ROBERT A. RESCORLA
ALLAN R. WAGNER

3 A Theory of Pavlovian Conditioning: Variations in the Effectiveness of Reinforcement and Nonreinforcement

In several recent papers (Rescorla, 1969; Wagner, 1969a, 1969b) we have entertained similar theories of Pavlovian conditioning. The separate statements have in fact differed more in the language of their expression than in their substance. The major intent of the present paper is to explicate a more precise version of the form of theory involved, and to indicate how it may be usefully applied to a variety of phenomena involving associative learning.

The impetus for a new theoretical model is not generally a new datum which clearly disconfirms existing theory. It is more likely to be the accumulation of a salient pattern of data, separate portions of which may be adequately handled by separate existing theories, but which appears to invite a more integrated theoretical account. Such, at least, is the better description of the background of the present work.

In the sections which follow we will first describe certain data from our laboratories which exemplify the kind of observations which have encouraged the present theorizing. The theory will then be presented in sufficient detail to show how it may be applied to experimental situations involving a variety of Pavlovian conditioning arrangements. Finally, we will briefly discuss the theory in relationship to more conventional approaches.

BACKGROUND

The background data pattern embraces a considerable range of phenomena. At the core, however, is a rather simple set of observations involving Pavlovian conditioning with compound CSs.

Suppose we have inferential knowledge concerning the "associative

strength" of some stimulus element A and of a second element X. This generally requires that we know certain things about the organism's history of experience with the separate cues, and something about the organism's behavior in their presence. We may know, for example, that A is a CS which has frequently been paired with a US, and which consistently elicits a sizeable CR. And we may know that X is a novel CS which neither elicits a CR nor inhibits the occurrence of otherwise elicited CRs. In this case, we would commonly attribute a high excitatory strength to A and a zero strength to X.

Suppose further then, that A and X cues which have been arranged to have special strength characteristics are presented concurrently and the AX compound is either reinforced by a US or is nonreinforced. What effect will such an AX trial, or a series of similar AX trials, have upon the behavioral influence, or "associative strength" of X alone? The answer, it appears, is that the effects will depend in a systematic fashion not only upon the current strength of X, but also upon the current strength of A, and hence upon the net strength of the AX compound. For example, if X has a relatively low excitatory value, a series of AX reinforced trials will increase the CR eliciting characteristic of X much more when A is arranged to have a relatively low excitatory value than when A is arranged to have a high excitatory value. Similarly, a series of AX nonreinforced trials will decrease the CR eliciting characteristic of X, or will increase the CR *inhibitory* characteristic of X much more if A is arranged to have a relatively high excitatory value than if A is arranged to have a low excitatory value.

Support for such generalizations may be drawn from a number of sources (e.g., Kamin, 1968; Egger & Miller, 1962; Konorski, 1948; Pavlov, 1927). But, it will be convenient to use several experiments from our laboratories to indicate the systematic variation involved. We will first illustrate the manner in which the effects of reinforcement appear to depend upon the net strength of the compound, and then the manner in which the effects of nonreinforcement appear also to depend upon the net strength of the compound.

Variation in the effects of reinforcement. Wagner and Saavedra (Wagner, 1969b) trained three groups of 20 rabbits in an eyelid conditioning situation in which the US was a 100 msec. 4.5-ma. shock to the area of the eye. In the reference condition (Group II) Ss received 200 conditioning trials in which a 1100 msec. compound CS, consisting of a flashing light (A) and a tone (X), always terminated with reinforcement. Two additional groups received an equal number of reinforced AX trials, but also received 200 trials with the A cue alone, irregularly interspersed among the compound trials. In Group I, A alone was always reinforced; in Group III, A alone was always nonreinforced.

Immediately following training all Ss received 16 reinforced test trials with X presented alone for the first time. In each of the three training conditions X had been experienced an equal number of times, and had always been followed by reinforcement. However, the conditions had always been followed by reinforcement. However, the conditions were designed to encourage different degrees of conditioning or "associative strength" to A with which X was always experienced in compound. In comparison to Group II which received only compound trials, reinforcing A alone in Group I should have increased the associative strength of A and hence of the AX compound during training, whereas nonreinforcing A alone in Group III should have decreased the strength of A and hence of the AX compound during training. The question was whether X would be differentially responded to in the three conditions as a function of this differential experience with A.

Figure 1 summarizes the percentage test trial responses of the three groups to the X element. Also included for comparison are the percentages of conditioned eyeblink responses to the AX compound and to the A element, where appropriate, during the immediately preceding block of training trials. As may be seen, relative to Condition II increasing the associative strength of A in Condition I decreased conditioned responding acquired by X, whereas depreciating the associative strength of A in Condition III increased conditioned responding acquired


Figure 1. Median percentage conditioned eyeblink responses to an AX compound and to the A and X elements alone, in three groups receiving either no training with A alone (II), training with A alone reinforced (I), or training with A alone nonreinforced (III), contemporaneous with AX reinforced.

by X. This ordering of the treatments is not only statistically reliable, but is very reproduceable in different situations. Wagner (1969b) has reported essentially identical results from similar comparisons involving Conditioned Emotional Response (CER) training or discriminated bar-press training with rats.

Rescorla, in a previously unpublished experiment, obtained similar effects, but in a situation in which the associative value of the A cue was manipulated prior to the start of AX training. Four groups of rats were first trained to bar-press on a VI schedule for food reinforcement. The several groups then received different Pavlovian conditioning treatments with a 2-min. tone CS (A) and a 0.5 sec. 1-ma. foot shock, while confined in a separate shock chamber. These treatments were designed to establish different behavior effects to A. For all groups 12 presentations of A occurred in each of 5 2-hour conditioning sessions. In Group .8—0, A was trained to elicit fear by presenting the shock with a probability of .8 during the CS but never in its absence. For a second group (Group 0—.8) shocks occurred with a frequency of .8 per 2-min. interval in the absence of A but the onset of A signalled a 4-min. period free from shocks. This procedure could be expected (e.g., Rescorla, 1969) to make A a conditioned inhibitor of fear. The remaining two groups were control groups in which the conditioning treatments could be expected to leave A relatively neutral. Thus, Group Control 0—.8 received the same number of shocks as Group 0—.8, and the same number of exposures to A, but the two were uncorrelated in time. Group Shock received the same schedule of shocks as Group 0—.8, but never experienced A.

Following this conditioning to A alone, all Ss received 8 trials in which a flashing light (X) was presented in conjunction with A and the compound was reinforced with shock on a 50% schedule. Finally, on each of four test days following this compound conditioning all Ss received 4 nonreinforced test presentations of X alone while bar-pressing.

Figure 2 summarizes the results of these extinction test sessions, in the form of mean suppression ratios (Annau & Kamin, 1961). This ratio yields a value of zero when the CS completely disrupts bar-pressing, and a value of .5 when bar-pressing behavior is unaffected by the CS. Thus, the lower the value indicated, the more effective was X alone. It should also be noted that the behavior of the two reference groups (Group Shock and Group Control, 0—.8) was virtually identical throughout testing, so that the two groups have been combined in Figure 2.

