

Discriminative Inhibition Is Specific to the Response–Reinforcer Association but Not to the Discriminative Stimulus

Charlotte Bonardi and Geoffrey Hall

In 4 experiments an instrumental contingency between a response and a reinforcer was introduced in the presence of a discriminative stimulus. Then a discriminative inhibitor (S^{Δ}) was established that signaled that the instrumental contingency would no longer operate in the presence of that discriminative stimulus, so that the S^{Δ} suppressed operant responding. The degree to which the S^{Δ} 's inhibitory properties transferred to different discriminative stimuli and different response–reinforcer associations was then explored. In Experiments 1 and 2 the S^{Δ} 's effects transferred perfectly to a 2nd discriminative stimulus, whereas the results of Experiments 3 and 4 were consistent with the hypothesis that the S^{Δ} 's inhibitory properties were specific to the original response–reinforcer association. The theoretical implications of these findings are discussed.

If an animal is rewarded for responding in the presence of a given stimulus, for example, a tone, it learns to confine its responding to the presence of this *discriminative stimulus* (S^d). Recent work has suggested that this property of the tone does not depend on a Pavlovian association with reinforcement, as was once thought (e.g., Holman & Mackintosh, 1981). Instead, it has been suggested that S^d 's set the occasion for responding (cf. Skinner, 1938) by facilitating retrieval or use of the response–reinforcer association.

It is also possible to create the inhibitory counterpart of an S^d ; if the animal is rewarded for responding in the presence of the tone but reward is omitted when the tone is presented in compound with a light, the animal will not respond when the light is present. The light is called a *discriminative inhibitor* (S^{Δ}). As with the tone, the light's power over responding does not appear to depend on any simple associative relationship with reinforcement (Bonardi, 1988a, 1988b; Goodall & Mackintosh, 1987). Although it remains unclear exactly how the light affects behavior, it follows that if the animal responds because the tone allows it access to the information represented by the response–reinforcer association, there are several points at which the light might be exerting its inhibitory properties (cf. Rescorla & Holland, 1977). One possibility is that the light could depress activation of the representation of the tone S^d . Alternatively, or in addition, it could affect the response–reinforcer association; for example, it could inhibit the individual components of the response–reinforcer association—either the representation of the response or that of the reinforcer, or even both of those representations; or, it could act on that particular response–reinforcer association in its entirety.

One way to discriminate among these various possibilities is to use transfer tests (cf. Brown & Jenkins, 1967; Hearst & Peterson, 1973). With the previous example, if, for instance, the light were acting on the response, it would affect performance of that response but not of any other. Experiments that have used this logic have provided some grounds for discriminating among the different accounts of S^{Δ} function. For example, it appears that the effects of the S^{Δ} are both response specific and reinforcer specific. Bonardi (1989) demonstrated that an S^{Δ} signaling the nonreinforcement of one particular response was better able to suppress performance of that response than of some other response; likewise, she found that an S^{Δ} signaling the omission of one particular reinforcer was better able to suppress a response rewarded with that reinforcer rather than with some other reinforcer (see also Colwill, 1991).

These observations suggest that the S^{Δ} does not act on only one of the components of the response–reinforcer association. The other three accounts, however, remain viable. It is obviously possible that the S^{Δ} might still independently inhibit both components of the response–reinforcer association. It is also possible that the S^{Δ} is acting on the response–reinforcer association as an independent unit, so that changing either the response or the reinforcer effects a change in that association and hence a loss of inhibitory control (cf. Jenkins, 1985). Finally, these results do not rule out the possibility that the S^{Δ} might also have an inhibitory effect on the S^d . More information is required if we are to decide among these three accounts of S^{Δ} function. The present experiments were designed to provide such information.

Charlotte Bonardi and Geoffrey Hall, Department of Psychology, University of York, York, United Kingdom.

This research was financed by a grant from the Science and Engineering Research Council. We thank Rob Honey for helpful discussion and Alan Willis for assistance in running the animals.

Correspondence concerning this article should be addressed to Charlotte Bonardi, Department of Psychology, University of York, Heslington, York YO1 5DD United Kingdom.

Experiment 1

In the first experiment, we examined the last of the aforementioned possibilities—the idea that the S^{Δ} might inhibit the representation of the S^d . We did so by looking at the degree to which the S^{Δ} 's properties transfer to a second S^d . If the S^{Δ} affects only the components of the response–

reinforcer association or the association in its entirety, then—all other things being equal—its inhibitory effects should be unaffected by the identity of the S^d with which it is presented. If the S^A inhibits the S^d , however, then it should not inhibit responding commanded by a second discriminative stimulus. Of course, this argument is slightly simplistic. Suppose that S^A s are S^d specific. Although at face value this account predicts that an S^A would show no transfer to a second S^d , this prediction ignores the role of stimulus generalization. To the extent that there is some generalization between the original S^d and the transfer S^d , the S^A might well be expected to have some effect on the transfer S^d , although this would be less than its effect on the original S^d . In other words, a theory according to which S^A s are S^d specific can predict some transfer of an S^A to a second S^d , but it predicts that this transfer will be incomplete.

Some studies have already examined whether an S^A established with one S^d can suppress responding in the presence of a second S^d and have found that the inhibitory properties of an S^A will in fact transfer in this way (e.g., Brown & Jenkins, 1967; Hearst & Peterson, 1973). This evidence is not conclusive, however. Hearst and Peterson, for example, did not assess whether the transfer they obtained was complete, so it is impossible to deduce from this study whether some component of the S^A 's effect was nevertheless S^d specific. Brown and Jenkins did include such a comparison and found transfer that was incomplete; but in their experiment, the two S^d s controlled different responses, so the failure to obtain perfect transfer may have been due to the response specificity of the S^A 's properties.

In the present experiment, all animals initially received identical discriminated operant training with two S^d s—a clicker and a noise. To provide evidence that the animals could discriminate these two stimuli, we arranged that the stimuli signaled the reinforcement of two different responses: Pressing a lever on the left of the food magazine was reinforced in the presence of the click, and pressing a lever on the right of the food magazine was reinforced in the presence of the noise. The subjects then received training in which two visual stimuli, light and dark, were established as S^A s; responding was no longer rewarded when the click was presented in compound with the light or when the noise was presented in compound with the dark. Once the subjects had mastered this discrimination they were divided into two groups and tested (see Table 1). In the test phase, all animals were rewarded for making a third response, the magazine response (operating the flap covering the food tray), during the click and the noise. During these sessions, for Group S (same) the click was always presented in compound with the light and the noise was always presented in compound with the dark, whereas for Group D (different) these stimuli combinations were reversed—the click was presented with the dark and the noise was presented with the light. If the inhibitory control over responding acquired by the light during discrimination training was specific to the click, and that acquired by the dark was specific to the noise, then animals in Group S should be slow to learn to make the magazine response during stimulus presentations in the test,

Table 1
Design of Experiment 1

Group	Stage 1	Test
S	Click (R1 → food)	Click + light (Rm → food)
	Click + light (R1 → 0)	
	Noise (R2 → food)	Noise + dark (Rm → food)
D	Noise + dark (R2 → 0)	
	Click (R1 → food)	Click + dark (Rm → food)
	Click + light (R1 → 0)	
	Noise (R2 → food)	Noise + light (Rm → food)
	Noise + dark (R2 → 0)	

Note. In Stage 1, responses were rewarded in the presence of either the click or the noise; responding was not rewarded when these stimuli were accompanied by one of the visual stimuli, light or dark (these were fully counterbalanced). The test response, operating the flap of the food tray, was rewarded during a compound of either click or noise with either light or dark. S = same; D = different; R1 and R2 = right- and left-lever responses (counterbalanced); Rm = operation-of-food-tray-flap response.

as the developing discriminative properties of the click and the noise would be inhibited by the presence of the light and the dark, respectively. No such retardation should be observed in Group D, for which, in this phase, the click and the noise were presented with the alternative inhibitory stimulus.

