

March 1993

Extending Continuous Versus Discontinuous Conditioned Stimuli Before Versus After Unconditioned Stimuli

Melody Albert

University of Nebraska-Lincoln

Sean Ricker

University of Nebraska-Lincoln

Rick A. Bevins

University of Nebraska-Lincoln, rbevins1@unl.edu

John J.B. Ayres

University of Nebraska-Lincoln

Follow this and additional works at: <http://digitalcommons.unl.edu/psychfacpub>



Part of the [Psychiatry and Psychology Commons](#)

Albert, Melody; Ricker, Sean; Bevins, Rick A.; and J.B. Ayres, John, "Extending Continuous Versus Discontinuous Conditioned Stimuli Before Versus After Unconditioned Stimuli" (1993). *Faculty Publications, Department of Psychology*. 207.

<http://digitalcommons.unl.edu/psychfacpub/207>

This Article is brought to you for free and open access by the Psychology, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Faculty Publications, Department of Psychology by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Extending Continuous Versus Discontinuous Conditioned Stimuli Before Versus After Unconditioned Stimuli

Melody Albert, Sean Ricker, Rick A. Bevins, and John J. B. Ayres

Conditioned suppression was used with rats to study the effects of extending conditioned stimuli (CSs) before versus after the delivery of unconditioned stimuli. These extensions are termed *B* and *A* extensions, respectively. Within-group designs were used to compare the effects of extending CSs when 2-min parts of those CSs were separated by temporal gaps of 6 min versus a separation of no gap. The results were as follows: (a) *B* extensions weakened conditioning more than did *A* extensions, with or without gaps; (b) under some conditions, this asymmetrical effect persisted with extended training; (c) gaps between 2-min parts of a *B* extension had no detectable effect; and (d) under some parameter values, gaps between 2-min parts of an *A* extension weakened conditioning significantly. These results are better predicted by the Sometimes Opponent-Process model (SOP; A. R. Wagner, 1981) than by the Rescorla-Wagner-Frey-Sears real-time model (J. J. B. Ayres, M. Albert, & J. C. Bombace, 1987).

Recently, researchers have begun to develop real-time models of Pavlovian conditioning (e.g., Ayres, Albert, & Bombace, 1987; Klopff, 1988; Moore & Stickney, 1980; Sutton & Barto, 1981, 1990; Wagner, 1981). Like the previous generation of models (e.g., Frey & Sears, 1978; Mackintosh, 1975; Pearce & Hall, 1980; Rescorla & Wagner, 1972), real-time models try to predict and explain the modulation of conditioning to target elements of a compound conditioned stimulus (CS) by other members of the compound. However, unlike those earlier models, real-time models also try to predict and explain the role of temporal factors, such as the interval between the onset of a CS and an unconditioned stimulus (US). In addition, real-time models allow a CS's associative value to change within a single CS presentation rather than only between trials.

In Figure 1, Panels 0, B, and A show conditioning procedures whose effects pose a challenge for real-time models. Panel 0 shows a standard, or baseline, procedure to which the other procedures are to be compared. In Panel 0, a relatively

brief CS coterminates with a US. In Panel B, the CS is extended before the US; we use the term *B extension* to refer to this procedure. In Panel A, the CS is extended after the US; we use the term *A extension* to refer to this procedure. Relative to the baseline condition in Panel 0, the 0 (*zero condition*), the *A* and *B* extensions weaken conditioning. However, of special interest is the finding that the *B* extension weakens conditioning more (Ayres & Albert, 1990; Ayres et al., 1987).

To explain this asymmetrical effect of *A* and *B* extensions, Ayres et al. (1987) made two assumptions. First, they viewed each CS as being composed of parts of equal duration. Second, they viewed the part that contained the US as the reinforced part and viewed all other parts as nonreinforced. Ayres et al. then used the Rescorla and Wagner (1972) learning rules to compute the change in conditioned value resulting from each part in turn. Finally, they used the dynamic attention rule of Frey and Sears (1978) to modify the Rescorla and Wagner salience parameter, α , after each part. For example, in Trial 1 of a *B* extension, salience weakened progressively across nonreinforced parts just as if each part were a CS-preexposure trial in a latent inhibition experiment (Lubow, 1973). By the time the final part was reinforced, the CS had lost so much salience that the effectiveness of reinforcement was greatly reduced. In contrast, in Trial 1 of an *A* extension, the reinforced part was not preceded by many nonreinforced parts, and salience was not so reduced. That is, reinforcement remained effective. Following US termination, the nonreinforced parts lost conditioned value but also progressively lost salience. This loss in salience reduced the effects of nonreinforcement and thus limited the losses in value. We call this model the *R-model*, short for the Rescorla-Wagner-Frey-Sears real-time model.

The *R-model* has not been tested extensively, but it has made at least two predictions that have received support. First, it predicted that the asymmetrical effects of *A* versus *B* extensions should be transient (see Ayres et al., 1987, Figure 8),

Melody Albert, Department of Psychology, Concordia College; Sean Ricker, Rick A. Bevins, and John J. B. Ayres, Department of Psychology, University of Massachusetts at Amherst.

This research was supported by Grants BNS 85-00277 and 87-23173 from the National Science Foundation and by a grant from the Joseph P. Healey Foundation to John J. B. Ayres.

Experiment 1A was conducted by Mary Hart and Jeff Ruell. Experiment 1B was conducted by Melody Albert. Experiment 2 was conducted by Sean Ricker. The replication experiment described in Footnote 2 was performed by Rick A. Bevins. The study was guided by John J. B. Ayres, who also prepared the manuscript.

We thank Allan R. Wagner, who supplied the computer program that we used to simulate the predictions of the sometimes opponent-process (SOP) model. We are extremely grateful to Kevin C. Clements for his work on the software we used to conduct these experiments.

Correspondence concerning this article should be addressed to John J. B. Ayres, Department of Psychology, Middlesex House, University of Massachusetts, Amherst, Massachusetts 01003.

and so far this seems to be the case (Ayres & Albert, 1990). Second, for both latent inhibition and extinction, it emphasizes the role of total CS exposure time instead of the number of CS exposure trials. The emphasis on total CS exposure time appears to be well placed both for latent inhibition (Albert & Ayres, 1989; Ayres, Philbin, Cassidy, Bellino, & Redlinger, 1992) and for extinction (Shipley, 1974).

Apart from its predictive value, the R-model has two important virtues. First, its heart is the Rescorla and Wagner model, which is widely known and understood. Second, it is simple enough to be easily implemented as a computer program that can simulate real-time experiments.

Although the R-model's simplicity is a virtue, it may also be a liability. For example, when the R-model is applied to successive CS parts, it ignores any gaps between them. In Figure 1, Panel BD shows a B extension with such gaps; we call this procedure a *BD (B discontinuous) procedure*. The R-model makes no distinction between Band BD procedures. Likewise, the R-model makes no distinction between the A and *AD (A discontinuous)* procedures of Panels A and AD, respectively. (For a similar view, see Gibbon and Balsam's [1981] treatment of whole-trial vs. split-trial procedures.) A computer-simulated rendering of the R-model's prediction about the effects of the procedures depicted in Figure 1 is shown in the top panel of Figure 2.