Although X was experienced, and had been followed by reinforcement an equal number of times in the several groups, X was not similarly effective during testing. In comparison to the performance of the reference groups, pretraining cue A to elicit fear in Group .8—0 decreased the acquisition of fear to X as a result of the AX reinforcements. In addition, pretraining cue A to be an inhibitor of fear in Group 0—.8 increased the acquisition of fear to X as a result of the AX reinforcements.

Figure 2. Mean suppression ratio for X following AX reinforced trials. Groups .8–0 and 0–.8 had prior excitatory and inhibitory training, respectively, to A, while the reference groups had prior treatment not expected to influence the associative strength of A.

Variation in the effects of nonreinforcement. A study conducted by Wagner, Saavedra, and Lehmann (Wagner, 1969b) was designed to evaluate whether nonreinforcement would also have different effects upon a stimulus element, depending upon the strength of the compound in which the element was imbedded. The study was conducted in eyelid conditioning and generally employed parameters similar to those in the earlier Wagner and Saavedra study.

Thirty-six rabbits were first conditioned to three separate stimulus elements, which will be referred to as A, B, and X. Over the course of two days training there were 224 A, 28 B, and 224 X trials, irregularly ordered, in which the respective cues were presented alone and reinforced.

The A and B cues, by virtue of the different numbers of reinforcements in their presence, were designed to have different associative strengths, i.e., A was designed to be a relatively strong cue, and B a relatively weak cue by the end of acquisition. For half of the Ss A was a flashing light and B was a vibration applied to Ss' chest. For the remaining Ss the nature of the cues was reversed.

The X cue, which for all Ss was a 3160 hz tone, was the element of special interest. Immediately following acquisition, Ss were assigned to one of two treatment conditions and administered 32 extinction trials, in which X was presented and nonreinforced. For half of the Ss X was presented during extinction in compound with Ss' A cue, while for the remaining 18 Ss it was presented in compound with the B cue.

On the 32 trials immediately following the extinction phase, X was again presented alone to all Ss and was reinforced. Comparison of Ss' responding during this reacquisition phase with the level of responding to X at the end of original acquisition, allowed a determination of the decremental effects suffered as a result of the intervening extinction, with either of the two compounds containing X.

Figure 3 represents the mean percentages conditioned eyeblink responses to the several CSs during the three phases of the experiment. The acquisition functions which summarize the responding of all 36 Ss prior to differential treatment, indicate that there was appreciable ac-


Figure 3. Mean percentage eyeblink responses during three training phases, involving acquisition to each of three separate component CSs, extinction with one of two compounds formed from the acquisition components, and reacquisition to the component common to the two extinction compounds. (From Wagner, 1969b.)

quisition to X, and, importantly, different amounts of acquisition to the A and B cues. Further evidence that A and B attained different associative strengths may be seen in the extinction phase panel of Figure 3. That group which received X in compound with the A cue responded more frequently during extinction than did that group which received X in compound with the presumably weaker, B cue.

The data of major interest, however, are depicted in the reacquisition functions which summarize the subsequent responding to X alone, in each of the two treatment groups. As is apparent, there was *less* responding to X following the AX extinction than following the BX extinction condition. That group in which the 32 nonreinforced exposures to X involved a relatively strong compound containing the A cue, experienced a significantly greater decrement in responding to X than did that group in which the same nonreinforced exposures to X involved a relatively weak compound containing the B cue.

Nonreinforcement may not only cause a CS to lose its tendency to elicit conditioned responses, but under appropriate circumstances may cause a CS to become "inhibitory," i.e., to act so as to decrease the likelihood of otherwise elicited CRs. The circumstances which are known to favor this occurrence are in fact consistent with the data from the previous study. That is, while simply nonreinforcing a previously neutral cue in isolation is unlikely to make that cue inhibitory, consistently nonreinforcing the same cue when in compound with an otherwise excitatory cue can result in a "conditioned inhibitor" (e.g., Konorski, 1948). This fact may be viewed as further indicating that the "decremental" effects of nonreinforcement are greater, the greater the net associative strength of all of the cues which precede the nonreinforcement.

To further evaluate this proposition, Wagner and Saavedra, in a previously unpublished experiment, only slightly modified the procedure of the Wagner, Saavedra, and Lehmann study referred to above. During an initial acquisition phase, cues A and B were again trained, as a result of differential numbers of reinforced trials (240 vs. 8), to have different associative strengths, and a third cue C, necessary for the test phase, was also highly trained (548 trials).

Following such training, a novel cue X was introduced in compound with either A or B for different groups of 20 Ss, and the compound was nonreinforced. Sixty-four such nonreinforced trials were irregularly alternated with a similar number of trials in which the cue paired with X continued to be presented alone and reinforced.

The X cue should have become a conditioned inhibitor as a result of either training schedule (e.g., Pavlov, 1927; Rescorla & LoLordo, 1965). The question was whether X would become more inhibitory as a result of being nonreinforced in compound with the stronger A cue, as compared to the weaker B cue. This was evaluated by returning the C cue,

Figure 4. Mean percentage conditioned eyelid responses, in evaluation of the conditioned inhibitory properties of X in two groups. In one group X had been nonreinforced in compound with a relatively excitatory cue, A, while in the other group X had been nonreinforced in compound with a less excitatory cue, B.

and determining in both groups the reduction in responding to C when in compound with X. This final test phase involved 16 reinforced presentations of C and of the CX compound.

For all Ss, C was a flashing light, X a vibratory stimulus, and A and B dissimilar auditory cues, the identification of the two as A and B counterbalanced within experimental groups. Conditioned responding observed during the initial training phases was appropriate to the experimental intention that A have a greater associative strength than B: prior to the introduction of X all Ss were responding at a higher level to their A then to their B cue, and a similar difference was continued in the performance of the separate groups subsequently receiving A reinforced vs. AX nonreinforced or B reinforced vs. BX nonreinforced.

Figure 4 presents the data of major interest from the final test phase. The two groups responded at the same high level to the C cue alone. The addition of the X cue, however, decreased this responding considerably (and reliably) more in the case of that group which had previously experienced the nonreinforcement of X in compound with the relatively strong A cue, then in the case of that group which had experienced the

GROUPS

Figure 5. Median suppression ratio to C alone and the CX compound. Stimulus X had received prior inhibitory training contrasted with different intensities of the US.

nonreinforcement of X in compound with the relatively weak B cue.

In all of the previous studies the associative strength of the cue with which X was eventually treated in compound was manipulated by varying the schedule of reinforcements and nonreinforcements with respect to that cue alone. There are, however, other variables which should influence the learning which would accrue to such cues alone and it might be expected that these variables would have an influence similar to that produced by varying the reinforcement schedule. For example, in the Wagner and Saavedra inhibition study above, it might have been as effective to bring A and B to differential strengths as a result of the same number of pairings with a US, but with a higher intensity US associated with A than with B.

Such reasoning gains support from an unpublished experiment by Rescorla. Following VI food-rewarded bar-press training, three groups of 8 rats received CER conditioning, with a 1200 Hz tone alone (A) and a compound (AX) composed of this tone plus a flashing light. In total, 45 A and 75 AX trials were irregularly distributed over 30 training days.

For all Ss the AX trials were consistently nonreinforced. The groups differed in the intensity of a .5 sec. shock US which they received on the A alone trials, being either 0-ma. (nonreinforcement), .5-ma., or 1-ma.

