THE QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY, 1998, *51B* (4), 349–362

Backward Sensory Preconditioning When Reinforcement Is Delayed

Jasper Ward-Robinson and Geoffrey Hall
University of York, U.K.

In two experiments rats received an initial phase of training in which two neutral stimuli were presented as a serial compound (A–X). In a second phase, A was established as a signal for a shock reinforcer, the shock being presented immediately after the termination of A (the *immediate* condition) or after a 5-sec interval (the *trace* condition). A final test phase showed that not only A but also X was capable of evoking conditioned suppression (a backward sensory preconditioning effect). The degree of suppression evoked by X was not correlated with that evoked by A. In both experiments the A trained with immediate reinforcement was more suppressive than that trained with the trace procedure, but in Experiment 1 the trace and immediate groups did not differ in the response they showed to X, and in Experiment 2 (which allowed a within-subject comparison) the trace procedure resulted in more suppression to X than did immediate conditioning of A. These results disconfirm the suggestion that the backward sensory preconditioning effect depends on the formation of an associative chain: X–A–shock. They are consistent with the proposal that the associatively activated representation of X is able to form a direct association with the reinforcer during A–shock training.

In standard demonstrations of sensory preconditioning subjects are first presented with paired stimuli (A and X), neither of which governs any marked response-eliciting properties. In a second phase of training, A is paired with an unconditioned stimulus (US) and acquires the ability to elicit a conditioned response (CR). Stimulus X is found also to govern conditioned responding. This effect may be obtained when the stimuli are presented either sequentially, with X preceding A (e.g. Hoffeld, Thompson, & Brogden, 1958; Prewitt, 1967; Rizley & Rescorla, 1972) or as a simultaneous compound (e.g. Brogden, 1939; Rescorla & Cunningham, 1978; Rescorla & Freberg, 1978) during the first phase of training. In either case the sensory preconditioning effect can be readily explained as being the consequence of the operation of an *associative chain*. Phase 1 is assumed to establish an excitatory association between the stimuli X and A; and Phase 2 training may establish an excitatory association between stimulus A and the US. The result of this learning would be that stimulus X would be able to elicit the CR, despite its

Requests for reprints and correspondence concerning this article should be sent to Jasper Ward-Robinson, Department of Psychology, University of York, York, YO1 5DD, United Kingdom. E-mail: jwrl@york.ac.uk

This work was supported by a grant from the Biotechnology and Biological Sciences Research Council. We thank C. Bonardi, N.J. Mackintosh, and R.R. Miller for their helpful comments.

never having been directly paired with the US: Its presentation is assumed to generate activity in a central representation of stimulus A, which, in turn, can produce the CR.

Sensory preconditioning has also been demonstrated when, during the first phase of training, stimulus A precedes stimulus X (Cole, Barnet, & Miller, 1995; Holland, 1981; Ward-Robinson & Hall, 1996). Some authors (e.g. Holland, 1981; Ward-Robinson & Hall, 1996) have suggested that this *backward* form of sensory preconditioning may depend on a quite different mechanism from that just described. They suggest that the formation of an A–X association in Phase 1 will allow A to activate the representation of X during Phase 2. This representation might then be able to enter into a direct association with the US that is presented during that phase (see also Konorski, 1967; Rescorla & Cunningham, 1978; Rescorla & Freberg, 1978).

The central feature of this account is the proposal that an associatively activated stimulus representation can be the subject of associative learning. If this suggestion could be confirmed it would be of importance not just for its application to backward sensory preconditioning, but because it has implications for associative learning theory generally (e.g. see Hall, 1996). The basic observation of backward sensory preconditioning is, however, insufficient to compel us to accept the reality of this mechanism as other analyses are available to explain the effect. In particular, the associative-chain account could be applied to backward sensory preconditioning if it were allowed that the A–X trials of Phase 1 resulted in backward excitatory conditioning, so that X became able to activate the representation of A. This interpretation may be considered implausible in that backward pairings are usually thought to produce inhibitory rather than excitatory learning (e.g. Siegel & Domjan, 1971; Wagner & Larew, 1985), but it remains a possibility (see, e.g. Spetch, Wilkie, & Pinel, 1981).

The experiments to be described in this article attempt to assess the viability of the associative-chain account of backward sensory preconditioning. One implication of this account is that the ability of stimulus X to evoke a CR will depend on the associative strength of stimulus A. Any procedure that reduces the strength of A should also restrict the magnitude of the CR governed by X. In our experiments we investigated the effects of interpolating a time interval between A and the US during Phase 2 of the backward sensory preconditioning paradigm. This trace conditioning procedure can be expected to produce relatively poor conditioned responding (see, e.g. Smith, 1968) and thus the associative-chain analysis predicts that it should support only a weak CR to stimulus X also.

