

Evidence for the hierarchical structure of instrumental learning

RUTH M. COLWILL and ROBERT A. RESCORLA
University of Pennsylvania, Philadelphia, Pennsylvania

In three experiments using rats, we examined the role of a discriminative stimulus (S) in governing the relation between a response (R) and an outcome (O) in an appetitive instrumental learning paradigm. In each experiment, we attempted to distinguish between a simple S-O association and a hierarchical relation in which S is associated with the R-O association. We used three variations on discriminative training procedures and three different assessment techniques for revealing the hierarchical structure. In Experiment 1, we employed a training procedure in which S signaled a change in the R-O relation but no change in the likelihood of O. Although such an arrangement should not produce an excitatory S-O association, it nevertheless generated an S that controlled responding and transferred that control to other responses. In Experiment 2, we used a discrimination procedure in which two Ss each had the same two Rs and Os occur in their presence but each S signaled that a different R-O combination would be in effect. This design provided the opportunity for equivalent pairwise associations among S, R, and O but unique hierarchical relations. The subjects learned the hierarchical structure, as revealed by the specific depressive effect of a subsequent lithium-chloride-induced devaluation of O on responding only in the presence of the S in which that response had led to that outcome. In Experiment 3, one S signaled two different R-O outcomes. Then, two new stimuli were presented with the original S; the R-O relations were retained in the presence of one of the added stimuli but were rearranged in the presence of the other. The added S came to control less responding when it was redundant with respect to the R-O relations than when it was informative. Although all of the results were of modest size and each has an alternative interpretation, together they provide converging evidence for the hierarchical role of S in controlling an R-O association.

In a typical instrumental learning situation, a response (R) produces a reinforcing outcome (O) in the presence of some discriminative stimulus (S). One consequence of such an arrangement is that the likelihood of the response increases in the presence of that stimulus. Three different accounts have been developed to explain the learned ability of the stimulus to increase the response. One simple possibility advocated by many traditional learning theories (e.g., Hull, 1943) is that an association forms between the stimulus and the instrumental response. According to this view, the reinforcing outcome serves a catalytic function; it promotes the development of the S-R association, thus enabling the stimulus to generate the instrumental response. A second possibility is suggested by classical two-process theorists (e.g., Rescorla & Solomon, 1967; Trapold & Overmier, 1972). According to their view, the occurrence of the reinforcer in the presence of the stimulus results in the development of a Pavlovian association between the stimulus and the outcome. That S-O association provides the stimulus with the power either

to motivate or to elicit the instrumental response. This is an elaboration of the classical r_g-s_g mechanism proposed by Hull and Spence (e.g., Spence, 1956). Finally, other theorists (e.g., Colwill & Rescorla, 1986; Mackintosh & Dickinson, 1979) have proposed that the stimulus becomes associated with the R-O relation (see also Skinner, 1938). According to this view, the stimulus operates on the R-O association, effectively setting the occasion upon which the response produces the outcome.

The intention of the experiments reported here was to provide some empirical evidence for the third, occasion-setting, account. In particular, the present experiments were aimed at distinguishing the contribution of a hierarchical relation between the stimulus and the R-O association from the contribution of a simpler S-O association. In these experiments, three variations on the standard instrumental learning paradigm and three different assessment techniques were used.

In Experiment 1, instrumental learning was studied under conditions designed to preclude the development of a Pavlovian S-O association but to allow the development of an S-(R-O) association. To assess the nature of the learning, the degree to which the stimulus controlled the performance of another response was measured in a transfer test. In Experiment 2, we employed a complex discrimination paradigm designed to encourage hierarchical encoding of the three elements, stimulus, response, and

This research was supported by National Science Foundation Grant BNS 83-08176. Correspondence concerning this article may be addressed to either Ruth M. Colwill, Department of Psychology, Brown University, Providence, RI 02912, or Robert A. Rescorla, Department of Psychology, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA 19104.

outcome. In this paradigm, no response or outcome was uniquely correlated with a stimulus; rather, each stimulus signaled which R-O combination obtained. The experiment assessed the degree of this learning by examining the impact of artificially devaluing the outcome. It is known that devaluing an outcome will specifically reduce the likelihood of responses with which it is associated (e.g., see Colwill & Rescorla, 1985). In Experiment 2, we asked whether each stimulus could activate a particular R-O association as revealed by specific depression of a particular response in its presence. In Experiment 3, we used a simple discrimination procedure in which a single stimulus signaled which outcome would follow a response. The control that stimulus had over the R-O association was assessed by its ability to block learning about another stimulus when they jointly signaled either the same or a different R-O relation.

EXPERIMENT 1

Colwill and Rescorla (1988) reported that a discriminative stimulus trained with an outcome would preferentially augment a new response subsequently trained with that same outcome. A new response trained with a different reinforcer was not affected by that stimulus. Such a result demonstrates that the stimulus provides fairly precise information about the reinforcer that is earned in its presence. Colwill and Rescorla offered several descriptions of the nature of this stimulus-reinforcer learning. One straightforward possibility is that the association is Pavlovian in nature. By virtue of the fact that the reinforcer occurs in the presence of the stimulus, a simple association may develop between the stimulus and the outcome. However, these circumstances also arrange for the stimulus to be present at a time when a particular relation obtains between the response and the outcome. Thus, a second possibility is that the stimulus is associated with the R-O relation rather than having separate associations with the individual elements. In the present experiment, we attempted to separate the hierarchical S-(R-O) view from a simple S-O account by training a stimulus under con-

ditions that are not conducive to the development of Pavlovian S-O associations. For this purpose, the outcome was programmed to occur with equal likelihood in the presence and absence of the stimulus, thus preventing the stimulus from signaling the occurrence of the outcome. However, the stimulus was arranged to signal a change in the R-O relation: in the absence of the stimulus, the outcome occurred freely, but in its presence, the response was required to earn the outcome. The question is whether, under these circumstances, the stimulus will nevertheless control the response and transfer that control to a new instrumental response trained with the same outcome.

The basic design is schematized in Figure 1. Rats were given discrimination training with two stimuli (light and noise), one response (nose poking), and two reinforcers (food pellets and liquid sucrose). In training sessions with one stimulus (S1), nose poking (Rc) was reinforced with pellets (O1) on a variable-interval (VI) schedule; in training sessions with the other stimulus (S2), nose poking was reinforced with sucrose (O2). In the periods between stimulus presentations, there were free deliveries of the reinforcer that was otherwise earned during the stimulus. The frequency of such intertrial reinforcers was controlled by a variable-time (VT) schedule whose value was equal to that of the VI operative during the stimulus, so that the stimulus bore no positive correlation to the reinforcer. Similar arrangements between a stimulus and an outcome are known to be effective in preventing the development of an excitatory Pavlovian association between the stimulus and the reinforcer in both appetitive Pavlovian conditioning (Durlach, 1983; Gamzu & Williams, 1973) and aversive Pavlovian conditioning (Rescorla, 1968). To ensure that the subject had general information about which reinforcer was available in a session, distinctive background cues (B1 and B2) were correlated with each outcome (and, consequently, with each stimulus). To achieve this, all sessions with the light stimulus contained an auditory background cue; all sessions with the noise stimulus contained a visual background cue. It is important to note that while this arrangement prevents the stimulus

Design of Experiment 1




Figure 1. Design of Experiment 1. (Rc, R1, and R2 are instrumental responses, nose poking, lever pressing, and chain pulling; S1 and S2 are discriminative stimuli, noise and steady light; O1 and O2 are reinforcers, sucrose and pellets; B1 and B2 are background stimuli, flashing light and tone.)

from signaling the outcome itself, it leaves the stimulus signaling the S-(R-O) relation. In the presence of the stimulus, reinforcers were response-contingent, but in the absence of the stimulus, responses and outcomes were uncorrelated.

