The Redundancy Effect in Human Causal Learning: Evidence Against a Comparator Theory Explanation

Tara Zaksaite (gintare.zaksaite@plymouth.ac.uk)
School of Psychology, Plymouth University, Plymouth, UK

Peter M. Jones (peter.m.jones@plymouth.ac.uk)
School of Psychology, Plymouth University, Plymouth, UK

Abstract

The blocking effect, canonical in the study of associative learning, is often explained as a failure of the blocked cue to become associated with the outcome. However, this perspective fails to explain recent findings that suggest learning about a blocked cue is superior to a different type of redundant cue. We report an experiment designed to test the proposal that blocking is not a failure of association, but a performance effect arising from a comparator process (Denniston, Savastano, & Miller, 2001). Participants received A+ AX+ BY+ CY- training containing a blocked cue X and another redundant cue Y, before rating outcome expectancies for individual cues. These ratings were inconsistent with the association-failure view. After subsequent A- Y+ training, participants rated cues again. Ratings in the second test were inconsistent with the comparator theory. Our data suggest that neither perspective is likely to provide a complete account of causal learning.

Keywords: associative learning; comparator theory; redundancy effect; blocking; cue competition

Introduction

In a typical causal learning task, participants are required to learn which cues cause an outcome. Many such tasks involve presentation of more than one cue on each trial, and this typically results in cue competition. That is, learning about a cue is dependent to some extent on accompanying cues. Probably the best-known example of cue competition is blocking (e.g. Dickinson, Shanks, & Evenden, 1984). In a blocking task, participants receive trials on which cue A is paired with an outcome (denoted A+) and trials on which A is presented alongside a second cue, X, and paired with the outcome (AX+). Blocking is said to have occurred if learning about X is restricted by the presence of A, relative to a control condition in which A+ trials are omitted. Learning about X is therefore influenced by the presence and associative history of A. This finding is analogous to classic demonstrations of blocking in nonhuman animals (e.g. Kamin, 1969).

Following the discovery of cue competition effects, Rescorla and Wagner (1972) outlined an elegant and much-cited model according to which an outcome will only support learning if it is surprising. Surprise is equivalent to prediction error, or the discrepancy between the outcome that is expected and the outcome that occurs. When an unexpected outcome occurs, the resulting prediction error enables the formation of an association between any cues present and the outcome. Critically however, outcome expectancy is based on all the cues that are present rather than individual cues. To illustrate this, consider the blocking effect. On AX+ trials, expectancy of the outcome is based on the extent to which it is predicted by both A and X. Because A is established as a predictor of the outcome on A+ trials, the outcome is expected on AX+ trials and little learning can take place. Learning about X is therefore 'blocked' by the presence of A. If A were not separately paired with the outcome, blocking would not occur. Informally, we can say that X is blocked because it is informationally redundant; it indicates no change in the outcome that is predicted by A. According to the Rescorla-Wagner model, this is operationalized as a failure by X to become associated with the outcome.

This description of blocking as a failure of association formation has been called into question by a recent result comparing learning about a blocked cue with another kind of redundant cue. Uengoer, Lotz, and Pearce (2013) compared learning about the blocked cue X with cue Y from a BY+ CY- discrimination. Here, the outcome was predicted by B and its absence was predicted by C. We refer to the common cue, Y, as an uncorrelated cue because it is paired with both the presence and the absence of the outcome. Uengoer et al. gave participants A+ AX+ BY+ CY- training, followed by test trials on which they were asked to rate outcome expectancy for each cue. The Rescorla-Wagner (1972) model predicts that learning about X should be blocked by A, as described above. The prediction for Y is perhaps more complex, but the model predicts that the strength of the association between Y and the outcome will increase overall during training. This results from the use of a combined prediction error in determining learning, as follows: On BY+ trials, the associations between B and Y and the outcome should strengthen. On CY- trials, the association between Y and the outcome should lead to expectation of the outcome, and its non-occurrence will in turn lead to decreases in the extent to which both C and Y predict the outcome. As a consequence, C should be established as an inhibitor of the outcome. This will enable Y to maintain its association with the outcome to some extent. Informally, we can say that Y could be a cause of the outcome if its absence on CY- trials is explained by the preventative status of C. The Rescorla-Wagner model, then, predicts that Y will become better associated with the outcome than will X. Contrary to this prediction, Uengoer et al.'s participants rated X as a more likely cause of the outcome than Y. This finding is known as the redundancy effect (for corresponding results using rats and pigeons, see Jones & Pearce, 2015; Pearce, Dopson, Haselgrove, & Esber, 2012).