Another real-time model that not only predicts asymmetrical effects of A and B extensions but also allows a role for gaps between CS parts is the sometimes opponent-process (SOP) model developed by Wagner and colleagues (Donegan & Wagner, 1987; Mazur & Wagner, 1982; Wagner, 1981). In SOP, external events, such as CSs and USs, are said to be represented internally at "memory nodes." A unique node is posited for each event. In each node, there are stimulus "ele-

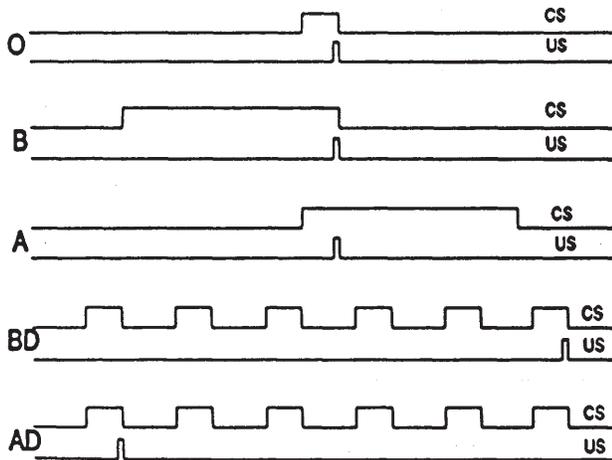


Figure 1. Typical Pavlovian forward delay conditioning procedure. (0 = the 0 [zero] procedure, a standard or baseline procedure against which the other procedures are to be compared. B = B extension; the conditioned stimulus [CS] is extended before the unconditioned stimulus [US]. A = A extension; the CS is extended after the US. BD = B discontinuous; the CS is extended before the US but is discontinuous. AD = A discontinuous; the CS is extended after the US but is discontinuous.)

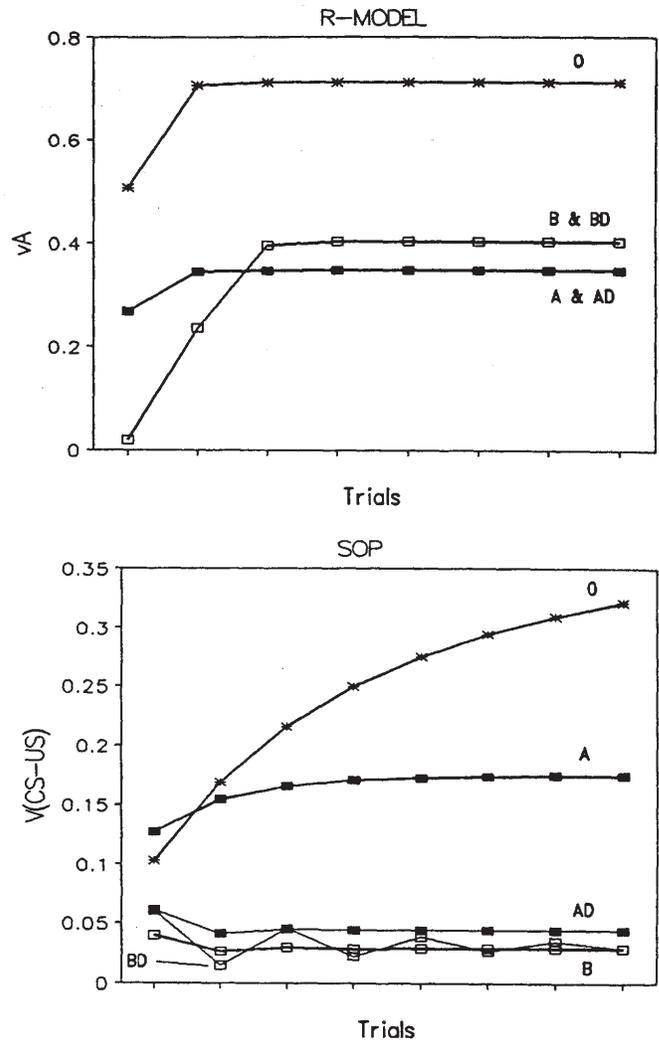


Figure 2. The Rescorla-Wagner-Frey-Sears real-time model's (R-model) predicted associative values (vA) resulting from 0, B, BD, A, and AD procedures, assuming that six 2-min parts of discontinuous conditioned stimuli (CSs) are separated by 6-min gaps (top panel) and the sometimes opponent-process model's (SOP) predictions (bottom panel). (The top panel is replotted from the top left panel of Figure 8 in Ayres, Albert, and Bombace [1987], and the reader is referred there for details, including parameter values used and the reason for their choice. SOP [e.g., Wagner, 1981] made "moment-to-moment" calculations beginning with CS onset and continuing for 10 min past CS termination. A moment was defined as 15 s. The values of SOP's free parameters were those used in previous work by Wagner and colleagues [e.g., Larew, 1986; Mazur & Wagner, 1982] and by Ayres and Albert [1990]. Thus, $CSp1 = 0.1$, $CSpd1 = 0.1$, $CSpd2 = 0.02$, $USp1 = 0.6$, $USpd1 = 0.1$, $USpd2 = 0.02$, $L+ = 0.1$, $L- = 0.02$. A = A extension; B = B extension; AD = A discontinuous; BD = B discontinuous; 0 = the zero procedure; US = unconditioned stimulus.)

ments." A given element can be in one of three "states." Before stimulus onset, all elements rest in an I state, a state of inactive, or long-term, memory. At stimulus onset, a fraction ($p1$) of the elements shift from I to A1, a state of central focus, or short-term memory. During each moment of the stimulus (a moment is an arbitrary unit of time similar to a part in the R-model), $p1$ of the elements in I go to A1, and a fraction ($pd1$) of the elements in A1 decay to A2. A2 is a state intermediate between A1 and I. Also, in each moment, a fraction ($pd2$) of the

elements in A2 decay back to I. In any moment, the change in the excitatory association between CS and US is found by multiplying an excitatory learning rate parameter, L^+ , by the product of the fraction of CS elements in A1 and US elements in A1. Similarly, the change in the inhibitory association between the CS and US is found by multiplying an inhibitory learning rate parameter, L^- , by the product of the fraction of CS elements in A1 and the fraction of US elements in A2. The current associative value of the CS, $V(\text{CS-US})$, is found by summing all these changes across all preceding moments. Once a CS has gained value, it can drive elements in the US node directly from I to A2. The fraction of US elements so driven is a function of the product of $V(\text{CS-US})$ and the fraction of CS elements in A1. Note that if the CS drives US elements to A2, the coincidence of CS elements in A1 and US elements in A2 fosters inhibitory CS-US associations, reducing $V(\text{CS-US})$. SOP constrains the fractions p_1 , pd_1 , and pd_2 and the learning rate parameters L^+ and L^- . Stimulus salience determines p_1 , so p_1 for the US is usually assumed to be much greater than p_1 for the CS. Also, the rate of decay from A2 to I is assumed to be much slower than the rate of decay from A1 to A2. Indeed, pd_2 is assumed to be five times smaller than pd_1 . Finally, L^- is assumed to be five times smaller than L^+ , and both are assumed to range from 0 to 1 (Wagner, 1981).

