

The Fate of Redundant Cues During Blocking and a Simple Discrimination

John M. Pearce and Jemma C. Dopson
Cardiff University

Mark Haselgrove
University of Nottingham

Guillem R. Esber
University of Maryland

In each of three experiments animals received blocking, A+ AX+, in which food was always presented after one stimulus, A, that was occasionally accompanied by another stimulus, X. They also received a simple discrimination, AX+ BX−, in which the presence and absence of food was signaled by two compounds that contained one unique cue, A or B, and one common cue, X. In each of these designs, X can be said to be redundant relative to A as a signal for food. Test trials at the end of training revealed that responding during X was stronger after blocking than after the simple discrimination. These results contradict predictions from theories of learning that assume changes in associative strength of a stimulus are determined by a global error term based on the outcome predicted by all the stimuli that are present for a conditioning trial. The results are interpreted, instead, by assuming either that animals store a memory of every trial to which they have been exposed, or that learning is governed by an error term based on the significance of individual stimuli.

Keywords: blocking, relative validity, appetitive conditioning, redundant cues

In any conditioning task some conditioned stimuli (conditional stimulus, CS) will be more informative about whether the unconditioned stimulus (US) will be delivered than others. Moreover, it appears that informative stimuli gain more control over behavior than those that are uninformative, or redundant. A classic demonstration of this effect was provided by Wagner, Logan, Haberlandt & Price (1968, see also Wagner, 1969) in an experiment with two groups. One group received what we shall refer to as a simple discrimination in which the US was presented after one compound CS but not another, with each compound containing a unique and a common cue, AX+ BX−. In this discrimination, A can be regarded as informative because it signals reliably whether the US will be delivered, whereas X can be regarded as redundant as it is less informative than A about the delivery of the US. Subsequent test trials revealed a strong conditioned response (CR) in the presence of A, and a weaker CR in the presence of X. The second group received the same two compounds but they were each paired with the US on half of the trials (AX+/- BX+/-). Here, X and A are equally informative about the delivery of the US, and subsequent test trials revealed a stronger CR during X than after the simple discrimination (see also: Pearce, Esber, George &

Haselgrove, 2008). Thus even though X was paired with the US according to the same schedule in both groups, it gained less influence over behavior when it was less informative than equally informative about the trial outcome, relative to the stimuli that accompanied it. This outcome of the experiment is referred to as the relative validity effect.

Another demonstration of how the informativeness of a stimulus determines the influence it has over behavior is blocking. In the classic demonstration of this effect by Kamin (1969), a blocking group of rats first received trials in which a shock was signaled by a noise. The same shock was then signaled by a noise-light compound. Given this training, the noise can be regarded as being informative about shock, as every delivery of this US was preceded by the noise, while the light can be regarded as being redundant as a cue for indicating whether the shock will occur. In keeping with the suggestion that redundant cues gain little control over behavior, a subsequent test with the light by itself revealed a weak CR. Moreover, this response was weaker than for a control group in which the light was made as informative as the noise about the delivery of shock by omitting the initial phase of training. A similar outcome was revealed by Wagner (1969), who trained rats to lever press for food. Two groups received reward for lever pressing during trials with a tone and light, while one of them also received, intermixed among these trials, reward for lever pressing during tone alone. Test trials with the light revealed stronger responding in the group that received only the trials with the tone-light compound. Thus blocking can be found when the two types of training trials are presented in separate phases, or when they are intermingled.

While it is generally acknowledged that redundant cues gain relatively little control over responding, rather little is known about the ultimate fate of a redundant cue. One possibility is that, irrespective of the manner in which redundancy is established, an

John M. Pearce and Jemma C. Dopson, School of Psychology, Cardiff University, Cardiff, United Kingdom; Mark Haselgrove, School of Psychology, University of Nottingham, Nottingham, United Kingdom; Guillem R. Esber, School of Medicine, Department of Anatomy and Neurobiology, University of Maryland.

The research reported in this article was supported by a BBSRC grant to John M. Pearce (# BB/H006176) at Cardiff University.

Correspondence concerning this article should be addressed to John M. Pearce, School of Psychology, Cardiff University, Cardiff, CF10 3AT United Kingdom. E-mail: pearcejm@cardiff.ac.uk

uninformative CS will, ultimately, exert no influence over behavior. Another possibility is that an uninformative CS retains a measure of influence over behavior, which is determined by the nature of the training it has received. In terms of blocking, for example, even though a cue might be redundant, it is nonetheless always followed by the US and thus may be regarded as having some informative value. In the case of the common cue from a simple discrimination, on the other hand, not only is it less informative about the trial outcome than the stimuli that accompany it, but by itself it is also unreliable as a predictor of when the US will occur. On this basis, therefore, it might be thought that information provided by the redundant cue from a simple discrimination is of less value than the redundant cue from a blocking experiment, with the consequence that the latter will elicit a stronger CR than the former. These proposals are admittedly speculative and informal, but they do not seem entirely implausible. It is thus of some interest to note that according to currently influential theories of learning, the opposite should be found. That is, the CR elicited by the redundant cue is predicted to be weaker after blocking than after a simple discrimination.

To make this point, we shall focus on the rule proposed by Rescorla and Wagner (1972), which stipulates that the change in the associative strength of a CS, ΔV , is given by Equation 1, where λ is the asymptote of conditioning supported by the US, V_{Σ} is the sum of the associative strengths of all the stimuli present on a trial, and α and β are learning rate parameters with values between 0 and 1.

$$\Delta V = \alpha \cdot \beta \cdot (\lambda - V_{\Sigma})$$

When this rule is applied to a blocking design, A+ AX+, then it follows that once A has gained asymptotic associative strength, there will be no opportunity for X to gain any associative strength, no matter how many compound trials are presented. Thus blocking should prevent the redundant cue from gaining any control at all over behavior. In the case of a simple discrimination, AX+ BX-, the trials with AX will permit X to gain a measure of associative strength. Although some of this strength will be lost on BX trials, as training progresses B will gain negative associative strength and prevent to some extent this loss. At asymptote, the theory predicts A will have high associative strength, X will have a modest level of excitatory strength, and B will have negative associative strength that matches in absolute magnitude that acquired by X. In terms of this analysis, therefore, the redundant cue, X, will gain more associative strength after a simple discrimination than after a blocking treatment. Other theories make similar predictions (e.g., Pearce, 1987, 1994), but to avoid repetition these are considered in the General Discussion where they can be considered in the light of the experimental findings.

To our knowledge, the effects of the two treatments just considered have not been compared in a single experiment. There are numerous experiments exploring the effects of blocking and the simple discrimination, but given the many differences in experimental design and the behavior that was recorded, there is little to be gained by comparing their findings. Suffice it to say that some experiments have found that training of the sort A+ AX+ results in relatively weak responding during tests with X (e.g., Kamin, 1969), whereas others have found it results in relatively strong responding during X (e.g., Good & Macphail, 1994). Likewise, Dopson, Esber, and Pearce (2010) found extremely weak respond-

ing during X after an AX+ BX- discrimination, whereas others have found a modest level (e.g., Wagner et al., 1968).

Experiment 1

A between-groups design was adopted for Experiment 1 using appetitive Pavlovian conditioning with rats (see Table 1). During Stage I a blocking group received trials in which food was signaled by CS A. The group then received AX+ trials in Stage II along with continued training with A+. This treatment, rather than the more conventional methodology of presenting A+ trials before but not intermingled with AX+ trials, was adopted to encourage conditioning with A to reach asymptote. It was also adopted because we have found it to be successful in obtaining reliable blocking with appetitive Pavlovian conditioning (Dopson, Pearce & Haselgrove, 2009). A discrimination group received an AX+ BX- discrimination in Stage II. To equate as closely as possible the experiences of the two groups, the discrimination group received C+ trials in Stage I, and the blocking group received C- trials in Stages I and II. The final session of Stage II included test trials with X presented by itself.

Method

Subjects. The subjects were 32 experimentally naïve male, Lister Hooded rats. Before the experiment they were gradually reduced to 80% of their free-feeding weights and were maintained at this level by being fed a restricted amount after each experimental session. The rats were housed in pairs in a light-proof room in which the lights were on for 14.5 hours each day. They were tested at the same time on successive days during the period when the lights were on in their holding room. At the start of the experiment they were assigned in equal numbers to the two groups.

