

## Occasion Setters: Specificity to the US and the CS–US Association

Charlotte Bonardi

*University of York, York, United Kingdom*

and

Jasper Ward-Robinson

*University of Cardiff, Cardiff, Wales, United Kingdom*

Two experiments using pigeon subjects examined whether the action of a positive occasion setter was specific to a particular CS–US association. In both experiments subjects were trained on a *switching* procedure, in which a diffuse feature, A, signaled that keylight x would be followed by one reinforcer, US1, and that a second keylight y would be followed a different reinforcer, US2. A second feature, B, signaled the converse arrangement, that x would be followed by US2 and y by US1. The intention was to establish whether the features allowed the birds to anticipate the specific keylight/US associations that they signaled. This was achieved using a blocking technique, by examining the ability of feature A to prevent learning about two test stimuli, S and D, presented in compound with x. On *same* trials A signaled stimulus x, which was followed, as in training, by US1; on these trials x was accompanied by S. On *different* trials x was signaled by A but followed by the other reinforcer, US2, and accompanied by the second test stimulus, D. If A signaled specific CS–US associations, then the US would be predicted in the former case but surprising in the latter, so that acquisition of associative strength by stimulus S would be blocked, but acquisition by stimulus D would not. In both experiments stimulus S supported less conditioned responding than stimulus D, which is consistent with this interpretation. The implications of these findings for theories of occasion setting are discussed. © 2001 Academic Press

*Key Words:* occasion setting; transfer; CS–US association.

Occasion setters are stimuli that signal that some conditioned stimulus (CS) will be followed by reinforcement, and as a result enhance responding

This work was supported by grants from the Biotechnology and Biological Sciences Research Council. We thank Geoffrey Hall and other members of the lab for much helpful discussion and Richard Dale and Iona Pinder for their care of the animals.

Address correspondence and reprint requests to Charlotte Bonardi, Department of Psychology, University of York, Heslington, York YO10 5DD, United Kingdom. E-mail: [cb2@york.ac.uk](mailto:cb2@york.ac.uk).



to that CS. They are intriguing because their behavioral effects are not obviously to be explained in terms of their own association with reinforcement (but see, e.g., Wagner & Brandon, 1989; Wilson & Pearce, 1989, 1990). A number of different suggestions have been made about the way in which these stimuli operate, and these generally agree in saying that the occasion setter acts on some part of the association between the CS and the unconditioned stimulus (US), enhancing the degree to which it can be retrieved or used. For example, it has been suggested that occasion setters facilitate use of the CS-US association in its entirety (e.g., Bonardi, 1989; Holland, 1983; see also Bouton, 1990); in contrast Rescorla (e.g., 1985) has argued that they act specifically on the US representation, lowering the threshold at which it can be activated. More recently Holland (e.g., 1989a) has argued that occasion setters will facilitate the activation of any association whose components, either separately or together, have themselves been the target of occasion setting.

These theories differ mainly in their predictions about the degree to which the occasion setter will affect responding to some different CS. For example, if the occasion setter acts on the CS-US association, its effects should be specific to *that particular* CS, whereas if it acts on the US representation, as Rescorla suggested, it should act with equal facility on *any* CS associated with the same US. And if Holland's (1989a) suggestion is correct, the occasion setter should act on any previously occasion-set CS (provided it is associated with a previously occasion-set US). Many experiments have attempted to discriminate between the various theories by examining these predictions. The results have been mixed (see Swartzentruber, 1995, for a review). Occasion setters frequently show quite substantial transfer to other CSs, suggesting that they are *not* CS-specific (although one could argue that a CS-specific occasion setter would show transfer provided there were generalization from the training CS to the transfer target). But the transfer that is seen is rarely complete (even to CSs that have themselves been the targets of occasion-setting discriminations), suggesting that occasion setters *are* CS-specific. In this case one could argue that a non-CS-specific occasion setter might fail to show perfect transfer because of generalization decrement produced by the novel feature-target combination. But some evidence against the latter possibility has been provided by Bonardi (1996), who demonstrated that imperfect transfer need not be explained in this way. If all instances of imperfect transfer are not to be explained in terms of generalization decrement, then one is forced to conclude that occasion setters are, at least in part, CS-specific.