EXPERIMENT 1

The backward sensory preconditioning procedure of Experiment 1 was similar to that used by Ward-Robinson and Hall (1996) and is summarized in the top section of Table 1. During Phase 1, rat subjects, responding on an instrumental baseline, were presented with two serial-compound stimuli, A–X and B–Y, which were composed of auditory and visual elements. In the second phase, stimulus A was paired with a shock reinforcer; stimulus B was presented non-reinforced. More conditioned responding to X than to Y in a subsequent test would provide evidence for backward sensory preconditioning. There were two groups of rats that differed only in the details of the conditioning

TABLE 1
Experimental Designs

Experiment	*Group*	*Phase 1*	*Phase 2*	*Test 1*	*Test 2*	*Test 3*
1	Immediate	A–X	A–sh	X	A	X–sh
		B–Y	B	Y	B	Y–sh
	Trace	A–X	A → sh	X	A	X–sh
		B–Y	B	Y	B	Y–sh
2		A–X	A–sh	X	A	
		B–Y	B → sh	Y	B	

Note: A and B represent white noise and a clicker; X and Y represent light and dark; sh = shock presentation; → indicates a 5-sec interval between the termination of a stimulus and shock presentation. In each stage, all animals in a given group experienced both of the trial types indicated.

procedure of Phase 2. Group immediate received conditioning in which A was followed immediately by the shock; group trace experienced an interval between A and the shock. It was anticipated that A would evoke a stronger CR in group immediate than in group trace. The question of central interest was whether this difference would also be reflected in the responding shown to X. This was examined first in a non-reinforced test (Test 1), which we have previously used to detect backward sensory preconditioning (Ward-Robinson & Hall, 1996). Test 1, however, revealed no backward sensory preconditioning, and consequently a shock-reinforced savings test was employed (Test 3), which, it was hoped, would be more sensitive to latent differences between the stimuli in their associative strength.

Method

Subjects

The subjects were 32 female hooded (Lister) rats whose mean ad lib weight was 171 g (range: 160–185 g). They were maintained at 85% of their ad lib weights by being fed a measured amount of food after each daily session. Fifteen of these rats had been used in a teaching demonstration of instrumental learning but were naive with respect to the stimuli and apparatus employed in the current experiment. The remaining subjects were fully naive. The rats were housed in pairs in a colony room illuminated from 0800–2000 h daily.

Apparatus

Four identical Skinner boxes, supplied by Campden Instruments Ltd., were used. Each was housed in a sound- and light-attenuating shell equipped with an exhaust fan serving to ventilate the chamber and generating a background noise level of 65 dB (Scale A). The boxes were equipped with a recessed food tray to which 45-mg pellets could be delivered. The tray was covered by a sprung, transparent plastic flap (6 cm high by 5 cm wide), which was hinged at the top. Pushing against this flap actuated a microswitch, and this was recorded as a response. The standard response levers were retracted throughout the experiment. The floor was made from stainless steel rods that

could be electrified by Campden Instruments Ltd. shock generators (model 521C) and shock scramblers (model 521S). The US was the delivery of a 0.5-sec shock. A loudspeaker mounted on the wall opposite the door was used to present a 10-Hz train of clicks and a white noise both at approximately 80 dB when measured against background noise produced by the exhaust fan. Illumination was provided by a 1.5-cm diameter, 2.8-W jewel light (rated for 24 V but operated at 15 V) mounted 14.5 cm above the base of the magazine tray. Interruption of this light constituted the dark stimulus. The ceiling of the box consisted of a sheet of translucent plastic and above this was positioned a 30-W strip-light. This was used to provide the stimulus referred to as light (see later). Events were controlled and recorded with a BBC microcomputer (Model B) using a version of BASIC.

Procedure

Pretraining. Initially the rats were given one 40-min session of magazine training in which food pellets were delivered to the recessed food tray according to a variable time (VT) 60-sec schedule. To retrieve these pellets, rats had to push the plastic flap that covered the recessed food tray. After rats had learned to retrieve pellets efficiently, pushing the plastic tray flap was trained as an instrumental response. Subjects were required to earn 25 food pellets according to a continuous reinforcement schedule in the second pretraining session and to respond on a variable interval (VI) 30-sec schedule in the next session. This and all subsequent sessions were 40 min in duration. Responding was maintained on a VI 60-sec baseline throughout the rest of the experiment.

Phase 1. Over the next six daily sessions, all subjects were presented with A–X and B–Y trials. Two trials of each type were presented during each session; the trial sequence was random but constrained to ensure that, over all six sessions, a given trial type was followed by a trial of the same type as often as by a trial of the other type. The organization of the Phase 1 treatment created four subgroups. For half of the subjects A was the clicker and B was the white noise; for the remaining subjects the roles of these stimuli were reversed. For half of the subjects in these two subgroups X and Y were, respectively, light and dark; for other animals the roles of these stimuli were reversed. As far as was possible, the non-naive subjects were distributed evenly among the four subgroups. Following Ward-Robinson and Hall (1996) the stimulus durations were 30 sec for A and B, and 1 sec for X and Y. Onset of the second stimulus on a trial immediately followed the termination of the first. The intertrial interval (ITI), measured from the termination of one stimulus to the onset of the next, was 462 sec.