Method

Subjects

The subjects were 16 naive male hooded Lister rats with a mean ad lib weight of 381 g (range: 363–413 g). Before the start of training the rats were reduced to 80% of their ad lib weights 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.

Apparatus

Four Campden Instruments operant chambers (Model 4104) were used. Each of the boxes had three walls of sheet aluminium, with a transparent plastic door as the fourth wall, and a translucent white Perspex ceiling. Each of the boxes contained a recessed food tray to which 45-mg mixed-composition food pellets could be delivered; this was situated in the center of one of the walls, adjacent to the door. The animals gained access to this food tray by means of a rectangular aperture 6 cm high × 5 cm wide. A transparent plastic flap of the same dimensions was attached by a hinge to the top of the entrance to the food tray. Pushing this flap inward from its vertical resting position allowed animals to enter the food tray. This movement activated a microswitch, and each closing of this switch was recorded as a single response. The flap automatically returned to its resting position when the rat removed its snout from the food tray. Each of the chambers was equipped with two retractable levers, which were mounted on the left and right of the food tray. Illumination was provided by a 2.8-W 15-V houselight, which was mounted in the front wall, directly above the food tray. A dark stimulus could be produced by turning off this houselight, and a light stimulus could be produced by turning on a 30-W striplight (rated for 240 V but operated at 100 V) that was situated above the Perspex ceiling. Two speakers were

mounted on the back wall through which an 83-dB, 10-Hz clicker and an 85-dB white noise (scale A) could be delivered from a Campden Instruments tone generator (Model 258) and a white noise generator (Model 530). The boxes were housed in sound- and light-attenuating shells; masking noise was provided by ventilating fans contained in these shells. The apparatus was controlled by BBC microcomputers (Model B) programmed in a version of BASIC.

Procedure

Preliminary training. In the first 30-min session the animals were trained to retrieve pellets from the food tray. The pellets were delivered according to a 60-s variable time (VT60) schedule. Animals who failed to eat all the pellets that were delivered were given an extra session. Both levers were retracted during this session.

The animals then were trained to press the two levers. For the rest of preliminary training and all preliminary variable ratio (VR) training, sessions were scheduled in duplicate pairs. The left lever was present during one session of each pair, and the right lever was present during the other. Reinforcement was contingent on responding on the available lever; in this way, we equated subjects' experience of responding on the two levers.

During the next two sessions the animals were consistently reinforced for responding, until a minimum of 75 responses had been made. All subjects then received a pair of stimulus preexposure sessions in each of which they experienced one presentation of each of the following stimuli: click, noise, noise + light, noise + dark, click + light, and click + dark. These presentations were designed to dissipate any unconditioned suppression to the various stimuli, thus allowing an uncontaminated assessment of their control over instrumental responding. Responding was reinforced throughout these sessions according to a variable interval (VI) 30-s schedule. In these sessions and throughout the rest of the experiment, each stimulus presentation was of 20-s duration and (unless stated otherwise) was preceded by a 20-s pretrial period. The entire 40-s period constituted a trial; the intertrial interval (ITI) in the preexposure sessions was variable, with a mean of 240 s (range: 180–300 s).

Preliminary VR training. At this point, animals were trained to confine their responding to the presence of the two discriminative stimuli, the click and the noise. Half of the animals were rewarded for responding on the left lever during the click and for responding on the right lever during the noise, and for the remaining animals these contingencies were reversed. Responding was without consequence in the absence of these stimuli. In the first three pairs of sessions of this stage there were twelve 20-s presentations of either the click or the noise, during which responding was rewarded according to a VR2 schedule. There was no pretrial period during these sessions; for the first pair of sessions the ITI was 11 s, for the second pair it was 21 s, and for the final pair it was 41 s. The next eight pairs of sessions each consisted of 12 trials, each preceded by a pretrial period; the ITI was variable, with a mean duration of 80 s (range: 50–110 s). In the first pair of these sessions, responding was rewarded on a VR2 schedule, the second and third pairs on a VR3 schedule, the fourth and fifth pairs on a VR4 schedule, the sixth pair on a VR5 schedule, and the seventh and eighth pairs on a VR6 schedule.

VR discrimination training. From this point until the start of transfer S^d training, both levers were always available at all times, and animals were explicitly trained to perform the correct response during each S^d . In each of the eight sessions of this stage, animals received 12 trials with the click and 12 trials with the noise. The

different stimuli were presented in a quasi-random order. The lever associated with reinforcement during these stimuli was the same as in the previous stage. In the first session, responding was rewarded on a VR4 schedule, the second on a VR5 schedule, and the final six sessions on a VR6 schedule.

Discrimination training. Twelve sessions of discrimination training followed, in which the two S^d 's were established. The first S^d signaled that responding would no longer be rewarded during the clicker, whereas the second signaled that responding would no longer be rewarded during the noise. As stated above, for half of the animals responses on the left lever were rewarded during the click, and responses on the right lever were rewarded during the noise, whereas for the remaining animals, the opposite was true. For half of each of these subgroups the light was the S^d for the click and the dark was the S^d for the noise, and for the remaining animals the converse was true. There were 24 trials in each session; of these, there were 9 with the noise and 9 with the click, during which responding was reinforced according to a VR6 schedule, as in the previous stage of training. There were also 3 presentations of each of the two stimulus compounds, during which no rewards were available. The different types of trial were presented in a quasi-random order.

Retardation test. Finally, the animals were given eight test sessions in which the effect of the S^d 's was evaluated. The animals were divided into two groups, Group S and Group D. The four counterbalanced subgroups described in the previous section were each divided into two, and half were allocated to one group and half to the other. No levers were present during this stage, and animals were trained to operate the magazine flap for food reinforcement during presentations of the click and the noise in compound with the S^d 's. For animals in Group S the click and the noise were each presented in compound with the S^d with which they had been paired during discrimination training; for subjects in Group D the click was presented in compound with the noise's S^d , and the noise was presented in compound with the click's S^d . There were 12 presentations of each of these 2 stimulus compounds per session. Responding was rewarded according to a fixed ratio-12 (FR12) schedule in the first session of this stage. However, few animals earned any reinforcers in this session, so in subsequent sessions the schedule was reduced to FR4 and remained at this value for the rest of the test sessions, all other aspects of which were identical to those of the discrimination training stage.

Results and Discussion

Inspection of the data revealed that the click and the noise acquired equally good excitatory control over responding, and that the light and the dark became equally good S^d 's; thus, we pooled response rates over click and noise trials and over compound trials with the light and the dark. A significance level of $p < .05$ was used throughout this and the subsequent experiments.

Discrimination training. Preliminary training and discrimination training proceeded normally. Response rate, in responses per minute, during the S^d 's in the last discrimination training session for Group S and Group D were 48.92 and 47.89, respectively; the corresponding pretrial response rates for these groups were 0.73 and 0.56, respectively. An analysis of variance (ANOVA) with group and S^d (responding during the S^d or during the pretrial period) as variables revealed a significant main effect of S^d , $F(1, 14) = 97.91$, which did not interact with group ($F < 1$). The main effect of group was not significant ($F < 1$).

As there was no difference between the two groups in responding during the S^d s, we evaluated the level of discrimination performance by calculating a ratio score of form $a/(a + b)$, where a was the rate of responding during compound trials and b was the rate during S^d alone trials. This yielded ratios of .110 for Group S and .170 for Group D. An ANOVA revealed that these scores did not differ ($F < 1$). Thus, there was no sign that the groups differed in their discrimination performance.

The rates of responding on the incorrect lever were low during this session—on S^d trials, the rates, in responses per minute, for Group S and Group D were 0.44 and 1.31, respectively; the corresponding rates on compound trials were 1.00 and 0.69. An ANOVA performed on these data with group and trial type (S^d alone or compound trials) as variables revealed no significant effects or interactions, largest $F(1, 14) = 1.38$.