Comparator theory

While the redundancy effect is not predicted by the Rescorla-Wagner (1972) model, it perhaps makes intuitive sense because X is consistently paired with the outcome. Y, on the other hand, is paired with the outcome only intermittently. Ignoring any effect of cue competition, we might expect X to become better associated with the outcome than Y. Accordingly, Uengoer et al. (2013) considered whether their results might be better accounted for by supposing that blocking occurs not because X fails to become associated with the outcome, but because of an additional process that acts during the test. According to comparator theory (Denniston, Savastano, & Miller, 2001), association formation is non-competitive and driven by an individual prediction error for each cue. Cue competition is then accounted for by a comparator process that operates at test to influence performance. This process compares the associative status of the target cue with that of any cues that have previously been presented alongside the target. This results in a decrease if companion cues have a strong association with the outcome, and an increase if the association is weak. In the case of X, outcome expectancy will be reduced because A is strongly associated with the outcome, and blocking will occur. This model also predicts the redundancy effect, because association formation is governed by each cue's relationship with the outcome. X should be better associated with the outcome than Y, because Y is only followed by the outcome on 50% of trials. The comparator theory therefore seems like a promising candidate for explaining both blocking and the redundancy effect.

However, two attempts have been made to test this account and both have cast doubt on its validity. Jones and Pearce (2015) conducted an experiment in which rats were given A+ AX+ BY+ CY- training, where each cue was an auditory or visual stimulus and the outcome was the delivery of a sucrose solution. Rats were subsequently tested in extinction with B, X, and Y. A larger response was elicited by X than Y, demonstrating the redundancy effect. Responding was also higher for B than for X. Jones and Pearce suggested that this was important, because it allowed a further test of the comparator theory. According to this theory, because B and X were both consistently paired with the outcome, they should have become associated with the outcome to the same extent. The larger response for B than for X at test must therefore have been the result of the comparator process. Because B had been presented alongside Y, which was only weakly associated with the outcome, the response to B was left largely intact. For X, however, the response was moderated because X had been trained alongside A, which was strongly associated with the outcome. To test this account, rats were given A- Y+ training. Following this, they were again tested with B and X. The comparator theory now predicts greater responding for X than for B, but the results closely resembled those from the first test. B elicited more responding than X despite revaluation of the comparator cues A and Y, apparently in contradiction of the theory. An objection may be raised, however, because of the nature of the outcome used in this experiment. Miller and Matute (1996) suggested that, once a target cue becomes associated with an outcome of motivational significance, the target cue itself acquires motivational significance. As a result, attempts to deflate responding to the target cue by further conditioning of an associate cue may be unsuccessful. In the experiment reported by Jones and Pearce, the appetitive outcome is likely to have had substantial motivational significance. It is therefore possible that responding to B was unaffected by Y+ training, not because the comparator theory is incorrect but because the manner in which it was tested was inadequate. Urushihara and Miller (2010) noted that such revaluation effects are difficult to observe in nonhuman animals because of the use of motivationally significant outcomes, but occur frequently in human causal learning.