Phase 2. Each of the two daily sessions of this phase included two trials, the first occurring 790 sec and the second 1610 sec into the session. One trial was with stimulus A and was shock reinforced; the other was with B and was non-reinforced. The shock intensity on reinforced trials was 1.0 mA. All animals received trials in the sequence A, B, B, A over the two sessions. Half of the subjects from each of the four subgroups created by stimulus counterbalancing were assigned to the immediate condition in which the shock occurred on the termination of stimulus A. The remainder were assigned to the trace condition and for these a 5-sec gap intervened between the termination of A and the occurrence of shock.

Phase 2 training resulted in a loss of baseline instrumental responding and accordingly all subjects next received six recovery sessions during which the VI 60 schedule was in operation but no other stimuli were programmed to occur.

Test Phase. Testing was designed to examine the extent to which stimuli A, B, X, and Y governed the suppression of instrumental responding. Test 1 consisted of three sessions each containing three non-reinforced presentations of stimuli X and Y. These were scheduled so that no more

than two trials of a given type occurred in succession. To allow a reasonable sample of behaviour to be obtained, the duration of each stimulus presentation was increased to 30 sec. The ITI was 322 sec. Test 2, which consisted of only a single session, was devoted to assessment of the suppression governed by A and B. Each stimulus was presented three times for 30 sec. Half of the subjects in each group received trials in the sequence A, B, B, A, B, A; the remainder received the sequence B, A, A, B, A, B. Again, the ITI was 322 sec. Test 3 consisted of two sessions each containing one reinforced presentation of X and one of Y; the stimulus duration remained at 30 sec. The reinforcer was a 0.2-mA shock. These trials occurred 790 sec and 1610 sec into the session. Half of the subjects in each subgroup received the sequence, X, Y, Y, X, whereas the remainder received the sequence Y, X, X, Y.

On each test, responding was recorded separately for each trial type and during the 30-sec stimulus-free period that preceded each trial. Suppression ratios were calculated for each stimulus on each test. These ratios took the form x/ (x+y), where x represents the rate of response pooled over all trials of a given type, and y the rate pooled over the corresponding prestimulus periods.

Results were statistically assessed using two-tailed tests and a Type I error rate of $p < .05$.

Results and Discussion

The rats readily learned to push the magazine flap in the pretraining phase. Baseline responding was not affected by the introduction of the trials of Phase 1 (see Table 1) but was disrupted to some extent by the shock-reinforced trials of Phase 2. The baseline response was reestablished during the recovery sessions, and by the sixth, and final, recovery day all animals responded at a sufficiently high rate to allow testing to proceed. Group immediate responded at a mean rate of 15.85 responses per min (rpm) on this session. The corresponding score for group trace was 19.93 rpm; these response rates did not differ reliably, $F(1, 30) = 2.49$.

The loss of baseline responding during Phase 2 made it difficult to assess the course of acquisition to stimulus A. Test 2, however, revealed that stimulus A produced a suppression of responding in both groups whereas stimulus B did not. More importantly, suppression was more profound in group immediate than in group trace. The upper part of Figure 1 shows, separately for the two groups, mean suppression ratios for A and B. An analysis of variance was conducted on these data, with group (immediate vs. trace) and stimulus (A vs. B) as the variables. There was no main effect of group, $F(1, 30) = 3.77$, but there was a significant effect of stimulus, $F(1, 30) = 78.07$, and a significant interaction between these variables, $F(1, 30) = 5.08$. An analysis of simple main effects showed that the groups differed in their response to stimulus A, $F(1, 60) = 8.79$, but not in their response to stimulus B ($F < 1$). The groups did not differ in their baseline response rates during this test. The rates, pooled over all prestimulus periods, were 20.40 rpm for group immediate and 24.94 rpm for group trace, $F(1, 30) = 2.11$.

The suppression governed by the stimuli X and Y was assessed in Tests 1 and 3. Test 1 failed to reveal any differences, either between the stimuli or between the groups. Group mean suppression ratios were .36 for X and .37 for Y in group immediate, and .33 for X and .32 for Y in group trace. An analysis of variance with group and stimulus type as the variables produced no significant effects; for the effect of group, $F(1, 30) = 1.79$, other Fs < 1. The groups did not differ in their baseline prestimulus rates. For group immediate