Following discrimination training, each animal was given concurrent free-operant training with two new responses (lever pressing and chain pulling), one of which (R1) was reinforced with pellets (O1) and the other (R2) with sucrose (O2). Finally, all animals were given extinction tests with the lever and chain. In one test, there were occasional presentations of S1 and S2; in the other test, B1 and B2 occurred. The question of principal interest was whether a stimulus would elevate a response with which it had never been trained but with which it shared a reinforcer. If transfer is mediated through an excitatory Pavlovian S-O association, there should be no selective enhancement. However, if the stimulus is associated with the R-O relation, that learning should still be intact and selective transfer might be observed. Such transfer might be based on the greater similarity between different R-O relations sharing the same outcome.

Method

Subjects

The subjects were 16 experimentally naive male Sprague-Dawley rats (Holtzman Co.) about 100 days old at the start of the experiment. They were maintained in individual cages at 80% of their free-feeding weights. Water was available on an ad-lib schedule.

Apparatus

The apparatus consisted of eight identical Skinner boxes measuring 22.9×20.3×20.3 cm. The two end walls of the chamber were aluminum; the side walls and the ceiling were clear Plexiglas. Each chamber had a recessed food magazine in the center of one end wall. A small metal cup measuring 1.25 cm in diameter and 1.5 cm deep was sunk in the floor of each food magazine. Located 2.5 cm from the left-hand wall of the food magazine was a lever, and 3 cm from the right-hand wall of the magazine was a chain suspended from a microswitch mounted on the roof of the chamber. The lever was mounted 9 cm above the grid floor; the distance between the end of the chain and the grid floor was 11 cm. Located 5.5 cm directly above the roof of the magazine was a nosepoke manipulandum, which consisted of a circular aperture, 2 cm in diameter and 1.3 cm deep. The back of this aperture was covered by a metal plate that operated a microswitch whenever it was depressed. The same model of microswitch (Unimax Switch Co., 2HBT-1) was activated by all three manipulanda. Access to the manipulanda could be denied by covering the lever with a metal plate, retracting the chain through an opening in the ceiling, and inserting a jeweled light cover into the aperture of the nosepoke manipulandum. The floor of the chamber was composed of 0.48-cm stainless steel rods spaced 1.9 cm apart. Each Skinner box was enclosed in a sound- and light-resistant shell. Mounted on the inside wall of each shell was a speaker that permitted presentation of a white noise (N), measuring approximately 76 dB re 20 μ N/m, which served as one discriminative stimulus. Next to the speaker was a 6-W light (L), the illumination of which served as a second discriminative stimulus. An additional speaker allowed presentation of an 1800-Hz tone measuring approximately 76 dB re 20 μ N/m against a background level of 62 dB. This auditory cue served as the background stimulus for the sessions containing the visual discriminative stimulus. An additional 6-W light, mounted close to

the grid floor on the outside of the operant chamber on the edge of the wall opposite the food magazine and flashed at a rate 2.5/sec, served as the background cue during the sessions containing the auditory discriminative stimulus. The outside roof of the shell supported a solenoid-operated gravity-feed system; one end was connected via plastic tubing to the cup in the food magazine and the other end was connected to a reservoir that contained the sucrose reinforcer. Experimental events were controlled and recorded automatically by relays and microprocessors located in an adjoining room.

Procedure

All subjects received one session of magazine training consisting of the delivery of 10 45-mg food pellets (Formula A, P. J. Noyes Co.) followed by 10 presentations of 0.3 ml of 8% sucrose liquid on a VT 60-sec schedule. In this and all subsequent phases, the delivery of a pellet was always signaled by a distinct click of the pellet dispenser; delivery of sucrose was accompanied by a 0.5-sec buzzer. During this phase, the response manipulanda were not accessible to the animals.

Following magazine training, all subjects received two sessions of continuous reinforcement training with the nosepoke manipulandum. For half the animals, each response was reinforced with pellets in the first session and with sucrose in the second session; for the remaining animals, nose poking was reinforced first with sucrose and then with pellets. Both sessions terminated after 30 reinforcers had been earned. Occasionally, an animal failed to make a response and was given an extra session of individualized shaping. Responding was then reinforced for two 20-min sessions on a VI 30-sec schedule and for 20 20-min sessions on a VI 60-sec schedule. Pellets served as reinforcers in half of these sessions and sucrose was the reinforcer in the remaining sessions. The order of training with the reinforcers was balanced across animals and followed a double-alternation sequence.

Discriminative stimulus training. Each session of discriminative stimulus training contained 16 presentations of either a 30-sec noise or a 30-sec light stimulus. Responding was reinforced on a VI 30-sec schedule with pellets during one of these stimuli and with sucrose during the other. For half the animals, responding was reinforced with pellets during the light and with sucrose during the noise; for the remaining animals, these contingencies were reversed. During the 60-sec intertrial interval (ITI), the reinforcer earned during the discriminative stimulus was delivered on a VT 30-sec schedule. Throughout the sessions with the light, the tone was continuously present; similarly, the flashing light was present throughout the sessions with the noise. There were 18 training sessions with each stimulus scheduled in a random order.

Target response training. The two target responses, lever pressing and chain pulling, were trained in separate sessions. Initially, each response was programmed to produce a reinforcer until 50 reinforcers had been earned. Each response was then reinforced on a VI 30-sec schedule for one 20-min session. Finally, both responses were made simultaneously available for four 20-min sessions. The two responses were reinforced on independent VI 60-sec schedules with the constraint that once a reinforcer had been made available, no other reinforcer could be set up until the first had been collected. For half the animals, lever pressing earned pellets and chain pulling earned sucrose; for the remaining animals, these reinforcement contingencies were reversed.

On the day before testing, the lever and chain were concurrently available for one 8-min session during which responding was never reinforced. The purpose of this extinction session was to obtain a level of responding that would be sensitive to the potential augmenting or depressing effects of the discriminative and the background stimuli.

Transfer tests. The first test session contained eight presentations each of S1 and S2—the light and the noise—in a counter-

balanced order with an ITI of 30 sec. Both the lever and the chain were available during this test but responding was never reinforced. The second test session was identical except that B1 and B2—the tone and the flashing light—were presented instead of S1 and S2. In each test, the question of interest was whether a stimulus would selectively augment the response with which it shared an outcome.

