


Provided for non-commercial research and educational use only.
Not for reproduction or distribution or commercial use.



This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>

Occasion setting is specific to the CS–US association

Charlotte Bonardi *

School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, UK

Received 18 May 2006; received in revised form 14 August 2006

Available online 10 October 2006

Abstract

In Experiment 1, rats were trained on a discrimination in which one occasion setter, A, signaled that one cue (conditioned stimulus, CS), x , would be followed by one outcome, p (unconditioned stimulus, US), and a second CS, y , by a different outcome, q ($x \rightarrow p$ and $y \rightarrow q$); a second occasion setter, B signalled the reverse CS–outcome relations ($x \rightarrow q$ and $y \rightarrow p$). In a subsequent stage, the animals were divided into two groups, and trained as before, except that both A and B were presented in compound with a novel occasion setter, C. For Group S (same) the CS–outcome relations following A and B were identical to those in the pretraining stage, whereas in Group D (different) they were reversed. In a subsequent test, stimulus C was shown to be a more effective occasion setter in Group D than in Group S. In Experiment 2, rats were trained on a negative occasion-setting discrimination in which CS x signaled outcome p , and y outcome q ; when x and y were signaled by the occasion setter A then no outcome followed ($x \rightarrow p$, $y \rightarrow q$, A: x -, and A: y -). In a subsequent stage, A was now trained as a positive occasion setter, signaling reinforcement of x and y . In Group S, x and y signaled the same outcomes as in the prior training stage (x -, y -, A: $x \rightarrow p$, and A: $y \rightarrow q$), whereas in Group D they signaled the opposite outcomes (x -, y -, A: $x \rightarrow q$, and A: $y \rightarrow p$); more efficient test performance was seen in the latter group. These results suggest that the each occasion setter conveyed information about the specific combination of CS and US paired in its presence (i.e., $x \rightarrow p$ and $y \rightarrow q$, or $x \rightarrow no p$ and $y \rightarrow no q$). These results are consistent with the suggestion that occasion setters operate, at least in part, on a specific CS–US association.

© 2006 Elsevier Inc. All rights reserved.

Keywords: Occasion Setting; Rats; CS-US association

* Fax: +44 1159 515324.

E-mail address: cmb@psychology.nottingham.ac.uk

An occasion setter is a stimulus that signals whether or not some other cue will be reinforced. For example, in a feature-positive occasion-setting discrimination, $F-t+ t-$, the animal will learn to respond to the target stimulus, t , more when it is preceded and thus signaled by the feature, F , than when it is presented alone. Such discriminations can be solved by means of Pavlovian conditioning: the feature is the best available predictor of the US, meaning that it should acquire associative strength, and hence boost responding to the target on reinforced $F-t$ trials through summation. But in some cases discrimination performance is independent of the feature's associative strength. In these cases, the feature is called an *occasion setter*.

Various theories have been proposed to explain how occasion setters control responding. One class of theories appeals to nonassociative processes to explain the phenomenon; for example, Rescorla (1985) suggested that a positive occasion setter lowers the activation threshold of the US representation, making it more easily activated by cues that are associated with it, such as the target CS. In contrast, Holland (1983, 1985) see also Bonardi (1989, 1998); Bouton (1990) proposed that the occasion setter facilitates the effectiveness of the entire CS–US association, perhaps by acting as an “and-gate” that controls the flow of activation from CS to US. Configural accounts, on the other hand, explain occasion setting solely in terms of known associative processes, by maintaining that the cue controlling behaviour is a configure of the feature and target cues which becomes directly associated with reinforcement. Discrimination is, nevertheless, independent of the feature's Pavlovian properties, because there is only limited generalization between the feature and the configural cue that controls responding (e.g., Brandon & Wagner, 1998; Brandon, Vogel, & Wagner, 2000; McLaren & Mackintosh, 2002; Pearce, 1987, 1994; Rescorla & Wagner, 1972; Wilson & Pearce, 1989, 1990). These configural theories will be considered in more detail below.

Much of the early work on occasion setting attempted to discriminate amongst these theories by concentrating on *CS-specificity*—the question of whether an occasion setter that controls responding to one target CS can also control responding to a different, transfer CS. The issue is critical: if the occasion setter acts on the US representation (Rescorla, 1985) then it should transfer to any CS that is associated with that US, whereas if it acts on the association between a specific CS and a specific US (Holland, 1983), then transfer to a different CS should never arise. Given the simplicity of these predictions, it is frustrating that the results actually obtained are far from clearcut (see Swartzentruber, 1995, for a review). Transfer is frequently obtained, especially to other “occasion-set” CSs (targets from other occasion-setting discriminations), but is typically *incomplete*—the effect of the occasion setter on responding to the transfer target is usually less profound than its effect on responding to the original target CS.

These findings do not seem to support either theory. The fact that transfer occurs at all contradicts the idea that occasion setters are CS-specific—*unless* one allows the possibility that there is some generalization between the training and transfer targets. To the extent that the animal fails to discriminate between these two stimuli, then even an occasion setter that is specific to one particular CS should transfer to some extent. Some authors argue that their stimulus parameters do not permit such generalization, but the possibility that it may have occurred is difficult to rule out, especially given that generalization between cues is affected not only by their physical characteristics but also by their training history (Honey & Hall, 1989, 1991). For example, an occasion-set cue generalizes more readily to another occasion-set cue than to a stimulus trained in some other way (Bonardi & Hall, 1993). If animals are especially poor at discriminating occasion-set cues, then this could explain why occasion setters transfer better to other occasion-set targets than to CSs that

have not been trained in this way. In short, the occurrence of transfer does not rule out the possibility that occasion setters are nevertheless CS-specific.

The observation that transfer is often incomplete is also theoretically ambiguous. At face value, incomplete transfer is consistent with theories predicting that occasion setters are CS-specific: because they explain transfer in terms of stimulus generalization, it follows that the occasion setter must be less able to elevate responding to the transfer target, which is only similar to the training target, than to the training target itself. But incomplete transfer could also be due to *generalization decrement*: the compound of occasion setter and transfer target is novel to the animal, and this alone could disrupt responding and so reduce the degree to which transfer is observed. This explanation is difficult to reject, as it is difficult to rule out the possibility that generalization decrement has occurred; however, it should be noted that Bonardi (1996) investigated the issue by controlling for this factor; in the task she employed she found no evidence that the imperfect transfer of occasion setting could be attributed to generalization decrement.

On balance, the evidence from transfer studies may be taken to support the notion that occasion setters are CS-specific. Theories predicting CS-specificity and hence absence of transfer can explain those instances of transfer that *do* occur in terms of stimulus generalization, whereas those theories that do not allow CS-specificity, and hence predict perfect transfer, cannot always explain why the transfer that is obtained is incomplete (Bonardi, 1996; see also Rescorla, 1991a, 1991b). The balance of evidence is therefore consistent with the idea that occasion setters act—at least in part—on the CS–US association.

If the bulk of the evidence supports the idea that occasion setters act on a specific CS–US association, then it is pertinent to note that this account (in common with Rescorla's (1985) proposal) must therefore predict that occasion setters should be *US-specific*—that an occasion setter signaling that a CS is paired with one particular US should not affect responding to a CS paired with a different US. Surprisingly, only three studies have addressed the issue (although see; Holland, 1983), and once again the evidence is equivocal (Bonardi & Ward-Robinson, 2001; Holland, 1989; Morell & Davidson, 2002). Holland (1989) reported that a feature (from a feature-negative discrimination) transferred perfectly to a CS paired with a different US, but only if the transfer US (and the transfer CS) had been occasion-set. Morell and Davidson (2002) partially replicated this finding in both positive and negative occasion-setting discriminations.