Retardation test. The results from the test sessions are shown in Figure 1, which shows response rates during stimulus presentations, and also, as responding in the absence of the stimuli appeared to differ between the two groups, during the pretrial periods. There was no sign that the compound stimuli were less effective at controlling responding in Group S than in Group D. On the contrary, responding was somewhat higher during stimulus presentations in Group S; however, as responding was also higher in the absence of the stimuli in this group, this did not seem to reflect better stimulus control. We performed an ANOVA on these data with group (S or D), stimulus (response rates during the stimuli or the pretrial period), and session as variables. The ANOVA revealed a significant effect of group, $F(1, 14) = 6.69$, suggesting that overall response rates were higher in Group S than in Group D. There was also a significant effect of session, $F(7, 98) = 37.85$; and of stimulus, $F(1, 14) = 115.28$; and a significant interaction between these two variables, $F(7, 98) = 35.19$. Presumably, this reflects the fact that responding during stimulus presentations increased over sessions, whereas responding in

pretrial periods did not. Most important, however, is that although the stimuli acquired control over responding, there was no sign that this differed in the two groups; nor did such a difference develop over sessions. This interpretation is supported by the statistics: Neither the interaction between group and stimulus nor the three-way interaction among group, stimulus, and sessions was significant ($F_s < 1$). Nothing else was significant ($F_s < 1$).

The results of this experiment provide no support for the idea that the S^A has a direct and specific effect on the S^d ; there was no indication that the S^d s acquired discriminative control over magazine entry more rapidly in Group D than in Group S. However, there are at least two reasons why this conclusion should be adopted with caution. First, the overall level of responding was higher in Group S than in Group D; although it is not clear why this was the case, it raises the possibility that whatever is responsible for this difference could be masking a real difference in the degree of stimulus control. A more serious criticism concerns the fact that, whereas in discrimination training the S^A s inhibited responding on one of the two levers, in the test phase the critical measure was their effect on a *third* response, magazine entry. This procedure could be problematic because, as we have noted, S^A s' effects are also response specific. This means that an S^A will affect performance of a second, transfer response only to the extent that there is generalization between this transfer response and the original response with which the S^A was trained. Consequently, one would not expect the S^A to have a large inhibitory effect on responding in this test procedure. Although it is still logically possible that the inhibition that is observed will be attenuated still further when the S^A is presented with the transfer S^d , in practical terms this might be difficult to detect if the S^A is able to only weakly suppress performance of this third response in the first place. We designed Experiment 2 to address these criticisms.

Experiment 2

The design of Experiment 2 is shown in Table 2. A single group of animals received discriminated operant training identical to that of the previous experiment, except that both click and noise signaled the reinforcement of the same lever-press response. An S^A was established for each of the S^d s, exactly as in the previous experiment. The animals then received a test in which responding during the S^d s was examined: The S^d s were presented alone, in compound with the S^A with which they had been trained, or in compound with the second, transfer S^A . We anticipated that each S^A would be able to inhibit responding during the S^d with which it had been trained. But if a component of that inhibition can be attributed to an effect of the S^A on the S^d itself, then the S^A should have less of an inhibitory effect on responding when in compound with the second, transfer S^d . With this procedure, we avoided the shortcomings of the previous experiment. First, conducting the test with the originally inhibited response ensured that sufficient inhibition would be obtained to allow detection of an attenuation




Figure 1. Mean rates of responding during stimulus presentations and during the pretrial period, for Group S and Group D in the eight test sessions of Experiment 1. S = same; D = different.

Table 2
Design of Experiment 2

Stage 1	Test 1	Test 2
Click (R → food) Click + light (R → 0)	Click (R)? Noise (R)? Click + light (R)? Noise + dark (R)? Noise + light (R)?	Click (Rm → food) Noise (Rm → food) Click + light (Rm → food) Noise + dark (Rm → food) Noise + light (Rm → food)
Noise (R → food) Noise + dark (R → 0)	Click + dark (R)?	Click + dark (Rm → food)

Note. In Stage 1, R was rewarded in the presence of the click and the noise; responding was not rewarded when these stimuli were accompanied by one of the visual stimuli, light or dark (these were fully counterbalanced). In Test 1, performance of R was examined during the click and the noise and during these stimuli in compound with the original and transfer S^ds. In Test 2, response Rm was rewarded during the same stimuli that were presented in Test 1. R = left-lever response; S^d = discriminative inhibitor; Rm = operation-of-food-tray-flap response.

of that inhibition. Second, as this experiment had a within-subjects design, the problem of potential group differences in overall response rate was avoided.

Finally, all animals received a second test, a within-subjects version of the one administered in Experiment 1. The levers were removed from the chambers, and all subjects received discriminated operant training with the magazine flap response. This response was rewarded during presentations of the S^ds when they were presented in compound with the S^Δs with which they had been trained and also when in compound with the transfer S^Δs. In contrast to Experiment 1, however, responding was also rewarded during presentations of the S^ds alone. The prediction was the same as in Experiment 1: If a component of the S^Δs inhibitory properties are directed at the original S^d, then animals should be slower to learn to make the magazine response during the S^d in compound with the original S^Δ than with the transfer S^Δ. This test procedure improved on that of Experiment 1 in two respects. First, because it was conducted within subjects, the results could not be complicated by between-group differences in response rate, as they were in the previous experiment. Second, we argued above that because S^Δs are response specific, they would not be expected to have a large effect on performance of the magazine response. One could, however, make a stronger argument; as no attempt was made to measure the extent to which inhibition did transfer to the magazine response in that experiment, it is possible that there was no transfer of inhibition at all, in which case a failure to observe an attenuation of inhibition would hardly be surprising. By measuring responding during the S^ds in the present experiment, we hoped to demonstrate that there was significant transfer of inhibition to the magazine response; this would not only invalidate such a criticism of the present experiment but also strengthen our interpretation of the test results of Experiment 1.

Method

Subjects

The subjects were 16 male hooded Lister rats with a mean ad lib weight of 431 g (range: 385–480 g) that had previously partici-

pated in a flavor aversion experiment but were otherwise naive; they were maintained as in the previous experiment.

Apparatus

The apparatus was identical to that used in the previous experiment, except that the click was 79 dB and the noise 78 dB, and only the left lever was used.

Procedure

All aspects of the procedure that are not specified were the same as those for Experiment 1.

Preliminary training. After magazine training, the animals were trained to press the left lever. They received one session in which they were consistently reinforced for responding and one preexposure session in which they experienced two presentations of the same stimuli to which they had been exposed in Experiment 1.

Preliminary VR training. In the previous experiment, sessions in this stage were scheduled in pairs, during each of which one of the two responses was reinforced during one of the S^ds. In the present experiment, as only one response was used, each of these pairs of sessions was amalgamated into a single session with twice the number of trials. Click and noise trials were presented in a quasi-random order. VR discrimination training was of course unnecessary, and therefore these sessions were omitted.

Discrimination training. Twelve sessions of discrimination training followed that were identical to those of the previous experiment except that only responding on the left lever was rewarded during click and noise presentations.

Summation test. The animals were then given four test sessions in which responding during the click and the noise was examined; these stimuli were presented alone, in compound with the S^Δ with which they had been trained, and in compound with the other, transfer S^Δ. In each session there were 24 trials, 4 each of the following: click, noise, click + light, noise + dark, noise + light, and click + dark. Two of each of these types of trial were test trials during which no reinforcers were delivered. In the remaining trials reinforcers were delivered independent of responding, according to a VT10-s schedule. By maintaining the availability of reinforcement, these trials were intended to minimize generalization decrement from discrimination training and so maintain responding throughout the test sessions. However, to the extent that the delivery of reinforcement itself might act as an S^d, we anticipated that responding during these reinforced trials would represent only

a contaminated measure of stimulus control over responding; consequently, responding during these trials was not further analyzed.

Discrimination training reminder. At this point, subjects were given four additional sessions to reestablish discrimination performance.