Results and Discussion

Over the course of discriminative stimulus training, responding increased substantially during S1 and S2 but declined during the ITI. Discrimination ratios for each session were calculated in this and the subsequent experiments by dividing the rate of responding during the stimulus by the combined rate of responding during the stimulus and the ITI. By the end of training, the mean discrimination ratio for the final session, collapsed across stimulus and reinforcer identity, was 0.76. There was no significant difference in performance associated with reinforcer identity. However, discrimination ratios were higher for the sessions with the visual stimulus than for the sessions with the auditory stimulus [0.83 and 0.69, respectively, Wilcoxon $T(15) = 4, p < .01$]. That difference was due both to a lower rate of responding during the ITI in sessions with the visual stimulus [2.9 vs. 4.2 responses/min, $T(16) = 22.5, p < .01$] and to a higher rate of responding during the visual stimulus [17.0 and 10.5 responses/min, $T(15) = 15, p < .05$].

Although the value of the VI schedule in effect during a discriminative stimulus was equal to that of the VT schedule operative in the ITI, not all VI reinforcers were actually earned. As a consequence, the rate of reinforcer delivery during the discriminative stimulus was typically below that in the ITI. For instance, on the final day of discrimination training with each stimulus, the mean rates of reinforcer delivery in the presence and absence of the stimuli were 1.25 and 2 reinforcers/min, respectively. The consequence is that the stimulus actually bore a substantial negative relation to the outcome.

These results suggest that a stimulus that fails to bear a positive relation to an outcome can nevertheless develop control over the instrumental response. This observation is in agreement with prior reports of discriminative control in multiple VI-VT schedules (e.g., Bersh & Lambert, 1975; Boakes, 1973; Huff, Sherman, & Cohn, 1975; Lattal & Maxey, 1971; Weisman & Ramsden, 1973). It suggests that the development of a strong S-O association is not essential to the operation of a discriminative stimulus.

Training of the target instrumental responses proceeded uneventfully, but there were differences among responses and reinforcers. In the final training session preceding extinction, the mean rate of responding was 8.2/min for pellets and 5.9/min for sucrose [$T(16) = 17, p < .05$]. There was also a difference in performance of the two responses such that the rate of lever pressing (8.0 responses/min) was reliably higher than that of chain pulling (6.0 responses/min) [$T(16) = 25, p < .05$]. During the 8-min extinction session, there was a gradual decrease in re-

sponse rates. Although there was no significant difference between the rates of lever pressing and chain pulling (8.5 and 6.7 responses/min, respectively), the likelihood of responses that had been previously reinforced by pellets continued to be higher than that of responses previously reinforced by sucrose [8.8 and 6.4 responses/min, respectively, $T(15) = 16, p < .05$].

The results of primary interest, from the test session with the target responses and the two discriminative stimuli, are displayed in Figure 2. Performance during a stimulus is shown separately for the response that had earned the same reinforcer as that earned during the stimulus (filled circles) and the response that had earned a different reinforcer from the one available during the stimulus (open circles). In addition, the level of responding in the ITI is displayed (dashed line). There are two points to be noticed about these results. First, the stimulus produced a small overall augmentation of both responses. Relative to the ITI rate, the stimulus augmented both the same-reinforcer [$T(16) = 7.5, p < .01$] and the different-reinforcer [$T(16) = 21.5, p < .05$] responses. Second, initially the stimulus showed a greater selective augmentation of the response with which it shared a reinforcer. This selective augmentation was lost during the test session. However, analysis of the first block of test trials revealed that the stimulus produced significantly more same-reinforcer responses than different-reinforcer responses [$T(16) = 28.5, p < .05$]. This data pattern was




Figure 2. Test results from Experiment 1. Mean responses/min in the intertrial interval and during stimuli trained with either the same outcome (filled circles) as the test response or a different outcome (open circles) from the test response.

observed for both reinforcers: for the sucrose reinforcer, the mean rates of same-reinforcer, different-reinforcer, and ITI responding were 7.3, 5.9, and 3.9/min, respectively; for the pellet reinforcer, the values were 8.6, 6.0, and 5.8/min, respectively.

These results suggest that a stimulus trained as a discriminative stimulus under conditions that prevent S-O learning will nevertheless transfer to new instrumental responses that share a common outcome. This finding indicates that the transfer effect is not mediated exclusively by S-O associations. Rather, this effect appears to be dependent on the stimulus being associated with the R-O relation.

However, it should be noted that the transfer effect observed here was substantially smaller and more transient than that seen in previous experiments (Colwill & Rescorla, 1988). Several factors may contribute to this reduced effect. First, discriminative control by the stimuli during training was not as strong as we have observed in prior experiments. This may reflect weaker S-(R-O) learning and thus more rapid extinction of that association during testing. This might result from the greater difficulty in discriminating between the R-O relations obtaining in the presence and absence of the stimulus in the present study. Second, S1 and S2 were necessarily tested in the absence of their background cues. The removal of those cues may have produced some general disruption in the processing of these stimuli and thus weakened their control over responding. Finally, it is likely that some of the transfer effect observed in prior experiments is indeed mediated by an S-O association. By eliminating that association, the present experiment revealed the contribution of an S-(R-O) association but reduced the overall level of transfer.

The present finding that different-reinforcer responses were also elevated by a discriminative stimulus has not been observed in the prior studies from this laboratory using more conventional discriminative stimulus training (Colwill & Rescorla, 1988). This suggests the possibility of a general augmenting effect of a stimulus trained in the present manner. However, two alternative accounts also seem plausible. First, the weaker transfer of the stimulus to a same-reinforcer response may permit the stimulus to better exhibit its ability to elevate a different-reinforcer response. Because the transfer test took place in a choice situation, the two responses are inherently in competition. Consequently, the stronger same-reinforcer transfer of prior studies may have masked a smaller different-reinforcer transfer. Second, generalization between the background and the discriminative stimuli may make some contribution to this transfer. Note that the stimulus relations were such that a background signaling one reinforcer shared a same stimulus modality with the discriminative stimulus that controlled a relation involving the other reinforcer. Consequently, any associative relations that developed between a background stimulus and an outcome could generalize to the discriminative stimulus that otherwise was related to the other reinforcer.

The results of testing with the background stimuli provide some support for this second possibility. Those stimuli showed a slight ability to augment responding in a way that was independent of reinforcer identification. During testing of the background stimuli, the mean overall rates of same-reinforcer, different-reinforcer, and ITI responding were 3.9, 3.9, and 3.0/min, respectively. Although there was no significant difference between the levels of same-reinforcer and different-reinforcer responding, both same-reinforcer [$T(16) = 19, p < .05$] and different-reinforcer responses [$T(15) = 6, p < .05$] were elevated relative to the ITI rate.