If occasion setters were US-specific, then they should never show transfer to other USs. Thus, the lack of transfer to nonoccasion-set USs implies US-specificity—but the finding that perfect transfer is obtained with occasion-set USs does not. This contradictory pattern of results led Holland (1989) to propose that an occasion setter can act on any CS and any US—*provided* they have been occasion-set. He maintained that this is why such good transfer is typically obtained with other occasion-set CSs, and why he only obtained transfer with occasion-set USs. This view implies that the representations of occasion-set CSs and occasion-set USs are in possession of some property that allows them to be influenced by occasion setters; however, according to this view the occasion setter is insensitive to the combinations in which these events are presented (Holland, 1989).

The observation that occasion setters show perfect transfer to other occasion-set USs thus challenges the proposal that occasion setters act on the CS–US association; however, these results may be open to alternative explanation. It is possible that generalization between USs, like that between CSs, is affected by their training history. If so, Holland's (1989) demonstration of selective transfer to other occasion-set USs might indicate that

their prior training had rendered them more similar to the original occasion-set USs, and that the transfer he observed was due to enhanced US–US generalization. This raises the possibility that, with more sensitive tests, transfer to occasion-set USs would be found to be incomplete, thus providing evidence for US-specificity even among occasion-set USs. A study by Bonardi and Ward-Robinson (2001) goes some way towards answering this question.

Bonardi and Ward-Robinson (2001) trained pigeons on a switching discrimination, with two occasion setters, A and B, two target CSs, x and y , and two USs, white and red lentils; A signaled that x would be followed by red lentils and y by white, whereas B signaled the opposite. In a subsequent phase, training continued as before, except that there were now two types of trials, *same* trials and *different* trials. On *same* trials, each occasion setter signaled the same CS–US combinations that it had signaled during training, and a novel stimulus S was presented in compound with the target CS. On *different* trials, the CS–US combinations were reversed, and a second novel stimulus, D , was presented in compound with the target CS.

If the animals had learned that the occasion setters signaled specific CS–US associations, so that (for example) after A, x would be followed by red lentils, then they should learn different amounts about stimuli S and D in the second phase. On *same* trials the association between x and red lentils would be fully activated by A, and the red lentils would be fully predicted, and hence unable to reinforce further learning; conditioning to S would therefore be “blocked”—that is, S would acquire little or no associative strength (e.g., Pearce & Hall, 1980; Rescorla & Wagner, 1972). But on *different* trials, although A signaled that x would be followed by red lentils, white lentils actually occurred; as these would be unpredicted, they would be able to support further learning about stimulus D , so that D should elicit more responding than S (e.g., Dickinson, Hall, & Mackintosh, 1976). But if occasion setters act on any CS or US that has been occasion-set (cf., Holland, 1989), then A is an occasion setter for both x and y , and for both red and white lentils. A should therefore be no less able to activate the $x \rightarrow$ white association than the $x \rightarrow$ red association. Thus, on both S and D trials the reinforcer would be well predicted, and learning about both S and D should be blocked.

The results provided preliminary evidence for the first of these predictions; the authors observed more responding to D than to S , suggesting that the animals had encoded which CS–US combination was signaled by each occasion setter (cf., Jenkins, 1985). To date this is the only report of such an effect, and in this respect one should probably note that the import of the findings were weakened by certain anomalies. Specifically, the experiment was counterbalanced so that for half the animals S was purple and D was a diamond, whereas for the other half of animals the reverse was the case. The statistical analysis of the test results with this counterbalancing factor and test stimulus (S or D) as factors revealed a significant interaction between them, which indicated that although overall animals responded more to D than to S , when S was the diamond animals actually responded non-significantly *more* to S than to D . The authors argued that this effect was due to an intrinsic tendency to respond to the diamond stimulus, which would work against seeing the blocking effect in the subgroup for which the blocked stimulus was the diamond. However, one could argue that it raises the possibility that the effects observed were a function of the choice of test stimulus, rather than a result of the fact that the birds had learned more about D than S . The present experiments thus aimed to reinforce and extend the generality of Bonardi and Ward-Robinson’s (2001) findings, by seeking more solid evidence for

occasion setters' specificity to the CS–US association in rat subjects, in both feature-positive and feature-negative discriminations.

Experiment 1

The first experiment employed a variation of the technique employed by [Bonardi and Ward-Robinson \(2001\)](#): blocking of the occasion setter, rather than blocking of the occasion-set target CS, was used as the tool for examining the specificity of the occasion-setter's action. Animals were trained with two features, A and B, two targets, x and y , and two reinforcers, sucrose pellets and groundnut oil (easily discriminable reinforcers, according to [Morell & Davidson \(2002\)](#); see [Table 1](#)). When feature A was the signal, x was reinforced with sucrose and y with oil, whereas when feature B was the signal, these reinforcement contingencies were reversed, x being reinforced with oil and y with sucrose.

After they had learned this discrimination the subjects were divided into two groups, Group S (same) and Group D (different). Training continued as before, but with two procedural changes. First, a third feature, C, was presented in compound with A, and with B; second, for Group S the reinforcement contingencies remained as during initial training, but for Group D they were reversed, such that when A was the signal, x was now reinforced by oil and y by sucrose, and while B was the signal the opposite was the case. The question of interest was the degree to which added feature C would come to control responding to the target CSs, x and y .

The logic behind this design relies on the observation that occasion setters are subject to blocking ([Bonardi, 1991, 1992](#); cf., [Ross & LoLordo, 1986](#) who first performed such an experiment, although see [LoLordo & Ross, 1990](#))—training a stimulus as an occasion setter for a particular CS–US association is impaired if that stimulus is trained in compound with another that has previously been established as an occasion setter for the same CS–US association. This is operationally analogous to Pavlovian blocking, in which pairings of a CS and a US result in poor conditioning if the CS is trained in compound with a second CS that has already been established as a signal for that US. According to many current theories, blocking is said to depend on the fact that the blocked stimulus is being paired with an outcome that is not surprising, because it is signaled by the pretrained CS (e.g.,

Table 1
Design of Experiment 1

Pretraining		Compound training		Test
<i>Group same</i>				
A: $x \rightarrow$ sucrose	$x \rightarrow$ nothing	AC: $x \rightarrow$ sucrose	$x \rightarrow$ nothing	
A: $y \rightarrow$ oil	$y \rightarrow$ nothing	AC: $y \rightarrow$ oil	$y \rightarrow$ nothing	C: $x \rightarrow$ nothing
B: $x \rightarrow$ oil	A \rightarrow nothing	BC: $x \rightarrow$ oil	A \rightarrow nothing	C: $y \rightarrow$ nothing
B: $y \rightarrow$ sucrose	B \rightarrow nothing	BC: $y \rightarrow$ sucrose	B \rightarrow nothing	
<i>Group different</i>				
A: $x \rightarrow$ sucrose	$x \rightarrow$ nothing	AC: $x \rightarrow$ oil	$x \rightarrow$ nothing	
A: $y \rightarrow$ oil	$y \rightarrow$ nothing	AC: $y \rightarrow$ sucrose	$y \rightarrow$ nothing	C: $x \rightarrow$ nothing
B: $x \rightarrow$ oil	A \rightarrow nothing	BC: $x \rightarrow$ sucrose	A \rightarrow nothing	C: $y \rightarrow$ nothing
B: $y \rightarrow$ sucrose	B \rightarrow nothing	BC: $y \rightarrow$ oil	B \rightarrow nothing	

Note. A and B refer to the illumination of the jewel lights, or to dark (achieved by turning off the dimmed house-light); x and y were a white noise and a clicker; C was a pulsed tone for all animals.

Pearce & Hall, 1980; Rescorla & Wagner, 1972). It follows that any change in the outcome that is present in the compound training stage will restore the USs surprise value, and thus attenuate the blocking effect, an effect that has been reported in a number of studies and is termed unblocking (e.g., Dickinson et al., 1976). By analogy, in the compound stage of the present experiment, if the pretrained feature is fully able to act as an occasion setter for the combination of CS and US that is presented, then blocking should occur, and the added feature should be curtailed in the occasion-setting power that it acquires.