SOP predicts the asymmetrical effects of A versus B extensions as follows. At CS onset, a fraction of the elements in the I state of the CS node shift from I to A1. In each ensuing moment, a fraction of elements remaining in I also go to A1. Meanwhile, elements in A1 begin decaying to A2. Because the rate of decay from A2 to I is slow, elements begin to pile up in A2, reducing the number available in I to be driven to A1. If the CS is extended long enough before the US, so many elements decay from A1 and pile up in A2 that few elements remain in A1 at US onset. The excitatory CS-US association is thus severely limited. In contrast, A extensions will not reduce conditioning as much. The US occurs before many CS elements accumulate in A2, leaving a large number in A1 at US onset when US elements go to A1. After US termination, CS elements continue to shift from I to A1 because the CS is still present. Even though US elements in A1 now start to decay to A2, many still remain in A1 at the same time that CS elements are in A1. Thus, $V(\text{CS-US})$ continues to grow during the first CS moments past US termination (as if backward excitatory conditioning were occurring). In later CS moments, US elements begin to pile up in A2, fostering inhibitory CS-US associations that reduce $V(\text{CS-US})$. It is difficult, however, for this process to offset fully the excitatory learning that occurred earlier when a large fraction of CS elements and US elements were in A1 both during and shortly after the US. This difficulty stems at least in part from the fact that the inhibitory learning parameter (L^-) is five times smaller than the excitatory parameter (L^+).

Unlike the R-model, SOP allows a role for gaps between CS parts. Furthermore, under the procedures we used in the present study, SOP predicts that the gaps between the parts of A extensions are especially important. Consider, for ex-

ample, the gap that immediately follows US termination in an AD extension (Figure 1, Panel AD). At the start of the gap, US elements still reside in A1. Although $V(\text{CS-US})$ increases after US termination to the extent that CS elements are also in A1, CS elements stop going to A1 from I during the gap. The number of CS elements in A1 is therefore less if a gap follows US termination than if it does not. In essence, a gap after US termination reduces the impact of backward excitatory conditioning on $V(\text{CS-US})$. The bottom panel of Figure 2 shows a computer simulation of the $V(\text{CS-US})$ that SOP predicts should result from each of the procedures depicted in Figure 1.

In this article, we describe three attempts to challenge the R-model's assumption that gaps between the parts of A and B extensions are of no consequence. A previous attempt by Albert (1987) to challenge this assumption provided no evidence against it. Using between-groups designs, Albert found no differences in the effects of B versus BD and A versus AD extensions. She used acquisition tests followed by extinction tests identical to those described in Ayres et al. (1987). The US intensity was 1.0 mA, and acquisition and extinction occurred rather rapidly. The transience of the suppression in her task may have made it insensitive to small differences among groups. Consequently, in the present study we used within-group designs, lowered the US intensity, and gave only one trial per day. We hoped these new procedures would slow the acquisition process and thereby prolong any differences that might emerge between conditions.

In Experiments 1A and 1B, we used within-group designs to contrast the effects of A and AD extensions. In addition, we included a within-group contrast of A and B conditions. This A versus B contrast was of interest for three reasons. First, it allowed us to evaluate the sensitivity of the within-group design to a previously demonstrated effect and perhaps to extend its generality. To date, only between-group procedures have been used to demonstrate the asymmetrical effect of A versus B extensions. Second, if the within-group contrast of A versus B should yield the expected effect, then the size of that effect would provide a baseline against which to compare the effects of A versus AD extensions. Third, including an A versus B contrast permitted us to ask whether the response to an A treatment is affected by the nature of a second treatment that the subject experiences. Thus, all of our subjects received an A treatment, but half received B as their second treatment, and half received AD as their second treatment.

In Experiment 2, we used a within-group design to contrast the effects of both A and AD extensions and B and BD extensions. In our procedures, neither the R-model nor SOP predicts large differences between B and BD extensions. However, SOP, but not the R-model, predicts that AD treatments should foster weaker excitatory conditioning than do A treatments (see Figure 2).

In all three experiments, we paired CSs with grid shock USs while hungry rats barpressed for food. Suppression of barpressing evoked by the CSs prior to US onset served to index conditioning.

Experiments 1A and 1B

Method

Subjects. The subjects for Experiment 1A were 32 experimentally naive male albino Holtzman rats, which were 90 days old on arrival. They were housed singly in suspended stainless steel mesh cages in a room lighted 24 hr per day. Water was always available in these home cages. The rats were given 7 days to adapt to the colony. Then over the next 7 days, they were reduced by restricted feeding to 80% of their normal body weights. The rats in Experiment 1B were similar to those in Experiment 1A but were not experimentally naive. In another laboratory, the Experiment 1B rats had, when 90 days old, received a one-trial conditioning procedure in which a 12-s white noise CS had been either forward paired, backward paired, or explicitly unpaired with a 4-s grid shock US of either 1 mA or 2 mA. Their assignment to the present conditions was counterbalanced for prior history. About a month after that service, the Experiment 1B rats were reduced to 80% of their free-feeding body weights. They then began Experiment 1B.

Apparatus. Eight Gerbrands operant boxes were housed in ventilated 0.61-m cubes of 12.7-mm plywood lined with acoustical tile. The inside dimensions of each box were $23.2 \times 20.3 \times 19.5$ cm. Each floor was made of 18 stainless steel rods, 2 mm in diameter, mounted 1.3 cm apart center to center. The end walls were aluminum. The side walls and lid were clear Plexiglas. Centered in one end wall was a standard Gerbrands bar, 5×1.5 cm, mounted 8 cm above the grid floor. In the lower left corner of this same wall was a $5 \times 5 \times 5.5$ -cm recessed dipper tray.

On the lid of each box were two 10-cm diameter speakers. One speaker provided a 1000-Hz 84-dB tone CS; the other speaker provided an 86-dB white noise CS. These stimuli were given against a background of about 67 dB created mainly by an exhaust fan. Sound intensity was measured with a General Radio model 1565-B sound-level meter set on the Cs scale with the microphone about 7 cm from the dipper tray. A light CS was provided by a 7.5-W 110-V white frosted bulb mounted on the rear wall of each housing cube. It could be presented against a dark background or against a background provided by a similar but red frosted bulb mounted just below the white bulb. Scrambled grid shocks (0.6 mA or 0.8 mA for 1 s) served as USs and were provided by eight Grason-Stadler shock sources (Models E1064GS and 700). Barpressing, the baseline response to be suppressed by CSs, was reinforced with 4-s access to a 0.1-ml dipper cup of 32% liquid sucrose. The CSs and USs were scheduled by an NCR pc-4 computer, which also recorded responses. Operant reinforcement schedules were controlled by solid state and electromechanical equipment. Control and recording gear was housed next door to the room with the boxes.