Operationally, these background stimuli are most similar to Pavlovian CSs. A particular reinforcer occurs freely in the presence of each background stimulus. Although previous experiments in our laboratory have found little ability of a Pavlovian conditioned stimulus to augment instrumental responses, such transfer has been reported by others (e.g., Baxter & Zamble, 1982; Kruse, Overmier, Konz, & Rokke, 1983; Lovibond, 1983). Moreover, some results suggest that such augmentation is more likely to be observed with Pavlovian conditioned stimuli that are of longer duration, such as the background stimuli of the present study (e.g., Brady, 1961; Henton & Brady, 1970; Kelly, 1973; Meltzer & Brahlek, 1970; Meltzer & Hamm, 1974, 1978). However, it should also be noted that these background stimuli were present not only during the VT but also during the VI portion of the session. Consequently, they also had the opportunity to have a response reinforced in their presence. As a result, they received not only Pavlovian but also instrumental discriminative stimulus training. It is possible that a portion of their transfer derives from that training.

In any case, the primary results of this experiment are the differential control of both a trained and a transfer response by a stimulus trained without a positive correlation with the outcome. That implies a contribution of a hierarchical S-(R-O) association to discriminative control.

EXPERIMENT 2

Experiment 1 extended, to training conditions under which S-O associations are unlikely, the observation that a discriminative stimulus would control the performance of a new instrumental response associated with the same reinforcer. This result implies that transfer depends in part on the stimulus having an association with the R-O relation rather than a simple association with the outcome. The purpose of Experiment 2 was to further explore this possibility by using a stimulus that provided information about unique R-O relations that could not be derived from simple S-O and S-R associations. This was accomplished by arranging for the same two responses and two reinforcers to be available in the presence of two different stimuli. However, the relations among the elements were such that each stimulus signaled the unique R-O combination that obtained in its presence.




Figure 3. Design of Experiment 2. (R1 and R2 are instrumental responses, lever pressing and chain pulling; S1 and S2 are discriminative stimuli, noise and light; O1 and O2 are reinforcers, sucrose and pellets; + and - indicate devalued or not.)

The design of the experiment is shown in Figure 3. Rats were trained to make one response, R1 (either leverpress or chain pull), for pellets (O1) and a different response, R2 (either chain pull or leverpress), for sucrose (O2) in the presence of one stimulus, S1 (a light or a noise); during presentations of a different stimulus, S2 (a noise or a light), the same responses were available but the reinforcers that followed them were switched. In this way, the stimuli signaled the same individual response and reinforcer elements but unique response-reinforcer relations. Then, one of the reinforcers was paired with a toxin (+) to reduce its value. Finally, the animals were given an extinction test with both the lever and the chain available and occasional presentations of S1 and S2.

It was expected that if the stimuli were associated with the response-reinforcer relations, responding would be selectively depressed only during the stimulus in whose presence that response had earned the now devalued reinforcer. Such conditional control could not emerge on the basis of simple S-O associations because both reinforcers occurred with equal frequency in the presence of both discriminative stimuli. Similarly, that control could not emerge on the basis of simple R-O associations that have previously been demonstrated to influence instrumental behavior (e.g., see Colwill & Rescorla, 1985).

Method

Subjects and Apparatus

The subjects were 16 experimentally naive male Sprague-Dawley rats (Holtzman Co.) about 90 days old at the start of the experiment. They were housed and maintained under conditions identical to those described in Experiment 1. The apparatus was the same as that used in Experiment 1, except that the nosepoke manipulandum was not employed.

Procedure

The procedure for magazine training was the same as that used in Experiment 1. Leverpress and chain-pull training began the following day. Initially, each response was programmed to produce a reinforcer. In Sessions 1 and 4, the animals were trained to lever press first for pellets and then for sucrose. In Sessions 2 and 3, chain pulling was reinforced first with pellets and then with sucrose. All sessions terminated after 25 reinforcers had been earned.

Discriminative stimulus training. Each session of discriminative stimulus training contained 16 30-sec presentations each of the light and noise stimuli. Responding in the presence of these stimuli

was reinforced on a VI 30-sec schedule. During one stimulus, lever pressing was followed by pellets and chain pulling by sucrose; during the other stimulus, the response-reinforcer relations were switched so that lever pressing earned sucrose and chain pulling earned pellets. The two combinations of stimulus and R-O relations were balanced across animals. Initially, only one manipulandum was available in a session. After each response had been trained separately for four sessions, both responses were made simultaneously available for 11 sessions. During these sessions, each response continued to be reinforced on a VI 30-sec schedule with one constraint: once one reinforcer had been made available, no other reinforcer could be made available until the first had been collected. The purpose of these procedures was to ensure that both responses would be sampled. The ITI was 15 sec for the first session of discriminative stimulus training with each response; for all subsequent training sessions, the ITI was 90 sec.

Aversion training. Flavor-aversion training was carried out in the operant chambers over five 2-day cycles. On the first day of each cycle, the to-be-conditioned reinforcer was delivered on a VT 60-sec schedule for 20 min or until the subject refused to consume it. At the end of the session or 5 min after consumption of the last reinforcer, the subjects were given a 0.5 ml/kg (i.p.) injection of 0.6 M lithium chloride (LiCl) and were then returned to their home cages. On the second day of each cycle, the other reinforcer was delivered on a VT 60-sec schedule for 20 min. At the end of this session, the subjects were simply returned to their home cages. For half the subjects, sucrose was paired with the LiCl toxin; for the remaining subjects, pellets were conditioned. Designation of the to-be-conditioned reinforcer was balanced across treatment conditions. Throughout this phase of the experiment, access to the response manipulanda was prevented. In addition, after the first two cycles, the doors of the sound-attenuating chambers were opened to permit observation of reinforcer consumption.

Extinction testing. All subjects received two tests with the lever and chain both present. There were eight presentations each of the noise and light stimuli with an ITI of 90 sec. Responding was not reinforced during these test sessions.

Finally, a consumption test was administered to assess the effectiveness of the aversion training. The subjects were placed in the operant chamber, and after 1 min a pellet was delivered. An observer recorded whether or not the reinforcer was consumed within 3 min of its delivery. This test was repeated with the sucrose reinforcer.

Results and Discussion

Discriminative responding developed without incident for all stimulus, response, and reinforcer combinations. The mean discrimination ratio calculated for the final session of discriminative stimulus training, collapsed across response, reinforcer, and stimulus identity was 0.82. There were no significant effects of those factors. During this final session, the mean rate of responding in the presence of the stimulus for the to-be-devalued reinforcer (9.9 responses/min) was not significantly different from responding for the not-to-be-devalued reinforcer (10.0 responses/min).