If the occasion setter acts on the CS–US association as a whole (cf., Bonardi, 1989; Holland, 1983), then changing that association by substituting the US should produce unblocking. Feature C should therefore acquire greater occasion-setting power in Group D, for whom the CS–US combinations are reversed during compound training, than in Group S, for whom they are maintained. But if, as argued by Holland (1989), the occasion setter acts independently on both CSs and both USs, then it should be insensitive to the combinations in which they are presented; there should, therefore, be no difference between the two groups in the amount of occasion-setting control acquired by C.

Method

Subjects

The subjects were 16 male hooded Lister rats with a mean ad lib. weight of 570 g (range = 485–620 g). They had previously participated in an experiment on flavor aversion learning but were naïve to the stimuli and procedures employed here. They were deprived to 80% of their ad lib weight before the start of the experiment and were maintained at this level for the rest of the experiment by being fed a restricted amount of food at the end of each session; they were housed in pairs in plastic tub cages with sawdust bedding. The colony rooms were lit from 8 am to 8 pm; the subjects were tested during the light portion of the cycle.

Apparatus

A set of four standard Skinner Boxes (supplied by Campden Instruments Ltd.) were used. Each box had three walls of sheet aluminum, a transparent plastic door as the fourth wall, a grid floor and a white translucent plastic ceiling. One of the walls adjacent to the door contained a recessed food tray covered by a transparent plastic flap, 6 cm high × 5 cm wide, that was hinged to the top of the opening to the food tray. Pushing this flap inward from its vertical resting position allowed subjects to gain access to the food tray. An inward movement of the flap actuated a microswitch, and each closing of the switch was recorded as a single response. The flap automatically returned to its resting position when the subject removed its snout from the food tray. The boxes were normally illuminated by a 2.8-W houselight, operated at 12 V, situated on the front wall directly above the food tray. Forty-five-mg sucrose pellets (Noyes, New Hampshire) could be delivered to the food tray, as could deliveries of groundnut oil, which were delivered from a reservoir outside the chamber with a peristaltic pump. During the experiment the reinforcers used were either the delivery of two sucrose pellets, or of .3 ml of groundnut oil (Sainsbury's, UK), delivered by operating the pump for 2 s. Retractable levers fitted alongside the food tray remained withdrawn throughout the experiment. Each box was housed in a sound- and light-attenuating shell. There were two visual stimuli and three auditory stimuli. One visual stimulus was provided by illuminating two, 2.8-W jewel lights, both of which were situated on the front

wall, one to the right of the food tray and one to the left. The second visual stimulus, dark, consisted of turning off the dimmed houselight. The three auditory stimuli, an 80-dB white noise, a 10-Hz 81-dB clicker, and an 80-dB, 1-kHz tone pulsed at 2 Hz, were provided by Campden instruments noise and tone generators, respectively, and delivered through a speaker mounted on the wall of the chamber. The boxes were controlled by a BBC micro-computer programmed in a version of BASIC.

Procedure

Initially, all animals were trained to retrieve sucrose pellets and oil from the food tray. Subjects first received six sessions of magazine training in which reinforcers were delivered on a VT-60s schedule. For three of these sessions the reinforcer was the delivery of two sucrose pellets, and for three it was the delivery of .3 ml groundnut oil. The first four sessions comprised ten reinforcer deliveries, and the final two sessions, five.

Then, all animals were trained on a positive-patterning switching discrimination. Each session consisted of 56 trials, eight of which were reinforced, and 48 nonreinforced. Reinforced trials consisted of a 10-s presentation of one of the visual feature stimuli, followed by a 10-s trace period during which no stimuli were presented, which in turn was followed by a 10-s presentation of one of the auditory target stimuli, the noise or the click. For half the animals, the light signaled pairings of the click and sucrose, and of the noise and oil, and the dark signaled the opposite, pairings of noise and sucrose, and of click and oil. For the remaining animals, the roles of the light and dark were reversed. There were two trials with each of the four feature–target combinations. The remaining nonreinforced trials comprised 12 presentations each of light, dark, noise and click. The different trials were presented in a semi-random order that was changed each day. In addition, each trial was preceded by a 10-s *pre-CS* period, during which responding was recorded separately. The intertrial interval was variable, with a mean of 65 s (range 50–80 s), resulting in a total session duration of 98 min.

Because after ten sessions, there was little sign of acquisition, the duration of the trace interval was reduced from 10 to 5 s, and training continued for a further ten sessions. Because discrimination was still poor, the number of reinforced trials was halved from eight to four, one of each feature–target combination, giving a total of 52 trials per session; training was then continued for another 16 sessions. Reducing the proportion of reinforced trials clearly improved discrimination performance, and also eventually resulted in a steady decline in response rates; accordingly at this point the session types with four and eight reinforced trials were intermixed (we have since routinely observed that discrimination performance is enhanced by reducing the proportion of reinforced trials in this type of discrimination; however, this is offset by the need to maintain a reasonable rate of responding, and also giving the animals sufficient experience of all the trial types required to allow them to learn the discrimination in the first place.) Sessions were arranged into four-session blocks, with two sessions of each type per block. This training régime was continued for a further 24 sessions.

Then the subjects were divided into two groups that were matched for discrimination performance in the previous stage. Training continued exactly as at the end of the previous stage, with blocks of four sessions, two of which included eight reinforced trials, and two of which included four. However, in contrast to the previous stage, every visual feature presentation was accompanied by the pulsed tone. For Group S the same reinforcement contingencies were in operation as in the previous stage, but for Group D the reinforcement

contingencies were reversed. This training continued for a further eight sessions. There followed a single test session, which was identical to the 56-trial training session except for the fact that the visual features and reinforcers were omitted on feature–target trials. The results of this test session were in the predicted direction, but the differences were small, so another 4-session training block was given, followed by a series of four test sessions including probe trials with the pulsed tone. These test sessions were also identical to the 56-trial training sessions, except that for one of each of the four types of reinforced trial the visual feature and reinforcement were omitted, thus providing four extinction test trials with the added feature.

Data treatment

In both this and the following experiment, flap-pressing was recorded separately during the various stimuli, the trace interval, and also during the 10-s pre-CS period that immediately preceded the onset of each trial. The measure of behavior was the rate of conditioned responding, flap-pressing, during the various types of CS. As the rates of pre-CS responding were low, there was no correction for pre-CS response rate. One rat failed to learn to drink oil from the foodtray and was consequently omitted from the experiment. A significance level of $p < .05$ was adopted in all the analyses that follow. Standard error terms for any means presented in the text are presented afterwards in brackets.

Results

Although the discrimination was difficult for the animals to acquire, by the end of training good discrimination performance had been achieved. During the last eight sessions of this stage animals responded consistently more to the various target stimuli when they were preceded by the feature stimuli than when they were presented alone, and discrimination performance was very similar in the two groups: the mean rate on reinforced and non-reinforced trials, respectively, was 15.43 (3.20) and 8.42 (1.27) rpm for Group S, and 16.52 (1.44) and 8.08 (1.16) rpm for Group D. An ANOVA with group (S or D) and trial type (feature + target or target alone) as factors revealed a main effect of trial type, $F(1, 13) = 35.51$, but no main effect of group or Group \times Trial type interaction, $F_s < 1$. The rates of responding during the pre-CS periods, the feature stimuli and the trace interval were low: the mean rates of responding for Groups S and D, respectively, over this period were .94 (.44) rpm and .55 (.15) rpm for pre-CS response rates, 2.42 (1.0) rpm and 1.86 (.61) rpm for responding during the feature stimuli, and 2.65 (1.37) rpm and 3.40 (1.05) rpm for responding during the trace interval. Analyses of variance revealed that the groups did not differ on any of the measures, $F_s < 1$.

