < .001$, $\eta_p^2 = 0.25$ (0.05–0.44) The effect of outcome, and the interaction (Stimulus × Outcome) were not significant, $Fs < 1$.

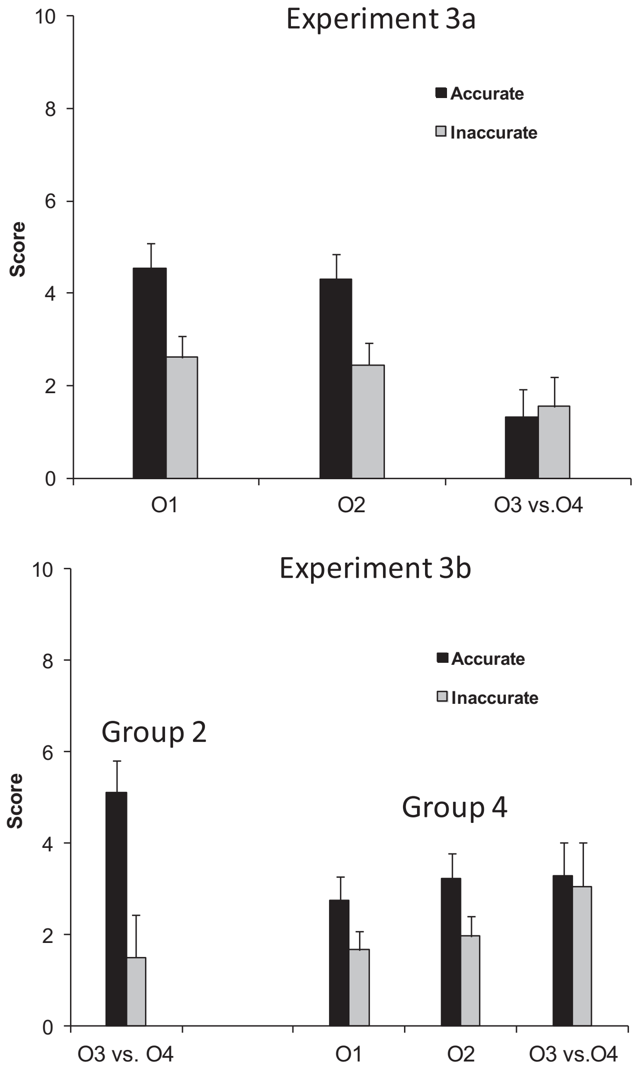

Figure 3. Results of the tests after nondifferential training in Experiments 3a (top panel) and 3b (lower panel). In Experiment 3a, all the participants were allowed to rate any of the outcomes experienced up to that point in the task (O1, O2, O3, O4). Three averaged scores for Accurate and Inaccurate predictors are presented separately: scores for O1, scores for O2, and a difference score (as described for the previous experiments) for O3 and O4. In Experiment 3b, participants in Group 4 were also allowed to rate the four outcomes; however, participants in Group 2 were allowed to rate only O3 and O4 (hence, for them, only the difference score for O3 and O4 is presented). Vertical bars represent standard errors of the mean.

The difference on this test between the accurate and inaccurate predictors is to be expected on the basis of standard accounts of association formation (e.g., Pearce & Hall, 1980; Rescorla & Wagner, 1972). For a compound made of accurate predictors (such as AC) each element would have become strongly linked to it associated outcome (A with 1 and C with 2; see Table 1) during Stage 1 training in which A was always and only followed by Outcome 1 and C was always and only followed by Outcome 2. For a compound made up of inaccurate predictors (such as VX; see Table 1), each element had been followed (on different trials) by both outcomes, and, on each trial, another and more reliable predictor of that outcome had been present (as, e.g., when AV precedes Outcome 1). This arrangement will restrict the growth of associative strength to the inaccurate predictor, and thus the likelihood of the compound generating a high rating on the final test will be reduced.