This pattern of results is at least consistent with theories proposing that occasion setters act on the CS-US association; but it is not the only interpretation. Such theories must also predict that occasion setters are to some extent US-specific, so that an occasion setter should not transfer fully to the original CS if it is subsequently paired with some other US. Without evidence for

any US specificity it remains possible that the occasion setter, rather than acting on the association, simply facilitates activation of the CS representation and has no effect on the US representation at all. To argue that occasion setters can act on the CS–US association, it is thus necessary to show that occasion setters show some US specificity as well as CS specificity, and yet, surprisingly, almost no experiments have examined this issue. The exception is a study by Holland (1989b), using negative occasion setting and two appetitive reinforcers, food and sucrose; using this procedure he was able to examine the specificity of the occasion setter to the *sensory* properties of the US. He reported that, provided that the transfer US had participated in an occasion-setting discrimination, his occasion setters transferred perfectly to the original CS after it had been paired with a different US. At face value Holland's finding provides no support for the view that occasion setters can be US-specific. However, the implications of this result for positive occasion setting depend on exactly what a negative occasion setter is supposed to do. For example, you could argue that both positive and negative occasion setters act on an excitatory target–US association — positive occasion setters as facilitators, negative occasion setters as inhibitors. As the target association is the same, then it is likely (although not inevitable) that both positive and negative occasion setters will be similar in terms of their specificity to the elements of that association. But an alternative possibility is that, in a negative occasion-setting procedure, the target CS forms an inhibitory association with the US, and the negative occasion setter is simply facilitating the use of this inhibitory association. According to this view, then, a positive occasion setter facilitates use of an *excitatory* CS–US association, and a negative occasion setter of an *inhibitory* CS–US association. Adopting this position would have implications for the interpretation of Holland's findings because there is little evidence that conditioned inhibitors convey precise information about the sensory properties of the US whose omission they signal (e.g., Pearce, Montgomery, & Dickinson, 1981). And if the target association encodes no sensory information about the US, it is unclear how the negative occasion setter that modulates that association should be able to do so. But there is far more reason to believe that Pavlovian excitatory CSs are associated with the sensory aspects of the USs they predict; thus it remains possible that although negative occasion setters are not specific to the sensory properties of the US, positive occasion setters might be. Our first aim was to address this question.

The second was to examine a further prediction of the view that occasion setters act on the CS–US association. Suppose it were demonstrated that the occasion setter can be not only CS-specific but also US-specific. Although this evidence is necessary for the view that the occasion setter acts on the CS–US association, it is not sufficient — it could be that an occasion setter has separate and independent actions on the *components* of an association, without being sensitive to the way in which these events are *combined*. In

order to sustain the view that an occasion setter acts on an association it is necessary to demonstrate not only that it is specific to the CS and to the US, but also that it is specific to their combination (Jenkins, 1985). The present experiments also addressed this question.

### EXPERIMENT 1

In this and the following experiment an occasion setter was established as a signal for two associations, one between CS1 and US1 and the other between CS2 and US2. The question was whether this occasion setter would be more effective with these associations than with associations formed by recombining two of the original events—CS1 and US2, for example (cf. Jenkins, 1985). A result of this type would allow several conclusions. First, it would make it clear that the occasion setter was not only CS-specific, but also US-specific; if it were not, its effects would not be attenuated when the US following the training CS was switched. Second, it would rule out the possibility that the occasion setter was only acting independently on CS and US representations—if it were, it should be insensitive to the combination in which these stimuli were presented. The only possible interpretation of such a result would therefore be that the occasion setter could act on the target CS–US combination, and so would be consistent with the idea that it was operating on the association between those events.

A within-subjects procedure was employed, the design of which is illustrated in Table 1. All birds received training on four occasion-setting discriminations with two diffuse feature stimuli, A and B, two keylight target CSs, x and y, and two USs, lentils of differing colors. Birds were trained on a switching task: in the presence of feature A, x was followed by red lentils

TABLE 1

STAGE 1		STAGE 2		
		Same	Different	
A (x → red)	x-	A (xS → red)	A (xD → white)	x-
A (y → white)	y-	A (yS → white)	A (yD → red)	y-
B (x → white)	x-	B (xS → white)	B (xD → red)	x-
B (y → red)	y-	B (yS → red)	B (yD → white)	y-

*Note.* A and B were presentations of a click and offset of the houselight, respectively; these feature stimuli were of 10-s duration and immediately preceded presentation of the target keylights x and y, which were green and a grid composed of three vertical and three horizontal lines. S and D were purple and diamond, counterbalanced, and immediately followed x and y presentations; all keylight presentations were of 5-s duration, as were the reinforcer presentations on reinforced trials. Red refers to the red lentil reinforcer, and white to the white lentil reinforcer. For half the animals red and white reinforcers were presented as shown above, and for the remainder the identity of these reinforcers was reversed. There were 32 training sessions and two test sessions.