The results of the 12 sessions of compound training are shown in Table 2; the data are presented in four-session blocks. Although discrimination performance appeared better in Group D than in Group S, this was not significant: an ANOVA with Group (S or D), trial type (signaled target versus target alone) and block revealed a significant main effect of trial type, $F(1, 13) = 22.17$, and of block, $F(2, 26) = 5.75$, and a significant interaction between these two factors, $F(2, 26) = 4.64$. However, neither the effect of group nor any of the interactions involving this factor were significant, largest $F(1, 13) = 1.04$. As in the previous stage, the mean rates of responding during the pre-CS periods, the feature, and the trace interval on reinforced trials remained low. The group mean rates for Groups S and D, respectively, were 1.03 (.34) and .71 (.27) rpm for pre-CS responding, 3.10 (1.03) and 3.45

Table 2

Group mean response rates during signaled and unsignaled presentations of the target CS in the three 4-session blocks of compound training in Experiment 1; figures below in italics represent corresponding standard errors

Block	1	2	3
Signaled/unsignaled			
Group S	15.02/7.14 <i>4.26/1.31</i>	14.77/6.86 <i>3.74/1.38</i>	9.87/5.25 <i>2.76/1.33</i>
Group D	18.98/8.30 <i>2.31/1.12</i>	17.25/7.82 <i>2.14/1.19</i>	15.16/7.74 <i>2.10/1.48</i>

(1.41) rpm for responding during the feature, and 3.50 (1.54) and 3.86 (.91) rpm for responding during the trace interval. ANOVAs revealed that the groups did not differ on any of these measures, $F_s < 1$.

The results from the single initial test session revealed little difference in discrimination performance between the two groups. The mean response rates for reinforced and nonreinforced trials, respectively, were 7.50 (2.25) and 4.88 (1.62) rpm for Group S, and 8.14 (1.50) and 6.11 (1.28) rpm for Group D; an ANOVA performed on these scores with group (S or D) and trial type (feature + target versus target alone) as factors revealed only a significant effect of trial type, $F(1, 13) = 6.79$; nothing else was significant, $F_s < 1$.

The group mean rates of responding during the target, when it was signaled by the added, pulsed tone, and when it was presented alone, for each probe test session are shown in Fig. 1. As the probe trials were conducted in extinction, and constituted half of the total number of reinforced trials in each of the test sessions, it was anticipated that extinction of C's control over responding to the CSs might occur over the course of these sessions, but the prediction was that C's control over responding would be weaker, and would extinguish more rapidly, in Group S than in Group D.

Because the first test session's reinforced probe trial data were lost for three animals in Group D, the data from the first session are omitted from the analysis that is presented below. The data for Group D presented in Fig. 1 for this session come from the remaining four animals (these four animals responded at rates similar to those of the remaining three subjects on session 2, suggesting that both subgroups were similarly representative of overall group performance). Overall response rates were slightly higher in Group D than in Group S; it is also clear that, despite a slight increase on session 2, the discrimination in Group S declined over the course of testing and had completely dissipated by session 4. Discrimination in Group D, on the other hand, was sustained, and if anything increased over the course of testing. An ANOVA performed on these data with Group (S or D), session (2–4) and trial type (feature + target and target alone) as factors revealed a significant three-way interaction between these factors, $F(2, 26) = 4.38$.

The source of this three-way interaction was further investigated. First, ANOVAs were performed on the data from each session, with group and trial type as factors. On session 2, this revealed a significant main effect of trial type, $F(1, 13) = 8.85$, but neither the main effect of group nor the interaction between these two factors was significant, $F < 1$ and $F(1, 13) = 3.44$, respectively. The same pattern was observed on session 3, with a significant main effect of trial type, $F(1, 13) = 9.14$, but no main effect of group or interaction, $F(1, 13) = 1.71$ and $F < 1$, respectively. On session 4, there was a significant main effect of group, $F(1, 13) = 5.32$; neither the main effect of trial type nor the interaction between these two factors was significant, $F(1, 13) = 2.67$ and 3.58, respectively.

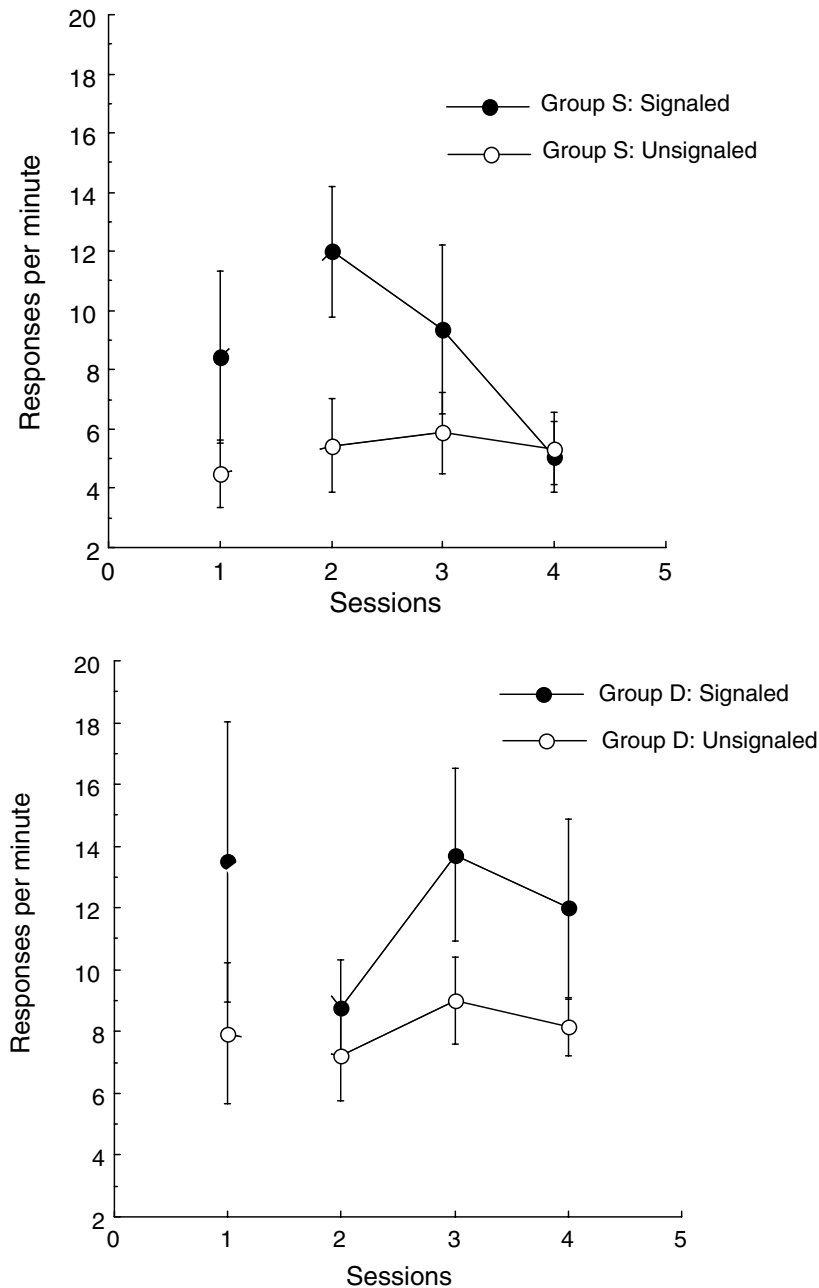


Fig. 1. Group mean response rates to signaled and unsignaled presentations of the target stimulus during the probe trials of the test sessions; data from Group S are shown in the top panel, and those from Group D in the bottom panel. The bars show standard errors.

These ANOVAs performed on the different sessions did not provide much encouragement to the prediction that Group D would show better discrimination performance than Group S. However, the results from the remaining tests did. Two separate ANOVAs with trial type (target + feature versus target alone) and sessions as factors were conducted for each group. In Group D, this revealed a significant effect of trial type, $F(1,6) = 7.02$; there was also an effect of sessions, $F(2,12) = 6.56$, but no interaction between these two factors, $F < 1$, showing that the slight increase in discrimination performance over testing in this group was not significant. The corresponding analysis for Group S revealed a main effect of trial type, $F(1,7) = 6.70$, no main effect of sessions, $F(2,14) = 2.45$, but, critically, a significant interaction between these two factors, $F(2,14) = 4.31$. Exploration of this interaction

with simple main effects analysis revealed that the discrimination was significant on session 2, $F(1, 7) = 9.05$; it was not significant on either of the remaining sessions, $F(1, 7) = 2.57$ and $F < 1$ for sessions 3 and 4, respectively. It also revealed that responding on reinforced trials declined significantly over the course of testing, $F(2, 14) = 4.64$, whereas on nonreinforced trials it did not, $F < 1$.