In order to evaluate the inhibitory effects of X in each of the three treatments it was necessary to train the CER to an additional cue (C). This cue was a 250 Hz tone, introduced for all Ss after the differential treatment phase and reinforced with a .5-ma. shock on a 50% reinforcement schedule. Testing was then accomplished by evaluating the degree of suppression produced when X was presented in compound with C, as compared to C alone. Two C and two CX trials, each nonreinforced, were presented on each of 6 consecutive test days.

Figure 5 presents the median suppression ratios during testing under the two cue conditions in the three groups. Again it should be noted that the smaller the ratio value the more effective was the CS in disrupting bar-pressing. As may be seen, C alone produced equivalent degrees of suppression in all three groups. The addition of X had little effect upon suppression in the 0-ma. group but increasingly interfered with suppression in the .5-ma. and 1-ma. groups. There was a clear and statistically reliable tendency for X to be a *more* effective conditioned inhibitor of the CER as a result of AX nonreinforcement, the more intense the US with which the A cue alone was paired.

THE BASIC THEORY

The generalization which applies to all of the results in the previous section is that the effect of a reinforcement or nonreinforcement in changing the associative strength of a stimulus depends upon the existing associative strength, not only of that stimulus, but also of other stimuli concurrently present. It appears that the changes in associative strength of a stimulus as a result of a trial can be well-predicted from the composite strength resulting from all stimuli present on that trial. If this composite strength is low, the ability of a reinforcement to produce increments in the strength of component stimuli will be high; if the composite strength is high, reinforcement will be relatively less effective. Similar generalizations appear to govern the effectiveness of a nonreinforced stimulus presentation. If the composite associative strength of a stimulus compound is high, then the degree to which a nonreinforced presentation will produce decrements in the associative strength of the components will be large; if the composite strength is low, the effect of a nonreinforcement will be reduced.

Certain similarities and differences between these generalizations and Hull's postulates for growth of $_sH_R$ will be readily recognized. The changes in associative strength are acknowledged to depend upon current levels of that strength. However, the statements above assert that

changes in the strength of a stimulus depend upon the *total* associative strength of the compound in which that stimulus appears, whereas for Hull only the strength of the component in question was relevant. It is just this dependence upon total associative strength which is central to the theory we wish to develop here.

There are a variety of theoretical languages in which this central idea can be expressed. One rather peripheralistic formulation has been suggested by Rescorla (1969). He proposed that the change in CR conditioned to a CS, as a result of a CS–US pairing may depend upon the discrepancy between the CR actually evoked on that trial and the maximum CR which the particular US will support. The CR occurring on a trial arises from all of the stimuli present on that trial, not simply the CS in question. Increments in conditioning may be assumed to occur when the actual CR evoked on a trial is smaller than the maximum which the ensuing US will support. Correspondingly, decrements result when the actual CR is larger than the maximum CR. Rescorla has particularly emphasized the implication that Pavlovian conditioned inhibition can be established to a CS by presenting it at a time when the actual CR is larger than the maximum CR which the subsequent US will support.

A somewhat different version of the central notion, that conditioning depends upon the associative strength of all stimuli occurring on a trial, has been suggested by Wagner (1969a, b). Wagner couched his proposal in terms of the changes in "signal value" of a cue, an associative construct meant to embrace both the incremental effects of reinforcement and the decremental effects of nonreinforcement. Specifically, the changes in signal value as a result of a trial were assumed to be linear functions of the composite signal value resulting from all stimuli present on that trial. Separate sets of such linear functions were suggested to be appropriate for the cases of reinforcement and nonreinforcement. The resultant signal value of the stimulus would presumably be reflected in the overt CR, although the specific relationship was not treated.

A less completely formulated version of these ideas has been suggested by Kamin (1968). Indeed, it was Kamin's notions concerning the "surprisingness" of a US that originally encouraged the formulations of Rescorla and Wagner. Attempting to account for his data on the so-called blocking effect, Kamin argued that conditioning will occur only when the US event is somehow "surprising" for the animal. Although the conditions which produce this surprise were not detailed, Kamin clearly intended that the surprise generated by a US be assumed to be reduced if that US is preceded by a CS which has previously been paired with it. Consequently, the surprise generated on a CS–US trial (and the resulting increment in conditioning to the CS) should depend upon the degree to which all stimuli present predict the US which occurs. It is not clear

from Kamin's formulation what should be the consequences for conditioned responding when the animal is variously "surprised" by the nonoccurrence of a US.

The central notion suggested here can also be phrased in somewhat more cognitive terms. One version might read: organisms only learn when events violate their expectations. Certain expectations are built up about the events following a stimulus complex; expectations initiated by that complex and its component stimuli are then only modified when consequent events disagree with the composite expectation.

A more precise formulation of the theory. It should be clear that these formulations all express the same core idea. They all generate essentially similar expectations with respect to the variable effects of reinforcement, as reported in the previous section. However, the ability of any of these formulations to make specific predictions is limited by their imprecise verbal nature. It has seemed profitable to us to ask whether, if we make more specific, formal assumptions around the central notion involved, we could expand the possibilities for experimental evaluation. In what follows we will attempt one such specification; the formulation follows most closely Wagner's (1969a, b) version of the theory.

As indicated above, one way to look at the central notion of this theory is as a modification of Hull's account of the growth of $_sH_R$. Similarly, one way to view the particular formalization to be proposed is as a modification of the mathematical model most closely related to the Hullian theory, the linear model. This model (e.g., Bush & Mosteller, 1955) specifies the changes in probability of a response as a result of a trial by the following equation:

$$\Delta p_n = \beta(\lambda - p_n).$$

where β is the learning rate parameter, p_n the probability of a response on trial n, and λ the asymptote of learning. The particular values of β and λ are determined by the US and CS events involved on the trial. Clearly the model incorporates the basic Hullian assumption that the increment (or decrement) in learning on each trial is dependent upon the amount already conditioned at the beginning of that trial as well as upon the final asymptote of learning which that US will support. Notice, however, that the model specifies the rules for growth in response probability while Hull's equations are for growth of habit strength, $_sH_R$.

The model we wish to propose constitutes a modification of the linear model in several ways. First, it describes the learning curves for strength of association, not response probability. In that sense it is more in line with the Hullian theory than is the linear model. Independent assumptions will necessarily have to be made about the mapping of

associative strengths into responding in any particular situation. Secondly, we will explicitly recognize that learning is tied to various external stimuli and discuss associative strength to various stimuli. In recognition of these two modifications, we will describe the model in terms of V_i, the strength of association to stimulus i.

It is also important to note that V_i will be allowed to take on both positive and negative values, corresponding roughly to conditioned excitation and conditioned inhibition. But, the most significant departure from the linear model is that when a stimulus compound, AX, is followed by a US, the changes in the strength to each of the component stimuli, A and X, will be taken to be a function of V_{AX}, i.e., the strength of the compound, rather than the strength of the respective components.