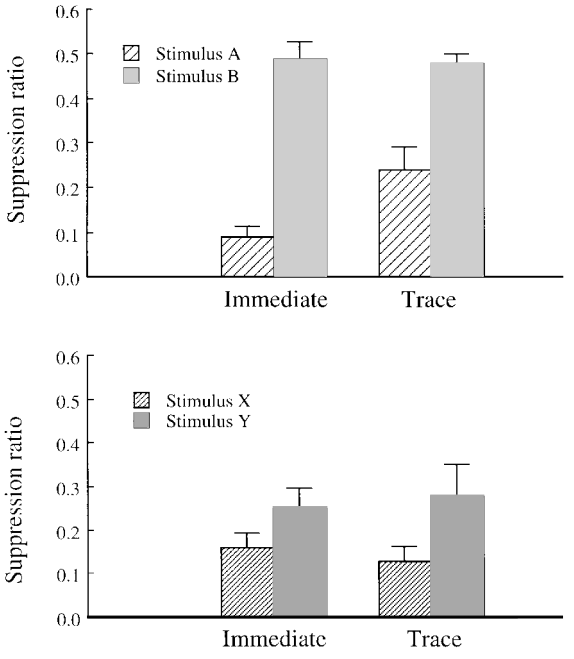

FIG. 1. Experiment 1: Group mean suppression ratios on Test 2 (non-reinforced presentations of A and B) and Test 3 (reinforced presentations of X and Y). In Phase 1 of training, X had been presented signalled by A and Y by B for all subjects. In Phase 2, A was presented followed immediately by shock (group immediate) or by shock after a 5-sec delay (group trace). Bars indicate the standard error of the mean.

this rate was 18.57 rpm; for group trace it was 22.44 rpm; $F(1, 30) = 2.08$. The backward sensory preconditioning effect reported by Ward-Robinson and Hall (1996) led us to expect more suppression to X than to Y on Test 1. We have no ready explanation for our failure to find a difference in this test—apart from the interposition of six sessions of baseline recovery training between Phase 2 and the test, the training procedure employed for group immediate was essentially identical to that used in Experiment 1 of Ward-Robinson and Hall. The final savings test (Test 3), however, yielded the expected results. The lower part of Figure 1 shows, separately for the two groups, the mean suppression scores for X and Y in this test. It is evident that X elicited more suppression than Y and that the groups did not differ in this respect. An analysis of variance with group and stimulus type as the variables yielded only a significant effect of stimulus type, $F(1, 30) = 4.98$; other Fs < 1. The mean baseline response rates based on responding shown in the prestimulus periods in Test 3 were 14.72 rpm for group immediate and 14.88 rpm for group trace. These rates did not differ significantly, $F < 1$.

It should be acknowledged that interpretation of the results of reinforced savings tests is complicated in that a treatment difference could reflect either differences in associative strength or differences in associability of the stimuli. In the present case it is possible that stimuli X and Y possessed identical initial levels of associative strength but that, for some reason, stimulus Y conditioned more slowly during the shock-reinforced trials of Test 3. One way this could occur would be if stimulus Y were latently inhibited (see e.g. Lubow,

1989); and indeed this may have occurred during its non-reinforced Phase 1 exposure. However, latent inhibition should also have occurred, and to an equal extent, to stimulus X. It is therefore unclear how a simple difference in latent inhibition could account for the differences in suppression between X and Y. Recently, however, Reed (1995) has suggested that a stimulus might show retarded conditioning as a consequence of its having been paired with some other latently inhibited stimulus. This raises the possibility that stimulus Y conditioned slowly during Test 3 because it had been paired with stimulus B, which was presented non-reinforced (and might therefore have become latently inhibited) during Phase 2. However, on the first trial of Test 3, the mean suppression ratio for stimulus X was .24 and the corresponding ratio for stimulus Y was .34. This difference (which was statistically reliable; by a sign-test $z = 2.17$) was present before either of the stimuli was paired with the shock and so cannot reflect differences in their associability. By elimination then, it is most likely that Test 3 provided an index of differences in the associative strength of stimuli X and Y.

The greater degree of suppression shown to X than to Y on Test 3 may thus be taken to constitute a demonstration of the backward sensory preconditioning effect. These stimuli differed only in that during Phase 1 X had been signalled by the subsequently reinforced A stimulus, whereas Y had been signalled by the non-reinforced B. The ability of X to evoke a CR, therefore, depends on the fact that its associate has been used to signal shock. It does not, however, appear to depend directly on the associative strength of the A stimulus—the two groups did not differ reliably in their response to X in spite of the fact that group immediate showed substantially more suppression to A than did group trace. In order to confirm this point, we carried out a further analysis of variance, which included for each group the scores for the reinforced stimulus A (Test 2) and for A's associated test stimulus X (Test 3). There was no main effect of group, $F(1, 30) = 2.31$, or of stimulus, $F < 1$, but the interaction between these two variables was significant, $F(1, 30) = 6.45$. An analysis of simple effects showed that the groups differed reliably in their response to A, $F(1, 60) = 8.37$, but not in their response to X, $F < 1$.