Finally, separate ANOVAs with group and session as factors were conducted on the data from reinforced and nonreinforced trials. The analysis of nonreinforced trials revealed no significant effects or interactions, largest $F(1, 13) = 2.35$; the analysis of responding on reinforced trials revealed no significant main effect of group or sessions, $F < 1$ and $F(2, 26) = 1.85$, but a significant interaction between these two factors, $F(2, 26) = 5.56$, $p < .01$. Exploration of this interaction with simple main effects revealed a significant difference between the groups on session 4, $F(1, 39) = 4.42$, but not on sessions 2 and 3, $F < 1$ and $F(1, 39) = 1.73$, respectively. There was also a significant main effect of session in Group S, $F(2, 26) = 5.26$, but not in Group D, $F(2, 26) = 2.35$.

The results of the preceding analyses may be taken to support the view that Group D showed performance on the test discrimination that was superior to Group S; discrimination performance in Group S was only significant on one of the test sessions, and extinguished rapidly, whereas the discrimination in Group D was maintained over all three sessions. Moreover, the groups differed in responding on reinforced, but not nonreinforced, trials. This pattern of findings is consistent with the prediction that C's control over target responding would be weaker in Group S than in Group D and would extinguish more quickly. This conclusion would, of course, be greatly strengthened if we had observed a significant interaction between group and trial type on sessions 3 and 4; however, such an effect was not observed.

As in all the other stages, the rates of responding during the pre-CS periods, and during the feature presentations and the trace intervals on reinforced trials, remained low. The mean rates for Groups S and D, respectively, were .90 (.41) and .91 (.30) rpm for pre-CS responding, 2.94 (1.05) and 4.00 (1.28) rpm for responding during the feature, and 2.75 (.65) and 3.29 (1.13) rpm for responding during the trace interval. ANOVAs performed on these scores revealed that the groups did not differ on any of the measures, $F_s < 1$.

Experiment 2

The results of Experiment 1 supported the contention that the occasion setters conveyed information about the combination of CSs and USs that were paired in their presence. Experiment 2 aimed to extend the generality of this result to a feature-negative preparation. Although designed as a companion experiment to the first study, it is important to note that there are certain design constraints in the feature-negative version of such a task. Specifically, the switching training employed in the previous experiment was as follows: $x \rightarrow$ nothing, $y \rightarrow$ nothing, A: $x \rightarrow p$, A: $y \rightarrow q$, B: $x \rightarrow q$, and B: $y \rightarrow p$. This allowed feature A to signal reinforcement of x with p , and of y with q , and for B to signal the opposite CS/outcome combinations. A feature-negative version of this discrimination would involve the features signalling that these CSs would be followed by the unexpected *omission* of these outcomes—A signaling that x would be followed by the unexpected omission of p , and so on. To do so would involve the following training: $x \rightarrow p$, $y \rightarrow q$, A: $x \rightarrow$ nothing, A: $y \rightarrow$ nothing, and also $x \rightarrow q$, $y \rightarrow p$, B:

Table 3
Design of Experiment 2

Pretraining		Reversal		
A: $x \rightarrow$ nothing	$x \rightarrow$ sucrose	Group S	A: $x \rightarrow$ sucrose	$x \rightarrow$ nothing
			A: $y \rightarrow$ oil	$y \rightarrow$ nothing
			A \rightarrow nothing	
A: $y \rightarrow$ nothing	$y \rightarrow$ oil	Group D	A: $x \rightarrow$ oil	$x \rightarrow$ nothing
			A: $y \rightarrow$ sucrose	$y \rightarrow$ nothing
			A \rightarrow nothing	

Note. A refers to the illumination of the jewel lights; x and y were a white noise and a clicker.

$x \rightarrow$ nothing, and B: $y \rightarrow$ nothing. Clearly, such training would make it impossible for the animal to learn that A and B signal the absence of specific CS–US combinations, because x (and y) must predict *both* p and q when presented alone, so that there is no specification of which outcomes are being omitted on A: x and B: x trials. For this reason, a simpler design was employed in the present experiment.

Animals were trained with one feature stimulus, A, two target stimuli, x and y , and two reinforcers, sucrose pellets and groundnut oil. When x and y were presented alone, x was followed by sucrose, and y by oil, but when they were signaled by A neither CS was followed by reinforcement (see Table 3). Then the animals were divided into two groups, S (same) and D (different). Both groups were trained on a positive-patterning task, effectively a reversal of the previously trained feature-negative discrimination, in that A now signaled the reinforcement of x and y , and x and y were nonreinforced when presented alone; A was also presented alone and nonreinforced, to reduce the impact of any Pavlovian excitatory conditioning to this stimulus on discrimination performance. For Group S, x was paired with sucrose and y with oil, exactly as they had been during the feature-negative training stage; for Group D this arrangement was reversed, such that x was now paired with oil and y with sucrose.

If during the initial stage of training feature A had come to signal the fact that x would *not* be followed by sucrose, then the reversal should be hard for animals in Group S, who must now learn exactly the opposite—that x is now to be followed by sucrose. Conversely, animals in Group D, who have to learn that x is now followed by oil in the reversal stage, experience no direct contradiction with what was learned in Stage 1. Thus, if occasion setters signal specific CS–US relations, Group S should perform less efficiently than Group D on the reversal task. But if the occasion setter merely acts independently on the individual CS and US components of the two occasion-set associations, x and y , the omission of oil, and the omission of sucrose, without being sensitive to the way in which these events are combined, then there would be no reason to predict a difference between the two groups in their ease of learning the reversal task.

Method

Subjects and apparatus

The subjects were 16 naïve male hooded Lister rats with a mean ad lib. weight of 604 g (range = 530–700 g), which were maintained at 80% of their ad lib weights, exactly as in Experiment 1. The apparatus was the same as in Experiment 1.

Procedure

Initially, all animals were trained to retrieve pellets and oil from the food tray. Subjects received sessions of magazine training in which 40 sucrose pellets were delivered according to a variable time (VT)-60s schedule, or three deliveries of .3 ml of oil were delivered according to a VT-240s schedule. Sessions were repeated until animals had collected all available reinforcements.

Then, all animals were given preliminary training in which 10-s presentations of each of the two excitatory stimuli, the click and the noise, were paired with reinforcement. For half the animals, the clicker was followed by sucrose and the noise by oil, and for the remaining animals this pattern was reversed. In order to maintain the same distribution of reinforced trials over the course of the session as would occur during the feature-negative training stage, the same control program was employed, but all the nonreinforced trials were omitted. Accordingly, each session theoretically comprised 56 trials programmed to occur with a mean ITI of 65 s (range 50–80 s); eight of these trials, four with the click and four with the noise, were reinforced; the remaining trials were dummy trials, on which no stimuli were programmed to occur. There were 12 sessions in this stage. There followed 48 sessions of feature-negative discrimination training, in which nonreinforced trials also occurred; these consisted of a 10-s presentation of the light, followed by a 5-s trace interval, in turn followed by a 10-s presentation of one of the auditory target stimuli, the noise or the click. There were 24 trials with each of the two nonreinforced feature–target combinations, and four of each of the two reinforced auditory stimuli, in each session. This type of session will be called type 1. As in the previous experiment, some adjustment was made to the number of reinforced trials per session over the course of this phase, in order to maximise observation of the discrimination that had been acquired. Thus after four sessions of type 1, animals received six sessions of type 2, which were the same as type 1 except for the addition of four more reinforced presentations of each of the two target stimuli (making 64 trials per session). The next four-session block (sessions 11–14) comprised three sessions of type 1 and one of type 2; the following block (sessions 15–18) comprised two of type 1, one of type 2, and one of type 3 (as type 1, but with only two of each of the reinforced types of trial, making 52 trials total per session). The next seven 4-session blocks (sessions 19–46) comprised three of type 1 and one of type 3, and one of the final two sessions was type 1 and the other type 3.