and y by white lentils, whereas in the presence of feature B the reverse was true. We then examined whether A and B told the animals anything about the specific CS-US combinations that they signaled by asking whether the birds could correctly anticipate which US would follow a particular target CS when it was signaled by a particular feature. As both reinforcers supported the same conditioned response, this would not be manifest in discrimination performance, and so an indirect, blocking technique was employed. The birds were given two types of trials, *same* trials and *different* trials; same trials were identical to the training trials, except that a further 5-s keylight stimulus, S, was inserted between the offset of x or y and the delivery of the reinforcer. Different trials were identical to same trials except that a different keylight, D, replaced S, and the USs were switched; thus if during training animals had received red lentils, then on these trials they received white and vice versa. S and D were novel, but during the course of this phase would presumably become associated with reinforcement and elicit conditioned responding. What was of interest was the relative speed with which this response would be acquired by the two stimuli. If the occasion setters act on specific associations, so that for example A acted as a signal that x would be followed by red lentils, and y by white, then on same trials the reinforcers would have been correctly signaled by x and y and so would have been unsurprising. S would therefore be paired with a predicted reinforcer, and so it would be difficult for these events to become associated, and responding to S would therefore be poor. On different trials, on the other hand, on being given A the animal would expect x to be followed by red lentils and so would be surprised when white lentils were presented. D would therefore be paired with a surprising reinforcer, learning about this stimulus would be good, and responding to it strong. In other words, if occasion setters act on specific associations, then one would predict more responding to D than to S. But if instead the occasion setters act independently on CS and US representations, then the occasion setters would be equally effective in allowing the target CSs to predict the reinforcer on same trials as on different trials, and one would not expect any difference in responding to S and D.

The rationale for this experiment depends on the assumption that unblocking may be produced by changing the sensory properties of the reinforcer, in this case the color of the grain reward. This perhaps requires some justification, as although unblocking is well documented after upshifts or downshifts in the *magnitude* of reinforcement (e.g., Dickinson, Hall, & Mackintosh, 1976), it is less well established that the same effect may be obtained after changes in the quality of the reinforcer (although see, e.g., Betts, Brandon, & Wagner, 1996; Rescorla, 1999). If the present experiment were to work, it could be taken as evidence for the existence of such a sensory unblocking effect. But although the current design relies on the existence of what is admittedly a rare phenomenon, if the anticipated result were to be obtained, it would be difficult to explain *without* assuming that the occasion

setters were giving the animals precise expectancies about the type of reinforcer that was about to be presented.

## Method

### *Subjects*

The subjects were 16 locally obtained pigeons maintained at 80% of their free-feeding body weights (mean *ad libitum* weight of 520 g; range = 475–585 g). They were housed in pairs in a colony room that was lit from 08:00 to 22:00 each day. They had served in a previous experiment using an autoshaping procedure, but were naive to all the stimuli used in the present experiment.

### *Apparatus*

Training was conducted in four standard three-key pigeon chambers supplied by Campden Instruments Ltd. Only the center key was used. An in-line projector behind the key allowed the following stimuli to be presented: white, green, grid, purple, and diamond. The colors were presented over the entire key (a 2.5-cm-diameter circle); purple was made by mixing red and blue. The other stimuli consisted of white figures against a dark background, presented in the center of the key. The diamond was 1 cm, and the grid was composed of three horizontal bars and three vertical bars superimposed on each other. A grain feeder positioned just above floor level and immediately below the center key was used to provide reinforcement; food delivery was accompanied by the illumination of a 2.8-W lightbulb situated inside the hopper. Except in the first autoshaping session, when standard grain was used, the feeder was either filled with red lentils (an orange color) or white lentils of the same size and shape (Alligator, Fishergate, York). Except when the stimulus “dark” was presented, each chamber was dimly illuminated by a 2.8-W bulb rated for 24 V, but operated at 15 V; this houselight was situated in the ceiling. A speaker mounted on the front wall could be used to deliver a 74-dB, 10-Hz clicker, produced by a Campden Instruments Ltd. tone generator. Each chamber was enclosed in a sound-attenuating shell; the operation of an extractor fan provided background masking noise. The equipment was controlled by a BBC (B) microcomputer programmed in a version of BASIC.

### *Procedure*

*Preliminary training.* The birds' previous experience of autoshaping meant that all would eat readily from the feeder and peck at illuminated response keys, so they received only a single reminder session of autoshaping at the start of the experiment. This consisted of 40 10-s presentations of a white center key, each of which was followed by a 5-s presentation of grain; the intertrial interval (ITI—trial offset to trial onset) was variable, with a

mean value of 45 s (range = 25–65 s). All subjects then received two extra hopper-training sessions in order to accustom them to the lentils. In each 20-min session 5-s hopper presentations were presented according to a variable time (VT) 60-s schedule; in one of these sessions red lentils were delivered, and in the other white. Birds that ate on less than 10 of these hopper presentations were given an extra session of training.