There also exists a test of whether the comparator theory can account for the redundancy effect in humans, reported by Uengoer et al. (2013). Since blocking is dependent on a comparison between X and A, it follows that revaluation of A should increase outcome expectancy for X. In one experiment, following initial A+ AX+ BY+ CY- training and subsequent individual cue testing, participants were given A- training and a further test. They found that outcome expectancy for X was equivalent for the two tests, contrary to the predictions of the comparator theory. This conclusion should be treated with caution, however. The crucial comparison is between outcome expectancy for X during the first and second tests. This means that the results are likely to have been contaminated to some extent by order effects. In the present paper, we report an experiment intended to provide a fairer test of the comparator theory. The experiment is conceptually similar to the Jones and Pearce experiment, except that it used human participants and a causal learning task. It therefore combines the better aspects of the existing evaluations of the comparator theory described above, while eliminating the shortcomings. The use of human participants should provide ideal conditions for observing revaluation effects and, because the adequacy of the comparator theory can be assessed by comparing B and X in the same test, the confounding effect of order present in the Uengoer et al. experiment is avoided.

A test of the comparator theory

The design of this experiment is summarized in Table 1. Stage 1 of the experiment was designed to establish the causal status of B, X, and Y. Each participant received four types of trial: A+, AX+, BY+, and CY-. Following Uengoer et al. (2013), training was embedded in a variant of the classic allergist task (Aitken, Larkin, & Dickinson, 2000).

On each trial, participants were shown one or two food pictures and asked to predict whether they would lead to stomach ache in a fictional patient, Mr. X. After participants made their predictions, they received feedback on whether stomachache did (+) or did not (-) occur. After the completion of Stage 1, a test stage was administered in which participants were shown the five individual food cues and asked to rate the likelihood of stomach ache for each food using a rating scale. These ratings served as the measure of outcome expectancy for each cue. We expected these ratings to resemble those obtained by Uengoer et al. That is, we expected ratings to be higher for X than for Y (the redundancy effect) and to be higher for B than for X. We also expected ratings to be high for A and low for Y. After this test, participants received further training in Stage 2. This training was designed to revalue A and Y, and consisted of A- and Y+ trials. Following this training, outcome expectancies were again measured in the same way as in the earlier test. If the comparator theory (Denniston, Savastano, & Miller, 2001) is correct, ratings for X should be higher than ratings for B in this test. Alternatively, if the outcome expectancy for B was higher than for X at Test 1 because of a difference in the strength of associations formed between these cues and the outcome during Stage 1, then ratings should still be higher for B than for X at Test 2.

Table 1: The design of the experiment.

Stage 1	Test 1	Stage 2	Test 2
A+	A	A-	A
AX+	B	Y+	B
BY+	C		C
CY-	X		X
	Y		Y
8 blocks	2 blocks	8 blocks	2 blocks

Method

Participants The participants were 50 Plymouth University undergraduate students studying Psychology. They received course credit for their participation in this experiment. They were aged 18-53 years (*M*=21.86, *SD*=7.1) and five were male.

Materials The experiment was run using computers attached to 22-inch monitors with a 1920 x 1080 resolution. The experiment was designed, cues presented and responses recorded, using E-prime 2.0 software (Psychology Software Tools, PA, US).

The cues were five images of foods on a white background, each measuring 300 x 300 pixels. The foods were: apple, cherry, grape, lemon and strawberry. Foods were randomly assigned to serve as each cue (A, B, C, X, Y) for each participant. Outcomes were stomach ache, signified by text and a sad face on a red background, and no stomachache, indicated by text and a happy face on a green background. Cues and outcomes were presented on a black background with white text. Participants responded using the mouse.

Procedure Each participant was initially asked to read on-screen instructions that were identical to those used by Uengoer et al. (2013). In the first stage of the experiment participants were presented with eight blocks of trials. Each of the four trial types (A+, AX+, BY+, CY-) were presented once per block, and were randomized within each block. Each trial started with the presentation of either one or two images of foods, below the phrase "The patient ate the following food(s):" The sentence "Which reaction do you expect?" was presented below the images. Participants responded by clicking one of two response buttons placed at the bottom of the screen. The left-hand button was labelled "No stomach ache", and the right-hand button was labelled "Stomach ache". As soon as the participant responded, the response buttons and the sentence above them were replaced by a statement and image showing the outcome of the trial. When the outcome was stomach ache, the statement was "The patient has stomach ache" and the picture of a sad face was shown. When the outcome was no stomach ache, the statement was "The patient has no stomach ache", and the picture of a happy face was shown. This feedback display remained on the screen for 3 s and was then followed by the next trial.