At this point, the subjects were divided into two groups that were matched for discrimination performance in the previous stage. Each reversal session comprised 24 presentations of each of the two auditory stimuli presented alone, and also 24 of the light, and 12 reinforced trials, six of each type, giving a total of 84 trials per session. During reinforced trials, the auditory stimuli were now signaled by the light, just as during the nonreinforced trials of the previous stage. For animals in Group S, the two target stimuli continued to signal the same reinforcer as they had in the first training stage, but for Group D this arrangement was reversed. There were 10 sessions in this stage. Although there was little sign of discrimination at this point, this was assumed to be due to the relatively high number of reinforced trials per session masking discrimination performance (see Experiment 1). Accordingly, at this point a test session was given, which was identical to the previous reversal training sessions except that there were only four reinforced trials with each combination, rather than six, giving a total of 80 trials.

Data treatment

The response measures and data analysis were the same as in the previous experiment.

Results

Acquisition of the discrimination was slow; two animals, one from each counterbalancing condition, were omitted as they were unable to master it. However, although the size of the discrimination was small by the end of training, the remaining animals were discriminating reliably; over the last eight training sessions, the mean rates of responding to the target on reinforced (unsignaled) and nonreinforced (signaled) trials, respectively, were 11.08 (2.26) and 10.25 (2.37) rpm for Group S, and 12.22 (1.84) and 11.80 (1.84) for Group D. An ANOVA with Group (S and D) and trial type (feature + target or target alone) as factors revealed a significant main effect of trial type, $F(1,12) = 5.36$; neither the effect of group nor the Group \times Trial type interaction were significant, $F_s < 1$. The rates of responding during the pre-CS periods, the feature stimuli and the trace interval were low: the rates for Groups S and D, respectively, were 1.13 (.28) rpm and 1.47 (.20) rpm for pre-CS responding, 1.65 (.45) rpm and 2.69 (.37) rpm for responding during the feature, and 1.67 (.49) rpm and 2.50 (.38) rpm for responding during the trace interval. ANOVAs performed on each of these scores, with group as factor, revealed that the groups did not differ on any of these measures, largest $F(1, 12) = 3.27$.

The results of the 10 test sessions are shown in Table 4, presented in two-session blocks; as noted above, there was little sign of learning in these data. This impression was supported by the results of an ANOVA with group, trial type and blocks as factors, which revealed no significant effects or interactions, largest $F(4,48) = 1.89$. The mean response rates in Groups S and D, respectively, were 4.03 (1.0) and 6.26 (1.03) rpm for the pre-CS periods, 4.22 (1.03) and 6.96 (1.37) for responding during the feature, and 3.74 (.97) and 6.75 (1.38) for responding during the trace interval. ANOVAs revealed that the groups did not differ on any of these measures, largest $F(1, 12) = 3.16$.

The results of the final, test session are shown in Fig. 2 and are clear: although the discrimination was still absent in Group S, it was now clearly present in Group D. This

Table 4

Group mean response rates to signaled and unsignaled presentations of the target stimuli in the five two-session blocks of the test of Experiment 2. Rates to the feature alone are shown below. Figures below in italics represent corresponding standard errors

Session	1	2	3	4	5
Signalled/unsignaled					
Group S	10.68/11.30 <i>1.96/1.98</i>	11.30/10.83 <i>1.50/1.81</i>	9.07/10.46 <i>1.17/1.62</i>	10.71/10.70 <i>1.74/1.57</i>	11.93/11.98 <i>1.93/1.98</i>
Group D	11.98/12.63 <i>2.65/2.16</i>	14.11/14.20 <i>1.63/1.53</i>	13.10/13.45 <i>1.86/1.81</i>	14.61/14.03 <i>2.52/2.60</i>	14.86/14.89 <i>1.89/1.96</i>
Feature					
Group S	2.75 <i>.63</i>	1.73 <i>.34</i>	1.95 <i>.53</i>	1.59 <i>.29</i>	2.86 <i>.75</i>
Group D	3.82 <i>.82</i>	3.13 <i>.71</i>	2.80 <i>.61</i>	3.46 <i>.78</i>	2.96 <i>.95</i>




Fig. 2. Group mean response rates to signaled and unsignaled presentations of the target stimulus in the test session of Experiment 2. The bars show standard errors.

impression was confirmed by the results of an ANOVA with group and trial type as factors; this revealed no significant main effect of group, $F(1, 12) = 3.01$, or of trial type, $F < 1$, but a significant interaction between these two factors, $F(1, 12) = 8.29$. Exploration of this interaction with simple main effects revealed a significant effect of trial type in Group D, $F(1, 24) = 6.51$, but not in Group S, $F(1, 24) = 2.31$. It also revealed that the groups differed on their response rates to reinforced but not to nonreinforced trials, $F(1, 24) = 4.93$ and 1.42, respectively. The mean pre-CS response rates on these sessions were 1.35 (.60) rpm for Group S and 1.31 (.42) for Group D; the corresponding rates for responding during the feature were 2.36 (.93) and 1.5 (.28) rpm, and for during the trace interval 1.29 (.89) and 3.43 (1.07) rpm for Groups S and D, respectively. None of these differences in rates was significant, $F < 1$, $F < 1$ and $F(1, 12) = 2.36$, respectively.

Discussion

These results suggest that the animals in Group D learned the reversal more easily than those in Group S. This result is consistent with the proposal that the occasion setter is acting on the specific CS–US relation that it signaled during initial training. There are, however, some issues associated with this result that deserve consideration. The first relates to the possibility that the feature was able to acquire Pavlovian associative strength during Stage 1 training; specifically, during feature-negative discrimination training, the light could have become a conditioned inhibitor for both oil and sucrose, and this could have retarded acquisition of excitatory associative strength for these two USs by the light during the reversal stage. If this had occurred, it would be expected to retard acquisition of the discrimination. But this could not explain the result we obtained, as there is no reason to suppose that this inhibition should be more profound in Group S than in Group D; there is therefore no reason to predict a difference between the two groups in their ability to learn the reversal task on the basis of this possibility.

Another potential explanation appeals to differences in learning about the light in Stage 2. In order to establish the light as a negative occasion setter for the associations between the clicker and one outcome, and between the noise and the other outcome, it was necessary for the clicker and the noise to be associated with different outcomes during the initial training stage. Thus, in Stage 2, on reinforced trials signaled by the light, the clicker and noise were being associated with their Stage 1 outcomes in Group S, but with the alternative outcomes in Group D. Being paired for the first time with reinforcement, albeit at some delay, the light would probably acquire some excitatory Pavlovian associative strength for the two outcomes in Stage 2. It is therefore possible that this conditioning to the light might have been blocked in Group S, for whom the outcomes were well predicted, relative to Group D, for whom they were not. To the extent that the final discrimination was the result of Pavlovian summation of the light's associative strength with that of the target CSs, this could in principle produce the result that was observed—greater elevation of target responding in Group D than in Group S. This interpretation might be considered unlikely, given the temporal delay between the light offset and outcome delivery, and the fact that the animals received many nonreinforced presentations of the feature throughout the reversal sessions, which would be expected to minimise any direct association between the feature and the outcomes. Nevertheless, to investigate the possibility further, responding during the feature in each test session was calculated and is presented in the lower panel of Table 4. Responding was low indicating that, as expected given the number of nonreinforced feature presentations delivered during the final discrimination phase, the feature commanded little associative strength. The level of responding was slightly higher in Group D than in Group S, but this effect was small and not statistically significant: an ANOVA with group and blocks as factors revealed no effect of group or of the Group \times Block interaction, $F(1,12) = 1.64$ and $F(4,48) = 1.66$, respectively; the effect of blocks was not significant, $F(4,48) = 2.29$. On the basis of these data, it seems there was no evidence of preferential Pavlovian conditioning to the light in Group D, and hence no indication that this factor could explain the results that we observed.