*Occasion-setting training.* There followed 32 sessions of occasion-setting training, in which all birds were trained on four feature-positive discriminations. The features, A and B, were the two diffuse stimuli, click and dark; the target CSs, x and y, were the grid and green keylights. All trials were composed of four consecutive 5-s bins. Reinforced trials began with a 10-s presentation of the feature, which occupied the first two 5-s bins; this was immediately followed by a 5-s presentation of the keylight, which was in turn immediately followed by a 5-s presentation of lentils. Nonreinforced trials consisted simply of a 5-s keylight presentation in the third 5-s bin of the trial. Each session consisted of 36 trials separated by a variable ITI of mean 80 s (range = 60–100 s). As only one type of lentil could be delivered in any session, animals received 2 sessions per day, 1 with each type of lentil; the order of session types over days was varied in an alternating double alternation sequence. For half of the subjects click-and-green and dark-and-grid were followed by red lentils, and click-and-grid and dark-and-green by white lentils; for the remaining subjects this arrangement of reinforcers was reversed; thus in each session there were 6 of each of the two types of reinforced trial that could be delivered in that session, and 12 nonreinforced trials with each of the target keylights. The different types of trial were presented in a quasirandom order.

*Extinction of test stimuli.* At this point the keylights that were to serve as the test stimuli underwent extinction to avoid contamination of test responding through generalization from the target CSs. All subjects received four sessions in which diamond and purple were presented in a semirandom order; the first, third, and fourth of these sessions comprised six presentations of each stimulus, and the second three. The ITI was variable with a mean of 110 s (range = 44–176 s).

*Test.* Two test sessions followed, one with each type of lentil. These were the same as the training sessions except in the following respects. Each session included three presentations of each of the four feature/target CS combinations. This meant that two of the four types of reinforced trials in each session would be followed by the same reinforcer as during training — same trials — and the remainder by the alternative reinforcer — different trials. Reinforced trials differed from those given in training in that a second 5-s keylight stimulus, either purple or diamond, was interposed between the offset of green or grid and the onset of lentil presentation. For half of the animals in each of the two counterbalanced subgroups purple was the stimulus interposed on same trials, and diamond on different trials, and for the remaining

subjects the converse was true. This is a slightly unusual test procedure in that in a typical blocking experiment the test stimuli are usually presented alone. However, given the sensitivity of pigeons to generalization decrement effects, we thought it would be more sensitive to measure responding to the test stimuli in the context in which they actually acquired associative strength, the compound training stage.

### Results and Discussion

Inspection of the training data revealed no systematic difference in responding to the green and the grid; in the last two training sessions the mean response rates for these two types of trial (in responses per minute, rpm) were 68.26 and 65.00 rpm, respectively. Nor was there an apparent difference in the ability of the two reinforcers to support conditioned responding, the corresponding means being 67.86 and 65.40 rpm for red and white lentils respectively. The birds responded slightly more to the keylights when they had been signaled by the click than by the dark—the mean response rates for these two types of trial in the same two sessions were 72.66 and 60.60 rpm respectively—but this factor appeared to have no influence on the effects of interest. For ease of presentation the response rates reported below are therefore pooled across these different factors. The birds seldom responded during the diffuse signals alone.

Although most birds learned to respond appropriately during discrimination training, in both this and the next experiment some were not discriminating reliably during the test phase. Consequently a performance criterion was introduced. In each test session the birds received several different types of reinforced trial (four in this experiment, eight in the next); as these only differed from training trials in what *followed*  $x$  and  $y$ , it was possible to use responding *during*  $x$  and  $y$  on these trials as an estimate of discrimination performance. Any bird that responded less on reinforced than on nonreinforced trials on more than 50% of the different types of reinforced trial, pooled over all test sessions, was rejected. In the present experiment application of this criterion resulted in the rejection of one bird (the only one which did not respond at all throughout the test); its data are not considered further. Finally, inspection of the data from both this and the next experiment revealed occasional differences in responding to the test stimuli, purple and diamond; analyses involving responding to these stimuli therefore included counterbalancing—subgroups for whom purple or diamond was stimulus  $S$ —as a factor. A significance level of  $P < .05$  was employed throughout this and the subsequent experiment.

#### *Occasion-Setting Training*

By the end of this stage the birds had learned all four discriminations, and were responding more to the keylights when they were preceded by the features than when they were presented alone. During the last pair of sessions



the mean response rates for reinforced and nonreinforced trials were 66.63 and 18.62 rpm respectively. An analysis of variance (ANOVA) performed on these data revealed that this difference was significant,  $F(1, 14) = 57.64$ .