Apparatus. Eight identically specified conditioning boxes (24.5 × 23.0 × 20.0 cm) were used. The boxes were housed in individual sound- and light-attenuating chambers. Three of the walls were constructed from aluminum, one wall was a transparent Perspex door, and the ceiling was translucent white Perspex. The floor was constructed from steel rods. There was a 5.0 × 6.0 cm recessed food magazine in the front wall of the box into which 45-mg food pellets (traditional formula, P. J. Noyes, Lancaster, NH) could be delivered. The base of the magazine was .5 cm above the floor. A clear Perspex flap, suspended from a hinge at the top, could be pushed back to gain access to the magazine. Three pairs of photodiode sensors were set into a rectangular frame in such a manner that horizontal beams, 5 mm in front of the closed flap, were located 10, 20, and 30 mm above the floor. Appropriate circuitry permitted a PC programmed in Visual Basic 6.0 to record the interruption of these beams and to control the experimental events using Whisker Multimedia software (Campden Instruments Ltd, Loughborough, U.K.). Two loudspeakers placed on the ceil-

Table 1
The Design of Experiment 1

Group	Stage 1	Stage 2	Test
Blocking	A+	A+ AX+ C-	X
Discrimination	C+	AX+ BX- C+	X

ing of the boxes delivered a 10-Hz clicker and white noise, both at an intensity of 76 dB. A 6-W bulb mounted above the ceiling could be illuminated to provide a constant visual stimulus. Three 6-W bulbs were mounted on the front wall of the chamber. One bulb was located directly above the center of the food magazine, at a height of 17 cm above the floor. The other bulbs were 5 cm to the left and right of this bulb and were 11 cm above the floor of the chamber. A flashing light was provided by illuminating the two side bulbs with an on-off frequency of 1 Hz. The overhead light served as Stimulus A for all rats. The flashing light served as Stimulus B in the discrimination group. For half the subjects in each group the clicker served as X and white noise as C, whereas these relationships were reversed for the remaining subjects.

Procedure. All rats initially received two 60-min sessions of magazine training. Within each session, a single food pellet was delivered into the magazine regularly once a minute for the first 30 min; the rat then remained in the conditioning chamber for a further 30 min.

In each of the eight sessions of Stage I, rats in the blocking group received 24 reinforced 10-s presentations of A, whereas rats in the discrimination group received 24 reinforced 10-s presentations of C. Reinforcement constituted the delivery of one pellet into the magazine immediately after the CS was turned off. In this and all other stages of the experiment the intertrial interval (inter-trial interval (ITI)) ranged from 80 s to 160 s, with a mean of 120 s. The first session of Stage II began on the day after the completion of Stage I. There were 24 trials in each of the 17 sessions of Stage II. There were eight presentations of each of the three different types of trial in every session. The blocking group received reinforced trials with A and AX and nonreinforced trials with C. The discrimination group received reinforced trials with AX and C, and nonreinforced trials with BX. For all groups, the sequence of reinforced and nonreinforced trials was random, with the constraint that no more than three trials of the same type could occur in succession.

Session 18 served as a test of the redundant cue, X. The first 12 trials of this session were conducted in the same manner as for the previous sessions. The next four trials consisted of 10-s presentations of X with the same ITI as for the previous trials.

Throughout the experiment, the amount of time spent in the magazine was recorded for every trial and for the 10 s that preceded every trial. The difference between these scores was then taken as measure of performance for each trial.

Results and Discussion

Both groups showed evidence of conditioning in the first stage of the experiment. During the final two sessions of this stage, the mean difference score for the trials with the light in the blocking group was 2.9 s, whereas for trials with the auditory stimulus in the discrimination group it was 5.5 sec. This difference between the groups, which was significant, $t(30) = 4.66$, is in keeping with the frequent finding from our laboratory that appetitive conditioning with an overhead light, as measured by magazine activity, is less effective than with auditory stimuli.

The mean difference scores during the three types of trial for Stage II of the experiment are shown in the left-hand panel of Figure 1 for the blocking group and in the right-hand panel for the discrimination group. The initial training with A in the blocking group resulted in consistently strong responding during A and AX from the outset of this stage, whereas responding during the nonreinforced trials with C was weak. As for the discrimination group, it soon acquired the AX+ BX- discrimination and responded at a consistently high rate during C. To compare the asymptotic performance of the two groups, individual mean difference scores were calculated for the final two sessions for the three types of trial. A two-way ANOVA with the factors of group and trial type (reinforced element, reinforced compound, and non-reinforced trial) revealed no significant effect of group, $F < 1$, but a significant effect of trial type, $F(2, 60) = 111.78$, $MSE = 1.57$,

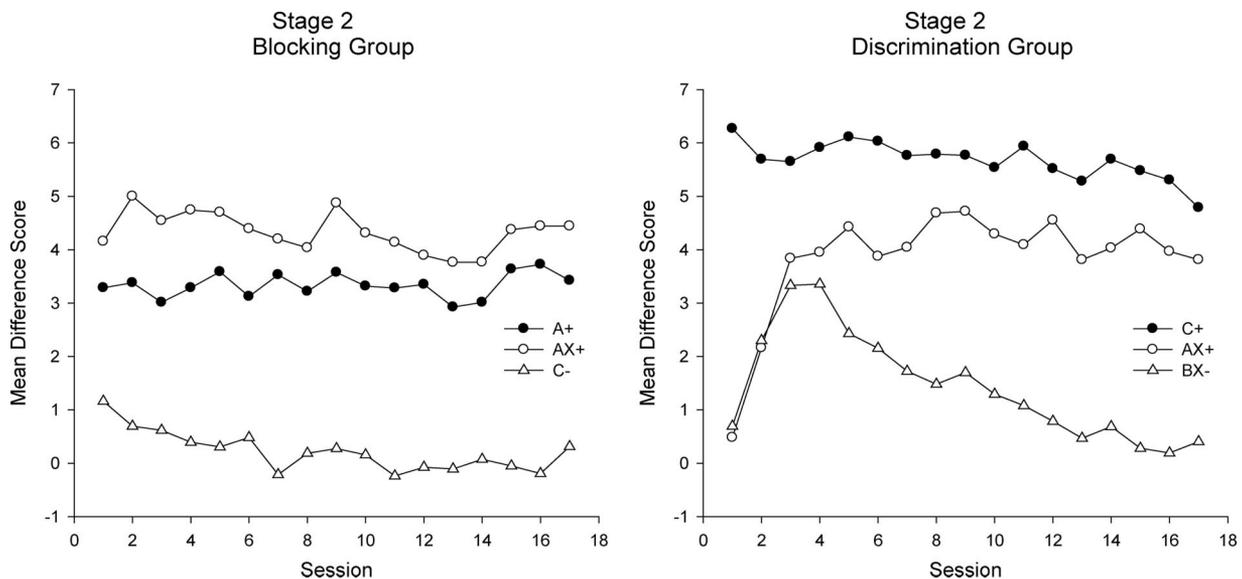


Figure 1. Group mean difference scores for Experiment 1 for the three trial types in Stage 2 with the blocking group (left-hand panel) and the discrimination group (right-hand panel).

and a significant Group \times Trial Type interaction, $F(2, 60) = 5.29$, $MSE = 1.57$. Further analysis of this interaction, using tests of simple main effects, revealed that the effect of trial type was significant in both groups, $F_s(2, 60) > 54.79$, $MSE = 1.57$. In addition, the mean difference scores were significantly greater during C+ trials in the discrimination group than during A+ in the blocking group, $F(1, 90) = 6.98$, $MSE = 2.48$. There was no effect of group for either of the other trial types, $F_s < 1$. Bonferroni-corrected pairwise comparisons revealed that the mean difference scores for the following comparisons in the blocking group were significant: A and AX, C and AX, and C and A, $t_s(15) > 3.60$. A similar analysis for the discrimination group revealed that the difference scores were significantly smaller during BX than during either C or AX, $t_s(15) > 6.77$, but the analysis of the difference scores during C and AX did not reach the criterion for significance, $t(15) = 2.55$, $p = .02$ (corrected p value for 6 comparisons = 0.008).