The top panel of Figure 3 also shows the results of central interest—the scores with respect to Outcomes 3 and 4. (All the participants rated at least one of these outcomes; when the rating of an outcome was absent it was computed as zero.) As before, we present a difference score (rating for correct minus rating for incorrect outcome). The critical result here is that no learned predictiveness effect was evident; that is, there was no significant difference between the accurate and inaccurate compounds in these scores, $t(42) = .263$, $p = .794$. Given the theoretical importance of this null result, we performed an additional analysis from a Bayesian approach in order to provide evidence that such a result is not a matter of the experiment being underpowered. We calculated a Bayes factor (BF) using the techniques developed by Wagenmakers (2007) as presented by Masson (2011), where the BF is estimated from the change in the Bayesian information criterion as $e^{.5(\Delta\mathrm{BIC})}$ with $\Delta\mathrm{BIC} = n \times \ln(SSeffect/(SSeffect + SStotal)) + \ln(n) \times DFeffect$. The BF captures the relative probabilities of null hypotheses and alternative hypotheses, with a factor of one signifying that each is equally likely. The BF captures the relative probabilities of null hypotheses and alternative hypotheses, with a factor of one signifying that each is equally likely. In this case, the analysis yielded a BF of 6.33, suggesting that the null hypothesis is over 6 times more likely than the alternative hypothesis for these data.

It should be noted that the general magnitude of the ratings was lower in this than in the previous experiments. Perhaps this is to be expected on the basis of the integration account of the basic effect. In the standard procedure, the links activating the representations of Outcomes 3 and 4 would be stronger (thus leading to enhanced responses) thanks to the extra amounts of associative strength borrowed from the links established with the Outcomes 1 and 2 previously. If allowing consideration of all four outcomes on test prevents the exploitation of this integrated structure then that borrowing effect might have been canceled resulting in weaker links and responses. It is also possible, of course, that the lowered test performance in comparison with previous experiments just reflects the unusual test procedure used—the absence of scales to rate during the test, and with it, the absence of the names of the outcomes to rate, might well have reduced the confidence of the participants in their responses. It is possible, then, that the null result obtained here simply reflects the insensitivity of a test generating such low scores. This problem is attenuated, however, in the results produced by the procedure used in Experiment 3b.

Experiment 3b. The lower panel of Figure 3 shows performance in the final test of Experiment 3b. The difference scores for Outcomes 3 and 4 for Group 2 (the group tested with just these outcomes) replicated the basic learned predictiveness effect, with higher scores for stimuli that had been accurate rather than inaccurate predictors during differential training; $t(33) = 3.01$, $p = .005$, $d = 0.72$ (0.22–1.23). The scores for Group 4 match those obtained in Experiment 3a. As before, accurate predictors had higher scores than inaccurate predictors in ratings of Outcomes 1 and 2. An ANOVA with stimulus type (accurate or inaccurate) and outcome (1 or 2) as the variables revealed a main effect of stimulus, $F(1, 31) = 6.84$, $p = .014$, $\eta_p^2 = 0.18$ (0.08–0.40), but no other significant effects (Fs < 1). There was no difference (i.e., no learned predictiveness effect) in the test for Outcomes 3 and 4, $t(31) = 0.22$, $p = .82$ (the decisions on null hypothesis rejection in the three latter t tests remained the same when controlling the false discovery rate, $q < .05$; Benjamini & Hochberg, 1995). Again, given the relevance of this null result for theoretical interpretation, we calculated a BF in the way previously described. In this case, the analysis yielded a BF of 5.51, suggesting that the null hypothesis is over five times more likely than the alternative hypothesis for these data. To strengthen the relevant comparison of Groups 2 and 4 in their ratings for Outcomes 3 and 4, an ANOVA, with group (two vs. four) and stimulus type (accurate vs. inaccurate) as the variables, was performed with these data. This analysis revealed a main effect of stimulus, $F(1, 64) = 5.68$, $p = .02$, $\eta_p^2 = 0.08$ (0.001–0.22), a nonsignificant effect of group, $F < 1$, but a significant Group × Stimulus interaction, $F(1, 64) = 4.36$, $p = .041$, $\eta_p^2 = 0.06$ (0.00–0.20). The presence of this interaction reinforces the conclusion that the learned predictiveness effect, evident in Group 2, was abolished in subjects given exactly the same training, differing only in that, on the test, they were permitted to rate not only Outcomes 3 and 4, but also the outcomes experienced in the first stage of training.