This pattern of results constitutes a problem for the attempt to explain backward sensory preconditioning in terms of the operation of an associative chain. If the ability of X to elicit a response depends on the chain X–A–US, then a treatment that alters the associative strength of A should modify the responding controlled by X. In this experiment the conditioning procedure used in Phase 2 resulted in a difference between the groups in their responding to A, but there was no difference in their responding to X. Discussion of possible interpretations of this finding will be postponed until after the next experiment has been described.

EXPERIMENT 2

The theoretically important conclusion to be drawn from Experiment 1 rests, in part, on a null result—the failure to find a difference between the groups in the size of the backward sensory preconditioning effect. A single null result cannot be decisive. Accordingly in this experiment we attempted a further examination of the effects of trace conditioning on the backward sensory preconditioning effect, making use of a within-subject comparison that

we hoped would provide a particularly sensitive measure of the response controlled by the test stimuli. This design has a further advantage. It is possible that the groups being compared in Experiment 1 differed in the degree to which background, contextual stimuli acquired associative strength during Phase 2 training. In particular, it might be expected that the context would acquire more strength in group trace than in group immediate. This raises the possibility that responding to stimulus X might—as predicted by the associative-chain analysis—have been less in group trace than in group immediate, but that the responding was boosted in group trace by extra excitation provided by the contextual stimuli. It is important to acknowledge that the reported failures to detect group differences in prestimulus response rates do not allow us to dismiss this possibility: It is possible that group differences that were ordinarily subthreshold were rendered detectable in the presence of the moderate excitation provided by stimulus X. The within-subjects design (summarized in the lower part of Table 1) is immune to this problem in that all subjects experience both immediate and trace conditioning in Phase 2. As in Experiment 1, all subjects received an initial phase of training consisting of presentations of two serial compounds, A–X and B–Y. In Phase 2, A was used as a CS for shock in the immediate conditioning procedure; for B the trace conditioning procedure was used. It was anticipated that A would come to elicit more suppression than would B; this was tested at the end of the experiment (Test 2). The question of interest was whether a difference in the strength of these CSs would also be evident in the responding shown to the stimuli with which they had been associated in Phase 1. This was assessed in Test 1, which consisted of separate presentations of X and Y.

Method

The subjects were 16 experimentally naive female hooded (Lister) rats with a mean ad lib weight of 216 g (range: 190–230 g). They were maintained as in Experiment 1. The apparatus was that used in Experiment 1.

Pretraining and Phase 1. Rats were first given pretraining in which they learned to retrieve pellets from the magazine before flap pushing was trained as an instrumental response. This was performed exactly as in Experiment 1. Phase 1 training followed this in which the stimuli A–X and B–Y were presented (see Table 1). Pretraining and Phase 1 training proceeded exactly as in Experiment 1.

Phase 2. Phase 2 of the experiment occurred over the following four days. During this phase animals received trials with stimuli A and B, reinforced with a 1.0-mA shock. The treatment of these stimuli differed, however, in that the shock occurred immediately after the termination of stimulus A, whereas a 5-sec trace interval was interpolated between B and the shock. One trial was given per day, occurring 462 sec into each 15-min session. Trials were organized so that half of the counterbalanced subgroups received the sequence A, B, B, A, whereas the remaining subjects received the sequence B, A, A, B. Four recovery days were given after Phase 2 training.

Testing. Test 1 assessed the suppression governed by X and Y. There were three 40-min sessions each consisting of three presentations of each of the stimuli. These were scheduled as in Test 1 of Experiment 1. The level of suppression governed by A and B was assessed in the following two

sessions (Test 2) each of which contained three presentations of A and three of B. In any details not specified here, the procedure was the same as that used in Experiment 1.

Results and Discussion

Baseline responding was acquired readily and was unaffected by the introduction of A–X and B–Y trials during Phase 1. Responding declined, however, during the reinforced trials of Phase 2 and, for one subject, was not re-established during the recovery sessions that followed this phase. The results that follow are for the remaining 15 subjects only.

Figure 2 shows group mean suppression ratios for X and Y (Test 1) and for A and B (Test 2). In each case responding was pooled over all presentations of a given stimulus type before the ratio was calculated. As might be expected, A and B (stimuli that had been directly paired with shock in Phase 2) evoked more suppression than did X and Y. A and B differed in that A—the stimulus associated with immediate reinforcement—evoked more suppression than did B—the stimulus that underwent trace conditioning. This pattern of suppression was not reflected in the responding governed by X and Y. Rather, X—the stimulus that had been associated with A in Phase 1—was less suppressive than Y—the stimulus that had been associated with B in Phase 1. A factorial analysis of variance was conducted on the data summarized in Figure 2. The variables were whether or not the stimulus had been directly reinforced (i.e. A and B vs. X and Y), and whether the stimulus had been associated (either directly or indirectly) with immediate or delayed reinforcement (i.e. A and X vs. B and Y). The first variable yielded a significant effect, $F(1, 14) = 39.61$, the second did not, $F(1, 14) = 1.19$; but there was a significant interaction, $F(1, 14) = 12.82$. Analysis of simple effects revealed a significant difference between A and B, $F(1, 14) = 9.88$, and a difference between X and Y that fell just short of the conventional level of significance, $F(1, 14) = 4.56$, $p < .052$. However, this difference was found to be reliable when subjected to a sign test, $z = 2.58$.