The results of primary interest are displayed in Figure 4. Responding during the test session is shown separately when the reinforcer earned during the stimulus had been devalued (filled circles) and when it had not been devalued (open circles). It is clear that responding was sensitive to the current value of the reinforcer and that sensitivity

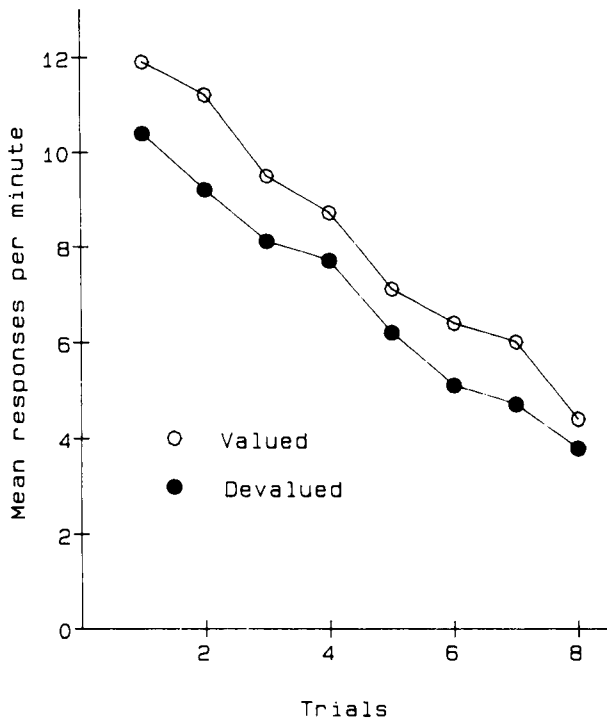


Figure 4. Test results from Experiment 2. Responding is shown when the stimulus predicted that the response would be followed by an outcome that had been devalued (filled circles) or not (open circles).

was controlled by the stimulus. When the stimulus signaled a relation between a response and the devalued reinforcer, that response was depressed; when the stimulus signaled that a valued outcome would follow the response, the level of responding was high. An overall analysis of the data, collapsed across response, reinforcer, and stimulus identity, revealed that there was significantly less responding when the stimulus signaled that the outcome would be the currently devalued reinforcer [$T(16) = 30$, $p = .05$]. Inspection of the response rates during the ITI periods revealed a decline from 3/min in the first block of two trials to 0.4/min in the final block of two trials.

The latency to make the initial response after stimulus onset was also analyzed. When the stimulus signaled that the reinforcer to follow a response had been devalued, the mean latency to make that response was longer (10.0 sec) than if the stimulus signaled a valued outcome for the response (8.9 sec). This difference, however, was not significant [$T(16) = 34.5$, $p < .10$]. On the second day of testing, essentially the same pattern of results for response rates and latency to respond was obtained. However, response rates were, in general, much lower and the differences did not reach statistical significance.

The data from the consumption test verified the success of the devaluation operation. No animal consumed the reinforcer that had been paired with the toxin, but all animals ate the nonpoisoned reinforcer.

These results support the conclusions of Experiment 1 that the animal encodes the relation between the response

and the outcome for that response as part of its learning about the stimulus. This relational knowledge cannot be deduced from any separate pairwise associations between the discriminative stimuli and the outcomes, since such S-O associations were arranged to be equivalent.

The discrimination required in this experiment can be viewed as a variant of that demanded in a simple two-choice discrimination procedure. In such a discrimination, each of two stimuli signals which of two responses will be reinforced. Such discriminations are routinely learned and can readily be explained by any of a variety of theoretical approaches. For instance, a classical S-R theory simply assumes that a reinforcer establishes an association between each stimulus and one of the responses. In the more complex discrimination studied here, both responses are reinforced in both stimuli. The S-R theory has no difficulty with the observation that each stimulus therefore is capable of evoking each response. However, the use of different reinforcers for the different S-R combinations allows one to reveal that the learned structure is substantially more complicated than such a theory implies. The selective effect of devaluation observed here suggests that the animal's learning involves something beyond any of the individual S-R, S-O, or R-O associations or combinations of those associations. Instead, it implies a true hierarchical organization of the three elements.

EXPERIMENT 3

The previous two experiments used two different techniques—transfer to a new response and reinforcer devaluation—to identify the presence of S-(R-O) associations in instrumental learning. In Experiment 3, we employed a third technique—blocking—that has commonly been used to assess the presence of S-O associations in Pavlovian conditioning. It has been well established that a stimulus will develop a reduced association with a paired reinforcer if the pairing takes place in the presence of another stimulus that already predicts that same reinforcer. In the present experiment, we exploited this fact to reveal further support for hierarchical associations in instrumental learning. The intention was to treat an R-O relation as a reinforcer itself capable of being paired with a signal. One stimulus was established as a signal for two different R-O relations. Then another stimulus was presented in conjunction with the first, with either the same or the opposite R-O relations in effect. It was expected that the added stimulus would develop better control over responding when it predicted novel R-O combinations than when it was redundant.

The design of this study is displayed in Figure 5. Rats were trained to make one response (either a leverpress or a chain pull) for pellets and a different response (either a chain pull or a leverpress) for sucrose in the presence of a noise (N). In a subsequent phase, N was presented in compound with two different lights, a steady light (L1) and a flashing light (L2). During one of the noise-light compounds, the operative R-O relations were the same

Design of Experiment 3

		Test
N: R1--01, R2--02	NL1: R1--01, R2--02	L1: R1 and R2
	NL2: R1--02, R2--01	
	NL1: R1--02, R2--01	L2: R1 and R2
	NL2: R1--01, R2--02	

Figure 5. Simplified design of Experiment 3. (R1 and R2 are instrumental responses, lever pressing and chain pulling; N, L1, and L2 are discriminative stimuli, noise, steady light, and flashing light; O1 and O2 are reinforcers, sucrose and pellets.)

as those previously signaled by N; during the other noise-light compound, the outcomes for the responses were switched so that a different pair of R-O relations were operative compared with those previously in effect during N. To assess the degree of discriminative control acquired by the lights, L1 and L2 were occasionally presented without N.

To the degree that instrumental learning involves S-(R-O) relations and N is capable of blocking that learning, there should be less responding in the presence of the light that accompanied no change in the R-O relations compared with the light introduced when those relations were shifted.

Method

Subjects and Apparatus

The subjects were 16 experimentally naive male Sprague-Dawley rats (Holtzman Co.) about 90 days old at the start of the experiment. They were housed and maintained under conditions identical to those described in Experiment 1. The apparatus was the same as that used in Experiment 1, except that the nosepoke manipulandum was not employed.

Procedure

The procedure for magazine training was the same as that used in Experiment 1. Following magazine training, all animals were trained to lever press and chain pull. For half the subjects, lever pressing was reinforced with pellets and chain pulling with sucrose; for the remaining subjects, these reinforcement contingencies were reversed. Each response was programmed to produce a reinforcer until 50 reinforcers had been earned. Lever pressing was trained in the first session and chain pulling was trained in the second session.

Discriminative stimulus training. Each session of discriminative stimulus training contained 32 30-sec presentations of a noise stimulus (N). Responding during N was reinforced on a VI 30-sec schedule with the reinforcer that had been used to train that response originally. Initially, only one manipulandum was available in any session. For the first session with each manipulandum, the ITI was 15 sec. It was increased to 90 sec for the next eight sessions, four sessions with the chain and four with the lever. After these separate training sessions, both manipulanda were made available for the remaining 16 sessions of discriminative stimulus training. Lever pressing and chain pulling were reinforced during the stimulus on

concurrent VI 60-sec schedules. To ensure continued sampling of both responses, a reinforcer could only be set up if the preceding reinforcer had been collected. Throughout these sessions, the ITI was 90 sec.