After the completion of Stage 1, Test 1 began. Here the participants were instructed to judge the probability with which specific foods will cause stomach ache in the absence of feedback. On each trial a single food was presented on the screen below the sentence "What is the probability that the food causes stomach ache?" Participants responded by clicking on an 11-point rating scale ranging from 0 (*Certainly not*) to 10 (*Very certain*). After participants chose a rating for each food, a blank screen was shown for 1 s. Each food that appeared in Stage 1 was presented twice, with the order of trials randomly determined for each participant. For each participant, the average of the two outcome expectancy ratings was calculated and used in subsequent analyses.

Participants then received further training in Stage 2. Training consisted of eight blocks of two trial types (A-, Y+) appearing once per block in a random order. The procedure for this stage was otherwise identical to Stage 1. Test 2 then measured final outcome expectancies, using the same procedure as Test 1.

Results

We applied an inclusion criterion of 60% correct predictions in Stage 1, commonly used in similar work (e. g. Le Pelley & McLaren, 2003). Four participants failed to meet this, and are excluded from all subsequent analyses. The remaining 46 participants learned readily, and made 98% correct responses during the final block of trials of Stage 1. These participants also made correct predictions on 98% of trials during the final block of Stage 2. Our analyses here focus on the critical test data.

Mean ratings from Tests 1 and 2 are shown in Figure 1. The pattern of results in Test 1 closely resembles those obtained by Uengoer et al. (2013), with higher ratings for X than for Y, and higher ratings for B than for X. For the comparator theory, the crucial comparison is between B and X at Test 2. As for Test 1, ratings for B were higher than for X. A two-way ANOVA with test and cue variables was conducted. This revealed a significant effect of test, $F(1, 45) = 6.59$, $p = .014$, $\eta_p^2 = .128$, a significant effect of cue, $F(4, 180) = 80.76$, $p < .001$, $\eta_p^2 = .642$, and a significant interaction, $F(4, 180) = 207.85$, $p < .001$, $\eta_p^2 = .822$. To explore this interaction, simple effects analyses were used to compare ratings from Tests 1 and 2 for each cue. Ratings differed between tests for A, $F(1, 45) = 431.15$, $p < .001$, $\eta_p^2 = .905$, for B, $F(1, 45) = 29.30$, $p < .001$, $\eta_p^2 = .394$, for C, $F(1, 45) = 23.28$, $p < .001$, $\eta_p^2 = .341$, and for Y, $F(1, 45) = 359.38$, $p < .001$, $\eta_p^2 = .889$. Ratings for X did not differ between tests, $F < 1$.

Separate analyses were conducted to test the most informative comparisons. Firstly, in order to check that the redundancy effect was obtained, we used a within-subjects t-test to compare ratings for X and Y at Test 1. Ratings for X were significantly higher than for Y, $t(45) = 7.58$, $p < .001$. Secondly, to confirm that the revaluation of A and Y was successful, we conducted a two-way ANOVA comparing ratings for A and Y in the two tests. We found an effect of cue, $F(1, 45) = 6.37$, $p = .0151$, $\eta_p^2 = .124$, an effect of test, $F(1, 45) = 5.12$, $p = .029$, $\eta_p^2 = .102$, and importantly, a significant interaction, $F(1, 45) = 559.83$, $p < .001$, $\eta_p^2 = .926$. Exploring this interaction, we found that ratings for A were higher than for Y at Test 1, $F(1, 45) = 559.62$, $p < .001$, $\eta_p^2 = .926$, but that ratings were higher for Y than for A at Test 2, $F(1, 45) = 112.41$, $p < .001$, $\eta_p^2 = .714$. The revaluation of A and Y was therefore successful. Thirdly, to test the predictions of the comparator theory, we conducted a similar two-way ANOVA to compare ratings for B and X at Test 1 and Test 2. We found an effect of cue, $F(1, 45) = 46.82$, $p < .001$, $\eta_p^2 = .510$, an effect of test, $F(1, 45) = 21.40$, $p < .001$, $\eta_p^2 = .322$, and a significant interaction, $F(1, 45) = 18.81$, $p < .001$, $\eta_p^2 = .295$. Exploring the interaction, we found that ratings for B were higher than for X at both Test 1, $F(1, 45) = 62.86$, $p < .001$, $\eta_p^2 = .583$, and Test 2, $F(1, 45) = 16.52$, $p < .001$, $\eta_p^2 = .268$. This disconfirms the predictions of the comparator theory. If outcome expectancies for B and X were determined by a combination of direct associations with the outcome and comparison with Y and A respectively, then ratings for X should have been higher than for B at Test 2. One notable feature of the data that might suggest some role for a comparator process is the change in ratings for B between the two tests. Participants rated B as a less likely cause of the outcome after Y+ training than they did before, which is consistent with the comparator theory. However, an opposite effect was observed for C. Ratings for C were higher at Test 2 than at Test 1, which is the opposite change to that predicted by the comparator theory. It therefore seems likely that these changes are not the result of a comparator process, but rather a general decrease in certainty at Test 2. Since A and Y had been revalued in Stage 2, some participants may have assumed that associations learned during Stage 1 were no longer reliable.