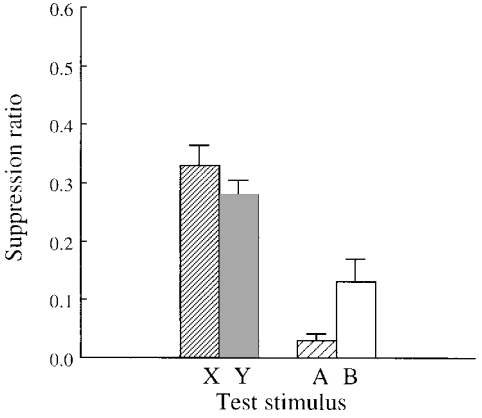

FIG. 2. Experiment 2: Group mean suppression ratios for Test 1 (presentations of X and Y) and Test 2 (presentations of A and B). In Phase 1 of training, X had been presented signalled by A and Y by B for all subjects. In Phase 1, A signalled immediate shock, and B signalled shock after a 5-sec delay. Bars indicate the standard error of the mean.

The baseline response rate was 16.74 rpm during Test 1 but was somewhat reduced, at 13.91 rpm, during Test 2. This difference, which we attribute to the fact that A and B were generally more fear-inducing than X and Y, proved to be statistically reliable, $F(1, 14) = 5.88$. The difference in baseline rates between Tests 1 and 2 means that a direct comparison between the outcomes of these two tests must be made with caution. Fortunately, however, the result of major theoretical interest is an interaction—more suppression to A than to B but less suppression to X than to Y. A change in overall level of responsiveness between the two tests could not be responsible for this pattern of results because, within each test, both stimuli would be affected equally.

In this experiment, as in Experiment 1, we found that the size of the backward sensory preconditioning effect was not directly related to the magnitude of the response acquired by the stimulus given first-order conditioning. Indeed in this experiment, the test stimulus (Y) whose partner (B) had been trained with a trace conditioning procedure proved to be more suppressive than stimulus X whose partner had been directly followed by shock in Phase 2. Although the equivalent difference was not seen in Experiment 1 (perhaps because the within-subject comparison used in this experiment provided a more sensitive measure) it is of interest that this result accords exactly with one previously reported by Cole et al. (1995). Consideration of this effect will be taken up in the General Discussion.

GENERAL DISCUSSION

In backward sensory preconditioning animals given A–X trials followed by training in which A signals the occurrence of a US are found to show a conditioned response to stimulus X. The associative-chain analysis of this effect proposes that stimulus X might control conditioned responding on test because (by way of backward excitatory conditioning in the first phase of training) it is able to activate the representation of A, which, in turn, can activate that of the reinforcer. The experiments reported here provide evidence against this interpretation. Experiment 1 showed that the capacity of stimulus A to produce conditioned responding could be reduced by using a trace-conditioning procedure with no detectable effect on the size of the backward sensory preconditioning effect. It is possible, of course, that trace conditioning did influence the magnitude of backward sensory preconditioning but that the test procedure we used was insufficiently sensitive to allow us to detect it. No such argument can be applied to Experiment 2, however, which showed that the backward sensory preconditioning effect can be enhanced by trace conditioning. These dissociations suggest that the effect does not derive from the operation of an associative chain. We turn now to other possible explanations.

First we should acknowledge the possibility that a modified version of the associative-chain principle might be capable of accommodating the results obtained here. Suppose that stimulus A produces a distinct set of after-effects (to be referred to as A′) and that these after-effects, like any other pattern of stimulation, are capable of entering into associations. The co-occurrence of A′ and stimulus X during Phase 1 training might thus be expected to establish an X–A′ association. If we assume also that the conditioning procedure of Phase 2 allows A′ to gain associative strength as a signal for the US, then test responding to X could be generated by the associative chain: X–A′–US. And finally, the assumption that A′ will be better able to gain strength in a trace conditioning procedure

than when reinforcement is immediate provides an explanation for the greater effectiveness of the former procedure in generating the sensory preconditioning effect.