When a compound, AX, is followed by US_1, the changes in associative strength of the respective components may be represented as:

$$\Delta V_A = \alpha_A \beta_1 (\lambda_1 - V_{AX})$$
$$\text{and}$$
$$\Delta V_X = \alpha_X \beta_1 (\lambda_1 - V_{AX}).$$

If AX is followed by a different valued US, i.e., US_2, which may include 0 or nonreinforcement, the changes in associative strength of the respective components may be represented as:

$$\Delta V_A = \alpha_A \beta_2 (\lambda_2 - V_{AX})$$
$$\text{and}$$
$$\Delta V_X = \alpha_X \beta_2 (\lambda_2 - V_{AX}).$$

As may be seen in the above equation, there are three sets of parameters which affect the magnitude of the changes involved. The alphas are learning rate parameters, each associated with one component stimulus, and are appropriately subscripted to indicate this identification. The value of alpha roughly represents stimulus salience and indicates our assumption that different stimuli may acquire associative strength at different rates despite equal reinforcement. The betas are learning rate parameters associated with the USs. The assignment of different beta values to different USs indicates our assumption that the rate of learning may depend upon the particular US employed. Alpha and beta values are confined to the unit interval, $0 \leq \alpha, \beta \leq 1$. Finally, the λ values represent the asymptotic level of associative strength which each US will support; presumably different USs will yield different asymptotic levels. Although λ is not formally bounded, changing the range of its permissible values simply shifts the scale on which we observe Vs.

In order to apply the model, two further specifications are needed. The associative strength of the compound, V_{AX}, must somehow be specified in terms of the strengths of the components. The simplest assumption, and the one we will make here, is $V_{AX} = V_A + V_X$. Notice

that although the Vs are in principle unbounded, in application the λ values set limits on the compound Vs.

Secondly, we need to provide some mapping of V values into behavior. We are not prepared to make detailed assumptions in this instance. In fact, we would assume that any such mapping would necessarily be peculiar to each experimental situation, and depend upon a large number of "performance" variables. For the analyses we wish to present in this paper, it will generally be sufficient simply to assume that the mapping of Vs into magnitude or probability of conditioned responding preserves their ordering. Stimulus compounds whose net V is negative would all be expected generally to map into a zero CR, but differential negative values could also be distinguished among by a variety of experimental procedures (Rescorla, 1969).

ELEMENTARY DEDUCTIONS FROM THE THEORY

Without making more specific assumptions about parameter values, certain general deductions can be made from the model. It should be clear that for the case of repeated reinforcement or nonreinforcement of a single cue, A, the equations reduce to essentially the linear model. For instance, as V_A increases with repeated reinforcement of A, the difference between V_A and λ will decrease. Consequently, increments in V_A will decrease and a negatively accelerated learning curve will result with an asymptote of λ. Similarly, if we assume that the λ value associated with nonreinforcement is lower than V_A, then a negatively accelerated extinction function is generated by repeated nonreinforcement of A.

Reinforcement of compound stimuli. But the more interesting cases result from compound stimuli, as in the experiments of the previous section. Consider first the case of reinforcement of an AX compound. The experiments of the previous section, together with those of Kamin (1968), indicate that prior conditioning of A reduces the degree to which reinforcement of an AX compound increments the associative strength of X. From the above equations it is clear that changes in V_X are governed by the difference between λ and the composite V_{AX}. The result of prior conditioning to A is that V_A, and thus V_{AX}, is large; hence the difference between λ and V_{AX} is reduced and the effectiveness of reinforcement correspondingly limited. Similarly, the prior establishment of A as an inhibitor, as in Rescorla's 0–.8 group means that V_A is negative. As a consequence, V_{AX} is reduced and the difference between λ and V_{AX} enlarged; thus X can be incremented proportionately more through reinforcement.

The arguments for the Wagner and Saavedra experiment are essentially similar; here A is not pretreated, but V_A is modified by inter-

spersing trials of A alone with reinforced AX trials. In the case where A is reinforced on those intermixed trials, again V_A will be large and result in an enlarged V_{AX}, thus limiting the amount of conditioning which can accrue to X on the compound trials. Early in conditioning, reinforcements will occur to AX while V_{AX} is still below asymptote and consequently V_X will increase initially. Nevertheless, V_X will eventually decrease to zero. Since V_A will increase toward λ, as a result of the A alone trials, V_{AX} will come to exceed λ. Notice that when this happens, the result of a *reinforced* AX trial is to *decrement* the associative strength of the components. As A and AX are both reinforced, increments to A will occur on the reinforced A trials and decrements to A and X on the reinforced compound trials. The result will be a transfer to A of whatever associative strength X may have initially acquired. It is an important characteristic of the model that even on reinforced trials, if V_{AX} exceeds λ, A and X will be decremented.

A similar account can be given of the results of nonreinforced presentations of A alone in the Wagner and Saavedra study. These presentations should lead to a reduction of V_A and hence a reduction of V_{AX}. This provides increased opportunity to condition X on the AX trials, as compared to a condition involving only reinforced AX trials.

Kamin (1968) has provided considerable additional data for the particular case in which AX reinforcement is preceded by a history of reinforcement of A. His experiments, carried out in a CER situation, indicate that with a high degree of prior conditioning to A, reinforcement of AX can be rendered almost completely ineffective in conditioning fear to X.

Several variations in the treatment of A, however, were found to attenuate the ability of A to "block" conditioning of X. For instance, as the number of prior conditioning trials to A was reduced, the ability of A to block the conditioning of X was lessened. Alternatively, if A was first highly conditioned but then extinguished, the extinction disrupted A's ability to block. Finally, if the intensity of A was decreased, blocking was reduced.

All of the latter manipulations might be expected to yield a lower V_A and thus a lower V_{AX} at the time of reinforcement of the compound. This deduction should require no elaboration in the cases where number of reinforced or extinguished trials to A alone was manipulated. In the case of decreased A intensity, it is only necessary to make the reasonable assumption that a lower V_A was attained prior to compound training as a result of a lower α associated with the weaker stimulus.

Kamin also reported that the nature of the US at the time AX is introduced is critical. For instance, if the prior conditioning of A is done with a 1-ma. shock and then AX is followed by a 1-ma. shock, a large interference with the conditioning of X is observed. But if the

AX compound is followed instead by a 4-ma. shock, considerably more conditioning to X results. This conditioning to X depends not simply upon the use of a high shock intensity, but requires an increase in shock intensity from the conditioning of A to that of AX. Alternatively, if A is followed by a single shock and then AX is followed by two shocks in close succession, similar conditioning to X results.

There is a natural way for the present model to handle these outcomes. There is evidence available (e.g., Annau & Kamin, 1961) that higher asymptotes of conditioning result from higher shock intensities. Thus it would not be unreasonable to assume that the λ value associated with a 4-ma. shock is larger than that associated with a 1-ma. shock. The result of increasing shock intensity when shifting from conditioning of A to conditioning of AX is that the potential for conditioning X is enhanced; consequently, the reduction of the blocking is not surprising. Notice that it is the increasing of the λ value between the two stages of the experiment, rather than simply having a larger λ throughout, that is critical to this prediction. A similar kind of reasoning might be applied to the case of shifting from a single shock following A to a double shock following AX.[1]

There are aspects of Kamin's data, however, with which the present model does not deal so well. For instance, Kamin found that the initial trial on which X is presented in conjunction with A generally yields less suppression than the previous A-alone trial. Furthermore, he provided evidence that most of the conditioning to X which occurs in his paradigm results from that first AX trial. The present theory provides no statement of performance axioms which might lead us to expect the introduction of X to interfere with suppression to A unless it is assumed that X is initially associated with a negative V. More problematic is the fact that while the theory predicts that the first AX trial will produce more X conditioning than any subsequent compound trial, it does not anticipate Kamin's claim that the first AX trial accounts for *all* of the conditioning to X. Further analysis of this problem must await additional data collection as well as the development of more detailed performance statements for the theory.