One other feature of the present design that deserves comment is that for Group S the target CSs were consistently paired with the same outcomes throughout the experiment, whereas in Group D they were paired with one outcome in Stage 1 and with the alternative outcome in Stage 2. It is possible that a CS associated with two outcomes elicits more (or less) responding than a stimulus associated with only one outcome (cf., [Watt & Honey, 1997](#)). This could introduce an overall difference in levels of responding between the two groups, and hence potentially affect interpretation of the results. But it should be noted that if such an effect occurred in the present experiment, it must have been extremely small; as described above, there was no significant difference between the two groups in responding to the target CSs.

General discussion

The results of these experiments support those reported by [Bonardi and Ward-Robinson \(2001\)](#), in providing evidence that an occasion setter provides information about the specific combination of CS and US that occurs in its presence. They extend the generality of [Bonardi and Ward-Robinson's \(2001\)](#) findings to a feature-negative and a positive-patterning preparation and to two different test procedures, blocking of the occasion setter (Experiment 1), and reversal learning (Experiment 2). The consistency of the results found

across all three experiments supports the idea that Bonardi and Ward-Robinson's (2001) findings were not attributable merely to the identities of the test stimuli they employed. The results thus constitute evidence that an occasion setter conveys information about the specific combination of CS and outcome that it predicts. They are not, on the other hand, consistent with Holland's (1989) suggestion that occasion setters act independently on the representations of any CSs and USs that have been subject to occasion setting and are hence insensitive to the combinations in which they are presented (cf., Jenkins, 1985). In this respect, the present findings parallel those from studies of instrumental discriminative control. In this field, a number of experiments have demonstrated that an instrumental discriminative stimulus conveys information about the specific response–reinforcer associations that it signals (e.g., Bonardi & Hall, 1994; Rescorla, 1990). This provides further support for the notion that the operation of discriminative stimuli and occasion setters obey essentially the same rules – appropriate, given that the term “occasion setting” was originally used in the context of instrumental learning (Skinner, 1938).

It was mentioned in the introduction that an alternative explanation of occasion setting may be couched in terms of configural learning (e.g., Brandon et al., 2000; McLaren & Mackintosh, 2002; Pearce, 1987, 1994; Rescorla & Wagner, 1972; Wilson & Pearce, 1989, 1990). According to configural theories, accurate performance on an occasion-setting discrimination is the result of Pavlovian conditioning to a configural cue produced by the compound presentation of the feature and target stimuli. Such an analysis could be provided for the results of the present experiments. For example, in Experiment 1, in the first stage of training each of the four possible feature–target combinations could have resulted in a configural cue that was associated with reinforcement, two being associated with one of the reinforcers (i.e., light + click → sucrose, dark + noise → sucrose) and two with the other (i.e., light + noise → oil, dark + click → oil). In the second stage of training, these combinations of configural cues and reinforcers were preserved in Group S, but not in Group D. If this were sufficient to produce an unblocking effect in Group D, then the new configural cue formed by the conjunction of the pretrained feature, the added feature and the target stimulus would acquire more associative strength in Group D than in Group S. However, this analysis relies on an assumption that may be questionable—that it is possible to produce unblocking of Pavlovian conditioning by changing a sensory property of the reinforcer, an effect that is not well established (e.g., Bakal, Johnson, & Rescorla, 1974; Ganesan & Pearce, 1988 although see, e.g., Betts, Brandon, & Wagner, 1996; Rescorla, 1999).

It might also be possible to construct an explanation of the results of Experiment 2 in terms of configural learning. In this study, animals were trained on a feature-negative discrimination, in which two target stimuli were nonreinforced when they were accompanied by the feature but reinforced when they were presented alone. The fact that the animals could learn this discrimination, and suppress their responding on compound trials, could be because they had learned that a configural cue of the feature and the target cue signaled the absence of reinforcement. If one were to reject the usual assumption that an inhibitor transfers its effects perfectly to CSs signaling outcomes of the same motivational valence, even when their sensory properties differ (e.g., Pearce, Montgomery, & Dickinson, 1981), then one might suggest that the Pavlovian inhibition accruing to these two feature–target compounds was *reinforcer-specific*, so that the configure of, for example, light + click signaled the absence of sucrose, and the light + noise configure the absence of oil. This would allow one to explain the pattern of results that was observed. To solve the reversal task, the

light+click configure would have to become associated with sucrose for animals in Group S, whereas for Group D it would have to become associated with oil. Thus, for Group S the reversal phase required them to learn the exact opposite of what they had learned in the first stage of training, whereas for Group D it did not; it follows that animals in Group S would find it harder to learn the reversal than those in Group D. Apart from the questionable nature of the assumptions underlying this explanation, one problem with this analysis is that it relies on the feature–target configure having greater Pavlovian associative strength in Group D than in Group S. One might expect at least some generalization from the configural cue to the feature or to the target alone, which should lead to higher responding to separate presentations of the feature and target stimuli in Group D than in Group S. No such effect was observed.

Whether or not one accepts these specific arguments, there are more general and convincing reasons for questioning a purely configural account of occasion setting. The first comes from some results reported by Morell and Davidson (2002). They established an occasion setter signaling that a target CS would be followed by a specific reinforcer, and then compared transfer to another target CS in two groups of animals; for one group the transfer target had been paired with the same outcome as had been employed in occasion-setting training, whereas for the other group it was paired with a different outcome. They observed substantially greater transfer of occasion setting when the outcomes were the same than when they were different. This result is, at face value, difficult for configural theories to explain, as the configural cue present at the transfer test was *the same for both groups*. The only way for configural theories to explain such a result would be to argue that the configural representation of a stimulus changes as a function of what it signals—that some US-specific element is activated whenever a predictor of that US is presented, and that this US-specific element mediates generalization between the cues.

A second reason to question a standard configural interpretation of occasion setting comes from a series of experiments conducted in this laboratory (Bonardi, submitted). Animals received feature-negative training in which two target stimuli, A and B, were paired with different outcomes (e.g., $A \rightarrow$ sucrose, $B \rightarrow$ oil); a feature stimulus, X, signaled the nonreinforcement of a simultaneous compound of A and B (i.e., $X \dots AB \rightarrow$ nothing). This training could have permitted the animal to learn that X signaled two things—that A would not be followed by sucrose, and that B would not be followed by oil; X, in other words, would give information about two specific cue/outcome combinations—A/no sucrose and B/no oil. After this training, the animals were divided into two groups and tested on a positive-patterning discrimination with X as the feature. For one group, Group S (same), X signaled that A would be followed by sucrose, and B by oil—in other words X signaled the exact opposite to what it had conveyed in the feature-negative phase; by implication, this task should be quite difficult for these animals. For Group D (different), this arrangement was reversed: X signaled that A would be followed by oil, and B by sucrose, so the information to be learned about X in this phase was not a direct contradiction of what had been learned in the previous phase, making this task relatively easy. In accord with this analysis, animals in Group D showed better performance than those in Group S in the test discrimination. This result cannot be explained by configural theories. Standard configural accounts all share the assumption that normal Pavlovian conditioning occurs to the configural cue that is activated on compound trials—and that the nature of the associative strength acquired by this cue is determined by the outcome on compound trials. But in the feature-negative phase of this experiment, the outcome on compound trials was the

unexpected absence of *both* oil and sucrose; this would have resulted in acquisition of inhibition for *both* outcomes by the configural cue—so that all information about the critical cue/outcome combinations (*A*/no sucrose and *B*/no oil) would have been lost. Thus, configural theories cannot explain the fact that information about these cue/outcome combinations appeared to have been retained during the test phase. This result adds to a growing body of evidence suggesting that configural learning cannot provide a complete explanation of occasion-setting effects (e.g., Honey & Ward-Robinson, 2001, 2002; Honey & Watt, 1998, 1999).