The prediction that trace conditioning should produce a particularly strong backward sensory preconditioning effect can also be derived from the *temporal coding hypothesis* of Cole et al. (1995). The hypothesis is that animals encode information about temporal relationships between the events they experience, and that they can integrate information acquired in separate phases of training. In Phase 1 of our procedure, therefore, they learn that X immediately follows A; in Phase 2 they may learn (if given trace conditioning) that A is followed, after a 5-sec interval, by the shock. Integrating these two items of information allows the conclusion that X precedes the shock, a predictive relationship that will, according to the hypothesis, result in the first of these events eliciting a conditioned response. It is a problem for this account, however, that it has no reason to predict backward sensory preconditioning when stimulus A is followed by an immediate shock during Phase 2. With this procedure, integrating the information acquired in the two phases of training leads to the conclusion that X and the shock share the same time of onset, and that X continues after the termination of the shock. Since X does not predict the occurrence of the shock it should not, according to the hypothesis, generate any conditioned responding. Indeed, the hypothesis appears to suggest that a stimulus that follows a US should acquire *inhibitory* properties (e.g. Barnet & Miller, 1996), not the results obtained here (or by Ward-Robinson & Hall, 1996).

The explanation offered by Ward-Robinson and Hall (1996; see also Holland, 1981, 1983) was that backward sensory preconditioning depends on the formation of a direct association between the associatively activated representation of X and the US presented during the reinforced phase of training. According to this account, a shock that occurs on the termination of A will be well placed to become associated with a representation of X that is likely to be activated during (and perhaps immediately after) the presentation of A itself. There is thus no problem in explaining the occurrence of the effect when reinforcement of A is immediate.

By taking into account the possible role of overshadowing, this interpretation can also provide an explanation for the fact that the size of the effect is enhanced by trace conditioning. In standard Pavlovian procedures, the conditioning that occurs to a given target stimulus will be restricted if that stimulus is trained in compound with some other salient stimulus. In Phase 2 of backward sensory preconditioning, the presence of stimulus A might be expected to overshadow acquisition by X. This will be particularly true when immediate reinforcement is used, as A will be physically present at the time the X-shock pairing occurs. With trace conditioning, however, stimulus A will have been absent for several seconds before the presentation of the shock. Perhaps the memory or trace of A might be able to overshadow learning about X, but this trace will presumably be less salient than stimulus A itself; we can expect, therefore, that overshadowing will be slight and that learning about X will occur readily with the trace conditioning procedure.

The notion that the associatively activated representation of a stimulus can gain direct associative strength can thus provide a coherent analysis of both the basic observation of backward sensory preconditioning and of its enhancement when reinforcement is delayed. A further advantage of this interpretation is that it can be incorporated into current standard models of associative learning without the need for a radical revision of their

basic tenets. It has been usual to assume (e.g. Wagner, 1981) that excitatory conditioning requires the representation of the CS to be activated directly (to be in what Wagner refers to as the A1 state). But, as Holland (1983) has pointed out, there is nothing to prevent a relaxation of this rule. The suggestion that an associatively activated representation (one in the A2 state) can also come to act as an excitatory CS when paired with a reinforcer allows the model to deal with the basic backward sensory preconditioning effect without compromising the other successful predictions of the model. It should be acknowledged, however, that certain studies of human associative learning (Dickinson & Burke, 1996; Van Hamme & Wasserman, 1994) have produced results suggesting that, in some circumstances, an associatively activated representation will acquire *inhibitory* rather than excitatory strength when activated along with a reinforcer. It is not clear what factors determine which effect will be obtained, but the need to accommodate both possible outcomes may yet require a more drastic overhaul of current associative theory.

The experiments reported here do not directly address the theoretical issues just raised. Their primary purpose was to assess the possibility that backward sensory preconditioning might depend on an associative chain mechanism, similar to that widely assumed to underlie orthodox sensory preconditioning. But our demonstration that the responding evoked by stimulus X is not directly correlated with that controlled by stimulus A disconfirms this interpretation and thus opens the way for further consideration of the conditions in which an associatively activated stimulus representation might acquire associative strength.

REFERENCES

Barnet, R.C., & Miller, R.R. (1996). Second-order excitation mediated by a backward conditioned inhibitor. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 279–296.

Brogden, W.J. (1939). Sensory pre-conditioning. *Journal of Experimental Psychology*, *25*, 323–332.

Cole, R.P., Barnet, R.C., & Miller, R.R. (1995). Temporal encoding in trace conditioning. *Animal Learning & Behavior*, *23*, 144–153.

Dickinson, A., & Burke, J. (1996). Within-compound associations mediate the retrospective revaluation of causality judgements. *Quarterly Journal of Experimental Psychology*, *49B*, 60–80.

Hall, G. (1996). Learning about associatively activated representations: Implications for acquired equivalence and perceptual learning. *Animal Learning & Behavior*, *24*, 233–255.

Hoffeld, D.R., Thompson, R.F., & Brogden, W.J. (1958). Effect of stimuli–time relations during preconditioning training upon the magnitude of sensory preconditioning. *Journal of Experimental Psychology*, *56*, 437–442.

Holland, P.C. (1981). Acquisition of representation-mediated conditioned food aversions. *Learning & Motivation*, *12*, 1–18.