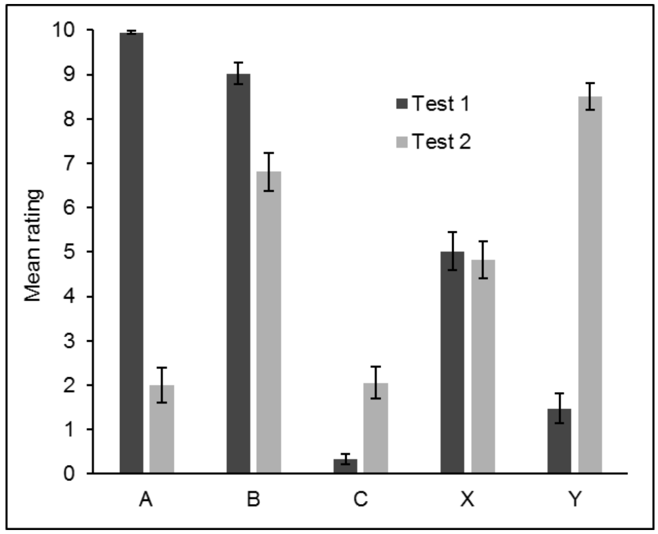

Figure 1: Mean ratings for Test 1 and Test 2, for each cue. Error bars show the standard error of the mean.

Discussion

The experiment reported here was designed to test an explanation of the redundancy effect based on the comparator theory (Denniston, Savastano, & Miller, 2001). Following A+ AX+ BY+ CY- training, participants were asked to rate the probability of the outcome occurring for each individual cue in Test 1. Ratings were higher for X than for Y; we therefore replicated the redundancy effect (Uengoer et al., 2013). This finding is consistent with the comparator theory, which states that the strength of the association formed between a cue and an outcome is determined by an individual (i.e. non-competitive) prediction error. Since X was consistently paired with the outcome and Y was not, it follows that X should have become better associated with the outcome than Y. Participants also gave higher ratings for B than for X during Test 1. Again, this is consistent with the comparator theory. Although the theory predicts that each of these cues will have become associated with the outcome to the same extent, it also states that outcome expectancies should have been moderated by the comparator process at test. Specifically, outcome expectancy for X should have been reduced because it had been trained alongside A, which was strongly associated with the outcome. Any reduction in outcome expectancy for B should have been smaller, because it had been trained alongside Y, which was only weakly associated with the outcome. However, the comparator theory is not consistent with the results of Test

2. Following Stage-2 A- Y+ training, participants again rated the probability of the outcome occurring for each cue. Ratings for B were again higher than for X. The comparator theory, however, predicts the opposite pattern of results. This is because, although the associations with the outcome should have remained unchanged for both B and X, the associative status of their comparator cues had changed. We therefore conclude that the comparator theory cannot account for our results. Of course, this conclusion relies on the assumption that Stage-2 training was successful in revaluation of A and Y. This is apparent in the higher ratings given for Y than for A at Test 2.