Compound discriminative stimulus training. The previous training was intended to establish N as a signal for particular relations between responses and outcomes. The intention of the present phase was to assess the degree to which N would block another stimulus that jointly signaled the same R-O relation.

All subjects received two sessions containing 32 30-sec compound presentations of N and a steady light (L1) with an ITI of 90 sec. Lever pressing and chain pulling continued to be reinforced on concurrent VI 60-sec schedules. For half the animals, the relation between responses and outcomes in the presence of the N-L1 compound was the same as that in effect during original training. Because N already signaled the operative response-reinforcer relations, it was expected that L1 would acquire relatively little control over behavior. To assess the degree of blocking, the R-O relations were switched in the presence of the N-L1 compound for the remaining animals. Thus, if lever pressing earned pellets and chain pulling earned sucrose during original training with N, then during N-L1 training, lever pressing earned sucrose and chain pulling earned pellets. In this control condition, L1 is informative about the particular response-reinforcer relations and should therefore develop a hierarchical association with those relations.

To assess learning about L1, both compound training sessions terminated with four presentations of L1 spaced 90 sec apart. Responding was not reinforced on these test trials.

To maximize detection of any potential blocking of control over response-reinforcer relations by N, compound training was repeated with a different stimulus, a flashing light (L2). However, the assignment of animals to the blocking and control treatments was reversed. Thus, the animals that had received the same response-reinforcer relations in N-L1 as in N (blocking condition) were trained with the opposite relations in the presence of the N-L2 compound (control condition). The animals that had different R-O relations in effect during N and N-L1 (control) experienced the R-O relations in N-L2 that had operated during N. There were two sessions of this compound training identical to the previous sessions. Each session terminated with four presentations of L2 with an ITI of 90 sec. Responding in the presence of L2 was never reinforced. These compound training sessions were preceded by four sessions of routine training with the noise.

After this second set of compound training sessions, the subjects received four additional retraining sessions with N. For the remainder of this experiment, they received intermixed sessions of training with L1 and L2. Each session consisted of 16 presenta-

tions of N alone and 16 presentations of N in compound with one of the two lights. In some sessions, the compound contained the steady light and in other sessions the compound contained the flashing light. The trial sequence was randomized with a 90-sec ITI. Both the lever and the chain were present and responding was reinforced on concurrent VI 60-sec schedules. The outcomes for responding were the same as those in effect during the previous compound trials. There were 16 sessions of this training, eight with the N-L1 compound and eight with the N-L2 compound. In this way, each animal was trained with two different lights in compound with the noise. One light, the control stimulus, was informative with respect to R-O relations; the other light, the blocked stimulus, was uninformative with respect to R-O relations.

To assess control by L1 and L2, four test trials with an ITI of 90 sec were scheduled immediately after the termination of Sessions 4, 8, and 16. Two of the test trials consisted of nonreinforced presentations of L1 and two consisted of nonreinforced presentations of L2. The order of testing was balanced across animals.

Results and Discussion

Discrimination training with the noise stimulus proceeded smoothly. Over the course of training, lever pressing and chain pulling occurred primarily in the presence of N and not in its absence. In the final session before compound training began, the mean overall discrimination ratio was 0.83. There were no significant effects of either response (lever pressing or chain pulling) or reinforcer type (pellets or sucrose) on discriminative performance. There was a slight disruption of discriminative control in the first session in which the steady light (L1) was presented in compound with the noise. However, performance quickly recovered and was unaffected by the presence of either L1 or L2 in any of the subsequent compound training sessions.

Figure 6 shows the results of the various test sessions. The left-hand pair of bars represents the data from the first four test sessions, scheduled at the end of separate N-L1 and N-L2 compound training. The right-hand bars show the data from the three test sessions conducted after intermixed compound training sessions. In both cases, responding has been collapsed across response and reinforcer type, but is presented separately for the blocked and control stimuli. Responding during the ITI is not displayed in the graph but it was consistently less than 1.2 responses/min throughout the tests.

After separate compound training, there was a slight indication of less responding to the blocked stimulus. However, that difference did not reach statistical significance. After more extensive intermixed training, differential responding did emerge. On those test sessions, there was reliably less responding in the presence of the blocked stimulus [$T(16) = 24.5, p < .05$]. The emergence of blocking over the course of testing is not surprising for two reasons. First, intermixing the original training trials and the blocking trials, which ensures that the blocker retains its original value, may be an especially powerful treatment (Wagner, 1969). Second, the overall level of responding was also lower during those test sessions, presumably because the animals were beginning to discriminate the nonreinforced test presentations of the

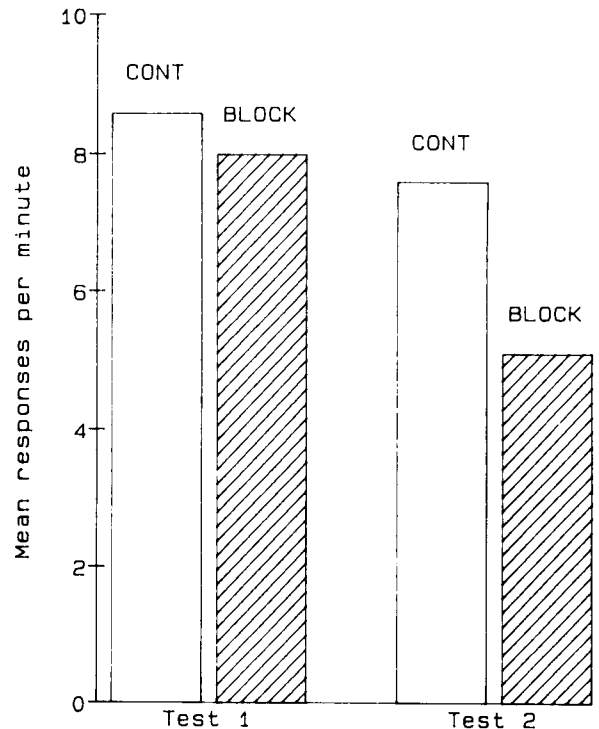


Figure 6. Test results from Experiment 3. Responding is shown during two visual stimuli, one of which (Block) was redundant with respect to the response-outcome relation and the other of which (Cont) was informative.

lights from the reinforced training presentations of the lights in compound with the noise.

These results indicate that when the response-reinforcer relations were predicted by one stimulus, the added stimulus acquired substantially less discriminative control. One interpretation of that difference is that instrumental learning involves associating the R-O relation with the discriminative stimulus. When that relation is already predicted by one stimulus, that stimulus is able to prevent or block an accompanying stimulus from developing its own association with that relation.