These results are consistent with the integration account (in whatever form, whether based on bottom-up or top-down processing). When the test procedure allows the participants to express separately their ratings of Outcomes 1 and 2, the learned predictiveness effect disappears and they respond according to the objective equality with which predictive and nonpredictive cues were paired with Outcomes 3 and 4 during the nondifferential training. These results are not to be expected on the basis of an interpretation in terms of learned changes in attention. Neither modulation of associability nor an account in terms of changes in attention for performance has reason to expect the effect to be abolished by making all outcomes available on test. It remains possible that these attentional processes contribute to the effect that is observed under standard testing conditions. If so, then it might be expected that some (perhaps small) effect should be found in the present experiments. Our results are not decisive on this matter—there is an indication of an effect from the test result for Outcomes 3 and 4 in Experiment 3b, but no sign of one in Experiment 3a.

General Discussion

Le Pelley and McLaren (2003) explained their demonstration of the learned predictiveness effect in terms of the theory of attention in conditioning proposed by Mackintosh (1975). This supposes that stimuli that have been good predictors of their consequences will have their associability enhanced, making them more readily learned about in a subsequent task. The results reported here show that a learned predictiveness effect can be obtained when effects depending on associability (when envisaged simply as a learning rate parameter) cannot be responsible. Our Experiment 2 reversed the usual order of the stages of the learned predictiveness procedure (as was used to demonstrate the basic effect in Experiment 1). It was found that training cues as accurate or inaccurate predictors could modify the effectiveness of these cues in controlling responses about information already acquired in a previous stage of training—that is, before differences in associability (if any) had been established.

One response to this finding is to note that associability is just one aspect of the set of processes labeled "attention" and to acknowledge the role of other aspects (Hall & Rodríguez, 2017, 2019; see also Esber & Haselgrove, 2011; Kruschke, 2001; George & Pearce, 2012; Le Pelley, 2004; Pearce & Mackintosh, 2010). Specifically, we have suggested (Hall & Rodríguez, 2017, 2019) that it is useful to consider also the role of stimulus salience, a property of the stimulus that will influence the vigor with which a cue will evoke a response or activate an associate. Salience is normally dependent on the physical intensity of the stimulus, but its value may change with experience. We have argued that the salience of a stimulus will normally decline with exposure to that stimulus but will be maintained when the stimulus is followed by a consistent consequence. Training a cue of as an accurate predictor will therefore maintain its salience and allow it to evoke a powerful response on test, whether that training preceded or followed acquisition of the information being tested. It may be added that although we have emphasized the role of salience as attention-for-performance, it may also influence the acquisition of associative strength—a cue high in salience (e.g., one that of high physical intensity) will be learned about more readily than one lower in salience. This allows an explanation of the finding reported by Le Pelley et al. (2009) in their Experiment 4. This experiment showed that a novel cue trained in compound with a cue that had previously been established as an accurate predictor (of some other outcome) was learned about less readily than when it was trained in compound with a cue previously established as a poor predictor. This restriction of learning about the novel target cue is an instance of overshadowing, and it is well established that the ability of a stimulus to overshadow depends on its salience or intensity (e.g., Kamin, 1969). The finding of Le Pelley et al. is thus to be expected on the basis of the hypothesis that accurate and inaccurate predictors differ in their effective salience.

In spite of these successes, this version of an attentional account runs into problems with the results of Experiment 3. Here it is shown that the learned predictiveness effect is abolished when subjects are permitted to express views about the relation between the cues and the outcomes that were used to establish these cues as good to poor predictors. There is no reason to expect this result on the basis of an attentional theory. It follows readily, however, from the suggestion that the basic effect depends on a tendency of subjects to integrate information acquired in the two phases of training, and that circumstances likely to restrict such integration will thus eliminate the effect. The present experiment thus adds to recent previous recent literature (e.g., Le Pelley, Mitchell, & Johnson, 2013; Mitchell et al., 2012; Shone, Harris, & Livesey, 2015) in suggesting an important role for nonattentional processes in generating the effect.