Procedure. The procedure of Experiment 1A began with 1 day of magazine training, 3 days of shaping, and 5 days of training on a variable interval (VI) 1-min schedule of reinforcement with a variable time 1-min limited hold. Beginning with Day 1 of VI training, all sessions were 90 min long. On Day 4 of VI training, each rat received two pretest trials with white noise (N) and two with white light (L). The order was NLLN. The noise was 86 dB. The light was the 7.5-W 110-V white frosted bulb turned on in an otherwise dark box. Both CSs were 2 min long.

After 5 days of VI training came 16 days of conditioning. On each day, 16 rats (Group A/AD) received either an A trial or an AD trial. The AD trials had 6-min gaps between 2-min parts of the 12-min CS. Thus, an AD trial covered a span of 42 min (six 2-min CS parts plus five 6-min gaps). The order of trial types across the 16 days was A AD AD A, AD A A AD, A AD AD A, AD A A AD. The CS always

began at the start of Min 45, and the grid shock US (0.6 mA for 1 s) came in the last 1 s of Min 46 (the last second of the first 2-min CS part). For 8 of the 16 rats, N was the CS on A trials, and L was the CS on AD trials. For the remaining rats, the reverse was true. These 16 rats thus provided a within-group contrast of A and AD extensions. Another 16 rats (Group A/B) were used to provide a within-group contrast of A and B extensions. For both groups of 16 rats, the A trials were identical. The order of trial types was also identical except that B trials replaced AD trials for Group A/B. On B trials, a 12-min CS coterminated with the 1-s 0.6-mA grid shock US; the US came in the last 1 s of Min 46. The modality of the trial type was again counterbalanced.

The procedure of Experiment 1B was identical to that of Experiment 1A with four exceptions: (a) a 1000-Hz 84-dB tone replaced the white noise CS; (b) the white light CS was turned on in a box already lit by the red bulb; (c) the US was 0.8 mA instead of 0.6 mA; and (d) there were 12 days of conditioning instead of 16 days.

Treatment of data. Suppression during various parts of CSs was indexed in terms of the Annau and Kamin (1961) suppression ratio, $D/(D + B)$, in which D denotes the number of barpresses during a 2-min part of the CS, and B denotes the number of responses in the 2-min period just before the CS. In Experiment 1B, 1% of the suppression ratios were 0/0. Following Annau and Kamin, we defined such ratios as 0. This definition is not entirely arbitrary, for our rats usually showed substantial CS-evoked suppression on the trial just before, just after, or both before and after a 0/0 trial. For all statistical contrasts, we adopted a two-tailed rejection region of .05.

Results

The results of Experiments 1A and 1B are shown in the top and bottom panels, respectively, of Figure 3. The left panels show the results for A versus B. The right panels show the results for A versus AD. The conditioning measure is based on only the first 2 min of the CS. This 2-min measure was chosen because, in the A and AD treatments, shock occurrence complicates behavior in later CS parts.

In both Experiments 1A and 1B, there was little suppression to the CS in pretesting, but suppression increased across conditioning trials. Also, suppression tended to be stronger, as expected, in Condition A than in Condition B. Not shown in the figure are modality effects that occurred in both studies. In Experiment 1A, light evoked more suppression than noise; in Experiment 1B, tone evoked more suppression than light. Therefore, in Experiment 1A, the difference between A and B was greater when A was light and B was noise than vice versa. In Experiment 1B, the difference between A and B was greater when A was tone and B was light than vice versa.

These summary statements were supported statistically. Modality \times Trial \times Treatment analyses of variance (ANOVAs) were conducted on the data of Group A/B in both Experiments 1A and 1B. Here, modality refers to whether A was noise or light in Experiment 1A or was tone or light in Experiment 1B. In Experiment 1A, the ANOVA found a significant effect of trial, $F(6, 84) = 14.21$, and significant interactions between trial and treatment, $F(6, 84) = 3.96$, and between modality and

¹ There are six, rather than seven, degrees of freedom for the trials because of the loss of data for Trial 6 in the B condition. No data from Trial 6 were included in the ANOVA.

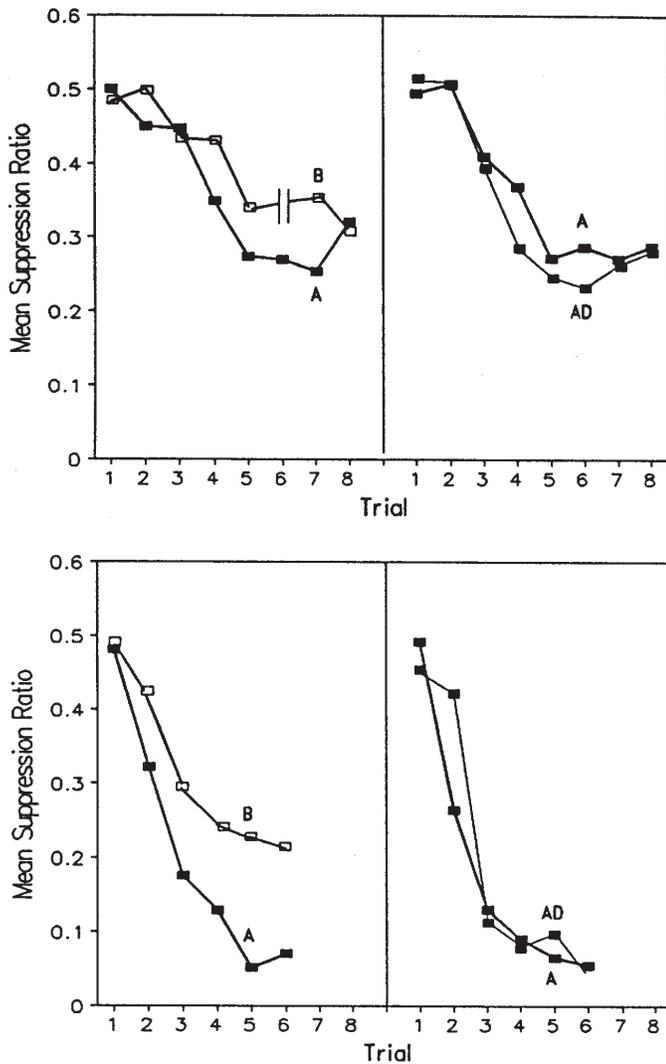


Figure 3. Acquisition of conditioned suppression at one trial per day when the A versus B conditions and A versus AD conditions were manipulated within subjects. (Top panels = Experiment 1A. Bottom panels = Experiment 1B. A = A extension; B = B extension; AD = A discontinuous.)

treatment, $F(1, 14) = 17.23$.¹ Of main interest was the Trial \times Treatment interaction, which showed that B evoked less suppression than A on some trials but not on other trials. In Experiment 1B, the Modality \times Trial \times Treatment ANOVA again found significant effects of trial, $F(5, 70) = 47.51$, trial \times treatment, $F(5, 70) = 2.92$, and modality \times treatment, $F(1, 14) = 11.82$. The main effect of treatment was also significant, $F(1, 14) = 34.48$, indicating more overall suppression to A than to B.