Although this experiment seems most naturally interpreted in terms of blocking of the hierarchical S-(R-O) association, it should be noted that it is possible to interpret the results in terms of differential S-O and S-R associations. The treatment of the control stimulus necessitated switching the R-O relations, with the result that the outcome following a response was inconsistent with the reinforcement history of that response. One consequence of this is that the outcomes occurring during the control stimulus may have been more surprising than those occurring during the blocked stimulus. It has been well documented that surprising outcomes are more effective reinforcers than are expected outcomes for both Pavlovian and instrumental learning (e.g., see Kamin, 1968, 1969; Mackintosh & Dickinson, 1979; St. Claire-Smith, 1979a, 1979b). Thus, it is possible that the difference observed between the control and the blocked stimulus may

derive from the control stimulus having stronger S-O associations and/or stronger S-R associations (see Rescorla & Holland, 1982).

GENERAL DISCUSSION

These experiments provide three different kinds of evidence for the involvement of hierarchical S-(R-O) associations in instrumental performance. In Experiment 1, a discriminative stimulus was trained under conditions that are highly unlikely to produce excitatory Pavlovian S-O associations. Yet that stimulus gained control over its own instrumental response and transferred that control to another response trained with the same outcome. In Experiment 2, a training procedure was used in which discriminative stimuli had equal opportunities to become associated with two responses and two outcomes while signaling unique R-O combinations. Subsequent devaluation of one outcome resulted in differential reduction of the likelihood of that response during the stimulus that signaled its relation to the devalued outcome. In Experiment 3, two different responses were arranged to be followed by unique outcomes in the presence of a common discriminative stimulus. A new stimulus added in compound with the original stimulus showed more control over responding when the R-O relations were switched than when they were kept constant, suggesting blocking based on the degree to which the stimulus signaled a particular relation.

Each of these results finds a natural account in terms of a hierarchical associative structure. But none is without an alternative interpretation. For instance, the interpretation of Experiment 1 depends on the assumption that the Pavlovian schedule arranged when the stimulus signals the transition from the outcome being response-contingent to being response-independent is inadequate to establish S-O associations. Although the treatment actually resulted in the stimulus having a negative correlation with the outcome, it is possible that a weak excitatory association developed between the stimulus and the outcome. This association might mediate the relatively small transfer effect obtained. However, it seems considerably less plausible that any such S-O association could support the original discrimination. Similarly, in Experiment 3, the fact that the outcome following a response changed for one target stimulus but not for the other might have allowed the outcome to be differentially effective. This, in turn, might have promoted either S-O or S-R associations that could have contributed to the reduced control by the blocked stimulus. Finally, Experiment 2 seems to admit of the least plausible alternative. There seems to be little possibility that the stimulus could be differentially associated with either of the responses or either of the outcomes. The very nature of the training paradigm appears to demand some sort of hierarchical organization. However, it is possible that the organization involves the organism using the joint occurrence of the stimulus and the response to signal a particular outcome, rather than using the stimu-

lus to signal the relation between the response and the outcome. As is the case with comparable alternatives in Pavlovian conditioning (see Holland, 1985; Jenkins, 1985; Rescorla, 1985), it is likely to prove very difficult to separate such an SR-O organization from an S-(R-O) structure.

It should also be noted that although each experiment points to the existence of a hierarchical association, the critical effects are, in each case, relatively small. One possibility is that the size of the effects is a result of the relatively complex designs required to separate the S-(R-O) learning from the simpler S-O learning. In order to attenuate the contribution of the S-O association, each experiment needed multiple stimuli, responses, and outcomes, often arranged in complex relations. To the degree that the animal fails fully to discriminate among the stimuli, to keep separate the responses, or to identify the outcomes as different, the critical results will be reduced in size. However, another possibility is that the S-(R-O) structure is only one contributor to the control by a discriminative stimulus. Procedures that emphasize that structure while reducing the contribution of the more elemental S-O, R-O, S-R associations may consequently exhibit weaker control.

The contrast between simple S-O associations and more hierarchical structures has arisen recently in Pavlovian conditioning experiments. Several recent results have suggested that sometimes a Pavlovian stimulus develops an occasion-setting power, in which it signals not the coming of another stimulus, but rather the presence of a relation between two other stimuli. For instance, Holland (1983) studied a feature-positive design in which rats received a tone (A) that was followed by food (+) on half of its occurrences. Which tone presentations terminated in food was signaled by a light (X) that preceded the tone. In such a procedure, X apparently does not develop a simple excitatory association with food; instead it signals that the A-food relation obtains. One result that supports this conclusion is that X does not control a conditioned response but instead enables A to evoke its own characteristic excitatory response. The X is not an elicitor but an occasion setter. Rescorla (1985) has reported similar results with pigeons.

These Pavlovian data have been given two different interpretations. One possibility is that X does develop an association with the food, but that the manner in which that association functions differs from that of a standard Pavlovian excitatory association. For instance, Rescorla (1985) has suggested that X takes on the power to lower the threshold for activation of the unconditioned stimulus (US) representation. Pavlovian excitors such as A have their effect by activating that representation; the role of X is to facilitate that activation. This possibility is parallel to an interpretation of the action of conditioned inhibition that results from the converse procedure in which X signals which presentations of A will not be followed by the US. Several authors (e.g., Konorski, 1948; Rescorla, 1979) have argued that in such a design, X functions to

raise the threshold for activation of the US representation, thereby interfering with A's conditioned response. This interpretation is hierarchical in the sense that X modifies the ability of A to act; but it does so by targeting one element of the A-US association, not by acting on the relation between A and the US. Another possibility, more truly hierarchical in character, has been advocated by Holland (1983). He argued that X is directly associated with the A-US association with the result that it can modulate the accessibility of that association.

A useful procedure for choosing between these alternatives is to test the ability of X to transfer to another conditioned stimulus (CS) also paired with the same US. If the action of X is primarily on the representation of the US, one might expect successful transfer. However, if X acts primarily on the A-US association, it might not transfer its effects to other CSs. In some Pavlovian preparations, that transfer is indeed substantial, providing support for the threshold-raising alternative (e.g., Rescorla, 1985). However, in other preparations, in which considerably less transfer has been observed (e.g., Holland, 1983), it seems plausible to believe that X acts on the A-US association. The transfer that is observed could then be attributed to generalization between A and other CSs paired with that US.

The transfer data obtained from instrumental training situations seems more compatible with the hierarchical alternative. Both here (Experiment 1) and elsewhere (Colwill & Rescorla, 1988), we have found that a discriminative stimulus for one response will transfer its action to another response trained with the same reinforcer. We have taken these data as indicating that a discriminative stimulus has some knowledge of the outcome that has occurred in its presence. However, the degree of control observed in the transfer test is well below that exercised over the original response. This suggests that the stimulus also has some knowledge of the response with which it has been trained. It thus becomes attractive to entertain the possibility that the stimulus is associated with an R-O unit. In that case, the observed transfer might be attributable to generalization from one R-O unit to another. The outcome-specificity of transfer of the stimulus precludes simple generalization from one response to another. However, it is easy to imagine that generalization might be based on the similarity of R-O units, where the degree of similarity would be enhanced by a shared outcome, O.