The mean difference scores for the four test trials combined with X were .31 s for discrimination group and 1.79 s for the blocking group. An analysis of individual mean difference scores revealed that this difference between the groups was significant, $t(30) = 2.47$. The mean duration of magazine activity during the 10 s before these trials was 0.68 s for the blocking group, and 1.07 s for the discrimination group. This difference between the groups was not significant, $t(30) = 1.13$.

The experiment has revealed the novel finding that the CR to a redundant cue is stronger after blocking than after a simple discrimination. This finding is opposite to that predicted by the Rescorla-Wagner (1972) theory. We shall see in the General Discussion that it is also inconsistent with predictions from at least one other theory of associative learning. Throughout both stages of the experiment, responding was stronger during the C+ trials in the discrimination group than the A+ trials in the blocking group. Could this difference have implications for the interpretation of the outcome of the test trials with X? One might argue, for example, that there was a general tendency for the discrimination group to respond more vigorously during any CS than the blocking group. If this were true, it would then follow that responding during the test trials with X would have been stronger in the discrimination than the blocking group, whereas the opposite outcome was found. It is also worthy of note that by the end of the second stage of the experiment there was no difference between the performance of the two groups during the trials with AX. Indeed, an inspection of Figure 1 reveals that responding to AX in these two groups was similar for the final 10 sessions. This finding demonstrates that the groups performed similarly when trained with the same stimuli and that the relatively low rate of responding during A in the blocking group was a consequence of the properties of the stimulus selected to perform this role—the overhead light. In view of this conclusion, it would seem that the poor performance during the training trials with A in the blocking group does not point to a theoretically uninteresting explanation for our results.

Rescorla and Durlach (1981) proposed that during Pavlovian conditioning with a compound CS, associations may develop between the elements of the compound. As a consequence, when one element is presented alone it may activate a representation of its partner which in turn may activate the representation of the US with which it has been paired and ultimately influence the strength of the CR that is performed. These proposals have important

implications for the present experiment. For example, such within-compound associations might have formed between A and X during the AX+ training trials in both groups. During the test stage, therefore, the presence of X might have retrieved a memory of A, and the strength of the CR during X would then be influenced, in part, by the associative strength of A. Because it is likely that the associative strength of A was greater in the blocking than the discrimination group, it would then follow that the response elicited by X would be stronger in the blocking than the discrimination group. The next experiment was designed to evaluate this explanation for the results of the present study.

Experiment 2

A summary of the design of the experiment can be seen in Table 2. During Stage I, a single group of rats received a blocking treatment of the form A+ AX+, and a simple discrimination of the form BY+ CY-. To test the reliability of the results from Experiment 1 a test was given with X and Y toward the end of Stage I. For the next stage of the experiment an A- C+ discrimination was given, followed by another test with X and Y. On the basis of the results from the previous experiment, it was expected that responding during the first test would be stronger during X than Y. If this outcome is solely attributable to the influence of within-compound associations then it is possible that, as a result of the second stage of training, responding during the second test will be weaker during X than Y. This follows because the presentation of X for the second test should excite a representation of A, which will have lost much of its excitatory influence during Stage II and thus add little to responding during the blocked cue. The test trials with Y, on the other hand, can be expected to activate a representation of B, which should have retained the excitatory influence it gained in Stage I, and of C, whose originally inhibitory influence should have been replaced by an excitatory influence as a result of Stage II. Thus any contribution of within-compound associations would be expected to enhance responding to Y to a greater extent than to X in the second test and produce the opposite outcome to that seen in the first test. From Table 2 it is evident that the experiment concluded with extinction trials with A and C. The purpose of these trials was to assess the effectiveness of the Stage-II training. In view of the two separate tests with X and Y, only a few trials were given in the first test so as not to jeopardize the outcome of the second test. Finally, the experiment took place in different test chambers, and with different visual stimuli, to the previous experiment. One reason for making this change was to avoid the problem posed in Experiment 1 by using a visual cue, the overhead light, which was less effective for conditioning than the other cues.

Table 2
The Design of Experiments 2 and 3

Stage 1	Test 1	Stage 2	Test 2	Extinction
A+ AX+ BY+ CY-	X & Y	A- C+	X & Y	A- C-

Note. The extinction stage was used only in Experiment 2.

Method

Subjects. The subjects were 30 experimentally naïve rats from the same stock and maintained in the same manner as for Experiment 1.

Apparatus. Eight identical conditioning chambers were used. The side walls and ceiling of each chamber were constructed from clear Perspex. Each wall had a height of 28 cm and a width of 30 cm. There was a grid floor positioned 5 cm above the base of the chamber. In the center of the back wall, there was a circular hole, diameter 3 cm, the center of which was 3 cm above the grid floor. The circular hole allowed access to a well into which sucrose solution (8% sugar, 92% water) was delivered. This area is henceforth referred to as the magazine. A peristaltic pump was located beneath each conditioning chamber, which delivered the sucrose solution via a plastic tube into the well. A PC with Whisker software, and programmed in Visual Basic 6.0, controlled the experimental events and recorded the duration of entries into the magazine from infrared sensors that were set into each chamber. Barriers were placed between the chambers to prevent the animals seeing each other. Auditory stimuli were delivered simultaneously to all chambers from a 5-ohm speaker located on the ceiling of each chamber. Visual stimuli appeared on two flat-screen thin film transistor (TFT) monitors that had a width of 33 cm and a height of 27 cm. These were placed side by side at an angle of 90°. The point where the screens met was 25 cm away from the center of the magazine entry hole in the back wall. The bottom edge of each screen was level with the floor of the chamber.

The auditory stimuli were a 10-Hz, 76-dB clicker and a 2-kHz, 70-dB tone. The visual stimuli were a striped pattern or a spotted pattern, which filled both of the computer screens, or the illumination of both screens to white. The spotted pattern consisted of a 13 × 11 matrix of black circles (diameter 1.5 cm) on a white background. The striped pattern consisted of 11 alternating black and white vertical stripes, each 3 cm wide, with a height of 27 cm.

The stimuli were counterbalanced such that for 16 subjects A, B, and C were the white, striped, and spotted patterns, respectively. For eight of these animals X was the tone and Y was the click, and vice versa for the other eight. For the remaining 14 rats A, B, and C were the spotted, striped, and white patterns, respectively. Again, for half of these animals X was the tone and Y was the click, and vice versa for the other half.

Procedure. All rats initially received two sessions of magazine training. During each of these 1-hr sessions, 0.2 ml of sucrose solution was delivered into the well over a period of 3 s, every minute for 30 min. The rat then remained in the conditioning chamber for a further 30 min. After magazine training, the animals were divided into the four counterbalancing groups.

During each of the 20 sessions of Stage I, all subjects received discrimination training with the four trial types shown in Table 2. There were 24 trials in each session, with each trial type presented six times. The following details applied to all stages of the experiment. Each stimulus was presented for 10 s, and the mean inter-trial interval was 120 s (range = 80–160 s). The trial types were presented in a random order with the constraint that no more than two reinforced or two nonreinforced trials could occur in succession. Sucrose solution was delivered on reinforced trials for 3 s immediately after the stimulus had been turned off.

The first test with X and Y was given during the 17th session of Stage-I training. During this session, all subjects received 24 training trials, as in the previous sessions. After the first 12 of these trials, subjects received two nonreinforced trials with each of X and Y presented in a random sequence.

On the day after Session 20, both groups received the first session of Stage II. During each of the five sessions of this stage subjects received 20 reinforced presentations of C and 20 nonreinforced presentations of A. The procedural details that have been omitted were the same as for the previous stage.

The second test with X and Y commenced on the day after Session 5 of Stage II. The session started with 12 trials of Stage-2 training followed by 6 trials with each of X and Y presented in a random sequence. The treatment with X and Y continued for another 12 trials, 6 with each stimulus, on the subsequent session. The final session of the experiment consisted of 12 extinction trials with A randomly intermixed among 12 extinction trials with C.