In the A versus AD conditions (right panels of Figure 3), A and AD evoked similar suppression except on Trial 2 of Experiment 1B (lower right panel). Analyses similar to those described above showed no differences between the A and AD treatments except in Experiment 1B, where the Treatment \times Trial interaction was significant, $F(5, 70) = 6.48$, and in which A evoked more suppression than AD on Trial 2, $F(1, 14) = 18.49$.

As noted in the introduction, the design of Experiments 1A and 1B permits us to ask if the response to an A exten-

sion is affected by the nature of another extension that the rat experiences. Comparing suppression to A in the left and right panels of Figure 3 suggests that A evoked similar suppression whether the other extension experienced was a B extension (left panels) or an AD extension (right panels). The ANOVAs on each trial failed to reveal any differences in responding to A as a function of the other extension type; in Experiment 1A, largest $F(1, 28) = 1.80$, and in Experiment 1B, largest $F(1, 28) = 1.14$.

Experiments 1A and 1B extend the finding of asymmetrical effects of B versus A extensions to a situation in which each rat receives a different treatment on alternate days (e.g., B on one day and A on another, or A on one day and AD on another). In general, the asymmetrical effects of B versus A extensions were much more robust in Experiment 1B, which used a 0.8-mA shock, than in Experiment 1A, which used a 0.6-mA shock. Similar results have been reported in between-groups designs by Ayres and Albert (1990).

For Group A/B, mean pre-CS response rates during conditioning in Experiment 1A were 22 and 23 responses per minute on A and on B trials, respectively. For Group A/AD, response rates were 20 responses per minute on both A and AD trials. In Experiment 1B for Group A/B, mean pre-CS rates during conditioning were 18 responses per minute on both A and B trials. For Group A/AD, response rates were 16 and 15 responses per minute on A and on AD trials, respectively. Thus, differences between treatments in terms of suppression ratios were not complicated by differences in pre-CS rates.

In Experiment 1A, suppression was uniform throughout the 12-min CSs over the first three trials. Thereafter, however, clear trends emerged. These trends are shown in Table 1. Under all conditions, suppression weakened throughout the CS. This weakening is especially remarkable in the B treatments because it is just the opposite of inhibition of delay. Here and elsewhere (Ayres & Albert, 1990; Ayres et al., 1987), we have been concerned that rats in the B condition might learn that the first 2-min part of the CS was nonreinforced and therefore might show inhibition of delay. If so, such learning might explain why B evoked weaker suppression in that period than did A. Table 1 reveals clearly that there was no evidence for that idea. Indeed, B evoked more suppression during its first 2-min part than in parts closer to reinforcement (for contrasting results, see Gaioni, 1982). Table 1 (Group A/AD) also shows that suppression in the second 2-min CS part was weaker in the A condition than in the AD condition. This difference probably reflects the activating effects of shock (e.g., Ayres & Vigorito, 1984; Fanselow, 1982). These activating effects are more likely in a CS part that immediately follows shock, as in the A condition, than in a CS part that occurs 6 min later, as in the AD condition. This difference in the suppression evoked by A and AD in the second CS part illustrates the complicating effect of shock on behavior in all but the first 2-min part of the CS. For this reason, we chose suppression in the first 2-min CS part as our measure of conditioning.

In the AD condition of Experiment 1A, there was little suppression in the gaps between CS parts. For example, over the

Table 1
Mean Suppression Ratios in Each 2-Min Part of a 12-Min Conditioned Stimulus (CS) Averaged Over the Last 2 Days of Conditioning in Experiments 1A and 1B and Averaged Over the Last 4 Days of Conditioning in Experiment 2

Condition	CS part					
	1	2	3	4	5	6
Experiment 1A						
A/B						
A	.29	.36	.37	.43	.43	.44
B	.33	.34	.36	.41	.45	.46
A/AD						
A	.28	.45	.44	.45	.48	.46
AD	.27	.35	.43	.48	.49	.49
Experiment 1B						
A/B						
A	.06	.25	.20	.26	.30	.31
B	.22	.26	.34	.37	.37	.39
A/AD						
A	.06	.21	.15	.25	.29	.35
AD	.07	.08	.16	.18	.20	.23
Experiment 2						
B/BD						
B	.46	.45	.46	.47	.48	.46
BD	.45	.45	.45	.43	.43	.42
A/AD						
A	.17	.42	.41	.45	.49	.48
AD	.28	.36	.44	.41	.46	.46

Note. D = discontinuous.

last two trials, the mean suppression ratio for any gap ranged from .48 to .52. In the AD condition of Experiment 1B, the range of means was .41 to .48.

Discussion

The within-group designs used in Experiments 1A and 1B were clearly sensitive to the difference between A and B effects but did not detect any differences between A and AD effects, with one exception. On Trial 2 in Experiment 1B, there was a significant difference in the direction predicted by SOP. Unfortunately, a difference on a single trial in the design we used is difficult to interpret. The problem is that 2 conditioning days were required to obtain a single estimate of conditioning to each of the two stimuli. When Stimulus A was given on Trial 2, it actually occurred on Conditioning Day 4. The rats had received reinforced occurrences of A, AD, and AD on Days 1, 2, and 3, respectively. We cannot tell whether the increase in suppression between Days 3 and 4 reflects a general learning process (e.g., learning that stimuli are followed by shock) or whether it reflects a true difference between the A and AD conditions. Experiment 2 provides the counterbalancing of treatment order necessary to reach a conclusion on this matter.

The present results extend the generality of the asymmetrical effects of A and B extensions in two ways. First, the earlier demonstrations (Ayres & Albert, 1990; Ayres et al., 1987) used between-groups designs as opposed to the within-group designs used in this research. Second, the earlier demonstrations used intermittent CSs (1 s on, 0.11 s off). We used non-intermittent CSs here. With regard to using within-group designs, it is interesting that the response to the A extension appeared to be unaffected by the nature of the other extension the rat received. Thus, rats given a B extension as their other stimulus and rats given an AD extension as their other stimulus responded very similarly to the A extension.

There was, however, one finding that did differ from previous results of CS extensions. In Experiments 1A and 1B, suppression weakened from beginning to end of B extensions. In contrast, Ayres et al. (1987) found roughly uniform suppression throughout B extensions, and Ayres and Albert (1990) found some evidence for weaker suppression at the beginning and end of the CS than in the middle of the CS. The tendency for suppression to weaken throughout the B extension here may have been due to the experience of the present rats with the A extension. Although previous reports did not describe suppression throughout the entire A extension, the suppression did weaken as it did here. Perhaps a rat's experience with an A extension affects the course of suppression within a B extension. The result is a pattern of responding opposite that of inhibition of delay.