Indeed, granting the associative structure of instrumental learning the form S-(R-O) permits an account of all the currently available evidence that has been taken as demonstrating the individual associations. Such a structure includes within it the R-O knowledge necessary to show the devaluation and transfer results that we have reported earlier (e.g., Colwill & Rescorla, 1985, 1986, 1988). It would also provide a means for the subject to encode the observation that the outcome occurs in the presence of a particular stimulus. And it would even contain the specific information about the S-R relation that is necessary to solve certain conditional discriminations

(see Mackintosh, 1983; Mackintosh & Dickinson, 1979). From this perspective, the task of understanding instrumental behavior becomes one of specifying in more detail the nature of the hierarchical S-(R-O) associative structure.

REFERENCES

- BAXTER, D. J., & ZAMBLE, E. (1982). Reinforcer and response specificity in appetitive transfer of control. *Animal Learning & Behavior*, *10*, 201-210.
- BERSH, P. J., & LAMBERT, J. V. (1975). The discriminative control of free-operant avoidance despite exposure to shock during the stimulus correlated with nonreinforcement. *Journal of the Experimental Analysis of Behavior*, *23*, 111-120.
- BOAKES, R. A. (1973). Response decrements produced by extinction and by response-independent reinforcement. *Journal of the Experimental Analysis of Behavior*, *19*, 293-302.
- BRADY, J. V. (1961). Emotional conditioning. In D. E. Sheer (Ed.), *Electrical stimulation of the brain* (pp. 413-430). Austin, TX: Hogg Foundation for Mental Health and the University of Texas.
- COLWILL, R. M., & RESCORLA, R. A. (1985). Post-conditioning devaluation of a reinforcer affects instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes*, *11*, 120-132.
- COLWILL, R. M., & RESCORLA, R. A. (1986). Associative structures in instrumental learning. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 20, pp. 55-104). New York: Academic Press.
- COLWILL, R. M., & RESCORLA, R. A. (1988). Associations between the discriminative stimulus and the reinforcer in instrumental learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 155-164.
- DURLACH, P. J. (1983). The effect of signaling intertrial USs in auto-shaping. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 374-389.
- GAMZU, E. R., & WILLIAMS, D. R. (1973). Associative factors underlying the pigeon's keypecking in auto-shaping procedures. *Journal of the Experimental Analysis of Behavior*, *19*, 225-232.
- HENTON, W. W., & BRADY, J. V. (1970). Operant acceleration during a pre-reward stimulus. *Journal of the Experimental Analysis of Behavior*, *13*, 205-209.
- HOLLAND, P. C. (1983). Occasion-setting in Pavlovian feature positive discriminations. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Discrimination processes* (Vol. 4, pp. 183-206). New York: Ballinger.
- HOLLAND, P. C. (1985). The nature of conditioned inhibition in serial and simultaneous feature negative discriminations. In R. R. Miller & N. E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 267-297). Hillsdale, NJ: Erlbaum.
- HUFF, R. C., SHERMAN, J. E., & COHN, M. (1975). Some effects of response-independent reinforcement on auditory generalization gradients. *Journal of the Experimental Analysis of Behavior*, *23*, 81-86.
- HULL, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.
- 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.
- KAMIN, L. J. (1968). Attention-like processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on predictability, behavior, and aversive stimulation* (pp. 9-33). Coral Gables, FL: University of Miami Press.
- KAMIN, L. J. (1969). Predictability, surprise, attention and conditioning. In B. Campbell & R. Church (Eds.), *Punishment and aversive behavior* (pp. 279-296). New York: Appleton-Century-Crofts.
- KELLY, D. D. (1973). Suppression of random ratio and acceleration of temporally spaced responding by the same prereward stimulus in monkeys. *Journal of the Experimental Analysis of Behavior*, *20*, 363-373.

- KONORSKI, J. (1948). *Conditioned reflexes and neuron organization*. Cambridge: Cambridge University 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 & Motivation*, **14**, 165-181.
- LATTAL, K. A., & MAXEY, G. C. (1971). Some effects of response independent reinforcers in multiple schedules. *Journal of the Experimental Analysis of Behavior*, **16**, 225-231.
- LOVIBOND, P. F. (1983). Facilitation of instrumental behavior by a Pavlovian appetitive conditioned stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, **9**, 225-247.
- MACKINTOSH, N. J. (1983). *Conditioning and associative learning*. Oxford: Oxford University Press.
- MACKINTOSH, N. J., & DICKINSON, A. (1979). Instrumental (Type II) conditioning. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation* (pp. 143-167). Hillsdale, NJ: Erlbaum.
- MELTZER, D., & BRAHLEK, J. A. (1970). Conditioned suppression and conditioned enhancement with the same positive UCS: An effect of CS duration. *Journal of the Experimental Analysis of Behavior*, **13**, 67-73.
- MELTZER, D., & HAMM, R. J. (1974). Conditioned enhancement as a function of schedule of reinforcement. *Bulletin of the Psychonomic Society*, **3**, 99-101.
- MELTZER, D., & HAMM, R. J. (1978). Differential conditioning of conditioned enhancement and positive conditioned suppression. *Bulletin of the Psychonomic Society*, **11**, 29-32.
- RESCORLA, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative & Physiological Psychology*, **66**, 1-5.
- RESCORLA, R. A. (1979). Conditioned inhibition and extinction. In A. Dickinson & R. A. Boakes (Eds.), *Mechanisms of learning and motivation* (pp. 83-110). Hillsdale, NJ: Erlbaum.
- RESCORLA, R. A. (1985). 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., & HOLLAND, P. C. (1982). Behavioral studies of associative learning in animals. *Annual Review of Psychology*, **33**, 265-308.
- RESCORLA, R. A., & SOLOMON, R. L. (1967). Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review*, **74**, 151-182.
- SKINNER, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- SPENCE, K. W. (1956). *Behavior theory and conditioning*. New Haven, CT: Yale University Press.
- ST. CLAIRE-SMITH, R. (1979a). The overshadowing and blocking of punishment. *Quarterly Journal of Experimental Psychology*, **4**, 51-61.
- ST. CLAIRE-SMITH, R. (1979b). The overshadowing of instrumental conditioning by a stimulus that predicts reinforcement better than the response. *Animal Learning & Behavior*, **7**, 224-228.
- TRAPOLD, M. A., & OVERMIER, J. B. (1972). The second learning process in instrumental learning. In A. A. Black & W. F. Prokasy (Eds.), *Classical conditioning: 2. Current research and theory* (pp. 427-452). New York: Appleton-Century-Crofts.
- WAGNER, A. R. (1969). Stimulus validity and stimulus selection in associative learning. In N. J. Mackintosh & W. K. Honig (Eds.), *Fundamental issues in associative learning* (pp. 90-122). Halifax, Nova Scotia: Dalhousie University Press.
- WEISMAN, R. G., & RAMSDEN, M. (1973). Discrimination of a response-independent component in a multiple schedule. *Journal of the Experimental Analysis of Behavior*, **19**, 65-73.

(Manuscript received January 18, 1989;
revision accepted for publication May 22, 1989.)