Holland, P.C. (1983). Representation-mediated overshadowing and potentiation of conditioned aversions. *Journal of Experimental Psychology: Animal Behavior Processes*, *9*, 1–13.

Konorski, J. (1967). *Integrative activity of the brain.* Chicago: University of Chicago Press.

Lubow, R.E. (1989). *Latent inhibition and conditioned attention theory.* Cambridge: Cambridge University Press.

Prewitt, E.P. (1967). Number of preconditioning trials in sensory preconditioning. *Journal of Comparative and Physiological Psychology*, *64*, 360–362.

Reed, P. (1995). Enhanced latent inhibition following compound preexposure. *Quarterly Journal of Experimental Psychology*, *48B*, 33–45.

Rescorla, R.A., & Cunningham, C.L. (1978). Within-compound flavor associations. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 267–275.

Rescorla, R.A., & Freberg, L. (1978). The extinction of within-compound flavor associations. *Learning & Motivation*, *4*, 411–427.

Rizley, R.C., & Rescorla, R.A. (1972). Associations in second-order conditioning and sensory preconditioning. *Journal of Comparative and Physiological Psychology*, *81*, 1–11.

Siegel, S., & Domjan, M. (1971). Backward conditioning as an inhibitory procedure. *Learning & Motivation*, *2*, 1–11.

Smith, M.C. (1968). CS–US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, *66*, 679–687.

Spetch, M.C., Wilkie, D.M., & Pinel, J.P.J. (1981). Backward conditioning: Reevaluation of the empirical evidence. *Psychological Bulletin*, *89*, 163–175.

Van Hamme, L.J., & Wasserman, E.A. (1994). Cue competition in causality judgments: The role of nonpresentation of stimulus elements. *Learning & Motivation*, *25*, 127–151.

Wagner, A.R. (1981). SOP: A model of automatic memory processing in animal behavior. In N.E. Spear & R.R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 5–47). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Wagner, A.R., & Larew, M.B. (1985). Opponent processes and Pavlovian inhibition. In R.R. Miller & N.E. Spear (Eds.), *Information processing in animals: Conditioned inhibition* (pp. 233–265). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Ward-Robinson, J., & Hall, G. (1996). Backward sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *22*, 395–404.

Manuscript received 14 July 1997
Accepted revision received 9 June 1998

Préconditionnement sensoriel rétroactif lorsque le renforcement est retardé

Dans deux expériences des rats on reçu une phase initiale de'entraînement durant laquelle deux stimuli neutres furent présentés en composé sériel (A–X). Dans une phase seconde, A signalait la présentation d'un choc électrique présenté immédiatement avant la fin de la présentation de A (la condition immédiate) ou après un interval de cinq secondes (la condition trace). La phase finale test a montré que non seulement A mais aussi X étaient capable d'évoquer un suppression conditionnée (un effet de préconditionnement sensoriel rétroactif). Le degré de suppression évoqué par X n'était pas corrélié à celui évoqué par A. Dans les deux expériences, le stimulus A entraîné avec unrenforcement immédiat était plus suppressant que celui entraîné avec une procédure trace, cependant la réponse évoquée par X était semblable dans les deux groupes, et dans l'expérience 2 (qui permettait une comparaison entre sujets) la procédure trace a produit le plus de suppression à X. Ces résultats opposent la suggestion que le préconditionnement sensoriel rétroactif est basé sur le développement d'une chaîne associative X–A–choc. Les résultats supportent la notion qu'une représentation de X activée associativement peut former une association directe avec le renforcement durant l'entraînement A-choc.

Precondicionamiento sensorial hacia atrás cuando se demora el reforzamiento

En dos experimentos unas ratas recibieron una primera fase de entrenamiento en la que dos estímulos neutros se presentaron como un compuesto serial (A–X). En una segunda fase, A se estableció como señal de una descarga, presentándose la descarga inmediatamente después de la finalización de A (condición inmediata) o después de un intervalo de 5 segundos (condición huella). Una última fase de prueba mostró que no sólo A sino también X fue capaz de provocar supresión condicionada (un efecto de precondicionamiento sensorial hacia atrás). El grado de supresión que provocó X no correlacionaba con el que provocó A. En ambos experimentos la A que durante el entrenamiento iba seguida de reforzamiento inmediato provocó una mayor supresión que la que iba seguida del procedimiento de huella, pero mientras en el Experimento 1 no hubo diferencias entre los grupos huella e inmediato en la respuesta que mostraron ante X, en el Experimento 2 (que permitía una comparación intrasujeto) hubo más supresión ante X en el procedimiento de huella que en el condicionamiento inmediato de A. Estos resultados no confirman la sugerencia de que el efecto de precondicionamiento sensorial hacia atrás depende de la formación de una cadena asociativa: X–A–descarga. Son conherentes con la propuesta de que la representación de X que se activa asociativamente puede formar una asociación directa con el reforzador durante el entrenamiento A-descarga.