Kamin has also investigated a phenomenon related to the blocking effect, so-called "overshadowing." If an AX compound is repeatedly reinforced and A is simply a more salient stimulus, little conditioning

---

[1] It might be noted in passing that just as the present model predicts a possible increment in associative strength of X when the reinforcement magnitude for AX is increased over that for X, so it predicts a decrement in associative value of X when the reinforcement magnitude for AX is *decreased* with respect to that for X. In that case, λ will be lowered and assuming that V_A has been made to approach the preshift λ_1, V_{AX} will be greater than the post-shift λ, resulting in decrements to both A and X. If X begins with no associative strength, this procedure might be expected to produce a conditioned inhibitor. Experiments are currently underway in our laboratories to investigate this possibility.

to X may result. Even though A is not reinforced more frequently than X, because it is a more salient stimulus it may overshadow X and interfere with development of associative strength to the latter cue. Kamin has demonstrated that the degree to which A overshadows X is dependent upon the relative intensities of the two stimuli, relatively more intense stimuli yielding more overshadowing. From the point of view of the present model, the effect of having A be a more salient stimulus is that it will have a larger alpha value. Hence when AX is reinforced, V_A will grow rapidly with respect to V_X. Consequently, the more salient A is with respect to X, the greater proportion of V_{AX} will be due to V_A and the more limited the conditioning to X. More precisely, it is expected that when V_A and V_X both begin conditioning at zero, that after any large number of conditioning trials with AX reinforced that V_X will be equal to $\dfrac{\alpha_X}{\alpha_A + \alpha_X} V_{AX}$.

Notice that, nevertheless, the more potent the US used, the higher should be λ and the more conditioning to AX and hence to X should result. Thus, overshadowing, measured in terms of absolute responding to X, should be attenuated by employing greater US magnitudes, a finding confirmed by Kamin (1968).

Nonreinforcement of compound stimuli. The preceding paragraphs indicate that the present model is capable of integrating a considerable amount of data on the effects of reinforcement, both from our own laboratories and those of others. We now turn to an account of some elementary effects of nonreinforcement. Because of the symmetry of the model, the arguments for nonreinforcement are analogous to those for reinforcement and can be presented briefly.

As pointed out above, if we assume that the λ value associated with a US of zero intensity is zero, then the model naturally generates a negatively accelerated extinction function. Considering the nonreinforced presentation of a compound stimulus, the changes in component associative strength should be dependent upon the total V of the·compound: the larger V_{AX}, the larger should be the decrement expected in both A and X as a result of a nonreinforced presentation of the compound. Consequently, any operation which enhances V_A and hence V_{AX} should result in a larger decrement to X as a result of nonreinforcement of AX. This prediction is consistent with the findings from the two eyeblink conditioning studies from Wagner's laboratory reported above, in which the number of prior reinforcements of A critically affected the decrementing of X. In fact, in the Wagner and Saavedra study where X was presumably introduced with a zero V, greater conditioned inhibition accrued to X when it was nonreinforced in conjunction with that stimulus which had been more frequently reinforced in the past. If

we assume that different intensities of the US result in different levels of conditioning, the conditioned inhibition results of the Rescorla experiment also fall into place. If A is followed by a more intense US and as a consequence V_A is larger, then nonreinforcement of X in the presence of A should result in greater conditioned inhibition to X.

There is an additional case of nonreinforcement which is of interest to consider even though no data are currently available. Suppose that stimulus A were pretreated so as to give it a negative V and that a novel cue, X, were then combined with it and the AX compound nonreinforced. The value of V_{AX} should then be negative, and assuming that the λ associated with nonreinforcement is zero, $\lambda - V_{AX}$ should be positive. Thus *nonreinforcing* X in conjunction with an inhibitor should give X *positive* associative strength. This prediction, however paradoxical, should further emphasize the kind of symmetry inherent in the model's treatment of nonreinforcement and reinforcement.

APPLICATION TO A PROBLEM IN DISCRIMINATION LEARNING

In the foregoing applications of the model, reasonably adequate predictions from the basic theory might well have been drawn without benefit of any quantitative formulations. The cue of interest was never presented in compound with more than a single additional stimulus, and it would generally have been possible and sufficient, in order to account for the data discussed, to specify that the associative strength of the latter stimulus had been manipulated to have various ordered degrees of excitatory or inhibitory value.

In other instances to which the theory should be applicable, however, it is not possible to proceed at such a level. The present section will attempt to illustrate one such instance, in which certain quantitative assumptions become critical.

The problem to be considered involves the differential reinforcement of stimulus compounds containing a so-called common cue. Suppose a compound CS, composed of experimentally isolatable cues A and X is consistently reinforced, while another compound CS composed of cues B and X is consistently nonreinforced. If the several components are adequately discriminable we should, of course, expect that differential associative strengths will be acquired, with V_{AX} approaching that asymptote appropriate to the US employed, and V_{BX} approaching that asymptote appropriate to nonreinforcement. But what should be the expected fate of V_X?

What makes this question especially interesting is that there is reason to believe (e.g., Wagner, Logan, Haberlandt, & Price, 1968) that a cue occupying the place of X in an AX, BX discrimination will come to be less responded·to alone than might be expected simply on the

basis of the schedule of reinforcement and nonreinforcement with which it is associated. And, several theories have been advanced (e.g., Restle, 1957; Sutherland, 1964; Mackintosh, 1965) which propose that such a common, or "irrelevant" cue will become "adapted," "neutralized," or "unattended-to," by virtue of the availability of more valid cues.

The present theory includes no such special "neutralization" assumption, but indicates only that the learning which accrues to X should be a function of the trial by trial strength of both AX and BX in relationship to the reinforcing events with which they are separately associated. But, should X then be expected to become neutral? Or, at least more "neutral" than in various comparison treatments not involving a discrimination?

In order to illustrate certain relatively general features expected in the course of discrimination learning when AX trials are alternated with BX trials, the former followed by reinforcement, and the latter by nonreinforcement, a sample learning run was computed. It should be recalled that the theory specifies that,

$$\Delta V_A = \alpha_A \beta_1 (\lambda_1 - V_{AX}),$$
$$\Delta V_X = \alpha_X \beta_1 (\lambda_1 - V_{AX})$$

when AX is reinforced, and,

$$\Delta V_B = \alpha_B \beta_2 (\lambda_2 - V_{BX}),$$
$$\Delta V_X = \alpha_X \beta_2 (\lambda_2 - V_{BX})$$

when BX is nonreinforced.

For purposes of this example, A, B, and X were assumed to begin with Vs of zero prior to the first learning trial, and were assumed to be equally salient ($\alpha_A = \alpha_B = \alpha_X = 1.0$). Reinforcement and nonreinforcement were taken to be associated with λs of 1.0 and 0 respectively, and with equal rate parameters ($\beta_1 = \beta_2 = .05$). Any set of parameter assumptions identifies a special case, but this set was intended to have some semblance to certain experimental arrangements (e.g., the initial Vs of zero might be relatively typical of untrained cues) and to avoid any assumptions that would beg justification (e.g., unequal αs or βs).