Our results are also difficult to reconcile with the model of learning proposed by Rescorla and Wagner (1972). Because it describes learning as being the result of a combined prediction error, X should have failed to become associated with the outcome and should have been rated as a less likely cause of the outcome than Y at Test 1. In other words, the Rescorla-Wagner model fails to account for the redundancy effect because it incorrectly predicts that learning about the blocked cue will be prevented. However, Vogel and Wagner (2017) have suggested a way in which the prediction of blocking can be modified to accommodate the redundancy effect. Their modification is based on the assumption that each cue shares some common features, denoted K. The training given in Stage 1 here could therefore be re-described as AK+ AXK+ BYK+ CYK-. K should become associated with the outcome, with two consequences that are relevant for interpreting the redundancy effect. Firstly, because K is present on CYK-trials, overexpectation of the outcome is increased and the weakening of the association between Y and the outcome on these trials is more substantial than when K is omitted. Secondly, When XK is presented at test, outcome expectancy is boosted by K; the model can therefore predict greater outcome expectancy for XK as a result of including the common features. Combination of these two changes allows the model to predict the redundancy effect. This version of the model also makes an interesting prediction regarding the effect of adding further trial types to Stage-1 training. Because the extent to which K becomes associated with the outcome is critical, adding extra trials on which the outcome does not occur (e.g. DK-) should reduce the influence of common features and eliminate the redundancy effect. This prediction remains untested. If it is correct, it would lend support to an account that provides a way to reconcile the Rescorla-Wagner model with the redundancy effect.

Another possibility is that learning is governed by quite different rules. Not all models of learning make such strong predictions about the restriction of learning about blocked cues. Pearce's (1987, 1994) configural model, for instance, predicts substantial outcome expectancy for blocked cues. According to this model, participants learn about configural representations that include all cues present on a given trial, rather than each cue entering into its own association with the outcome. In the case of blocking, participants would come to associate A with the outcome on A+ trials, and to associate AX with the outcome on AX+ trials. Outcome expectancy for X alone would then be determined by generalization from AX, based on their similarity. Outcome expectancy for X would therefore be weaker than for AX, but considerably stronger than it would have been without any training. However, Pearce at al. (2012) note that the theory is unable to predict the redundancy effect because it predicts that outcome expectancy for Y will be higher still. As with the Rescorla-Wagner (1972) model, it is possible that some modification of the configural model would alter this prediction, but it is not clear at present what that modification might be.

Whether cues are learned about individually or as configurations, the redundancy effect might be accommodated if we suppose that the amount of attention paid to blocked and uncorrelated cues changes during training. For instance, it is commonly assumed (Le Pelley, 2004; Mackintosh, 1975) that cues are processed to the extent that they have predictive value. Since blocked and uncorrelated cues are both redundant, we might expect the amount of attention they are paid to be reduced. In order to explain the redundancy effect, however, we need to propose that this reduction in attention differs in magnitude for blocked and uncorrelated cues. If we suppose that participants learn quickly that Y is irrelevant during BY+ CY- training, then we might expect substantial decreases in the amount of attention paid to Y and a weak association between Y and the outcome as a result. Attention to X, on the other hand, might be maintained for longer, allowing a stronger association to form between X and the outcome. In an attempt to evaluate this claim empirically, Jones and Zaksaite (2017) monitored participants' eye gaze during A+ AX+ BY+ CY- training. The duration of eye gaze for each cue has been used extensively as a measure of overt attention in learning tasks (e.g. Beesley & Le Pelley, 2011). Jones and Zaksaite found that participants spent more time looking at Y than at X, but that this was likely to have been a consequence of differing trial durations. When X and Y were presented on the screen together in a subsequent stage of training, gaze was equivalent for each. This experiment therefore failed to provide any evidence that the amount of attention paid to blocked and uncorrelated cues differs.