Retardation test. Finally, the animals were given eight further test sessions. These were identical to the test sessions administered in Experiment 1, except that the animals received six different types of trial: click, noise, click + light, noise + dark, click + dark, and noise + light; operation of the magazine flap was rewarded according to an FR4 schedule during all stimulus presentations throughout these sessions.

Results and Discussion

We pooled response rates over click and noise trials and over compound trials with the light and the dark, as in Experiment 1.

Discrimination training

Preliminary training and discrimination training proceeded normally. The response rate during the S^d s in the last discrimination training session was 53.77 responses per minute; the corresponding pretrial response rate for this session was 3.11 responses per minute. Discrimination performance was evaluated using a ratio score as in the previous experiment; this yielded a ratio of .15, confirming that the S^A s had acquired substantial inhibitory control over responding.

Summation Test

Responding during the S^d s was maintained during the test sessions; the mean response rate was 66.80 responses per minute. Suppression ratios were computed for each animal in each test session—one for the S^d s in compound with the S^A s with which the animals had been trained and a second for the S^d s in compound with the transfer S^A s. The resulting scores were pooled across sessions to produce two scores for each animal. These data are shown in Figure 2. It is clear that, although the S^A s suppressed responding, these stimuli were no less effective when they were presented in compound with the transfer S^d s than when they were presented with the S^d s with which they had been trained. This interpretation was supported by the results of an ANOVA that compared the suppression ratios during the original and the transfer compounds ($F < 1$).

Discrimination Training Reminder

The animals continued to perform accurately during these sessions; the mean suppression ratio during the final session of this stage was .20.

Retardation Test

The rates of responding in the four test sessions, during the S^d s alone and during the S^d s in compound with the




Figure 2. Mean suppression ratios in the summation test sessions of Experiment 2. $S^d = S^d$ = discriminative stimulus; S^A delta = S^A = discriminative inhibitor.

original and the transfer S^A s, are shown in Table 3. Although the differences are numerically small, responding during the S^d s was reduced by addition of the S^A s, and the transfer S^A was no less effective than the original S^A at producing this inhibitory effect. This interpretation was supported by the results of an ANOVA with session and trial type (S^d , S^d + original S^A , or S^d + transfer S^A) as variables, which revealed a significant effect of session, $F(3, 45) = 21.65$, and of trial type, $F(2, 30) = 7.25$; but no interaction between these two variables, $F(6, 90) = 1.68$. Subsequent examination of the significant main effect of trial type with a Newman-Keuls test revealed that responding was significantly higher during the S^d s when they were presented alone than when they were presented in compound with either the original or the transfer S^A s, $p < .01$; however, responding during these two types of compound trial did not differ. Pretrial responding remained low during these sessions—5.63, 6.07, 7.86, and 6.43 responses per minute on the first, second, third, and fourth test sessions, respectively.

The results of this experiment confirm those of Experiment 1 in suggesting that the effect of an S^A is not specific to a particular S^d . Moreover, the results of Experiment 2 improve on those of Experiment 1 in a number of respects. First, the within-subjects design ensured that the results were uncontaminated by the between-group differences in overall response rate that were observed in Experiment 1. Second, the retardation test results not only replicate the

Table 3
Response Rates (in Responses per Minute) of the Retardation Test of Experiment 2

Stimulus	Session			
	1	2	3	4
S^d	18.09	30.02	33.12	37.24
S^d + original S^A	15.97	28.78	32.04	37.00
S^d + transfer S^A	15.73	30.38	31.48	36.54

Note. S^d = discriminative stimulus; S^A = discriminative inhibitor.

findings of Experiment 1 but also improve on them. Although in Experiment 1 there was no direct evidence that the S^A s had an inhibitory effect on performance of the magazine entry response, the present experiment provided such evidence. Despite the fact that the S^A s had a demonstrable inhibitory effect on performance during the S^d s, however, there was no sign that this effect was diminished when the S^A s were presented with the transfer S^d s. One could still criticize these data, however. First, one could argue that although the inhibitory effect of the S^A s was highly significant, it was nevertheless too small to allow us to discern any attenuation of that effect. Second, one could argue that because no untreated control stimulus was included in the experiment, there was no assurance that the inhibition that was observed was not due to generalization decrement induced in the S^d s. These criticisms cannot be leveled at the results of the summation test, however. In this test, the S^A s were tested with the response-reinforcer association with which they had been trained. This meant that they exerted a substantial inhibitory effect on responding, so there should have been no difficulty in observing an attenuation of that effect. Moreover, there is good evidence from other experiments that the inhibitory effect of S^A s generated with such procedures is not to be attributed to generalization decrement (Bonardi, 1988a, 1988b). Nevertheless, it was again demonstrated that the efficacy of the S^A s did not vary as a function of the S^d s with which they were compounded.

These results are difficult to reconcile with the idea that the inhibitory properties of an S^A are specific to the S^d with which it is trained. Two hypotheses remain, (a) that the S^A independently inhibits the representations of both response and reinforcer, or (b) that it acts on the response-reinforcer association in its entirety. The latter account predicts that the properties of an S^A will be sensitive to specific response-reinforcer combinations (cf. Jenkins, 1985), whereas the former does not. In Experiments 3 and 4, we exploited these differing predictions to discriminate between the two accounts.

Experiment 3

In Experiment 3, all animals were initially taught that performance of response R1 would be rewarded with a particular reinforcer, Rf1, in the presence of a discriminative stimulus. Then, an S^A was established that signaled that R1 would no longer be rewarded with Rf1. The effect of this S^A on responding produced by the original R1-Rf1 association was then compared with its effect on responding produced by three other associations: the one between R1 and a different reinforcer, Rf2; the one between a different response, R2, and the original reinforcer Rf1; and the one between a different response and a different reinforcer, R2 and Rf2. On the basis of the results reported by Bonardi (1989) and Colwill (1991), we anticipated that changing either the response (R2-Rf1) or the reinforcer (R1-Rf2) would result in a decrement of the S^A 's inhibitory properties relative to its effect on the original R1-Rf1 association. The question of interest was the effect it would have on the

R2-Rf2 association. If the S^A acts independently on both the response and the reinforcer, then changing both should have an additive effect and should result in a greater decrement of the S^A 's inhibitory power than changing either one. According to this account, the S^A should be significantly less effective at inhibiting responding that is based on the R2-Rf2 association than inhibiting responding that is based on either the R1-Rf2 or the R2-Rf1 associations. If the S^A acts on the R1-Rf1 association in its entirety, however, then it is immaterial whether one, the other, or both components of that association are changed. Changing both components should have no greater impact on the S^A 's inhibitory properties than should changing just one, because in each case the association is no longer the same, and the S^A should lose its power accordingly. In other words, this account predicts that the S^A , although effective with the R1-Rf1 association, should be equally bad at suppressing responding based on the other three response-reinforcer associations.

In this experiment, all animals initially received identical discriminated operant training in which performance of response R1 (responding on either the left or the right lever) was rewarded with a particular reinforcer Rf1 (food or sucrose pellets) during presentations of a click. They then received training in which the light was established as an S^A ; when the click was presented in compound with the light, responding was no longer reinforced. Once the subjects had mastered this discrimination they were divided into four groups (see Table 4). All groups received further discriminated operant training with a second S^d , a noise. The nature of the response-reinforcer association operative during noise presentations varied among the four groups. Subjects in Group SS (same-same) were, as before, rewarded with Rf1 for performing R1; those in Group DS (different-same) were also rewarded with Rf1, but for mak-

Table 4
Design of Experiment 3

Group	Stage 1	Stage 2	Test
SS	Click (R1 → Rf1) Click + light (R1 → 0)	Noise (R1-Rf1)	Noise (R1)? Noise + light (R1)?
SD	Click (R1 → Rf1) Click + light (R1 → 0)	Noise (R1-Rf2)	Noise (R1)? Noise + light (R1)?
DS	Click (R1 → Rf1) Click + light (R1 → 0)	Noise (R2-Rf1)	Noise (R2)? Noise + light (R2)?
DD	Click (R1 → Rf1) Click + light (R1 → 0)	Noise (R2-Rf2)	Noise (R2)? Noise + light (R2)?