Results

The group mean difference scores for each trial type on every session of Stage I are shown in the left-hand panel of Figure 2. By the end of this stage the duration of magazine activity was at a similar high level during each of the three reinforced trial types, whereas this activity was hardly affected by the nonreinforced trials. A one-way ANOVA of individual mean durations of responding during the two sessions before the first test with X and Y revealed a significant difference among the four trial types, $F(3, 87) = 27.56$, $MSE = 1.30$. Bonferroni-corrected pairwise comparisons for the difference scores during A and AX, A and BY, and AX and BY revealed no significant differences, $t_s(29) < 1.08$, but the difference scores for each of these three trial types differed significantly from those during the nonreinforced trials with CY, $t_s(29) > 5.69$.

The group mean difference scores for the first set of test trials with X and Y can be seen in the left-hand panel of Figure 3. As predicted, on the basis of the previous results, the response to X was stronger than to Y. Analysis of individual difference scores for the two stimuli revealed a significant difference with a one-tailed t test, $t(29) = 1.79$, $p = .04$.

The center panel of Figure 2 shows the effects of the training in which C was paired with sucrose for the first time, and A was presented in extinction for the first time. Although the initially stronger response to A than C soon disappeared, this treatment was relatively ineffective at reversing the original pattern of responding. In support of this observation an analysis of individual mean difference scores for the final two sessions combined revealed no difference during the two trial types, $t(29) = 1.65$.

The group mean difference scores for the second test with X and Y are shown in the right-hand panel of Figure 3, from which it is evident that the response during X was again stronger than during Y. A similar analysis, but with a two-tailed test, to that for the first test with X and Y revealed a significant difference, $t(29) = 2.09$.

The results from the extinction trials with A and C in the final phase of the experiment can be seen in two-trial blocks in the right-hand panel of Figure 2. Throughout this phase performance in the presence of A and C was very similar. In support of this observation, a two-way ANOVA of individual mean difference scores with the factors of trial type (A or C) and block, revealed

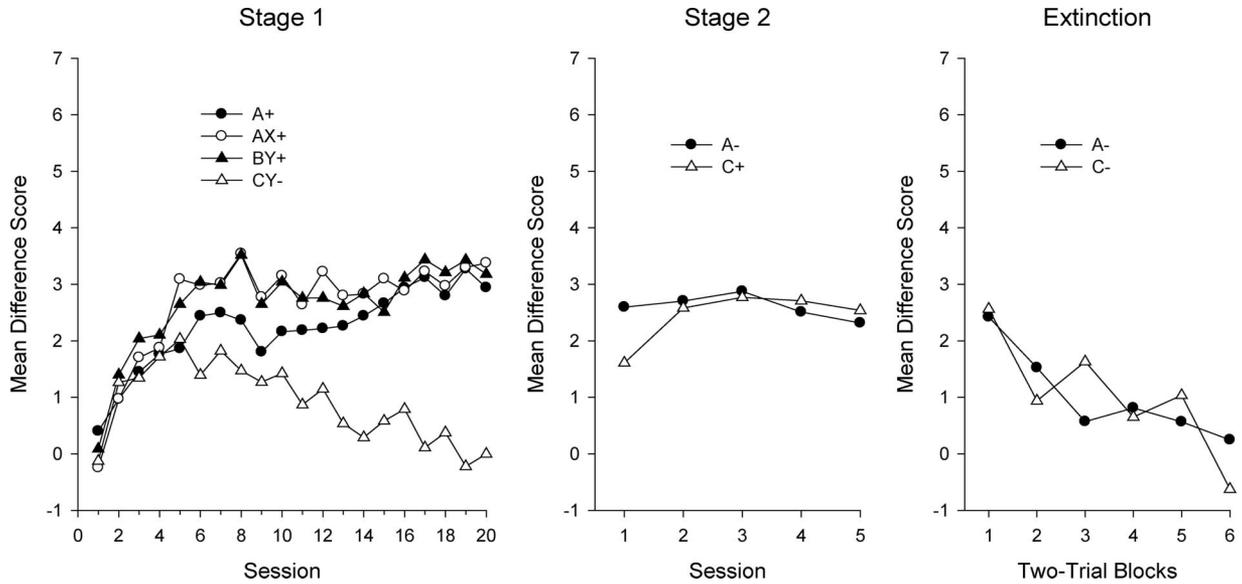


Figure 2. Group mean difference scores for Experiment 2 for the four types of trial in Stage1 (left-hand panel), for the reversal training during Stage 2 (center panel), and for the final extinction stage (right-hand panel).

that there was a significant effect of block, $F(5, 145) = 8.83$, $MSE = 5.19$, but no significant effect of trial type, $F < 1$, and no significant Trial Type \times Block interaction, $F(5, 145) = 1.47$, $MSE = 5.12$.

The experiment has again shown that the redundant cue from a blocking paradigm, A+ AX+, gains more associative strength than its counterpart, Y, from a simple discrimination, BY+ CY-. In essence, the present results constitute a within-group demonstration of the effect revealed by Experiment 1 using a between-groups methodology, but they go beyond the earlier findings in an

important way. They show that the difference between responding to X and Y can be found even when there is no difference between the associative strengths of A and C. Such a finding makes it hard to attribute the difference in responding to X and Y to the influence of any between-cue associations that might have been formed during the initial stage of training. Indeed, given that Y was also paired with the excitatory B, any analysis of the results in terms of between-cue associations would lead to the prediction that during the final test with the two stimuli, responding during Y will be stronger than during X.

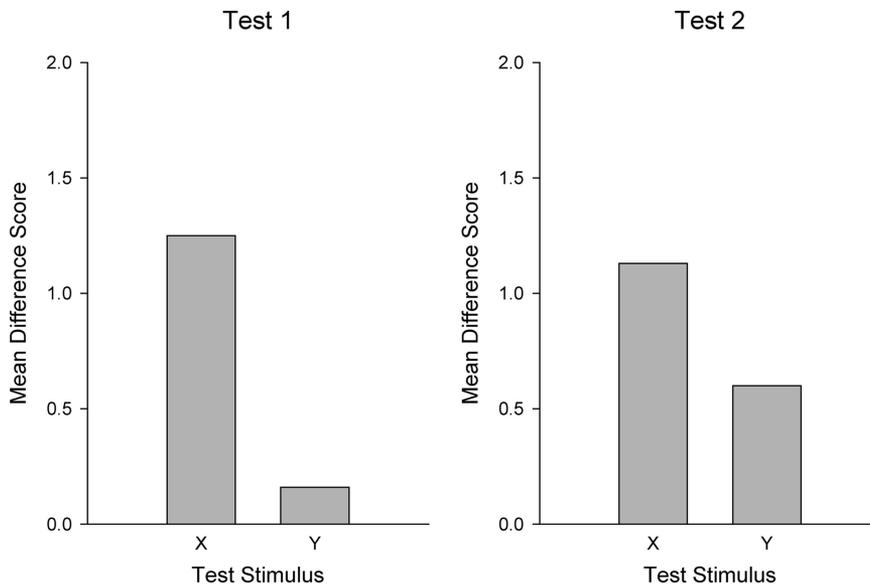


Figure 3. Group mean difference scores for Experiment 2 for the first (left-hand panel) and the second (right-hand panel) set of test trials with the redundant cue from blocking, X, and the simple discrimination, Y.

The A- C+ training in Stage II was intended to result in the associative strength of C being greater than of A. There was a hint of this being the case by the end of this stage, but it was not statistically significant. Even so, the failure to reverse the significance of A and C does not invalidate the conclusions we wish to draw from this study. Providing there is no evidence that the associative strength of A was greater than of C at the time of the final test with X and Y, then it is unlikely that the superior performance to X than Y in the final test was attributable to the influence of between-cue associations. The results from the extinction trials with A and C point forcefully to the conclusion that their associative strengths were similar when the second test with X and Y took place.

Experiment 3

The results from the two experiments above are novel and, as argued in the Introduction, they may pose a problem to current theories of associative learning. Before taking the theoretical significance of our findings too seriously, however, it is worth ensuring that they are not an artifact of a particular method of testing with a particular species. If this were the case, then the theoretical implications of the results from Experiments 1 and 2 would be slight. In view of this concern, the final experiment was based on the design of Experiment 2, but it was conducted using autoshaping with pigeons. Experiments with pigeons have already revealed effective blocking (e.g., Schreurs & Westbrook, 1982). They have also shown that the redundant cue of an AY+ BY- discrimination gains rather little associative strength (e.g., Dopson et al., 2010). To our knowledge, however, the effects of these different methods of training on the associative properties of the redundant stimuli in pigeons have never been compared in a single study. If it could be shown with pigeons that the blocking treatment again results in the redundant cue gaining more associative strength than the simple discrimination, then such an outcome would confirm both the reliability and generality of the findings revealed in Experiments 1 and 2.