The left panel of Figure 6 depicts the mean V_{AX} and mean V_{BX} computed over successive blocks of 4 trials. The right panel of Figure 6 depicts the corresponding V values of the separate components. Several features of the plotted functions should be noted. With all Vs beginning at zero, V_{AX} and V_{BX} both increase over the early trials. Although the rate parameters β_1 and β_2 associated with reinforcement and nonreinforcement were set equal, $\lambda_1 - V_{AX}$ is initially large, whereas $\lambda_2 - V_{BX}$ is concurrently small. As a consequence V_X as well as V_A increase rapidly in relationship to a slower decline in V_B. Only when the absolute value of $\beta_1(\lambda_1 - V_{AX})$ becomes equal to $\beta_2(\lambda_2 - V_{BX})$ does V_X cease to grow

Figure 6. Mean Vs computed from a sample learning run in which AX reinforced trials were alternated with BX nonreinforced trials. The left panel depicts the compound theoretical values, the right panel the corresponding component values.

and does the rate of decrease in V_B equal the rate of increase in V_A. As trials continue, V_{AX} and V_{BX} necessarily approach λ_1 and λ_2, so that in this example, the terminal value of V_A is sufficiently positive so that $V_A + V_X = 1$, and the value of V_B is sufficiently negative that $V_B + V_X = 0$.

It is unnecessary for present purposes to completely detail the way in which this picture changes as all the values of the model are manipulated. For example, whereas V_{AX} and V_{BX} will always approach λ_1 and λ_2, the terminal values of V_A, V_B, and V_X in relationship to λ_1 and λ_2 will depend appreciably upon their initial starting values. But, for the conclusions we wish to draw, we can restrict ourselves without danger to instances like that exemplified, in which all Vs begin at zero.

It is worth noting, however, the effects of variations in α and β. For instance, if β_1 is made larger than β_2, the initial positive course of V_{BX}, associated with the nonreinforced compound, would attain a higher absolute level and be more protracted. Similarly, V_X would initially rise to a higher value and then fall. However, for all values of β_1 and β_2

greater than zero, the *asymptotic* levels of V_A, V_X, and V_B as well as V_{AX} and V_{BX} are independent of the βs.

Variation in the relative saliences of A, B, and X has a more marked effect. For example, if α_X is made larger relative to α_A and α_B, the formation of the discrimination in V_{AX} and V_{BX} is not only slowed, but the terminal values of V_A, V_B, and V_X are modified. It can be shown that under the conditions, $\lambda_1 = 1$, and $\lambda_2 = 0$, and where $\alpha_A = \alpha_B$, the asymptotic value of V_X is equal to $\dfrac{\alpha_X}{\alpha_A + 2\alpha_X}$.

The picture of discrimination learning in V_{AX} and V_{BX} shown in Figure 6 resembles many empirical functions. Especially satisfying is the initial rise predicted in the strength of the nonreinforced compound, a phenomenon which is frequently observed (e.g., Wagner, 1968). It is evident, however, that the theory does *not* generally predict that the associative strength of a common cue such as X will attain a value of zero. Rather as indicated above, X should attain some associative strength depending upon its relative salience in comparison to that of the discriminative cues. It seems, in fact, that an important expectation from the model for the case described is that X will have some positive value, such that discriminative cues, in the nonreinforced compound must become *inhibitory* in order for the strength of the latter compound to approach zero.

Still, it may be more informative to ask whether X should be more "neutral" as a result of being imbedded in such a discrimination than as a result of other treatments involving the same associated schedule of reinforcement.

We may simply declare, that the strength of X should, according to the model, generally be less following discrimination training than following an identical partial reinforcement schedule of X in isolation. But, granted the discussion thus far, this is a rather uninteresting comparison, since it can largely be viewed in terms of the overshadowing that occurs when a cue is trained in compound with other cues, as compared to being trained in isolation.

A more interesting empirical comparison was provided by Wagner, Logan, Haberlandt, and Price (1968). In both CER and eyelid conditioning, these investigators compared responding to an isolatable common cue occupying the place of X in an AX, BX discrimination, with the responding to a similar cue, experienced in a "pseudodiscrimination" treatment in which AX and BX were both partially reinforced on a 50% schedule. Although X in the two treatments was associated during training with the same reinforcement schedule, and in compound with the same cues, it was much more responded to when tested alone following pseudodiscrimination as compared to discrimination training.

We have already specified the asymptotic value of V_X to be expected in an AX, BX discrimination in which each compound is experienced on half of the trials. It will thus be useful to compare this expected value with the comparable value expected following a pseudodiscrimination procedure in which half of each of the AX and BX trials are reinforced and half nonreinforced, in the manner of Wagner, et al. (1968).

It can be shown that according to the model, the asymptotic value approached by a partially reinforced compound should be equal to

$$\frac{\pi\beta_1\lambda_1 - (\pi-1)\,\beta_2\lambda_2}{\pi\beta_1 - (\pi-1)\,\beta_2}$$

where π is the proportion of reinforced trials. Adopting, as we have previously, the assumption that $\lambda_1 = 1.0$ and $\lambda_2 = 0$, then in the case of a 50% reinforcement schedule this becomes $\dfrac{\beta_1}{\beta_1 + \beta_2}$

This quantity thus expresses the theoretical asymptotic value of V_{AX} and V_{BX} in the pseudodiscrimination case under consideration.

Of course $V_{AX} = V_A + V_X$ and $V_{BX} = V_B + V_X$ and it can be demonstrated that in the instance in which X is presented on all trials, and A and B each on half the trials that,

$$V_X = \frac{2\alpha_X}{\alpha_A + 2\alpha_X}\,V_{AX} = \frac{2\alpha_X}{\alpha_B + 2\alpha_X}V_{BX}, \text{ assuming that } \alpha_A = \alpha_B.$$

Thus the asymptotic V_X attained under the pseudodiscrimination procedure should be:

$$\frac{2\alpha_X}{\alpha_A + 2\alpha_X}\left(\frac{\beta_1}{\beta_1 + \beta_2}\right).$$

An appreciation of the relative V_X expected in the discrimination and pseudodiscrimination treatments is most evident if we now express the asymptotic V_X in the pseudodiscrimination condition, minus the asymptotic V_X in the discrimination condition, which difference is:

$$\frac{\alpha_X}{\alpha_A + \alpha_X}\left[2\left(\frac{\beta_1}{\beta_1 + \beta_2}\right) - 1\right].$$

Wagner (1969b) has noted that the present form of the theory can account for the greater responding to X alone under a pseudodiscrimination as compared to a discrimination treatment, only if one is willing to specify certain quantitative differences in the effects of reinforced and nonreinforced trials. The mathematical expression above makes this evident in terms of the current model. Only when the rate parameter

associated with reinforcement (β_1) is greater than the rate parameter associated with nonreinforcement (β_2) will the quantity expressed be positive, i.e., will V_X be greater in the pseudodiscrimination than in the discrimination treatment. The above expression also indicates that the *magnitude* of this effect will depend upon the cue saliences, so that any difference between the two conditions will be augmented as α_X approaches 1.0 and α_A becomes small.