### *Extinction Phase*


During the extinction phase responding to purple and diamond, which were to be the test stimuli, was low; over all the extinction sessions the mean rate of responding to the stimulus that was to serve as S was 23.20 rpm, and the corresponding rate to D was 20.53 rpm. An ANOVA with stimulus (S or D) and counterbalancing (subgroups for whom S was purple or diamond) as factors revealed no main effect of stimulus,  $F < 1$ ; there was an effect of counterbalancing,  $F(1, 13) = 4.70$ , but this did not interact with the effect of stimulus,  $F(1, 13) = 1.53$ .

### *Test*

The mean rates of responding to the test stimuli S and D during the two test sessions are shown in Fig. 1. Acquisition to S and D occurred rapidly, and as predicted the birds responded more to D than to S. This impression was confirmed by statistical analysis; an ANOVA with counterbalancing and stimulus (S or D) as factors revealed a main effect of stimulus,  $F(1, 13) = 5.63$ , no effect of counterbalancing,  $F(1, 13) = 4.26$ , and no interaction between these factors,  $F < 1$ .

The results of this experiment are consistent with the hypothesis that each occasion setter came to signal the specific CS-US associations that followed it during training. This allowed the animal to anticipate the type of reinforcer that was to be delivered after each feature-target CS combination, so that when a different reinforcer was presented, this was a sufficiently surprising event to allow unblocking and hence enhanced acquisition to stimulus D.

There is, however, an alternative explanation of these results. This arises



**FIG. 1.** Mean rates of responding during S and D stimuli during the two test sessions of Experiment 1.

from the fact that during training it was only possible to deliver one type of reinforcer in any session. It is possible that the anticipation, or aftereffects, of the reinforcer delivered in a particular session became part of the effective occasion setter, along with the feature, that predicted reinforcement of the target CS. So, for example, the compound of click and the anticipation of red lentils might have come to predict reinforcement of green during training (and click with the anticipation of white lentils, reinforcement of grid). In the test the click/green combination would for the first time be presented in a white lentil session to constitute a different trial. But the effective feature that signaled reinforcement of green, click plus anticipation of red, would have been altered on these trials, as the anticipation of red lentils would have been replaced by anticipation of white. Thus the occasion setter would lose its effectiveness, so that green would be only poor at predicting the reinforcer, and this would allow unblocking to occur. This critical feature of this analysis is that it is not necessary to suppose that the occasion setter tells the animal anything about *which* reinforcer follows the target CS; the same result would be generated if the animal knew only that food was about to occur. The present results would thus be far more compelling if one could reject this interpretation. One way of doing this would be to replicate the results of the previous experiment, but using a procedure in which both reinforcers could be delivered in the same session. This was the aim of the second experiment.

## EXPERIMENT 2

Experiment 2 employed the same logic as that used in Experiment 1; the critical difference was that the two reinforcer types were delivered in the same session. This was achieved by using standard grain as the reinforcer, but replacing the white hopper light that usually accompanied reinforcement delivery by LED lights of two different colors, red and green, so that the grain could be illuminated differently on different trials. In this way it was possible to have, as in the last experiment, two reinforcers differing in color, but these two reinforcer types could be delivered within a session.

### Method

#### *Subjects*

The subjects were 16 locally obtained pigeons maintained at 80% of their free-feeding body weights (mean *ad libitum* weight of 459 g; range = 395–573 g). They were housed as in the previous experiment. They had had previous experience of the autoshaping procedure, but were naive to the stimuli employed in the present experiment.

#### *Apparatus*

The apparatus was the same as that used in the previous experiment, except in the following respects. The keylight stimuli used in this experiment were

white, orange, purple (a compound of red and blue), horizontal, and diamond. The colors were presented over the entire key, and the horizontal was composed of three horizontal bars. Except when the flashing light stimulus and LED-lit reinforcers were presented, each chamber was dimly illuminated by the houselight. The flashing light was produced by operating the houselight alternately at 24 V (i.e., bright) for 0.5 s and then at 0 V (i.e., dark) for 0.5 s. When the feeder was operated the grain inside the hopper could be illuminated either with the normal, white hopper light or with one of two LED lights, one red and one green; the dim houselight was turned off during LED-lit grain delivery. A speaker mounted on the front wall could be used to deliver an 80-dB white noise, produced by a Campden Instruments Ltd noise generator.

### *Procedure*

Any aspect of the procedure that is not mentioned was the same as in the previous experiment.