Method

Subjects. The subjects were 24 adult homing pigeons who had previously completed an autoshaping experiment in which they were required to discriminate between black and white patterns. The pigeons were housed in pairs and had free access to water and grit in their home cages. They were maintained at 80% of their free-feeding weights by being fed a restricted amount of food after each experimental session. They were maintained in a lightproof room in which the lights were on for 14.5 hr each day.

Apparatus. The experimental apparatus consisted of eight pigeon test chambers (30.0 × 33.0 × 35.0 cm). Each contained an 8.3-cm high × 6.3-cm wide clear acrylic panel, which was hinged at the top. Pecks on the panel were detected by a reed relay that was operated whenever a magnet attached to its lower edge was displaced by a distance greater than 1 mm. The midpoint of the panel was 24 cm above the floor of the chamber. A color TFT TV with a 15.5 × 8.7-cm screen was located 2.0 cm behind the acrylic panel. Food was delivered by operating a grain feeder (Colbourn Instruments, Lehigh Valley, PA) with an opening that measured 4.6 cm × 5.4 cm located in the same wall as the response key. The

midpoint of the opening was 9.0 cm above the chamber floor and 7.0 cm to the left of the midline of the wall. The feeder was illuminated whenever grain was made available. The chambers were permanently lit during all experimental sessions by a 2.8-W bulb, operated at 24 V, located in the center of the ceiling of the chamber. A PC (Research Machines, Abingdon, U.K.) running Windows XP was used for the control of events, recording of responses, and generation of stimuli on the TV screens. The computer was programmed in Visual Basic 6.0, and the interface with the experimental apparatus was controlled by Whisker software (Campden Instruments Ltd, Loughborough, U.K.).

Stimuli. The stimuli were presented on the TV screen, which was otherwise black. Compound stimuli consisted of a rectangle, with height 2.7 cm and width 4.2 cm, which was divided along the diagonal from the bottom left corner to the top right to form two triangles. The top triangle represented either A, B, or C, and the bottom triangle represented X or Y. The center of the rectangle was located at the center of the TV screen. For trials on which A appeared alone, only the top triangle was presented on the screen, located in the same position as this triangle on compound trials. Each triangle was filled with one of five plain colors: yellow, red, green, cyan, and blue. The stimuli were counterbalanced in the following way. For all 24 birds, C was yellow. For 12 birds, A and B were green and cyan, respectively. For half of these birds X was red and Y was blue, and vice versa for the other half. For the remaining 12 birds, A and B were cyan and green, respectively. Again, for half of these birds X was red and Y was blue, and vice versa for the other half.

Procedure. All pigeons initially received two sessions of autoshaping, in which a white circle with a diameter of 2.4 cm was presented in the middle of the TV screen for 10 s. There were 40 trials in a session. The mean ITI was 60 s (range = 40–80 s), and food was made available in the hopper for 4 s whenever the white circle was removed from the TV screen. During each of the 10 sessions of Stage I, subjects received the four trial types shown in Table 2. There were 48 trials in each session, with each trial type being presented 12 times. The following details applied to all stages of the experiment. Each stimulus was presented for 10 s, and the mean ITI was 60 s (range = 40–80 s). The trial types were presented in a random order with the constraint that no more than four reinforced trials could occur in succession. Food was presented for 4 s as soon as the pattern signaling its delivery was turned off.

The first test with X and Y took place in Session 8 of Stage I. During this session, all subjects again received 48 training trials. Intermixed with these trials were four nonreinforced test trials with each of X and Y. The test trials were presented in a random sequence, and spaced throughout the session, with an average of six training trials appearing in between them. On the day after Session 10 of training, both groups received the first session of Stage II. During each of the two sessions of this stage subjects received 24 reinforced presentations of C and 24 nonreinforced presentations of A. The trial types were presented in a random order with the constraint that no more than two trials of the same kind could occur in succession. The procedural details that have been omitted were the same as for the previous stage.

On the day after the completion of Stage II, subjects received a second test with X and Y. The first 24 trials of this final session consisted of Stage-2 training, followed by 24 nonreinforced trials

in which X and Y were presented an equal number of times in a random sequence (with the constraint that no more than two trials with either X or Y could occur in succession).

Results

The rate of pecking at the Perspex panel in front of the TV screen was recorded on every trial throughout the experiment. The left-hand panel of Figure 4 shows the group mean rates of pecking that were recorded during the four trial types for the eight sessions of Stage I. The results were entirely as expected. By the end of this stage responding during the three reinforced trials occurred a high rate, whereas it occurred at a low rate during the nonreinforced presentations of CY. A one-way ANOVA of individual mean responses per min for the two sessions before the first test with X and Y revealed a significant effect of trial type, $F(3, 69) = 51.93$, $MSE = 1342.69$. Bonferroni-corrected pairwise comparisons revealed no significant differences between trial Types A, AX and BY, $t_s(29) < 1.34$, but responding during each of these trial types differed from that during CY, $t_s(29) > 6.97$.

During the first test with X and Y, responding was found to be faster during X than Y. The results from this test can be seen in the left-hand panel of Figure 5, which shows the group mean rates of responding during the four test trials with each of the two stimuli. Analysis of individual mean rates of responding revealed a significant difference between the two stimuli, $t(23) = 3.74$.

The right-hand panel of Figure 4 depicts the group mean response rates, in half-session blocks, for discrimination training with A- and C+ in Stage II. It is evident that the training was successful in reversing the significance of A and C, with the result that by the end of this stage responding was more rapid during C than A. Analysis of the results for the second session revealed that this difference was significant, $t(23) = 3.66$.

The right-hand panel of Figure 5 depicts the mean rates of responding during the second test with X and Y. Once again, responding was more rapid during X than Y. Analysis of individual mean rates of responding for the 12 trials combined with each stimulus confirmed that this difference was significant, $t(23) = 2.33$.

In contrast to the previous experiment, the reversal training in Stage II was completely successful and resulted in stronger responding during C, which had previously signaled the absence of food, than A which had previously been paired with food. Despite this reversal in the significance of A and C, the final test with X and Y revealed the same qualitative pattern that was observed during the first test. Responding to X was again stronger than to Y. In view of these results, it is hard to see how the stronger responding to the redundant cue after the blocking treatment than the simple discrimination can be explained by appeal to the influence of within-compound associations. In the discussion that follows we consider alternative explanations for this outcome.

General Discussion

Three experiments have shown that a blocking treatment, A+ AX+, results in the redundant cue, X, eliciting a stronger CR than when this stimulus takes part in a simple discrimination, AX+ BX-. Both the blocking treatment and the simple discrimination involve pairing the redundant cue with other stimuli, which raises the possibility that between-cue associations were responsible for the results that have been described. For reasons that have already been presented, the results from Experiments 2 and 3 indicate quite forcefully that such associations were not responsible for the outcome of our experiments. Instead, a more likely explanation for the results is that the associative strength of the redundant cue, that