Note. In Stage 1, R1 was rewarded in the presence of the click with Rf1; responding was not rewarded when the clicker was accompanied by the light. In Stage 2, one of the two responses was rewarded with one of the two reinforcers in the presence of the noise. In the Test phase, performance of the response trained in Stage 2 was compared during the noise and the noise-light compound. SS = same-same; SD = same-different; DS = different-same; DD = different-different; R1 and R2 = right- and left-lever responses (counterbalanced); Rf1 and Rf2 = food and sucrose reinforcers (counterbalanced).

ing the other response, R2. Subjects in Groups SD and DD were rewarded with the second reinforcer, Rf2, Group SD for performing R1, and Group DD for performing R2. Finally, the animals were given a summation test in which the effect of the light on responding during the noise was evaluated. We predicted that the light would be more suppressive in Group SS than in the other three groups, but the critical group was Group DD. If the light's inhibitory power were exerted separately on response and reinforcer representations, then it should be less effective in Group DD than in either Group SD or Group DS; but if the light acts on the response-reinforcer association as a whole, then there should be no difference among these three groups.

Method

Subjects

The experiment was run in two replications, each using 16 naive male hooded Lister rats. The subjects in the first replication had a mean ad lib weight of 315 g (range: 300–350 g) and those in the second replication had a mean ad lib weight of 329 g (range: 307–357); all animals were maintained as in the previous experiment.

Apparatus

The apparatus consisted of two Campden Instruments operant chambers identical to those used in Experiment 1 except in the following respects. Each of the boxes had a ceiling of sheet aluminium and was equipped with two pellet dispensers, one of which delivered 45-mg mixed-composition food pellets, the other of which delivered 45-mg sucrose pellets. Both types of pellet were delivered into the same recessed food tray. Illumination was provided by a 2.8-W 15-V houselight, which was mounted in the front wall, directly above the food tray. A light stimulus was provided by turning on three 2.8-W 24-V jewel lights, which were placed above and to the left and right of the food tray. Two speakers were mounted on the back wall through which an 80-dB, 10-Hz click and an 85-dB white noise (scale A) could be delivered from a Campden Instruments tone generator (Model 258) and white noise generator (Model 530).

Procedure

All aspects of the procedure that are not specified were the same as those of Experiment 2.

Preliminary training. In the first 30-min session the animals were trained to retrieve food and sucrose pellets from the food tray. A VT60-s schedule, in which a pair of pellets, one food and one sucrose, were delivered on an average of once per minute, operated throughout the session.

From this point on, one of the two levers was always available. For the rest of preliminary training and throughout discrimination training this was always the same lever; a response on this lever was designated R1. For half of the animals, R1 was made on the left lever and for the other half it was made on the right. R1 was always rewarded with the same reinforcer, Rf1, during this period. For half of the animals for whom R1 was the left lever and for half of those for whom it was the right, Rf1 was food; for the remaining subjects, Rf1 was sucrose.

In the next session animals were consistently reinforced for responding, and then all subjects received a stimulus preexposure session identical to that of Experiment 2.

Preliminary VR training. In this stage the animals were trained to respond only during the discriminative stimulus, the click. Rf1 delivery was contingent on performance of response R1 according to a VR schedule during click presentations. In the first session of this stage there were thirty-six 20-s presentations of the click during which responding was rewarded according to a VR2 schedule; there was no pretrial period in this session. For the first 12 trials the ITI was 11 s, for the second 12 trials it was 21 s, and for the final 12 trials it was 41 s. The next six sessions each consisted of 24 trials, each with a 20-s pretrial period and an ITI of mean duration 80 s (range: 50–110 s). The first of these sessions consisted of 12 VR2 trials followed by 12 VR3 trials, the second session consisted of 12 VR3 trials followed by 12 VR4 trials, the third session consisted of 12 VR4 trials followed by 12 VR5 trials, and the fourth session consisted of 12 VR5 trials followed by 12 VR6 trials. The final two sessions of this stage consisted entirely of VR6 trials.

Discrimination training. Eight sessions followed, in which the light was established as an S^A , signaling that performance of R1 would no longer be followed by Rf1. Each session consisted of 24 trials, 18 with the click alone and 6 with the click and the light in a simultaneous compound.

Transfer S^d training. In this stage, all animals were given discriminated operant training with a new, transfer S^d , the noise. The animals were divided into four groups, which differed in the identity of the response and the reinforcer used in this stage. The counterbalancing was such that for half of each group R1 was a left lever response and for the other half it was a right lever response. For half of each of these subgroups Rf1 was food and for the other half it was sucrose. Groups SS and SD continued to be rewarded for making R1; Group SS was rewarded with Rf1, but Group SD was rewarded with the other reinforcer, Rf2. For Groups DS and DD the lever corresponding to R1 was retracted and replaced with the alternative lever, on which animals could make R2. Subjects in Group DS were rewarded as before with Rf1, whereas subjects in Group DD were rewarded with the alternative reinforcer, Rf2. In other respects, this stage was identical to the preliminary VR training stage, except that it commenced with two 36-trial sessions rather than with one.

Discrimination training reminder. Before the test, subjects were given reminder sessions of discrimination training to ensure that they were still performing accurately. In the first replication there were two sessions in this stage. In the second replication one animal did not discriminate in the second session, so, in an attempt to rectify this, all subjects received a third session. However, the subject in question (in Group SS) still failed to discriminate in this extra session and thus was dropped from the experiment. This was the only respect in which the replications differed.

Summation test. Finally, the animals were given five test sessions in which the effect of the light on responding during the transfer S^d was evaluated. The manipulanda present during these sessions were those from transfer S^d training—the lever corresponding to R1 for Groups SS and SD and that corresponding to R2 for Groups DS and DD. Similarly, when reinforcers were delivered, they were also those used during transfer S^d training—Rf1 for Groups SS and DS, Rf2 for Groups SD and DD. In each session there were 24 trials, 12 with the noise and 12 with the noise-light compound. Six of each of these types of trial were test trials during which no reinforcers were delivered, and in the remainder of the trials reinforcers were delivered independent of

responding; these sessions were otherwise identical to those of Experiment 2.

Results and Discussion

Transfer S^d Training

Subjects readily learned to make the appropriate response during the noise in this stage; this stimulus acquired good discriminative control over responding that did not appear to differ among the groups. In the last session the response rates, in responses per minute, for Group SS, Group SD, Group DS and Group DD were 30.07, 31.86, 32.97, and 30.00, respectively. The corresponding pretrial response rates for this session for Group SS, Group SD, Group DS, and Group DD were 3.45, 2.14, 1.69, and 1.99, respectively. A factorial ANOVA performed on these data with response, reinforcer, and S^d (whether the noise was present or absent) as variables revealed a significant main effect of S^d , $F(1, 27) = 115.92$; no other effects or interactions were significant ($F_s < 1$). Thus there was no evidence for any differences among the groups in responding during the transfer S^d , thus allowing it to provide an unbiased baseline of responding from which to estimate the efficacy of the light.

Discrimination Training Reminder

Discrimination training proceeded normally, and this performance was maintained in the reminder sessions. All subjects responded more in the presence of the click than in its absence, and this did not differ among the groups. The response rates during the click in the last session of this stage, in responses per minute, for Group SS, Group SD, Group DS, and Group DD were 35.67, 52.92, 36.72, and 30.08, respectively. The corresponding pretrial response rates for this session, for Group SS, Group SD, Group DS, and Group DD were 2.26, 3.46, 2.15, and 1.21, respectively. A factorial ANOVA with response (whether the response was changed between training and test), reinforcer (whether the reinforcer was changed between training and test) and S^d (whether the S^d was present or absent) as variables revealed a significant main effect of S^d , $F(1, 27) = 136.54$; no other effects or interactions were significant, largest $F(1, 27) = 3.99$.