Experiment 1A also found some evidence that the asymmetrical effects of A versus B extensions were transient. Thus, in Figure 3, the A and B curves converge on the last trial. The R-model predicts that the effects should be transient; this prediction has received support in a between-groups design (Ayres & Albert, 1990). No evidence for transience was found in Experiment 1B, but that experiment may not have had enough trials to allow a good test of the prediction.

Experiment 2

In Experiment 2, one group of rats received both A and AD extensions, and a second group received both B and BD extensions. The design provided within-groups comparisons of A versus AD and B versus BD and permitted between-groups comparisons of B versus A and BD versus AD. In contrast to Experiments 1A and 1B, CS modality and order of treatment were fully counterbalanced. We also attempted to equate the saliences of CSs in different modalities so as to eliminate the complicating Modality \times Treatment interactions of Experiments 1A and 1B.

Method

The subjects were 32 experimentally naive male albino Holtzman rats, which were 90 days old on arrival. After 1 week in the colony, they were reduced by restricted feeding over 7 days to 80% of their normal body weights. All apparatus and housing conditions were as before. The rats were magazine trained, shaped, and VI trained as in Experiment 1A. On Day 4 of VI training, the rats were pretested with

the CSs to be used later. The light CS was that of Experiment 1B (i.e., the white light superimposed on the red-light background). The tone (T) CS was also that of Experiment 1B but was reduced in intensity from 84 to 80 dB in an effort to match its salience to that of the light. Pretest trials were 2 min long and were given in the order TLLT. Trials began at the start of Mins 20, 30, 50, and 70.

Pavlovian conditioning began on the day following Day 5 of VI training and lasted 20 days. On each day, 16 rats (Group A/AD) received either an A trial or an AD trial. The AD trials had 6-min gaps between 2-min parts of the 12-min CS. The order of trial types across the 20 days was A AD AD A, AD A A AD, A AD AD A, AD A A AD, A AD AD A for 8 of the rats. For the other 8 rats, AD replaced A and vice versa. Also, for 8 of the 16 rats, L was the CS on A trials, and T was the CS on AD trials. For the other 8 rats, the reverse was true. In short, CS modality and treatment order were fully counterbalanced. On each day, the A or AD CS began at the start of Min 30, and the grid shock US (0.8 mA for 1 s) came in the last 1 s of Min 31 (the last second of the first 2-min CS part).

Another 16 rats (Group B/BD) received on each conditioning day either a B trial or a BD trial. The BD trials had 6-min gaps between 2-min parts of the 12-min CS. The CS modality and treatment order were counterbalanced as in Group A/AD. On each day, the B or BD CS began at the start of Min 30, and the 0.8-mA shock US came in the last 1 s of the last 2-min part of the CS (Min 41 on B trials and Min 71 on BD trials).

The rats were run in four squads of 8. In each squad, 4 rats received A or AD trials and 4 received B or BD trials. In Squads 1 and 3, rats in Boxes 1 through 4 received A or AD trials, while rats in Boxes 5 through 8 received B or BD trials. The reverse was true for Squads 2 and 4. In this way, potential box effects and time-of-day effects were equated across treatments.

Results

There was little suppression to tone or light during pretesting. Averaged over both pretest trials, the mean suppression ratio for all 32 rats was .48 for tone and .47 for light.

Figure 4 shows acquisition for Groups A/AD and B/BD. Once again, the suppression ratio is based on only the first 2-min part of the 12-min CS. Suppression increased across trials more in A and AD treatments than in B and BD treatments. Suppression was weaker for AD than for A. Near the end of training (Trials 9 and 10), suppression weakened markedly in the A and AD treatments (a postasymptotic performance decrement). Statistical analyses supported these descriptions.

The data were analyzed with an A-B \times C-D \times Trial ANOVA, in which A-B refers to whether the extension was an A or B type and C-D refers to whether the extension was continuous or discontinuous. The ANOVA revealed a significant effect of A-B, $F(1, 30) = 63.31$, denoting more suppression to A and AD than to B and BD. The trial effect and the Trial \times A-B interaction were also significant, $F_s(9, 270) > 15.64$, meaning that suppression increased across trials and did so more for A and AD treatments than for B and BD treatments. There was also a main effect of C-D, $F(1, 30) = 7.91$, denoting stronger suppression in continuous versus discontinuous extensions. The A-B \times C-D interaction was also significant, $F(1, 30) = 6.03$, meaning that the effect of gaps was larger for A extensions than for B extensions. Separate ANOVAs on the A/AD data and the B/BD data confirmed this interpretation. These

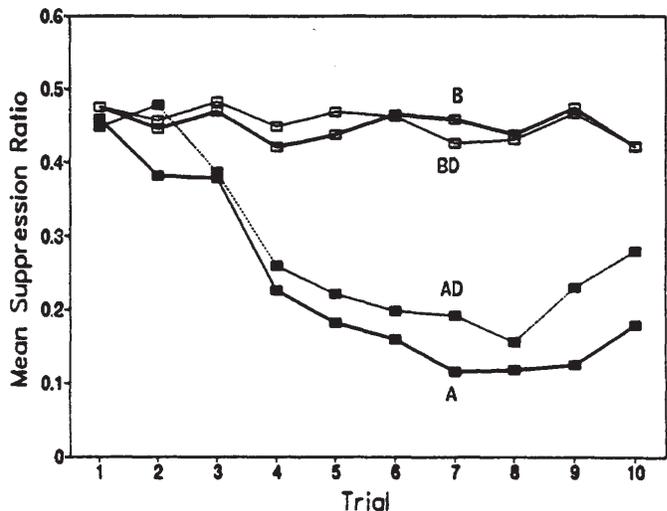


Figure 4. Acquisition of conditioned suppression at one trial per day under B, BD, A, and AD conditions. (Group B/BD received both B and BD treatments. Group A/AD received both A and AD treatments. A = A extension; B = B extension; AD = A discontinuous; BD = B discontinuous.)

ANOVAs found a significant difference between A and AD, $F(1, 15) = 9.30$, but not between B and BD ($F < 1$). Because the difference between A and AD seemed especially large on Trials 9 and 10, when the postasymptotic performance decrement was occurring, we were concerned that the A versus AD difference might somehow be tainted, and so we reanalyzed the A/AD data, omitting those trials. The difference between A and AD was still significant, $F(1, 15) = 7.27$.

Attempts to equate the salience of the light and tone were apparently successful. Separate ANOVAs carried out on the A/AD data and the B/BD data found no effect of modality ($F_s < 1$).

Differences in suppression between treatments were not complicated by differences in pre-CS rates. Those rates were 33 responses per minute under both A and AD conditions and were 26 and 28 responses per minute under B and BD conditions, respectively.

The bottom part of Table 1 shows that suppression tended to weaken across CS parts for A and AD treatments but not for B and BD treatments. Even so, there was little suggestion of inhibition of delay in the B and BD treatments. The table also shows that even though A evoked more suppression than AD in the first 2-min CS part (as already seen in Figure 4), it evoked less suppression in the second CS part. Once again, the weaker suppression to A than to AD in the second CS part presumably reflects the activating effects of shock. As shown in Table 1, this effect occurred in all three experiments. As mentioned earlier, it illustrates why we chose suppression in the first 2-min CS part as our index of conditioning.