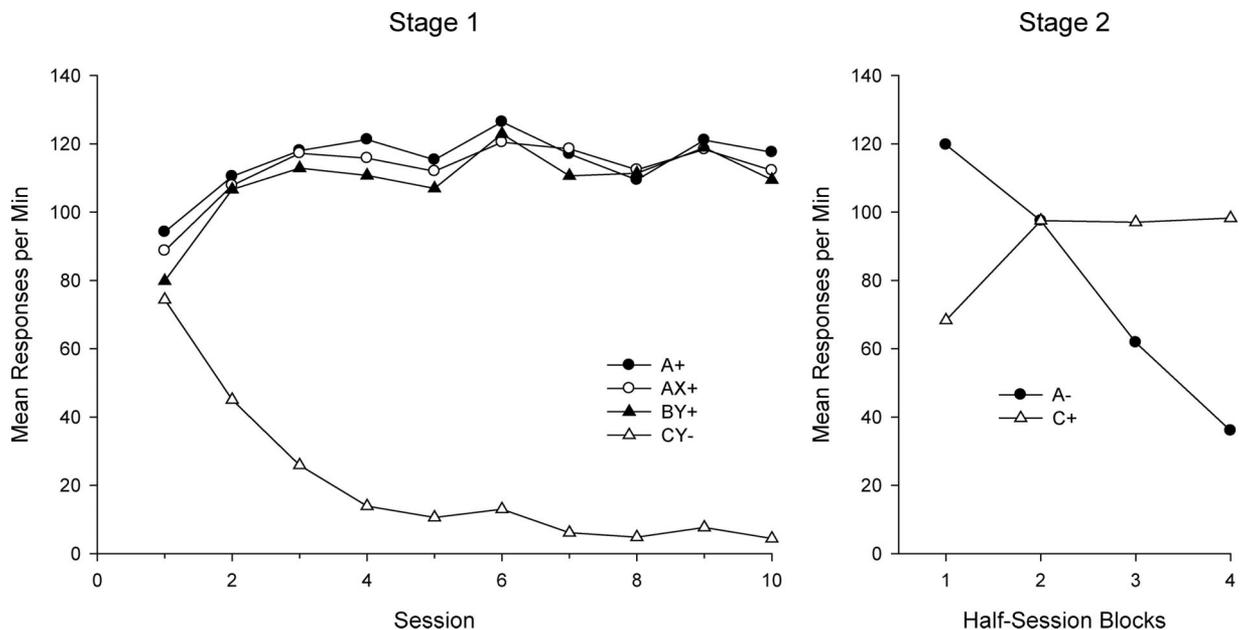


Figure 4. Group mean difference scores for Experiment 3 for the four types of trial in Stage 1 (left-hand panel) and during the reversal training during Stage 2 (right-hand panel).

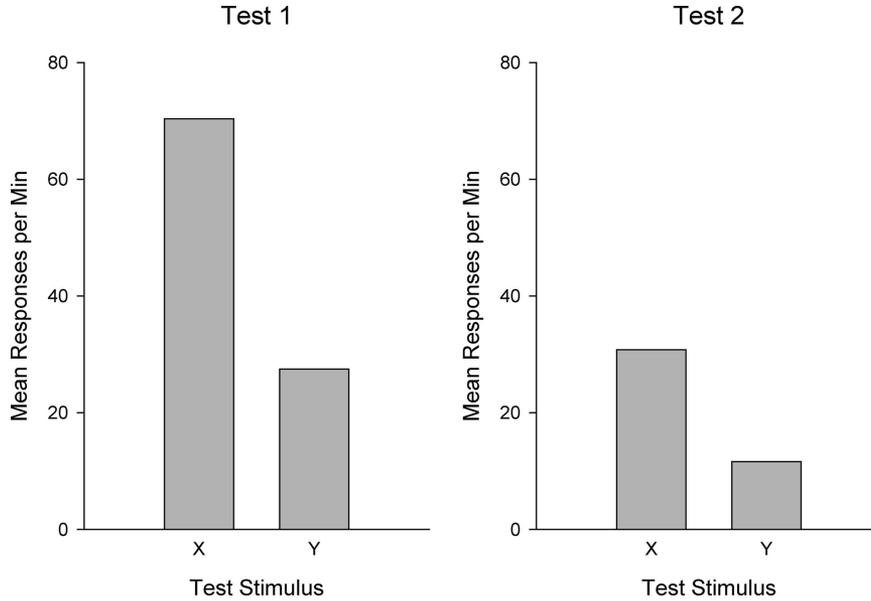


Figure 5. Group mean difference scores for Experiment 3 for the first (left-hand panel) and the second (right-hand panel) set of test trials with the redundant cue from blocking, X, and the simple discrimination, Y.

is, the strength of its association with food, was greater after blocking than the simple discrimination.

The foregoing conclusion does not fit comfortably with predictions that can be derived from what might be regarded as the standard account of associative learning, the Rescorla-Wagner (1972) theory. As noted in the Introduction, this theory predicts that A+ AX+ training will result in A acquiring asymptotic associative strength, and X acquiring none. On the other hand, an AX+ BX- discrimination will result in A gaining the lion's share of the available associative strength, but some strength will be left for X to acquire. Thus the theory predicts the opposite outcome to that found in each of the experiments. This prediction is based on the assumption that conditioning had reached asymptote by the time of the test trials with the redun-

dant cue. To determine the predictions made by the Rescorla-Wagner theory before asymptote has been reached a series of computer simulations was conducted, based on Equation 1 and the design of Stage I for Experiments 2 and 3, A+ AX+ BY+ CY-. The main purpose of the simulations was to explore the effect of assigning different values to the learning rate parameter β for trials in which an increment, β_E , and a decrement, β_I , in associative strength is predicted. The value of the learning rate parameter α was set at .2 and the value of λ was set at 100 and 0, respectively, for reinforced and nonreinforced trials. Figure 6 portrays representative graphs for simulations when β_E was greater than β_I (left-hand panel), when β was the same on all trials (center panel), and when β_E was less than β_I (right-hand panel). For the first two of these conditions the associative

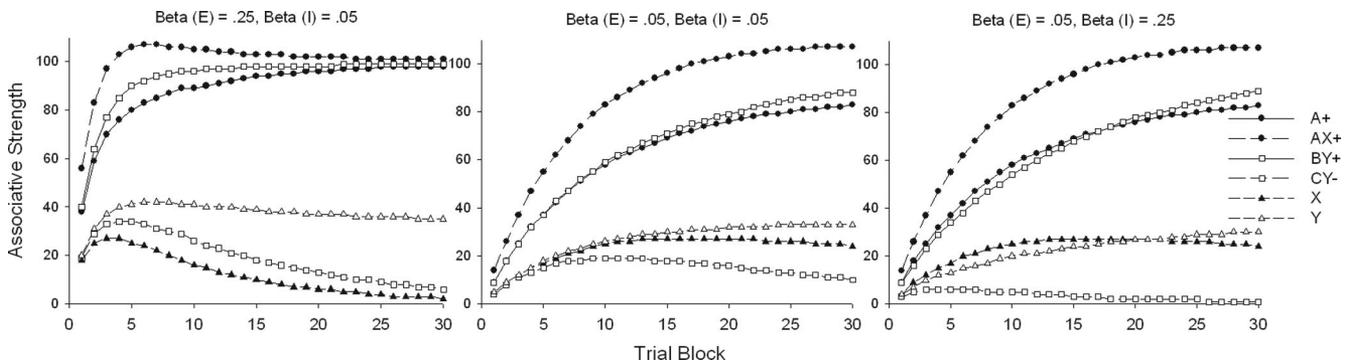


Figure 6. The associative strength predicted by the Rescorla-Wagner (1972) rule for every fifth trial of the training used in Experiments 1 and 2. The graphs depict the predicted course of acquisition of associative strength during a blocking treatment, A+ AX+, and a simple discrimination, BY+ CY-, with various values of β for trials in which there is a growth (E) or a decrease (I) in associative strength. The graphs also depict the predicted acquisition of associative strength to the redundant cues X and Y.

strength of the redundant cue, X, from blocking is predicted to be either the same as, or less than that of the redundant cue, Y, from the simple discrimination. That is, the predictions are inconsistent with our findings. A more complex pattern of results is predicted for the third condition. At first, the associative strength of Y is predicted to be greater than of X, which is what we found, but this relationship reverses with extended training.