As the response rates during the click did not differ among the groups, we evaluated discrimination performance by calculating a ratio score. This yielded ratios of .132 for Group SS, .180 for Group SD, .094 for Group DS, and .134 for Group DD. A factorial ANOVA performed on these data, with response and reinforcer as variables, revealed that the groups did not differ, largest $F(1, 27) = 1.19$. Thus there was no evidence for any difference among the groups in the efficacy of the S^A at this stage, suggesting that any group differences in the test trials could be safely attributed to differences in the response–reinforcer contingency operating during the transfer S^d .

Test

Response rates, in responses per minute, to the noise were pooled over the test trials for all five test sessions and were similar among the four groups. The mean rates for Group SS, Group SD, Group DS, and Group DD were 29.96, 29.67, 28.85, and 18.99, respectively, and the corresponding pretrial response rates were 1.83, 6.05, 1.50, and 2.60, respectively. A factorial ANOVA performed on these data with response, reinforcer, and S^d as variables revealed a significant main effect of S^d , $F(1, 27) = 134.28$; no other effects or interactions were significant, largest $F(1, 27) = 3.22$.

As there was no evidence that the groups differed in their rates of responding during the noise alone, we evaluated the inhibitory power of the light using ratio scores—a ratio was computed for each animal for each test session, and the resulting scores were pooled to produce a single ratio for each subject. The resulting data are shown in Figure 3. It is clear that there was more suppression in Group SS than in any of the other three groups, as predicted. Moreover, the light did not appear to be more effective in Group DD than in Groups SD and DS; although there was more suppression in Group DD than in Group SD, there was slightly less than in Group DS. However, a factorial ANOVA with response and reinforcer as variables revealed a main effect of response, $F(1, 27) = 12.01$; but no effect of reinforcer, $F(1, 27) = 1.14$, and no interaction between these two variables ($F < 1$).

The results of this analysis suggest that changing the response had an effect on the S^A 's inhibitory properties, whereas changing the reinforcer did not. If this were the case, then it is not surprising that changing both response and reinforcer did not have a greater effect than changing either one. However, other experiments have succeeded in finding an attenuation of an S^A 's effects when the reinforcer is changed (Bonardi, 1989; Colwill, 1991). One possible explanation for this discrepancy is that, in contrast to these other studies, the two reinforcers were experienced in successive phases of training in the present experiment. By minimizing the animals' opportunity to compare the two,




Figure 3. Mean suppression ratios for the four groups in the test sessions of Experiment 3. SS = same–same; DS = different–same; SD = same–different; DD = different–different.

this might have made the change of reinforcer less salient for them and resulted in less of an effect on their behavior. But there did, in numerical terms, seem to be some effect of reinforcer change; the data in Figure 3 suggest that Group SD was less suppressed than Group SS. To investigate this apparent difference further, we performed a second ANOVA, with group as a variable; this revealed a significant main effect of group, $F(3, 27) = 4.63$. Newman-Keuls tests revealed that, although Group SD did not differ from Group SS, it did not differ from the other groups either. Group SS differed from both Groups DS and DD, which did not differ from each other.

The results of this second ANOVA provide partial support for the hypothesis that the S^A acts on the entire response-reinforcer association. Groups SD, DS, and DD did not differ, suggesting that the light was equally poor at suppressing instrumental responding in these three groups, as predicted. Indeed, the light in Group DS was, if anything, slightly less suppressive than that in Group DD. But although they did not differ statistically, Group SD was numerically less suppressed than Group DD; moreover, Group SD did not differ from Group SS.

Two points should be made concerning the present experiment. The first is simply that ceiling effects may have made it difficult to observe greater attenuation of inhibition in Group DD than in Group DS. The second point is more subtle. The design of the experiment was such that the critical difference among the four groups was the extent to which the response-reinforcer contingency was altered between S^A training and transfer S^d training. However, the design ensured that this change in contingency was confounded with whether the response or reinforcer components of the transfer S^d contingency were novel, because changes in the response-reinforcer contingency were achieved by introducing novel elements into that contingency. Thus, for Group SS both response and reinforcer were familiar, for Group DD both were novel, and for the remaining groups one was novel and one familiar. Although

it is not clear how this might account for the results we observed, it would clearly be more compelling to replicate this finding under conditions in which no such confound is present.

Experiment 4

The design of Experiment 4 was based on a suggestion made by Jenkins (1985; see also Holland, 1989). All animals were initially trained to perform two responses, R1 and R2, for two reinforcers, Rf1 and Rf2 (see Table 5). R1 was always reinforced with Rf1, and R2 was always reinforced with Rf2. Then, an S^A was established, which signaled both that R1 would no longer be followed by Rf1, and that R2 would no longer be followed by Rf2. Finally, the animals were divided into two groups and were trained with a new, transfer S^d . Group S animals were rewarded for performing R1 with Rf1 or for performing R2 with Rf2, exactly as they had been during initial training. Group D animals were trained with responses and reinforcers in different combinations, so that R1 was now reinforced with Rf2, or R2 with Rf1. The effect of the S^A on responding during the transfer S^d was then assessed in a summation test, exactly as in Experiment 2.

During initial training the light signaled that R1 would no longer be followed by Rf1 and that R2 would no longer be followed by Rf2. So, if the S^A 's inhibitory power is the product of two separate effects—one on the response and one on the reinforcer—then the light will have acquired inhibitory power over both R1 and R2, and both Rf1 and Rf2 in this stage. Thus, it is immaterial how these responses and reinforcers are combined during the transfer S^d s—the light should be equally good at suppressing responding in both groups. But if the S^A acts on the entire response-reinforcer association, the prediction is rather different. According to this account, the light specifically inhibits the association between R1 and Rf1 and the association be-

Table 5
Design of Experiment 4

Group	Stage 1	Stage 2	Test
S	Click (R1 → Rf1) Click + light (R1 → 0)	Pulsed tone (R1-Rf1) or Pulsed tone (R2-Rf2)	Pulsed tone (R1)? Pulsed tone + light (R1)? or Pulsed tone (R2)? Pulsed tone + light (R2)?
	Tone (R2 → Rf2) Tone + light (R2 → 0)	Pulsed tone (R1-Rf2) or Pulsed tone (R2-Rf1)	Pulsed tone (R1)? Pulsed tone + light (R1)? or Pulsed tone (R2)? Pulsed tone + light (R2)?
D	Click (R1 → Rf1) Click + light (R1 → 0)	Pulsed tone (R1-Rf2) or Pulsed tone (R2-Rf1)	Pulsed tone (R1)? Pulsed tone + light (R1)? or Pulsed tone (R2)? Pulsed tone + light (R2)?
	Tone (R2 → Rf2) Tone + light (R2 → 0)	Pulsed tone (R1-Rf2) or Pulsed tone (R2-Rf1)	Pulsed tone (R1)? Pulsed tone + light (R1)? or Pulsed tone (R2)? Pulsed tone + light (R2)?

Note. In Stage 1, R1 or R2 was rewarded in the presence of the click or the tone with either Rf1 or Rf2; responding was not rewarded when either the clicker or the tone was accompanied by the light. In Stage 2, one of the two responses was rewarded with one of the two reinforcers in the presence of the pulsed tone. In the test phase, performance of the response trained in Stage 2 was compared during the pulsed tone and the pulsed tone-light compound. S = same; D = different; R1 and R2 = right- and left-lever responses (counterbalanced); Rf1 and Rf2 = food and sucrose reinforcers (counterbalanced).

tween R2 and Rf2; therefore, there is no reason why it should have any inhibitory effect on responding based on the new associations between R1 and Rf2 and R2 and Rf1 that are signaled by the transfer S^d in Group D. Thus, the prediction of this account is that there should be more suppression in Group S than in Group D.

As well as providing further support for the interpretation given of the results of Experiment 2, the present design avoids the problem with that experiment—that the differences between the groups' treatment was confounded with the novelty of the elements of the response-reinforcer contingency used during transfer S^d training. A difference between Group S and Group D in the present experiment could not be explained in this way, because both groups had equal experience of both the responses and the reinforcers that were used in the final test.