Mean suppression ratios in the gaps for discontinuous CSs over the last 4 days of conditioning ranged from .46 to .51 in the AD condition and from .46 to .49 in the BD condition. Suppression tended to be weaker during gaps than during CSs.

Discussion

The results of Experiment 2 are more consistent with SOP (Wagner, 1981) than they are with the R-model (Ayres & Albert, 1990; Ayres et al., 1987). SOP clearly predicts weaker conditioning in AD than in A extensions (compare Figures 2 and 4). The R-model, in contrast, incorrectly predicts no differences between A and AD (but see below). The R-model also predicts that the asymmetrical effects of A and B extensions should be transient. SOP does not. Figure 4 suggests that the effects are not transient. It appears that the A and B curves would not converge with continued training unless the postasymptotic performance deficit were so strong as to allow the A curve to approach the B curve. This is not the kind of convergence that the R-model predicts. The top panel of Figure 2 shows that the B curve should catch up with the A curve after A reaches asymptote, which is certainly not suggested by the present data. Ayres and Albert presented evidence for transience. Reinspection of their data in the light of the present findings suggests that a postasymptotic performance decrement may have contributed to convergence of A and B curves in both of their experiments (see also Experiment 1A here) and that a ceiling effect on suppression in their A treatment may have additionally played a role in their Experiment 2. Because of the complications of ceiling effects and postasymptotic performance decrements, it may be that the clearest test of transience would come from blocking or second-order conditioning experiments. That is, would A and B late in training have similar or different abilities to block conditioning to a third CS or to serve as reinforcers in the second-order conditioning of a third CS?

Although SOP correctly predicts weaker conditioning in AD than in A, it errs in predicting that conditioning in AD should resemble that in B and BD. In this instance, the R-model seems to do a better job. It clearly predicts, at least preasymptotically, stronger conditioning in AD than in B and BD. In computer simulations performed after collecting the present data, we have attempted to vary SOP's free parameters in several ways to make SOP predict large differences between AD and B and BD. We have had no success. SOP continued to predict similarly weak learning in AD, B, and BD.

Our results indicate that A versus AD effects are relatively weak and difficult to detect. Albert (1987) could not detect these effects with between-groups designs, nor did we detect them in Experiment 1A with a within-group design. We may have detected a short-lived effect in Experiment 1B, but it was confounded with the order of treatments across trials. Yet, the effect was quite clear in Experiment 2. Factors responsible for the clarity presumably include (a) complete counterbalancing of treatments across trials, (b) selection of equally salient CSs so as to remove complicating Modality \times Treatment interactions, (c) choice of parameters that fostered moderate levels of suppression, thereby avoiding both floor and ceiling effects, and (d) use of a within-group design.

General Discussion

Using within-group designs, our experiments have replicated the asymmetrical effects of B versus A extensions (Ayres & Albert, 1990; Ayres et al., 1987) and have shown a similar asymmetrical effect for discontinuous B and A extensions as well (i.e., weaker suppression to BD than to AD). Neither the within-groups design of Experiment 2 nor the between-groups designs used by Albert (1987) detected an effect of gaps in a B extension. This finding, then, is quite similar to the results of a series of studies of latent inhibition (Albert & Ayres, 1989; Ayres et al., 1992) and extinction (Shipley, 1974), suggesting that total CS exposure time and not number of CS exposure trials is the crucial parameter. Although we found no effect of gaps in a B extension, Experiment 2 very clearly, and Experiment 1B less clearly, showed evidence for a role for gaps in an A extension. The gaps weakened conditioned suppression to the part of the CS that preceded shock. This effect is anticipated by SOP (Wagner, 1981) but not by the R-model (Ayres et al., 1987; but see below).

According to SOP, two factors (more backward excitatory conditioning to A and more backward inhibitory conditioning to AD) foster a stronger $V(\text{CS}-\text{US})$ under the A procedure than under the AD procedure. Under both procedures, US onset drives US elements to A1 from where they gradually decay to A2. They then begin piling up in A2 over many moments, then even more gradually decay from A2 to I. At any moment, the number of US elements in A1 is identical for both procedures, and so is the number in A2. Therefore, the main difference between the two procedures is the number of CS elements in A1 at any moment. In the moments after US termination, there are many more CS elements in A1 for the A procedure than for the AD procedure. The reason is that the CS is physically present during those moments only in the A procedure. Thus, although the CS in both procedures gains excitatory value during many of the post-US moments, the gain is much larger in the A procedure. This gain may be viewed as a backward excitatory conditioning-like effect. As time continues, the number of US elements in A2 increases to the point where the conjunction of CS elements in A1 and US elements in A2 fosters inhibitory conditioning, reducing $V(\text{CS}-\text{US})$. Again, the crucial difference between A and AD procedures is the number of CS elements in A1. The greater that number, the greater the losses in $V(\text{CS}-\text{US})$. Under the A procedure, that number declines steadily over the remaining moments of the trial. However, under the AD procedure, the number declines in the gaps, then increases during subsequent CS parts. These increases favor relatively large losses in $V(\text{CS}-\text{US})$.

2 We have recently replicated the A versus AD difference reported in Experiment 2. We used a within-groups design with 32 rats. We counterbalanced for order effects and modality as in Experiment 2. The effect was extremely small. Averaged over 10 trials, the mean suppression ratio to A was .19 and that to AD was .22. Still, with 32 rats in a within-groups design, the difference was significant at $p < .01$.

These losses may be viewed as reflecting backward inhibitory conditioning. In short, the CS in the AD procedure benefits less from backward excitatory conditioning and suffers more from backward inhibitory conditioning.

It is possible, of course, that the A versus AD difference is not due to either backward excitatory or inhibitory conditioning. Another factor that might produce the difference is extinction to CS onsets. There are more CS onsets in the AD treatment than in the A treatment. If extinction occurs to these onsets, then that extinction should generalize to the first CS onset of each trial. Such generalized extinction would thus result in weaker suppression in the first CS part of an AD treatment than in the first CS part of an A treatment. This same generalized extinction, however, should also foster weaker suppression in BD treatments than in B treatments. We found no evidence for such an effect in our Experiment 2, and Albert (1987) found no such evidence in her experiments. Unfortunately, in our Experiment 2, a difference between BD and B could have been obscured by a floor effect. Suppression in our B condition was so weak that it could not be further weakened in the BD condition. Albert (1987) may have failed to detect a difference between B and BD because of the insensitivity of her between-groups designs.