*Preliminary training.* The birds received a single reminder session of autoshaping, consisting of 20 10-s presentations of a white center key, each of which was followed by a 5-s standard grain presentation, illuminated with the normal, white hopper light; the ITI was variable, with a mean value of 105 s (range = 95–125 s). They then received two 30-min hopper-training sessions to accustom them to the illuminated grain, as in the previous experiment; in one session the reinforcer was red-lit grain, and in the other it was green-lit grain.

*Occasion-setting training.* There followed 30 sessions of occasion-setting training, in which all birds were trained on four feature-positive discriminations; all birds received two sessions per day. The features, A and B, were the two diffuse stimuli, noise and flashing light; the target CSs, x and y, were the horizontal and orange keylights. Each session consisted of 12 non-reinforced trials and 6 reinforced trials with each keylight. Half of each of these types of reinforced trial were signaled by the noise, and the other half by the flashing light. For half the animals the reinforcer was red-lit grain when orange was signaled by the flashing light and when horizontal was signaled by the noise and green-lit food when orange was signaled by the noise and horizontal by the flashing light; for the remaining subjects this arrangement of reinforcers was reversed.

*Extinction of test stimuli.* All subjects received four sessions, each consisting of six presentations of diamond and six of purple, presented in a semirandom order. The variable ITI was the same as that used during occasion-setting training.

*Test.* Eight test sessions followed, each consisting of 24 trials. These were the same as the training sessions except in the following respects. Each session comprised eight nonreinforced presentations of each of the target stimuli and two presentations of each of the four feature–target CS combinations.

A second keylight stimulus, either purple or diamond, was interposed between the offset of orange or horizontal and the onset of grain presentation on reinforced trials. On half of each of the four types of reinforced trial the same reinforcer was delivered as had occurred during training; these were designated same trials. On the remaining reinforced trials, different trials, the alternative reinforcer was presented. For half of the animals in each of the two counterbalanced subgroups purple was the stimulus interposed on same trials, and diamond on different trials, and for the remaining subjects the converse was true.

### Results and Discussion

All aspects of the analysis not mentioned here were the same as in the previous experiment. Inspection of the data revealed that there were no systematic differences in responding to the keylights according to whether they had been signaled by the flashing light or by the noise; in the last two training sessions the mean response rates for these two types of trial were 82.16 and 80.32 rpm, respectively. Similarly, there was little difference in responding to the keylights according to whether they were reinforced by the red-lit or green-lit grain; the corresponding mean response rates for the same two sessions were 80.24 and 82.24 rpm, respectively. Consequently the response rates reported below are pooled across these different trial types. One bird had made no responses after six sessions of occasion setting training and was omitted from the experiment at this point.

Application of the performance criterion during the test phase resulted in the rejection of three birds from the experiment. Two of these performed considerably worse than the criterion, responding incorrectly on more than 70% of occasions, and the third made only 19 responses during the entire course of the test. These three birds responded very little to the two added stimuli during the test; the mean total responses to these stimuli for the remaining birds was 299 (range = 106–520); for the three rejected birds the corresponding totals were 19, 7, and 6.

#### *Occasion-Setting Training*

By the end of this stage the birds had learned all four discriminations and were responding more to the keylights when they were preceded by the features than when they were presented alone. During the last pair of sessions the birds were responding reliably more on reinforced than on nonreinforced trials; there was also a slight tendency for the birds to respond more on trials with horizontal than on those with orange; the mean response rates for reinforced and nonreinforced trials were, respectively, 95.46 and 5.69 rpm for horizontal and 67.00 and 4.58 rpm for orange. An analysis of variance performed on these data with keylight (orange or horizontal) and discrimination (reinforced or nonreinforced) revealed a significant main effect of keylight,  $F(1, 12) = 19.45$ , of discrimination,  $F(1, 12) = 44.67$ , and a significant


interaction between these two factors,  $F(1, 12) = 33.6$ ; simple main effects analysis revealed that the discrimination was significant with both keylights, smallest  $F(1, 12) = 30.02$ .

### *Extinction Phase*

During the extinction phase responding to purple and diamond, which were to be the test stimuli, was low; over all the extinction sessions the mean rate of responding to stimulus S was 9.62 rpm, and the corresponding rate to D was 10.12 rpm. An ANOVA with stimulus (S or D) and counterbalancing (subgroups for whom S was purple or diamond) as factors revealed no main effect of stimulus,  $F < 1$ ; nothing else was significant, largest  $F(1, 11) = 1.25$ .