Method

Subjects and Apparatus

The subjects were 16 naive male hooded Lister rats with a mean ad lib weight of 355 g (range: 331–394 g). They were housed in pairs and maintained exactly as in Experiment 1. The stimuli were an 87-dB, 1-kHz tone, an 87-dB, 10-Hz click, an 84-dB white noise, and an 83-dB 2-kHz tone pulsed at 1 Hz, light (as in Experiment 2) and dark (offset of the houselight). All other aspects of the apparatus were identical to those of Experiment 3.

Procedure

All aspects of the procedure that are not specified were the same as those in Experiment 1.

Preliminary training. After magazine training, the animals were trained to respond on the two levers. As in Experiment 1, sessions in preliminary training and preliminary VR training were scheduled in duplicate pairs, one with each lever. For half of the animals, responding on the left lever was followed by food, and responding on the right was followed by sucrose. For the remaining subjects, these contingencies were reversed; these reinforcement arrangements were maintained until the start of transfer S^d training.

In the first pair of sessions, animals were consistently reinforced for responding, and in each of the second pair they experienced one presentation each of the click, tone, light, and dark, and compound presentations of noise + tone, noise + click, noise + light, and noise + dark.

Preliminary VR training. In this stage, half of the animals were rewarded for responding on the left lever during the tone and for responding on the right lever during the click; for the remaining animals, these contingencies were reversed. Half the animals in each of these subgroups were rewarded for responding on the left lever with food, and for responding on the right with sucrose; for the other half of each subgroup, these contingencies were reversed. These sessions were identical to the preliminary VR training sessions of Experiment 1 except that, first, the first three sessions (those with no pretrial period) were amalgamated into a single 36-trial session with each lever, and second, animals received only one pair of sessions with a VR3 schedule and one with a VR4 schedule.

VR discrimination training. Animals then received eight sessions of VR discrimination training, identical to those of Experiment 1.

Discrimination training. Twelve sessions of discrimination training followed. In the first session, the noise was used as an S^A . This session consisted of 24 trials—9 with the clicker alone, 9 with the tone alone, 3 with the noise + tone compound, and 3 with the click + tone compound. However, in this session animals did not respond on the appropriate lever during compound trials, suggesting that the added noise induced too much generalization decrement for the animals to be able to identify the click and the tone. Therefore, we discontinued training with the noise. Animals were given two more sessions of VR discrimination training with a VR6 schedule, and then returned to discrimination training exactly as before, except that the noise was replaced by the light.

Transfer S^d training. At this point, the animals were given training with the transfer S^d , the pulsed tone. The animals were divided into two groups; for half of each group responding on the left lever had been reinforced during the tone and responding on the right lever had been reinforced during the click, and for the remaining subjects the opposite had been true. Likewise, for half of these subgroups, responding on the left lever had been rewarded with food and responding on the right with sucrose, and for the remaining animals the opposite had been true.

From the start of this stage until the end of the experiment, only one lever was available. For half of each of the four counterbalanced subgroups in each experimental group, responding on the left lever was reinforced during this stage and for the remaining animals responding on the right lever was reinforced. The reinforcer during this stage depended on the group the subject was in. Thus, of the animals trained to make R1 in this stage, those in Group S were rewarded with Rf1 and those in Group D were rewarded with Rf2. Likewise, of the animals trained to make R2 during this stage, those in Group S were rewarded with Rf2 and those in Group D were rewarded with Rf1.

The first session consisted of thirty-six 20-s presentations of the pulsed tone and was identical to the first session of preliminary VR training in all other respects. The next six sessions each consisted of 24 trials; the first of these sessions consisted of 12 VR2 trials followed by 12 VR3 trials, the second session of 12 VR3 trials followed by 12 VR4 trials, the third session of 12 VR4 trials followed by 12 VR5 trials, and the fourth session of 12 VR5 trials followed by 12 VR6 trials. The final three sessions of this stage consisted entirely of VR6 trials. In other respects these sessions were identical to the remaining sessions of preliminary training.

Discrimination training reminder. Subjects were then given two reminder sessions of discrimination training to ensure discrimination performance was maintained.

Summation test. Finally, the animals received four test sessions identical to those of Experiment 1 except that the noise was replaced by the pulsed tone.

Results and Discussion

Inspection of the data revealed that response rates during the click and the tone were very similar throughout the experiment, and so were pooled in all the analyses below.

Transfer S^d Training

Preliminary training and discrimination training proceeded normally. In the last session of transfer S^d training there was no sign of any difference between the groups; in both groups the pulsed tone maintained a similarly strong control over responding. The response rates, in responses

per minute, during the pulsed tone for Group S and Group D were 42.83 and 49.50, respectively, and the corresponding pretrial rates for these groups were 0.22 and 1.27 respectively. An ANOVA with group and S^d (whether the pulsed tone was present or absent) as variables revealed a significant main effect of S^d , $F(1, 14) = 88.29$, that did not interact with group ($F < 1$). The main effect of group was not significant ($F < 1$).

Discrimination Training

Response rates, in responses per minute, during the last reminder session were 45.39 for Group S and 41.81 for Group D during the S^d s and 0.21 for Group S and 0.65 for Group D during the corresponding pretrial periods. An ANOVA performed on these data with group and S^d (whether the S^d was present or absent) as variables revealed a significant main effect of S^d , $F(1, 14) = 128.98$, which did not interact with group ($F < 1$); the main effect of group was not significant ($F < 1$).

As before, we evaluated discrimination performance with a ratio measure. The resulting scores were .080 for Group S and .078 for Group D; an ANOVA revealed that these scores did not differ ($F < 1$). Thus, there was no sign that the groups differed in their discrimination performance.

The rates of responding on the incorrect lever were low during this session. On S^d trials the response rates, in responses per minute, were 2.29 for Group S and 2.19 for Group D, and, on compound trials, 1.00 for Group S and 2.19 for Group D. An ANOVA on these data with group and trial type (S^d or S^d + light trials) as variables revealed no significant effects or interactions ($F_s < 1$).

Test

We pooled response rates to the pulsed tone over the four test sessions, exactly as in Experiment 2. Inspection of these data revealed that response rates were higher on the right lever than on the left lever, and that this appeared to influence discrimination performance. As it seemed likely that this variable would contribute to the variance in the data, lever (left or right) was included as a variable in analysis of these test data. The response rate, in responses per minute during the pulsed tone were 33.86 for Group S (40.97 on the right lever and 26.75 on the left) and 31.43 for Group D (35.22 on the right lever and 27.63 on the left); the corresponding pretrial rates were 0.20 for Group S and 0.88 for Group D. An ANOVA with group, lever, and S^d (whether the pulsed tone was present or absent) as variables revealed a significant main effect of S^d , $F(1, 12) = 112.21$; which did not interact with group ($F < 1$); the main effect of lever was not significant, $F(1, 12) = 3.05$; nor was the interaction between lever and S^d , $F(1, 12) = 3.42$. Nothing else was significant ($F_s < 1$). Thus, there was no evidence that the groups differed in their rates of responding during the pulsed tone.

We evaluated the inhibitory power of the light using the ratio measure, exactly as in Experiment 2. One of the

animals in each group failed to respond at all during the pulsed tone in one of the four test sessions, making it impossible to evaluate the effect of the light. The score for this session was thus replaced by that subject's average score from the preceding session and the subsequent session. The resulting data are shown in Figure 4; it is clear that there was more suppression in Group S than in Group D. This was supported by the results of an ANOVA with group and lever as variables, which revealed a significant main effect of group, $F(1, 12) = 6.60$. There was also a main effect of lever, $F(1, 12) = 9.61$; this seemed to reflect the fact that ratios were lower for animals tested on the right lever—.129 for animals in Group S and .195 for those in Group D, as opposed to .218 for Group S and .369 for Group D for subjects that had been tested on the left lever. It seems clear, however, that the effect was present in both these subgroups; this was confirmed by the fact that the Group \times Lever interaction was not significant, $F(1, 12) = 1.02$. The results of this experiment provide further support for the hypothesis that the S^A acts on the response-reinforcer association as a whole rather than acting separately on its components. If the S^A inhibited the representations of the response and the reinforcer independently, there would have been no difference between the two groups during the test.