The results from Experiments 2 and 3 can thus be explained by the Rescorla-Wagner (1972) theory if it is assumed, first, that the current values of associative strength were some distance from asymptote when testing took place and, second, that the value of β is greater for nonreinforced than reinforced trials. Unfortunately, it is difficult to know how close the effects of conditioning were to asymptote by the time testing took place in Experiments 2 and 3, which makes it difficult to say whether the first of these assumptions is justified. There are, however, two reasons for questioning whether the second assumption is justified. First, Rescorla and Wagner (1972) show that their theory is able to explain the relative validity effect described by Wagner et al. (1968), which was mentioned in the Introduction, but only if it is assumed that the value of β_E is greater than β_I . They further note that this assumption is necessary if other formal models are to be successful (e.g., Bush & Mosteller, 1955; Lovejoy, 1966). Second, Rescorla (2002) has shown, using appetitive conditioning with rats and pigeons, that the rate of change in associative strength is greater on a reinforced than a nonreinforced trial, which can only be understood in terms of the Rescorla-Wagner theory if β_E is greater than β_I . There are, therefore, both theoretical and empirical grounds for believing the simulation in the left-hand panel of Figure 6 is the most appropriate of the three for deriving predictions from the Rescorla-Wagner theory concerning Experiments 2 and 3. If this conclusion is correct, then the findings from the last two experiments pose a challenge to the theory, even if it is accepted that conditioning was not close to asymptote when testing took place. Of course, in Experiment 1 steps were taken to ensure the associative strength of A had reached asymptote for the test trials. On this basis, the results from the first experiment are hard to reconcile with the Rescorla-Wagner theory, no matter what values are assigned to β_E and β_I .

The blocking treatment in all three experiments left the redundant cue with a measure of associative strength. In the case of Experiment 3, the first test revealed that this measure can be quite considerable. Perhaps, therefore, an explanation for the present results should be sought in a theory that predicts the effects of blocking will not be as profound as predicted by Rescorla and Wagner (1972). One such theory is the configural theory of associative learning proposed by Pearce (1987, 1994). In the case of conditioning with A+ AX+, the theory assumes that representations of A and of AX will each enter into associations with the US. The strength of the CR during AX (or A) is then determined by its own associative strength together with that generalizing to it from A (or AX). A computer simulation based on equations put forward by Pearce (1994), with β_E set at .2 and β_I set at .1, indicated that at asymptote the associative strength of both A and AX, after A+ AX+ training, will be $.66\lambda$, and that half of this strength will generalize from one pattern to the other to ensure the overall response to each of

them is of asymptotic magnitude, λ . If X should then be presented by itself for the first time, generalization from AX will enable it to elicit a CR of magnitude $.33\lambda$. The theory therefore predicts the failure to observe complete blocking in each of the experiments. Additional computer simulations were conducted to determine the predictions made by the theory concerning the fate of the redundant cue from an AX+ BX- discrimination. The ultimate associative strength predicted for AX is 1.07λ , and for BX it is $-.27\lambda$. Generalization between the two compounds will then result in the former eliciting a CR of magnitude λ , and the latter no CR at all. Crucially, a test trial with X is then predicted to result in a CR of $.4\lambda$, which is stronger than for the redundant cue from the blocking treatment. This failure to predict the correct outcome to the present experiment also occurs if the test trials are assumed to take place before the asymptote of conditioning has been reached, provided that β_E is equal to or greater than β_I . However, when β_I is greater than β_E then, like the Rescorla-Wagner theory, responding is initially predicted to be stronger during the redundant cue from blocking than from a simple discrimination. This effect is predicted to be small and to reverse with extended training. Of course, we have already seen there are good reasons for questioning whether β_I was greater than β_E for our experiments.

Both of the theories that have been considered assume that the current associative strength of a stimulus contains no information about how the strength was acquired. Associative strength is said to be path independent. As a consequence, if a stimulus has associative strength of, say, $.5\lambda$, then it is impossible to know whether this value represents the half way stage in the process of continuous reinforcement, or extinction, or whether it was arrived at through a partial reinforcement schedule. The independence of path assumption has not prevented the theories from successfully predicting a wide range of phenomena, but it is conceivable that it is responsible for their failure to explain the present results. The obvious difference between the two methods of training is that with the blocking treatment the redundant cue is always paired with food, whereas with the simple discrimination it is paired with food on half the trials. Perhaps the animal is able to retrieve this information at the time of testing and is therefore more likely to respond vigorously during the redundant cue from the blocking treatment than the simple discrimination. Such an explanation may be in keeping with one's intuitions about why animals performed in the way that they did on the test trials, but it is not without its problems.

The first problem is to explain how an animal knows whether a cue has been paired with food on a continuous or an intermittent schedule. One possibility is that a memory of each trial is stored as an independent episode, associated with the outcome (Babb & Crystal, 2006; Clayton & Dickinson, 1998; Heinemann & Chase, 1990). Such a record would then provide the animal with sufficient information to identify the reinforcement schedule associated with a particular stimulus. This proposal then leads to the next problem, which is to specify how retrieved memories of individual episodes relate to the strength of responding during a test with a particular stimulus. Until a set of rules can be identified that govern how individual memories influence behavior, it will be very difficult to test the suggestion that memories of individual training trials were responsible for the present results.

A further problem is that of relating an episodic model of associative learning to the standard models which have dominated theorizing for more than 40 years. The great strength of these models derives from the breadth of the phenomena they are able to explain in a formal manner. By way of example, they can explain blocking (Kamin, 1969), overshadowing (Kamin, 1969), conditioned inhibition (Rescorla, 1969), and perhaps of particular relevance to the present discussion, a phenomenon referred to as superconditioning, or supernormal conditioning. An example of this last effect is provided by Pearce and Redhead (1995). Rats first received an appetitive A+ AX- discrimination which was followed by trials in which AX was paired with food. Test trials after this treatment revealed that it resulted in A acquiring greater than asymptotic strength. This outcome follows in a straightforward manner from the proposals of Rescorla and Wagner (1972), and of Pearce (1994), but it is not clear how an episodic model of associative learning might explain superconditioning. At the time of testing with A, subjects may retrieve memories of their previous experiences with this stimulus—A+, AX-, and AX+—but there does not seem to be a good reason why the combination of these memories should result in supernormal responding during A.

Another feature of the two theories that have been considered is the assumption that changes in associative strength of individual stimuli are determined by a global error term, which represents the discrepancy between the associative properties of all the stimuli present on a trial and the magnitude of the US. An alternative is to assume learning is governed by an individual error term which is determined by the discrepancy between the associative strength of the stimulus under consideration and the magnitude of the US (e.g., Bush & Mosteller, 1951; Hebb, 1949; Spence, 1940). As the redundant cue in a blocking treatment is always paired with the US, and its counterpart in a simple discrimination is intermittently paired with food, it follows from this assumption that the associative strength of the former will ultimately be greater than of the latter. Thus, the assumption that changes in associative strength are determined by an individual error term is able to explain our results, but by itself it is unable to explain such well-established phenomena as blocking and overshadowing. Denniston, Savastano, and Miller (2001, see also Stout & Miller, 2007), however, have shown that these effects can be explained if learning is governed by an individual error term, by incorporating this term into what they call the extended comparator hypothesis.

In addition to assuming that learning is governed by an individual error term, the hypothesis asserts that during a test trial with a single stimulus, a comparator process will reduce its associative strength by subtracting from it a proportion of the associative strength of any stimulus that accompanied it during conditioning. The first of these assumptions leads to the prediction that the associative strength of the redundant cue in the blocking treatment, X, will be greater than of its counterpart in the simple discrimination, Y. Moreover, because during training both X and Y are paired with a stimulus that has repeatedly been paired with food, the comparator process will weaken responding to both stimuli to the same extent. The combination of these influences will then lead to stronger responding during the test with the redundant cue from blocking than from the simple discrimination.

Despite this success in explaining our results, there are number of problems with the extended comparator hypothesis that are

pertinent to the present discussion. One problem is posed by the superconditioning study of Pearce and Redhead (1995) described above. The initial A+ AX- discrimination is predicted to result in A by itself eliciting a strong CR, in part, because it is paired with a stimulus, X, that is associated with the absence of the US. One effect of the subsequent AX+ trials will be to increase the associative strength of X and the extended comparator hypothesis then predicts the strength of the response to A will decline. In fact, the opposite outcome was observed. Another problem is that recent studies have failed to confirm predictions made by the extended comparator hypothesis concerning blocking (Esber, Pearce, & Haselgrove, 2009; Dopson, Pearce & Haselgrove, 2009). At the very least, it would seem wise to keep an open mind about the merits of the explanation offered by the extended comparator hypothesis for how redundant cues acquire their control over behavior.