### *Test*

The mean rates of responding to S and D during each of the eight test sessions are shown in Fig. 2. It is clear that, although on the first session animals responded more to S than to D, this pattern reversed on session 2, and animals continued to respond more to D than to S for the remainder of testing. We have no explanation as to why the pattern of responding should



**FIG. 2.** Mean rates of responding during S and D stimuli during the eight test sessions of Experiment 2.

have been reversed on session 1, although it suggests that the effect we observed could not be attributed to preexisting, subthreshold differences in responding to the test stimuli. This description of the data was largely confirmed by statistical analysis; an ANOVA with counterbalancing, stimulus (S or D), and session as factors revealed a main effect of stimulus,  $F(1, 11) = 7.11$ . Nothing else was significant, largest  $F(7, 77) = 1.59$ , except for the interaction between stimulus and counterbalancing,  $F(1, 11) = 14.85$ . Examining the interaction further with an analysis of simple main effects revealed that none of the effects were significant, largest  $F(1, 11) = 2.16$ . The interaction seemed to arise because, although responding to S was lower than responding to D in animals for whom S was purple, mean response rates 91.89 (S) and 118.14 rpm (D), respectively, this effect was slightly reversed in animals for whom S was diamond, corresponding means 99.24 (S) and 94.44 (D) rpm, respectively. Presumably, an overall tendency to respond at a higher rate to the diamond (perhaps because of its similarity to the white key used in autoshaping training) obscured the tendency to respond *less* to this stimulus when it acted as S. This interpretation is supported by the fact that when the S/D comparison is made for the same physical stimulus, responding was less to S than to D in both cases [91.89(S)/94.44(D) for purple and 99.24(S)/118.14(D) for diamond].

### GENERAL DISCUSSION

The first aim of the present experiments was to examine whether occasion setters show US specificity: does the occasion setter that signals reinforcement of a target CS tell the animal anything about the nature of that reinforcement? As we saw above, although essentially all theories of occasion setting predict that the occasion setter's effects should be specific to the training US, almost no experiments have examined this question, and what evidence there is suggests that, at least in the case of negative occasion setting, the occasion setter does not necessarily give any information about the training US at all (Holland, 1989b). This, in combination with the fact that occasion setters are at least in part specific to the training CS, is consistent with the view that occasion setters do no more than modulate the activation threshold of the CS representation. The results of the present experiments allow us to reject this interpretation. In this positive occasion-setting preparation the occasion setter apparently allowed the animal to anticipate quite precise characteristics of the reinforcer that followed presentation of the target CS. It seems that positive occasion setters show not only CS specificity but also US specificity.

The second purpose of these studies was to address the possibility that occasion setters affect representations of both CS and US, but that these effects are independent, so that an occasion setter would be equally effective with any *combination* of CS and US, provided both had previously been the targets of occasion setting (cf. Holland, 1989a). This possibility was directly

tested in the present experiments and found to be incorrect. The implication is that the effects of an occasion setter show specificity to the particular combination of CS and US that it signaled during training. Of course, the present results do not preclude the possibility that occasion setters *also* have separate effects on the CS and/or the US representations. For example, the present arguments have concentrated on the fact that occasion setters are in part CS-specific. But as was noted in the introduction, although it might not be complete, substantial transfer across CSs is often observed. Although one could try to argue that occasion setters are completely CS-specific and that all such transfer is the result of target–target generalization, it is also possible that there is a non-CS-specific component of occasion-setter action. But even if this were so, the present data suggest that some part of an occasion setter's action is specific to the CS–US combination. This is consistent with those theories that argue that the occasion setter facilitates operation of the entire CS–US association (e.g., Bonardi, 1989; Holland, 1983; see also Bouton, 1990).

This finding is also easily accommodated by theories that attribute occasion setting to a type of configural learning, in which a configural cue, resulting from the combination of the feature and the target, predicts a particular reinforcer (e.g., Wagner & Brandon, 1997; Wilson & Pearce, 1989, 1990; and see also Swartzentruber, 1998; but see Bonardi, 1996). But there are other reasons for rejecting this interpretation. For example, it has been demonstrated that occasion setters show blocking: If a feature signals the reinforcement of a target CS in compound with a second feature that has already served as an occasion setter for that CS, the added feature will acquire relatively poor control over conditioned responding (e.g., Bonardi, 1991; Bonardi & Hall, 1993). In other words, acquisition of occasion-setting properties by the added feature is blocked by the presence of a pretrained occasion setter. This result is perfectly consistent with an analysis according to which a configure of the pretrained feature and the target acquires associative strength, which blocks substantial further acquisition when the second feature is added. According to configural theories, such as that proposed by Pearce (1987, 1994), this will ensure that at test the configure of the added feature and the target CS will have little associative strength and hence command little conditioned responding. But this account anticipates a similar blocking effect when a *nonexplicit* training procedure, in which the target is never presented alone in the absence of the feature, is employed. The fact that blocking is not observed under these conditions (Bonardi & Hall, 1993) is therefore a problem for this interpretation. Another line of evidence that is inconsistent with the configural account comes from a recent report by Honey and Watt (1998); they trained animals on a switching design in which two occasion setters, A and B, signaled one pair of target–US relations (i.e.,  $x \rightarrow a$ ,  $y \rightarrow b$ ) and a second pair of occasion setters, C and D, the converse set of target–US relations (i.e.,  $x \rightarrow b$ ,  $y \rightarrow a$ ). Subsequent pairing of A with