General Discussion

The results of the present experiments allow a more precise specification of the way in which an S^A acts. The results of Experiments 1 and 2 suggested that S^A s do not act on the representation of the S^d , whereas Experiments 3 and 4's results lent support to the idea that they act on the response-reinforcer association as a whole. These findings have implications for possible theoretical accounts of S^A function.

It appears that S^A s do not depend on simple associations with the presence or absence of reinforcement for their control over instrumental responding (e.g., Bonardi, 1988a, 1988b; Goodall & Mackintosh, 1987). In this respect, they bear a resemblance to Pavlovian occasion setters. Operationally speaking, these are the Pavlovian counterparts of




Figure 4. Mean suppression ratios for the two groups in the test sessions of Experiment 4. S = same; D = different.

discriminative stimuli and inhibitors. Just as an S^d signals that a response will be reinforced and an S^{Δ} signals that it will not be reinforced, so a positive occasion setter signals that a Pavlovian conditioned stimulus (CS) will be reinforced, and a negative occasion setter signals that reinforcement of that CS will be omitted. Moreover, occasion setters acquire the ability to modulate Pavlovian conditioned responding exactly as S^d s and S^{Δ} s modulate instrumental responding. Finally, both positive and negative occasion setters, in the same way as S^d s and S^{Δ} s, do not depend on simple classically conditioned associations for their control over behavior (e.g., Holland, 1984; Rescorla, 1985; Ross & Holland, 1981). Given these striking parallels, it seems likely that the same theory will govern both classes of stimuli.

Occasion setting has attracted more attention in recent years than has discriminative control, and as a result, a number of different theories of occasion setting have been specified, all of which may be applied to discriminative stimuli and inhibitors. For example, Rescorla (1985) suggested that occasion setters exert their effects by modulating the degree to which the reinforcer representation may be activated. In contrast, Holland (1983) argued that the occasion setter acts on the entire CS–reinforcer association, modulating the degree to which it can be used or retrieved (see also Bonardi, 1989; Bouton, 1990). The results of the present experiments are clearly more compatible with the second of these accounts in that they imply that the entire association is modulated, rather than one of its components.

Bonardi (1989) elaborated Holland's (1983) account by specifying that the way in which occasion setters acquire the ability to modulate CS–unconditioned stimulus (US) associations is through a process of associative learning. Bonardi argued that the CS–US pairing signaled by the occasion setter may be regarded as a complex event that can enter into associations in exactly the same way as can CS and US alone. She went on to suggest that the occasion setter's properties arise through the occasion setter becoming associated with the CS–US pairing. In these terms, a positive occasion setter can be regarded as a conditioned excitor for the CS–US pairing, and a negative occasion setter can be regarded as a conditioned inhibitor for that same event. This account has been explicitly applied to discriminative control (e.g., Bonardi, 1988a; Mackintosh, 1983; Rescorla, 1990a, 1990b), with some success.

The results of the present experiments agree to some extent with this view. First, the idea that the S^{Δ} is specific to the response–reinforcer association is clearly consistent with the notion that it is acting as a conditioned inhibitor for that event. But there is little evidence that Pavlovian conditioned inhibitors convey precise information about the USs whose omission they signal (but see, e.g., Kruse, Overmier, Konz, & Rokke, 1983). For example, an inhibitor for one type of shock will transfer readily to a second type of shock (Pearce, Montgomery, & Dickinson, 1981). Indeed, one major theory of inhibitory learning asserts that a Pavlovian inhibitor conveys no information about the sensory properties of the reinforcer whose omission it predicts; instead, the inhibitor relies on being associated

with the affective properties of reinforcer omission for its properties (Dickinson & Dearing, 1979). In contrast, in the present experiments the S^{Δ} conveyed a great deal of information about the response–reinforcer contingency whose omission it signaled. It remains to be seen whether this discrepancy is the product of insensitivity in the Pavlovian procedures or whether it implies a more fundamental theoretical inconsistency.

There still are certain things that remain unspecified with this analogy between S^{Δ} s and Pavlovian inhibitors. For example, Konorski (1948, 1967) argued that a Pavlovian inhibitor could act in two ways: It could be associated with the reinforcer representation through an inhibitory link, or it could be involved in an excitatory association with a representation of reinforcer absence. The same argument could be applied to S^{Δ} s: Are they involved in an inhibitory association with the response–reinforcer association, or does the animal have a representation for the fact that the response is no longer followed by food, and it is with this that the S^{Δ} becomes associated? Clearly, further work is required to decide among these possibilities.

Whatever the outcome of such inquiries, investigation of the parallels between simple associative learning and the formation of occasion setters and discriminative stimuli might be a fruitful strategy, not only because it might give us information about how such modulatory stimuli operate, but also because it could give us a deeper understanding of association formation itself.

References

- Bonardi, C. (1988a). Associative explanations of discriminative inhibition effects. *Quarterly Journal of Experimental Psychology*, *40B*, 63–82.
- Bonardi, C. (1988b). Mechanisms of inhibitory discriminative control. *Animal Learning & Behavior*, *16*, 445–450.
- Bonardi, C. (1989). Inhibitory discriminative control is specific to both the response and the reinforcer. *Quarterly Journal of Experimental Psychology*, *41B*, 225–242.
- 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.
- Brown, P. L., & Jenkins, H. M. (1967). Conditioned inhibition and excitation in operant discrimination learning. *Journal of Experimental Psychology*, *75*, 255–266.
- Colwill, R. M. (1991). Negative discriminative stimuli provide information about the identity of omitted response-contingent outcomes. *Animal Learning & Behavior*, *19*, 326–336.
- Dickinson, A., & Dearing, M. F. (1979). Appetitive–aversive interactions and inhibitory processes. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation* (pp. 203–231). Hillsdale, NJ: Erlbaum.
- Goodall, G., & Mackintosh, N. J. (1987). Analysis of the Pavlovian properties of signals for punishment. *Quarterly Journal of Experimental Psychology*, *39B*, 1–21.
- Hearst, E., & Peterson, G. B. (1973). Transfer of conditioned excitation and inhibition from one operant response to another. *Journal of Experimental Psychology*, *99*, 360–368.
- 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. (1984). Differential effects of reinforcement of an inhibitory feature after serial and simultaneous feature negative discrimination training. *Journal of Experimental Psychology: Animal Behavior Processes*, *10*, 461–475.
- 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.
- Holman, J. G., & Mackintosh, N. J. (1981). The control of appetitive instrumental responding does not depend on classical conditioning to the discriminative stimulus. *Quarterly Journal of Experimental Psychology*, *33B*, 21–31.
- 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.
- Konorski, J. (1948). *Conditioned reflexes and neuron organization*. Cambridge, England: Cambridge University Press.
- Konorski, J. (1967). *Integrative activity of the brain*. Chicago: University of Chicago Press.
- Kruse, J. M., Overmier, J. B., Konz, W. A., & Rokke, E. (1983). Pavlovian conditioned stimulus effects upon instrumental choice behavior are reinforcer specific. *Learning and Motivation*, *14*, 165–181.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford, England: Clarendon Press.
- 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. (1990a). 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. (1990b). The role of information about the response–outcome relation in instrumental discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*, 262–270.
- Rescorla, R. A., & Holland, P. C. (1977). Associations in Pavlovian conditioned inhibition. *Learning and Motivation*, *8*, 429–447.
- Ross, R. T., & Holland, P. C. (1981). Conditioning of simultaneous and serial feature-positive discriminations. *Animal Learning & Behavior*, *9*, 293–303.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.

Received July 8, 1993

Revision received January 10, 1994

Accepted March 4, 1994 ■