A long-standing theoretical claim is that animals pay more attention to cues that are relevant rather than redundant predictors of the outcome of a trial (e.g., Sutherland & Mackintosh, 1971; Mackintosh, 1975a). To take this claim a step further, it is possible that redundant cues in some tasks, such as simple discriminations, receive even less attention than those in other tasks, such as blocking. Although this proposal might form the basis of an explanation for our results, it is hard to justify on the basis of available evidence. Experiments have shown that blocking (e.g., Mackintosh, 1975b) and a simple discrimination (e.g., Dopson, Esber, & Pearce, 2010) can result in rather little attention being paid to the redundant cue but, to our knowledge, there is no evidence that the former treatment results in the redundant cue receiving more attention than the latter treatment. It is also hard to justify the foregoing proposal theoretically. The theory of attention that may be best suited to explaining the present results is the one proposed by Mackintosh (1975a), because it assumes that changes in associative strength are determined by an individual error term. On this basis the theory then predicts that X from a blocking treatment, A+ AX+, will gain more associative strength than X from a simple discrimination, AX+ BX-, as the former is continuously reinforced and the latter is partially reinforced. However, the effectiveness of conditioning will also be determined by the amount of attention paid to these stimuli. Mackintosh (1975a) argued that the rate of decline in attention to a redundant cue is determined by how poorly it predicts the outcome of a trial, relative to the stimulus that accompanies it. Given the independent pairing of the blocking cue with the US in each of the three experiments, it then seems likely that attention to X during A+ AX+ trials will decline more rapidly than to Y during AY+ BY- trials. Unfortunately, it is not clear from the theory how these differences in attention paid to X and Y will interact with their different histories of reinforcement in order to determine the strength of responding in their presence.

By relying on individual error terms, as LePelley (2004) points out, the theory of Mackintosh (1975a) is unable to provide a satisfactory account for the development of conditioned inhibition. In view of this shortcoming, LePelley developed a hybrid model of attention which is based on the principles of Mackintosh (1975a) but which incorporates a global error term and, for good measure, the principles of attention advocated by Pearce and Hall (1980). For a related hybrid model see Pearce and Mackintosh (2010).

Although these hybrid theories can account for a wide range of attentional effects, because they incorporate a global error term they again predict that the redundant cue from blocking will eventually end up with less associative strength than its counterpart from a simple discrimination. A similar problem arises for other theories of attention that incorporate a global error term (e.g., Esber & Haselgrove, 2011; Pearce & Hall, 1980). Changes in attention may well be responsible for the results reported here but, for the present, there does not seem to be a suitable theory to explain how these changes take place.

Findings such as blocking, relative validity, and superconditioning provide compelling support for currently popular associative models of learning (e.g., Rescorla & Wagner, 1972; Pearce, 1994) and it would be reckless to abandon these theories, together with their assumption that learning is governed by a global error term, on the basis of the present results. On the other hand, the inability of these theories to explain our results contrasts with the ease with which they can be understood by assuming either that animals remember individual training episodes or that learning is governed by an individual error term. We have just seen, however, that neither of these assumptions can explain many of the phenomena readily accounted for by the assumption that learning is governed by a global error term. Perhaps, therefore, animals have more than one mechanism of learning, and that a satisfactory account of association formation will need to make reference to two, or more, of the processes considered above.

References

- Babb, S. J., & Crystal, J. D. (2006). Discrimination of what, when and where is not based on time of day. *Learning & Behavior*, *34*, 124–130. doi:10.3758/BF03193188
- Bush, R. R., & Mosteller, F. (1951). A model for stimulus generalization and discrimination. *Psychological Review*, *58*, 413–423. doi:10.1037/h0054576
- Bush, R. R., & Mosteller, F. (1955). *Stochastic models for learning*. New York, NY: Wiley.
- Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*, 272–274. doi:10.1038/26216
- 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* (pp. 65–117). Hillsdale, NJ: Erlbaum.
- Dopson, J. C., Esber, G. R., & Pearce, J. M. (2010). Differences in the associability of relevant and irrelevant stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, *36*, 258–267. doi:10.1037/a0016588
- Dopson, J. C., Pearce, J. M., & Haselgrove, M. (2009). Failure of retrospective revaluation to influence blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 473–484. doi:10.1037/a0014907
- Esber, G. R., & Haselgrove, M. (2011). Reconciling the influence of predictiveness and uncertainty on stimulus salience: A model of attention in associative learning. *Proceedings of the Royal Society, B*, *278*, 2553–2561.
- Esber, G. R., Pearce, J. M., & Haselgrove, M. (2009). Enhancement of responding to A after A+/AX+ training: Challenges for a comparator theory of learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 485–497. doi:10.1037/a0014908
- Good, M., & Macphail, E. M. (1994). Hippocampal lesions in pigeons (*Columba livia*) disrupt reinforced preexposure but not overshadowing or blocking. *The Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, *47*, 263–291.
- Hebb, D. O. (1949). *The organization of behavior*. New York, NY: Wiley.
- Heinemann, E., & Chase, S. (1990). A quantitative model for pattern recognition. In M. L. Commons & R. J. Herrnstein (Eds.), *Computational and clinical approaches to pattern recognition and concept formation* (Vol. 9, pp. 109–226). Hillsdale, NJ: Erlbaum.
- Kamin, L. J. (1969). Selective association and conditioning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental Issues in Associative Learning* (pp. 42–64). Halifax, Canada: Dalhousie University Press.
- LePelley, 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*, 192–243.
- Lovejoy, E. (1966). Analysis of the overlearning reversal effect. *Psychological Review*, *73*, 87–103. doi:10.1037/h0022687
- Mackintosh, N. J. (1975a). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298. doi:10.1037/h0076778
- Mackintosh, N. J. (1975b). Blocking of conditioned suppression: Role of the first compound trial. *Journal of Experimental Psychology: Animal Behavior Processes*, *1*, 335–345. doi:10.1037/0097-7403.1.4.335
- Pearce, J. M., Esber, G. R., George, D. N., & Haselgrove, M. (2008). The nature of discrimination learning in pigeons. *Learning & Behavior*, *36*, 188–199. doi:10.3758/LB.36.3.188
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychological Review*, *87*, 532–552. doi: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. 11–39). Oxford, UK: Oxford University Press.
- Pearce, J. M., & Redhead, E. S. (1995). Supernormal conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *21*, 155–165. doi:10.1037/0097-7403.21.2.155
- Pearce, J. M. (1987). A model of stimulus generalization for Pavlovian conditioning. *Psychological Review*, *94*, 61–73. doi:10.1037/0033-295X.94.1.61
- Pearce, J. M. (1994). Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*, *101*, 587–607. doi:10.1037/0033-295X.101.4.587
- Rescorla, R. A., & Durlach, P. (1981). Within-event learning in Pavlovian conditioning. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 81–111). Hillsdale, NJ: Lawrence Erlbaum.
- 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.
- Rescorla, R. A. (1969). Pavlovian conditioned inhibition. *Psychological Bulletin*, *72*, 77–94. doi:10.1037/h0027760
- Rescorla, R. A. (2002). Comparison of the rates of associative change during acquisition and extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*, 406–415. doi:10.1037/0097-7403.28.4.406
- Schreurs, B. G., & Westbrook, R. F. (1982). The effects of changes in the CS-US interval during compound conditioning upon an otherwise blocked element. *The Quarterly Journal of Experimental Psychology B*, *34*, 19–30.
- Spence, K. W. (1940). Continuous versus non-continuous interpretations of discrimination learning. *Psychological Review*, *47*, 271–288. doi:10.1037/h0054336
- Stout, S. C., & Miller, R. R. (2007). Sometimes-competing retrieval

- (SOCR): A formalization of the comparator hypothesis. *Psychological Review*, 114, 759–783. doi:10.1037/0033-295X.114.3.759
- Sutherland, N. S., & Mackintosh, N. J. (1971). *Mechanisms of animal discrimination learning*. New York, NY: Academic Press.
- Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, 76, 171–180. doi:10.1037/h0025414
- Wagner, A. R. (1969). Stimulus validity and stimulus selection in associative learning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 90–122). Halifax, Canada: Dalhousie University Press.

Received November 7, 2011

Revision received January 13, 2012

Accepted January 25, 2012 ■