Although we have, up to this point, presented the attentional and the integration accounts of learned predictiveness as rivals or alternatives, it is worth considering the possibility that both may be operating in our experiments. They may be seen as complementary processes with a common end—that of optimizing the use of the processing resources of the cognitive system. We can try to illustrate this notion with an everyday example. Consider how we learn to play a new game that involves elements present in another that we have played before. A strategy that will facilitate learning is to take as a starting point knowledge acquired playing the first game. For example, in learning to play poker with dice after playing poker with cards, you will not need to start from zero. You will soon discover that the symbols on dice are substitutes for, or functionally equivalent, to those on the cards, the new game will be recognized as a sort of "poker" and the known rules will be transferred speeding learning of the commonalities (and the differences) of the two versions. This integration process (whether it be driven via bottom-up or top-down processing) will coexist with and interact with attentional processes. For example, there is no doubt that aces are important events in any version of poker which makes them salient stimuli. This high salience will make that stimulus able to command attention, but it will not supply information about its specific role—that information would be encoded in the associative links on which integration will occur.

This general analysis is supported by the results of our Experiments 2 and 3 and by those from several other related studies. If the effect depends on considering the two pairs of outcomes as equivalent, then changing the nature of these pairs should attenuate the effect, which is the result obtained by Le Pelley et al. (2005). In addition, explicit instructions indicating that the arrangements for the two stages are not equivalent should preclude the appearance of the effect. Supporting this, Mitchell et al. (2012; see also Shone, Harris, & Livesey, 2015) found that when participants were explicitly instructed that the accurate predictors in the first stage would be irrelevant in the next stage, the effect was reversed. This would be expected if the procedure led subjects to apply to the accurate predictors the causal structure learned for the inaccurate predictors, and vice versa. The analysis might even be extended to accommodate the result of Le Pelley et al. (2009; Experiment 4) that we explained previously in associative terms as an instance of overshadowing. When during nondifferential training, the familiar cues of the previous stage are presented in compound with novel cues, the use of an integrated causal structure should lead to the participants to infer that a novel cue presented in compound with an accurate predictor, must be an inaccurate predictor, and a novel cue presented in compound with an inaccurate predictor must be an accurate predictor.

Clearly it is now necessary to specify in detail how attentional/associative processes interact with these postulated integration processes. If a formalization can be achieved it is appropriate to ask whether the mechanism proposed also operates in nonhuman animals when trained on discrimination tasks of the sort used here. We acknowledge, however, that a failure to find evidence of integration processes in nonhuman animals would not necessarily have major theoretical significance. It is possible that the primary role that integration processes seem to play in studies of learned predictiveness in people could be a consequence of specific features of the experimental procedure used for human subjects. In particular, presentation of the stimuli is preceded by instructions introducing background information and participants will assume that this information must be used in some way; they are aware that they are subject to an experimental procedure, and that this procedure has been designed by the same researcher who provided the instructions. These features of human causal learning procedures are not present in parallel studies, conducted with similar theoretical purposes, but using nonhuman animals as subjects. Differences in the results from human and nonhuman animal procedures might be due not to fundamental differences in their psychological mechanisms but to the different conditions under which these mechanisms are being tested.

References

Beesley, T., Nguyen, K. P., Pearson, D., & Le Pelley, M. E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *The Quarterly Journal of Experimental Psychology, 68,* 2175–2199. http://dx.doi.org/10.1080/17470218.2015.1009919

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society B, 57,* 289–300.