In conclusion, the pattern of results reported here, and in the experiments reported by Bonardi and Ward-Robinson (2001), is consistent with the idea that occasion setters are, at least in part, specific to the particular combination of events comprising the CS–US association. This specificity is consistent with a number of accounts of occasion setting. For example, some argue that occasion setting is the product of associative learning (e.g., Bonardi, 1989, 1991, 1998; cf., Mackintosh, 1983); one potential mechanism for such an account is that each CS–US association has linked to it a hidden unit, activation of which facilitates the ability of the CS to activate the representation of the US; the occasion setter could become associated with this hidden unit, so that presentation of the occasion setter would activate the hidden unit and hence facilitate use of the CS–US association. Others have suggested that a hidden unit links the occasion setter and the target with the representation of reinforcement (e.g., Honey and Watt, 1989). Which, if any, of these suggestions turns out to be correct must form the topic of further research.

Acknowledgments

I thank Dómhnall Jennings for helpful comments on the manuscript, and Jo Temperon, Stuart Morley and Richard Wood for technical assistance. Correspondence regarding the paper should be addressed to: Charlotte Bonardi, School of Psychology, University of Nottingham, University Park, Nottingham NG7 2RD, England.

References

- Bakal, C. W., Johnson, R. D., & Rescorla, R. A. (1974). The effect of change in US quality on the blocking effect. *Pavlovian Journal*, 9, 97–103.
- Betts, S. L., Brandon, S. E., & Wagner, A. R. (1996). Dissociation of the blocking of conditioned eyeblink and conditioned fear following a shift in US locus. *Animal Learning and Behavior*, 24, 459–470.
- Bonardi, C. (1989). Inhibitory discriminative control is specific to both the response and the reinforcer. *Quarterly Journal of Experimental Psychology*, 41B, 225–242.
- Bonardi, C. (1991). Blocking of occasion setting in feature-positive discriminations. *Quarterly Journal of Experimental Psychology*, 43B, 431–448.
- Bonardi, C. (1992). Occasion setting without feature-positive discrimination training. *Learning and Motivation*, 23, 343–367.
- Bonardi, C. (1996). Transfer of occasion setting: the role of generalization decrement. *Animal Learning and Behavior*, 24, 277–289.
- Bonardi, C. (1998). Conditional learning: An associative analysis. In P. C. Holland & N. A. Schmajuk (Eds.), *Associative learning and cognition in animals: Occasion setting*. Washington, DC: APA.
- Bonardi, C., & Hall, G. (1993). Occasion-setting training renders stimuli more similar: Acquired equivalence between the targets of feature-positive discriminations. *Quarterly Journal of Experimental Psychology*, 47B, 63–81.
- Bonardi, C., & Hall, G. (1994). Discriminative inhibition is specific to the response–reinforcer association but not to the discriminative stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, 20, 278–291.

- Bonardi, C., & Ward-Robinson, J. (2001). Occasion setters: Specificity to the US and the CS-US association. *Learning and Motivation*, 32, 349–366.
- Bouton, M. E. (1990). Context and retrieval in extinction and in other examples of interference in simple associative learning. In L. W. Dachowski & C. F. Flaherty (Eds.), *Current topics in animal learning: Brain, emotion and cognition* (pp. 25–53). Hillsdale, NJ: Erlbaum.
- Brandon, S. E., Vogel, E. H., & Wagner, A. R. (2000). A componential view of configural cues in generalization and discrimination in Pavlovian conditioning. *Behavioral Brain Research*, 110, 67–72.
- Brandon, S. E., & Wagner, A. R. (1998). Occasion setting: Influence of conditioned emotional responses and configural cues. In P. C. Holland & N. A. Schmajuk (Eds.), *Associative learning and cognition in animals: Occasion setting*. Washington, DC: APA.
- Dickinson, A., Hall, G., & Mackintosh, N. J. (1976). Surprise and the attenuation of blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, 4, 313–322.
- Ganesan, R., & Pearce, J. M. (1988). Effect of changing the unconditioned stimulus on appetitive blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, 14, 280–291.
- Holland, P. C. (1983). Occasion-setting in Pavlovian feature positive discriminations. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Discrimination processes* (Vol. 4, pp. 183–206). New York: Ballinger.
- Holland, P. C. (1985). The nature of conditioned inhibition in serial and simultaneous feature negative discriminations. In P. C. Holland & N. A. Schmajuk (Eds.), *Associative learning and cognition in animals: Occasion setting*. Washington, DC: APA.
- Holland, P. C. (1989). Transfer of negative occasion setting and conditioned inhibition across conditioned and unconditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 311–328.
- Honey, R. C., & Hall, G. (1989). Acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 338–346.
- Honey, R. C., & Hall, G. (1991). Acquired equivalence and distinctiveness of cues using a sensory-preconditioning procedure. *Quarterly Journal of Experimental Psychology*, 43B, 121–135.
- Honey, R. C., & Ward-Robinson, J. (2001). Transfer between contextual conditional discriminations: An examination of how stimulus conjunctions are represented. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 196–208.
- Honey, R. C., & Ward-Robinson, J. (2002). Acquired equivalence and distinctiveness of cues: I. Exploring a neural network approach. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 378–387.
- Honey, R. C., & Watt, A. (1998). Acquired relational equivalence: Implications for the nature of associative structures. *Journal of Experimental Psychology: Animal Behavior Processes*, 24, 325–334.
- Honey, R. C., & Watt, A. (1999). Acquired relational equivalence between contexts and features. *Journal of Experimental Psychology: Animal Behavior Processes*, 24, 325–334.
- Jenkins, H. M. (1985). Conditioned inhibition of keypecking in the pigeon. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 327–353). Hillsdale, NJ: Erlbaum.
- LoLordo, V. M., & Ross, R. T. (1990). Retraction of Ross and LoLordo's findings concerning blocking in serial feature-positive discriminations. *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 402–406.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford: Clarendon Press.
- McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representation: II. Generalization and discrimination. *Animal Learning and Behavior*, 30, 177–200.
- Morell, J. R., & Davidson, T. L. (2002). Transfer across unconditioned stimuli in serial feature discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes*, 28, 83–96.
- Pearce, J. M. (1987). A model of stimulus generalization for 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., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87, 532–552.
- Pearce, J. M., Montgomery, A., & Dickinson, A. (1981). Contralateral transfer of inhibitory and excitatory eyelid conditioning in the rabbit. *Quarterly Journal of Experimental Psychology*, 33B, 45–61.
- Rescorla, R. A. (1985). Conditioned inhibition and facilitation. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned Inhibition* (pp. 299–326). Hillsdale, NJ: Erlbaum.
- Rescorla, R. A. (1990). Evidence for an association between the discriminative stimulus and the response-outcome association in instrumental learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 16, 326–334.

- Rescorla, R. A. (1991a). Combinations of modulators trained with the same and different target stimuli. *Animal Learning and Behavior*, *19*, 355–360.
- Rescorla, R. A. (1991b). Transfer of inhibition and facilitation mediated by the original target stimulus. *Animal Learning and Behavior*, *19*, 65–70.
- Rescorla, R. A. (1999). Learning about qualitatively different outcomes during a blocking procedure. *Animal Learning and Behavior*, *27*, 140–151.
- 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* (pp. 64–99). New York: Appleton-Century-Crofts.
- Ross, R. T., & LoLordo, V. M. (1986). Blocking during serial feature-positive discriminations: Associative versus occasion-setting functions. *Journal of Experimental Psychology: Animal Behavior Processes*, *12*, 315–324.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Swartzentruber, D. E. (1995). Modulatory mechanisms in Pavlovian conditioning. *Animal Learning and Behavior*, *23*, 123–143.
- Watt, A., & Honey, R. C. (1997). Combining CSs associated with the same or different USs. *Quarterly Journal of Experimental Psychology*, *50B*, 350–367.
- Wilson, P. N., & Pearce, J. M. (1989). A role for stimulus generalisation in conditional discrimination learning. *Quarterly Journal of Experimental Psychology*, *41B*, 243–273.
- Wilson, P. N., & Pearce, J. M. (1990). Selective transfer of responding in conditional discriminations. *Quarterly Journal of Experimental Psychology*, *42B*, 41–58.