The robustness of the effect demonstrated by Wagner, Logan, Haberlandt, and Price, would suggest that β_1 is considerably larger than β_2 in the situations which they employed. How adequate this assumption will otherwise turn out to be remains to be determined. It is worth noting, however, that a similar assumption has frequently been deemed necessary in the application of related models to other data areas (e.g., Bush & Mosteller, 1955; Lovejoy, 1966). The advantage of the present analysis is that it makes clear the kinds of additional quantitative assumptions which are necessary in order to account for the differences in strength of the common cue in these different treatments.[2]

APPLICATION TO BACKGROUND STIMULI AS COMPONENTS

Although the present model is stated in terms of increments and decrements in the associative strength of component stimuli as a result

2 Because of the summation assumption, it may appear that the present model will be inappropriate for certain cases of discrimination learning. For instance, there is some evidence that organisms are capable of learning to respond to a compound stimulus while withholding their response to its components, and vice versa (Woodbury, 1943). Such discriminations seem to call for an appeal to some special characteristic of the compound not present in its components. However, Mr. Donald Rightmer has suggested to us a way of conceptualizing such discriminations without recourse to "configuring." Consider a compound composed of two component stimuli. These components, although readily discriminable from each other, will nevertheless contain some common properties. To indicate this, we may describe the components as AX and BX and the compound formed from their combinations as ABX. If we apply the model to a case in which this compound is consistently reinforced and these components nonreinforced, it correctly predicts learning of that discrimination. In this case V_{ABX} approaches λ_1 while V_{AX} and V_{BX} both approach λ_2, as a result of V_A and V_B both approaching $\lambda_1-\lambda_2$, while V_X approaches $2\lambda_2-\lambda_1$. Much of the burden of the learning rests with the common parts of the component stimuli and consequently the rate of learning will depend upon the assumed salience of X. A similar result occurs for the case of reinforcement of AX and BX and the nonreinforcement of ABX. We mention these examples only to indicate that at least one kind of evidence commonly cited in criticism of summation notions is compatable with the present model.

A common approach to discrimination learning (e.g., Estes & Burke, 1953) is to assume, however, not only that any pair of CSs can be theoretically conceptualized as being composed of unique and common cues, but that the discriminability of the CS depends only upon the relative weights of the two sets of elements. While this latter manner of analysis might appear especially congenial to the present theory, it presents sufficient difficulties that we would prefer not to commit ourselves to this more general strategy. A consideration of the alternatives would, unfortunately, take us beyond the scope of the present paper.

of reinforcement and nonreinforcement of compounds, it also has implications for situations not obviously involving compound stimuli. One interesting such application is to data recently collected by Rescorla (1969) pointing to the importance of CS-US correlations in Pavlovian fear conditioning. Consider a situation in which an animal receives brief, intense electric shocks randomly distributed in time. Suppose further, that tonal stimuli are presented irregularly without regard to the occurrence of the shocks, i.e., in such a way that shocks may occur in both the presence and absence of the CS, and there is no correlation between the CS and shock. The question of interest is to what degree the tones will acquire associative strength.

A typical experiment asking this question employed a CER procedure with rats (Rescorla, 1968, Experiment 1). Three groups of animals received tone CSs and shock USs. Group 1 received the tones and shocks in random relation to each other, with shocks occurring both in the presence and absense of the CS. Group II received the identical treatment except that all shocks programmed to occur in the absence of the CS were omitted. Notice that these two groups received the same number of shocks during the CS; they differed only in that the first group also received shocks at other times. Finally, a third group received the same reduced number of shocks as Group II, but those shocks were distributed randomly in time, in the manner of Group I. When these stimuli were subsequently presented while the rats bar-pressed for food reward, only Group II showed fear of the CS. The two groups for which the CS and shock were independent showed no measurable fear conditioning.

This result suggests that the correlation of the CS and US, in addition to the number of reinforced CS presentations, is an important determinant of fear conditioning. One way to describe this correlation is in terms of the probability of occurrence of the US in the presence and absence of the CS. When the two events are positively correlated, then the probability of shock is higher during the CS than in its absence; when they are uncorrelated, then those probabilities are equal. Furthermore, this description suggests a third case: when the probability of the US is higher in the absence of the CS than in its presence, the two events are negatively correlated. In a series of experiments, Rescorla (1969) has accumulated evidence that these relative probabilities are important in determining the amount of conditioning obtained. According to that evidence, when the probability of shock is higher during the CS than in its absence, the CS becomes a conditioned elicitor of fear; when the CS signals a period which is relatively free from otherwise probable shocks, it becomes a conditioned inhibitor of fear. Finally, when the probabilities of shock are equal in the presence or absence of the CS, little or no conditioning of either sort occurs.

Somehow the organism appears to evaluate the probability with

which shocks occur both in the presence and in the absence of the CS, and it is the relation between these two probabilities which determines the amount of fear conditioning observed to the CS. The organism is apparently behaving as a relatively complex probability comparitor. What we wish to suggest here is that the present model may provide one way of understanding how the animal can be sensitive to such subtle relations via a relatively simple process.

The important point to notice for this analysis is that the CS occurs against a background of uncontrolled stimuli. To speak of shocks occurring in the absence of the CS is to say that they occur in the presence of situational stimuli arising from the experimental environment. Although these stimuli are not explicitly manipulated by the experimenter, they nevertheless can be expected to influence the animal. Thus, one way to think about the occurrence of the CS is as an event transforming the background stimulus, A, into background-plus-CS, AX. The present model, of course, has been designed to account for the conditioning of X when it appears in such a compound, as a function of the treatment of A elsewhere.

In order to exemplify the application of the model to this particular case, the experimental session was taken to be divisible into time segments the length of the CS duration. Each segment containing the CS is thus treated as an AX "trial" and each segment not containing the CS as an A "trial." It is possible then to specify the sequence of reinforcement and nonreinforcement over each of the two kinds of trials.

Sample learning runs were computed from the model with schedules of background alone (A) and background-plus-CS (AX) as might be the case in experiments such as those of Rescorla (1968). Figure 7 shows the results of one such application of the model. This figure describes the V value of the CS over trials, as a function of different shock probabilities in the presence and absence of the CS. The first digit labeling each curve indicates shock probability during the CS; the second the probability in the absence of the CS. The particular parameter selections used in arriving at the functions plotted were as follows: The CS was assumed to be present 1/5 of the time according to an irregular sequence and to have a salience 5 times that of the background ($\alpha_A = .1$, $\alpha_X = .5$); the λ values associated with reinforcement and nonreinforcement were taken to be 1 and 0, respectively, while the rate parameter associated with reinforcement was set at twice that associated with nonreinforcement ($\beta_1 = .1$, $\beta_2 = .05$).

The asymptotic values of the functions represented in Figure 7 are in general agreement with Rescorla's data in that they are clearly ordered by the relative probability of shock in the presence and absence of the CS. In addition, positive V values are associated with positive correlations between the CS and shock; negative V values are associated with negative correlations. Furthermore, the magnitude of the correlation

Figure 7. Predicted strength of association of X as a function of different US probabilities in the presence and absence of X. The first number next to each curve indicates the probability of the US during the CS; the second, the probability in the absence of the CS.

may be seen to be important, with stronger correlations generating Vs more removed from zero. Finally, the asymptotic V of a CS uncorrelated with the occurrence of shock is zero.