Using our A versus AD contrast in a within-group design, it is possible to experimentally compare the importance of extinction to multiple CS onsets as opposed to the backward excitatory and inhibitory conditioning factors discussed above. For example, using A and AD procedures, one could reinforce the first 2-min CS part in some sessions and present the five remaining nonreinforced CS parts in other sessions. These five remaining parts would then be too far away from the US for the A treatment to benefit from backward excitatory conditioning or for the AD treatment to suffer more from backward inhibitory conditioning. As in the present research, the treatments would be equated in terms of the total CS exposure time but would differ in terms of the number of CS onsets. If that number is important, then suppression to the first CS part should once again be weaker to AD than to A. If this result occurred, then it would not only challenge SOP's account of the difference between the effects of A and AD extensions; it would also reopen the broader question of the role of number of CS onsets in extinction and latent inhibition. Both extinction (ShIPLEY, 1974) and latent inhibition (Albert & Ayres, 1989; Ayres et al., 1992) have been shown to depend not on the number of CS onsets but rather on the total CS exposure time. However, the studies supporting this conclusion all used between-groups designs. Perhaps studies using within-group designs would reveal a previously undiscovered role for number of onsets.

Besides the generalization of extinction from multiple CS onsets, there is another factor that might create the difference in the effects of A and AD extensions. A number of researchers (e.g., Ayres & Vigorito, 1984) have suggested that temporal cues that follow widely spaced shocks become conditioned inhibitors of fear. Because it adopts the learning equations of the Rescorla and Wagner model, the R-model holds (a) that non-

reinforcing a cue in compound with an excitor should make that cue inhibitory and (b) that nonreinforcing a target excitor in the presence of an inhibitor should protect the target from extinction. Thus, the R-model predicts that postshock temporal cues in our A and AD treatments should become inhibitory and should offer some protection from extinction to the A and AD extensions. The question is whether such protection is better for A extensions than for AD extensions. In principle, the R-model predicts more protection for an A extension. The reason is that the salience of the postulated temporal cues should decrease as a function of the duration of their nonreinforced exposure. They are thus more salient when nonreinforced in compound with CS parts from an A extension than from an AD extension. This greater salience allows them to become more inhibitory in A procedures than in AD procedures. The more inhibitory they become, the more protection from extinction they offer. We sought to test these ideas in simulations that used the R-model and two stimuli, the target CS and a hypothetical postshock temporal cue. We chose values of the R-model's free parameters identical to those in Figure 2 (see also Ayres et al., 1987) and equated the starting salience of the postshock temporal cue to that of the target CS. Not knowing the duration of the postshock cue, we assigned it different durations in different simulations. We chose durations of 2, 10, and 40 min. In all three simulations, the R-model correctly predicted the direction of our A versus AD effect; however, the predicted difference in associative values was so small as to be apparent only beyond the third decimal place! Part of the reason for such a small difference is that the temporal cue could not become strongly inhibitory until after the target CS had become strongly excitatory. By that time, the temporal cue had been rendered nonsalient by nonreinforcement. By assuming a slightly higher starting salience for the temporal cue (on the grounds that it was produced by shock), we could enhance the predicted difference between A and AD, but it was still detectable only beyond the second decimal place. Nevertheless, the fact that the R-model can be forced to predict any difference between A and AD treatments is interesting and at-tests perhaps to the resilience of this simple model.

References

- Albert, M. L. (1987). *Extension of continuous versus discontinuous conditioned stimuli before versus after the unconditioned stimulus*. Unpublished master's thesis, University of Massachusetts at Amherst.
- Albert, M. L., & Ayres, J. J. B. (1989). With number of preexposures constant latent inhibition increases with preexposure CS duration or total CS exposure. *Learning and Motivation, 20*, 278–294.
- Annau, Z., & Kamin, L. J. (1961). The conditioned emotional response as a function of intensity of the US. *Journal of Comparative and Physiological Psychology, 54*, 428–432.
- Ayres, J. J. B., & Albert, M. (1990). Extending conditioned stimuli before vs. after unconditioned stimuli: Convergence of effect over trials. *Learning and Motivation, 21*, 399–414.

- Ayres, J. J. B., Albert, M., & Bombace, J. C. (1987). Extending conditioned stimuli before versus after unconditioned stimuli: Implications for real-time models of conditioning. *Journal of Experimental Psychology: Animal Behavior Processes*, *13*, 168–181.
- Ayres, J. J. B., Philbin, D., Cassidy, S., Bellino, L., & Redlinger, E. (1992). Some parameters of latent inhibition. *Learning and Motivation*, *23*, 268–287.
- Ayres, J. J. B., & Vigorito, M. (1984). Posttrial effects of presenting vs. omitting expected shock USs in the conditioned suppression procedure: Concurrent measurement of barpress suppression and freezing. *Animal Learning & Behavior*, *12*, 73–78.
- Donegan, N. H., & Wagner, A. R. (1987). Conditioned diminution and facilitation of the UR: A sometimes opponent-process interpretation. In I. Gormezano, W. F. Prokasy, & R. F. Thompson (Eds.), *Classical conditioning: III: Behavioral, neurophysiological, and neurochemical studies in the rabbit* (pp. 339–369). Hillsdale, NJ: Erlbaum.
- Fanselow, M. S. (1982). The postshock activity burst. *Animal Learning & Behavior*, *10*, 448–454.
- Frey, P. W., & Sears, R. J. (1978). Model of conditioning incorporating the Rescorla-Wagner associative axiom, a dynamic attention process, and a catastrophe rule. *Psychological Review*, *85*, 321–340.
- Gaioni, S. J. (1982). Blocking and nonsimultaneous compounds: Comparison of responding during compound conditioning and testing. *Pavlovian Journal of Biological Sciences*, *17*, 16–29.
- Gibbon, J., & Balsam, P. (1981). Spreading association in time. In C. M. Locurto, H. S. Terrace, & J. Gibbon (Eds.), *Autoshaping and conditioning theory* (pp. 219–253). San Diego, CA: Academic Press.
- Klopf, A. H. (1988). A neuronal model of classical conditioning. *Psychobiology*, *16*, 85–125.
- Larew, M. B. (1986). *Inhibitory learning in initial Pavlovian backward conditioning trials*. Unpublished doctoral dissertation, Yale University.
- Lubow, R. E. (1973). Latent inhibition. *Psychological Bulletin*, *79*, 398–407.
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*, 276–298.
- Mazur, J. E., & Wagner, A. R. (1982). An episodic model of associative learning. In M. L. Commons, R. J. Herrnstein, & A. R. Wagner (Eds.), *Quantitative analyses of behavior: Acquisition* (Vol. 3, pp. 3–39). Cambridge, MA: Ballinger.
- Moore, J. W., & Stickney, K. J. (1980). Formation of attentional-associative networks in real time: Role of the hippocampus and implications for conditioning. *Physiological Psychology*, *8*, 207–217.
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*, 532–552.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: II. Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Shipley, R. H. (1974). Extinction of conditioned fear in rats as a function of several parameters of CS exposure. *Journal of Comparative and Physiological Psychology*, *87*, 699–707.
- Sutton, R. S., & Barto, A. G. (1981). Toward a modern theory of adaptive networks: Expectation and prediction. *Psychological Review*, *88*, 135–170.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. W. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: MIT Press.
- 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: Erlbaum.

Submitted September 1992; revised January 1993; accepted January 25, 1993.