In addition to associative accounts of blocking, others (e.g. Lovibond, Been, Mitchell, Bouton, & Frohardt, 2003) have argued that blocking is the result of inferential reasoning. According to this view, blocking occurs because participants do not have independent evidence that the blocked cue causes the outcome (i.e. training trials on which X is presented without A). However, since participants also lack evidence that the blocked cue does *not* cause the outcome, they should be uncertain about the causal status of X and blocking should be relatively weak. This uncertainty might be enhanced because the magnitude of the outcome is fixed, meaning that compound presentation of two causal cues would lead to the same outcome as either cue alone. Lovibond et al. provided support for this position by showing that blocking is enhanced when the magnitude of the outcome varies in accordance with the number of causes present, allowing participants to infer that the blocked cue is not a cause of the outcome. In light of this account, we should consider whether the intermediate ratings for X in the present experiment were the result of an intermediate level of learning, or of uncertainty about its causal status. An unpublished experiment from our laboratory suggests that this might be a promising approach. In addition to rating the probability of the outcome for each cue, participants rated their confidence in these judgments. Confidence ratings were lower for blocked than for uncorrelated cues, suggesting that the redundancy effect might be due at least in part to uncertainty about X.

Conclusion

We have considered theories that account for learning by using individual and combined prediction errors. While combined prediction error models (e.g. Rescorla & Wagner, 1972) are difficult to reconcile with the redundancy effect, individual prediction error does not result in cue competition effects such as blocking, unless an additional process is invoked. We tested a theory that includes such a process (Denniston, Savastano, & Miller, 2001), but found that it was not consistent with the results of Test 2. We suggest two lines of future enquiry. Firstly, data should be collected that evaluate the predictions arising from Vogel and Wagner's (2017) addition of common features to simulations of the Rescorla-Wagner model. While the cues used in the present experiment are likely to have shared some common features, we cannot currently evaluate the claim that learning about these features enables the redundancy effect to occur. Secondly, since neither combined nor individual prediction errors seem capable of producing our results, attempts should be made to evaluate some combination of the two. In particular, any models containing both kinds of prediction error should be tested against the idea that cue competition occurs because of inferential reasoning processes.

References

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

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

Denniston, J. C., Savastano, H, I., & Miller, R. R. (2001). The extended comparator hypothesis: Learning by contiguity, responding by relative strength. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories*. Mahwah, NJ: Erlbaum.

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

Jones, P. M., & Pearce, J. M. (2015). The fate of redundant cues: Further analysis of the redundancy effect. *Learning and Behavior, 43*, 72-82.

Jones, P. M., & Zaksaite, T. (2017). The redundancy effect in human causal learning: no evidence for changes in selective attention. Manuscript submitted for publication.

Kamin, L. J. (1969). Selective attention and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning*. Halifax, Nova Scotia: Dalhousie University Press.

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

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

Lovibond, P. E., Been, S. L., Mitchell, C. J., Bouton, M. E., & Frohardt, R. (2003). Forward and backward blocking of causal judgment is enhanced by additivity of effect magnitude. *Memory and Cognition, 31*, 133-142.

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

Miller, R. R., & Matute, H. (1996). Biological significance in forward and backward blocking: Resolution of a discrepancy between animal conditioning and human causal judgement. *Journal of Experimental Psychology: General*, *125*, 370-386.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review, 94,* 61-73.

Pearce, J. M. (1994). Similarity and discrimination: a selective review and a connectionist model. *Psychological Review, 101*, 587-607.

Pearce, J. M., Dopson, J. C., Haselgrove, M., & Esber, G. R. (2012). The fate of redundant cues during blocking and a simple discrimination. *Journal of Experimental Psychology: Animal Behavior Processes, 38,* 167-179.

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 theory and research*. New York, NY: Appleton-Century-Crofts.

Uengoer, M., Lotz, A., & Pearce, J. M. (2013). The fate of redundant cues in human predictive learning. *Journal of Experimental Psychology: Animal Behavior Processes, 39*, 323-333.

Urishihara, K., & Miller, R. R. (2010). Backward blocking in first-order conditioning. *Journal of Experimental Psychology: Animal Learning and Cognition*, *36*, 281-295.

Vogel, E. H., & Wagner, A. R. (2017). A theoretical note on the interpretation of the "redundancy effect" in associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition, 43*, 119-125.