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. http://dx.doi.org/10.3758/BF03196351

Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: A model of attention in associative learning. *Proceedings Biological Sciences, 278,* 2553–2561. http://dx.doi.org/10.1098/rspb.2011.0836

George, D. N., & Pearce, J. M. (2012). A configural theory of attention and associative learning. *Learning & Behavior, 40,* 241–254. http://dx.doi.org/10.3758/s13420-012-0078-2

Griffiths, O., Johnson, A. M., & Mitchell, C. J. (2011). Negative transfer in human associative learning. *Psychological Science, 22,* 1198–1204. http://dx.doi.org/10.1177/0956797611419305

Hall, G., & Pearce, J. M. (1979). Latent inhibition of a CS during CS-U.S. pairings. *Journal of Experimental Psychology: Animal Behavior Processes, 5,* 31–42. http://dx.doi.org/10.1037/0097-7403.5.1.31

Hall, G., & Pearce, J. M. (1982). Restoring the associability of a preexposed CS by a surprising event. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 34,* 127–140. http://dx.doi.org/10.1080/14640748208400881

Hall, G., & Rodríguez, G. (2010). Attentional learning. In C. Mitchell & M. Le Pelley (Eds.), *Attention and associative learning* (pp. 41–70). Oxford, UK: Oxford University Press.

Hall, G., & Rodríguez, G. (2017). Habituation and conditioning: Salience change in associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 43,* 48–61. http://dx.doi.org/10.1037/xan0000129

Hall, G., & Rodríguez, G. (2019). Attention to perceive, to learn, and to respond. *Quarterly Journal of Experimental Psychology, 72,* 335–345.

Holland, P. C. (1997). Brain mechanisms for changes in processing of conditioned stimuli in Pavlovian conditioning: Implications for behavior theory. *Animal Learning & Behavior, 25,* 373–399. http://dx.doi.org/10.3758/BF03209846

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, NY: Appleton-Century-Crofts.

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology, 45,* 812–863. http://dx.doi.org/10.1006/jmps.2000.1354

Lawrence, D. H. (1949). Acquired distinctiveness of cues; transfer between discrimination on the basis of familiarity with the stimulus. *Journal of Experimental Psychology, 39,* 770–784. http://dx.doi.org/10.1037/h0058097

Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 57,* 193–243. http://dx.doi.org/10.1080/02724990344000141

Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology, 56,* 68–79. http://dx.doi.org/10.1080/02724990244000179

Le Pelley, M. E., Mitchell, C. J., Beesley, T., George, D. N., & Wills, A. J. (2016). Attention and associative learning in humans: An integrative review. *Psychological Bulletin, 142,* 1111–1140. http://dx.doi.org/10.1037/bul0000064

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

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. http://dx.doi.org/10.1037/0097-7403.31.2.226

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

Lochmann, T., & Wills, A. J. (2003). Predictive history in an allergy prediction task. In F. Schmalhofer, R. M. Young, & G. Katz (Eds.), *Proceedings of EuroCodSci 03: The European Cognitive Science Congress* (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. http://dx.doi.org/10.1037/h0076778

Masson, M. E. J. (2011). A tutorial on a practical Bayesian alternative to null-hypothesis significance testing. *Behavior Research Methods, 43,* 679–690. http://dx.doi.org/10.3758/s13428-010-0049-5

Mitchell, C. J., Griffiths, O., Seetoo, J., & Lovibond, P. F. (2012). Attentional mechanisms in learned predictiveness. *Journal of Experimental Psychology: Animal Behavior Processes, 38,* 191–202. http://dx.doi.org/10.1037/a0027385

Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87,* 532–552. http://dx.doi.org/10.1037/0033-295X.87.6.532

Pearce, J. M., & Mackintosh, N. J. (2010). Two theories of attention: A review and a possible integration. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 11–40). Oxford: New York: Oxford University Press.

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

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84,* 127–190. http://dx.doi.org/10.1037/0033-295X.84.2.127

Shone, L. T., Harris, I. M., & Livesey, E. J. (2015). Automaticity and cognitive control in the learned predictiveness effect. *Journal of Experimental Psychology: Animal Learning and Cognition, 41,* 18–31. http://dx.doi.org/10.1037/xan0000047

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review, 14,* 779–804. http://dx.doi.org/10.3758/BF03194105

Received October 17, 2017
Revision received January 17, 2019
Accepted January 27, 2019 ■