shock resulted in more fear to B than to D; this result is impossible to explain in terms of a feature–target configure predicting the US.

The present findings are also broadly in agreement with analogous experiments in instrumental learning. The term occasion setting was, of course, originally coined to describe the way in which a discriminative stimulus allows the animal to access information about the response–reinforcer association that operates in its presence (Skinner, 1938). One suggestion as to how such stimuli might operate was made by Mackintosh (1983), who argued that the properties of a discriminative stimulus derive from an association between this stimulus and the response–reinforcer relationship. A clear implication of this view is that the discriminative stimulus, or occasion setter, should be specific to the particular association that operates in its presence, and Rescorla (e.g., 1990) has provided some elegant evidence in support of this view. This work has marked parallels with that being described here. There is an obvious operational analogy between the relationship of a discriminative stimulus to a response–reinforcer association and that of an occasion setter to a stimulus–reinforcer association. And Mackintosh's suggestion that a discriminative stimulus's powers are the result of an association between it and the response–reinforcer association it controls has been applied, with some success, as an explanation of Pavlovian occasion setting (e.g., Bonardi, 1991, 1996, 1998). This perspective predicts, just as in the instrumental case, that an occasion setter's action should be specific to the particular stimulus–reinforcer association whose operation it signals, and the present results may be taken as evidence in support of this view. This adds support to the general idea that Pavlovian and instrumental learning arise from the same learning mechanism by demonstrating that the hierarchical control of both types of learning appears to obey the same rules.

Finally, although we have emphasized an interpretation of the present results in terms of unblocking produced by a change in the sensory properties of the reinforcer, no independent evidence was provided in these experiments that any blocking had in fact occurred. And there are other possible mechanisms for the differences we observed. For example, it has been observed in rats that greater summation is observed when two stimuli signal different USs than when they both predict the same US (Watt & Honey, 1997). Although the source of this effect is unclear, it could explain the present results in that on some trials both the trained and added keylights signaled the same US, whereas on different trials they signaled different USs. To the extent that there was summation between the response tendencies elicited by the training keylights and the S and D keylights that followed, this could be manifest as differential responding during S and D. But the important point is that both the unblocking and the differential summation accounts require that the animals were able to correctly anticipate the reinforcers that were to be presented on the various types of trial. The most parsimonious conclu-



sion is that the occasion setters signaled the specific stimulus–reinforcer associations that held in their presence.

## REFERENCES

- 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. (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: American Psychological Association.
- Bonardi, C., & Hall, G. (1993). A search for blocking of occasion setting using a nonexplicit training procedure. *Learning and Motivation*, **25**, 105–125.
- 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.
- Dickinson, A., Hall, G., & Mackintosh, N. J. (1976). Surprise and the attenuation of blocking. *Journal of Experimental Psychology: Animal Behavior Processes*, **4**, 313–322.
- 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. (1989a). Acquisition and transfer of conditional discrimination performance. *Journal of Experimental Psychology: Animal Behavior Processes*, **15**, 154–165.
- Holland, P. C. (1989b). 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., & Watt, A. (1998). Acquired relational equivalence: Implications for the nature of associative structures. *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.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford, UK: Clarendon Press.
- 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., 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. (1999). Learning about qualitatively different outcomes during a blocking procedure. *Animal Learning and Behavior*, **27**, 140–151.
- 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.
- Swartzentruber, D. E. (1998). Perspectives on modulation: Modulator and target-focused views. In P. C. Holland & N. A. Schmajuk (Eds.), *Associative learning and cognition in animals: Occasion setting*. Washington, DC: American Psychological Association.
- Wagner, A. R., & Brandon, S. E. (1989). Evolution of a structured connectionist model of Pavlovian conditioning (AESOP). In S. B. Klein & R. R. Mowrer (Eds.), *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (pp. 149–189). Hillsdale, NJ: Erlbaum.
- 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.

Received October 17, 2000, revised January 24, 2001

Published online May 21, 2001