In fact, it is possible to arrive at a relatively simple expression describing the asymptote of conditioning ($\dot{V}_X$) for these various treatments. The equation for the asymptotic value of V_A is:

$$\dot{V}_A = \frac{\pi_A \beta_1}{\pi_A \beta_1 - (1-\pi_A) \beta_2}.$$

Similarly, the equation for the asymptotic value of V_{AX} is:

$$\dot{V}_{AX} = \frac{\pi_{AX} \beta_1}{\pi_{AX} \beta_1 - (1-\pi_{AX}) \beta_2}.$$

Thus, one may arrive at the asymptotic value of V_X as:

$$\dot{V}_X = \dot{V}_{AX} - \dot{V}_A$$

It may be seen that $\dot{V}_X$ will depend only upon the probability of reinforcement in the presence of X (π_{AX}) and in the absence of X (π_A), as well as upon the rate parameters associated with reinforcement and nonreinforcement. Since the latter parameters are constants in the two equations for $\dot{V}_{AX}$ and $\dot{V}_A$ it should be evident how $\dot{V}_X$ will vary with the relative probabilities of reinforcement. When π_{AX} is greater than π_A, $\dot{V}_X$ will have a positive value, as π_{AX} becomes equal to π_A, $\dot{V}_X$ will approach zero, and finally when π_A is greater than π_{AX}, $\dot{V}_X$ will become negative.

Although we will discuss below the effects of other variables upon V_X prior to asymptote, none of these influence the final product of learning. In particular it is worth pointing out that in this instance the initial Vs, at the begining of any of the probability treatments, leave no permanent effect; whatever the starting V_A and V_X, a given treatment will eventually yield asymptote values, as specified, appropriate to that treatment. For example, should V_X first be incremented to some high value prior to a .8—.8 "extinction" treatment, the asymptotic value of X would still be zero according to the model.

Whatever the overall shock probability, if AX and A are reinforced with equal probability, the value of X will approach zero. Notice, however, that this zero level of conditioning for X occurs against different levels of conditioning to A which is dependent upon the overall shock probability.

Random control procedure. It is of interest to consider further the case in which shock probability is equal in the presence and absence of the CS. This is the "truly random" control treatment suggested by Rescorla (1967) as a procedure against which to evaluate the effects of CS-US contingencies. As noted above, asymptotically the predictions from the model agree with Rescorla's findings of little conditioning using such a procedure. However, the model suggests that early in conditioning such a CS may in fact show an initial rise in V, followed by a return to the zero point.

The basis for this initial rise in V lies in the rates of conditioning of A and AX. Early in conditioning, reinforcement of the AX compound occurs while V_A and hence V_{AX} is relatively low. Consequently, V_X can receive a sizable increment. As V_A increases due to shocks in the absence of the CS, V_{AX} approaches its asymptotic value and the increments in X correspondingly decrease in size. Furthermore, when V_A approaches its asymptote, V_{AX} will begin to exceed its asymptote due to the contribution of V_X. When this occurs, V_{AX} will be decremented. What follows is

a period in which V_A is incremented from shocks in the absence of the CS and V_{AX} is decremented because is exceeds its asymptote; effectively this produces a redistribution of the associative strength of AX among the components A and X. The process stops when V_A and V_{AX} are equal, i.e., V_X is zero.

This account makes clear that the degree of initial conditioning of the CS predicted by the model in the random procedure is a function of the relative conditioning rates of A and AX. There are a variety of experimental manipulations and parameters of the model which consequently should influence the magnitude and duration of this rise. Two experimental conditions are especially important. The first is the overall probability of shock in both the presence and absence of the CS. As the overall shock probability increases, the magnitude of the initial conditioning to X is greater, and the approach to the final zero asymptote is slower. It is interesting to note, however, that the stage of training at which the maximum V_X is reached remains the same. A second influential manipulation is the proportion of the total session during which the CS is present. In Figure 7, the CS was assumed to be present 1/5 of the time, as it was in fact in Rescorla's experiments. However, it can be shown that as the proportion of the session during which the CS is present is increased, the magnitude of the initial rise in V of the CS is increased; furthermore, the peak magnitude occurs earlier in conditioning, and the attainment of the final asymptote is retarded.

In addition to these experimental manipulations, two parameters of the model are important in determining the magnitude of the initial conditioning of X. As might be expected, one of these is the relative stimulus salience assumed for the background and the CS. According to the model, as the relative salience of the CS increases, the initial positive value taken on by a random CS is enlarged and its duration prolonged. Finally, the assumed relative importance of reinforcement and nonreinforcement is also relevant. As the rate parameters associated with reinforcement is assumed to be progressively larger than the rate parameter associated with nonreinforcement, the magnitude of this initial rise would be expected to increase.

Since most of the experiments employing this procedure have only assessed conditioning to the CS after extended training, there is relatively little direct evidence bearing on the details of these predictions from the model. However, because the manipulations which are predicted to affect the magnitude of the initial rise are also predicted to prolong its presence, some studies employing extended training might yet reflect the effects in question. One interesting example is a recent dissertation carried out at McMaster by Kremer. Kremer (1968) reported nonzero terminal levels of fear following a random procedure. His conditions were similar to those of Rescorla except that he used a more

salient CS (white noise vs. 720 Hz tone) and the CS was present a larger proportion of the session (1/3 or 1/2 vs. 1/5). It should be noted that these modifications of Rescorla's procedure are ones which according to the present model should enhance and prolong the initial positive conditioning of the CS. Furthermore, Kremer found that more frequent alternation between CS and non-CS periods produced more fear conditioning in the random procedure. It seems likely that more frequent alternation of these periods would also enhance CS salience, in which case the model again predicts a prolongation of positive phase of the CS. Consequently, if we assume that Kremer's data were collected prior to asymptotic levels of conditioning, his results would generally fit with the model.

Some comment should be made about the consequences of these predictions for the suitability of the truly random procedure as a "control" treatment in Pavlovian conditioning. This procedure was designed as a control for a particular operation, namely the establishment of a contingency between CS and US. On the assumption that contingencies are important in conditioning, this procedure gives a baseline of no correlation with which to make comparisons. This application of the procedure remains indifferent to the assignment of any particular theoretical learning value to the CS. What the present account attempts is a theoretical understanding of the results of arranging various degrees of correlation between a CS and US. And according to that account, although 'conditioning results are ordered by degree of correlation at every stage of learning and although asymptotically the random procedure does attain an associative value of zero, nevertheless there are stages at which it has positive value. But other theoretical accounts of this procedure are also possible and all would leave equally unaffected the status of the random treatment as a *procedural* control.

Positive CS–US correlations. From Figure 7 it is clear that throughout learning the degree of excitatory conditioning varies with the magnitude of the correlation between the CS and US; furthermore, positive correlations always yield positive Vs. However, the learning curves predicted from this account of positive correlated situations differ from typical learning curves in that they are not all monotonic. When shocks are delivered both in the presence and absence of the CS but with a higher probability during the CS, the V associated with the CS may sometimes attain a value early in conditioning which exceeds its final asymptotic value. The reason for this is similar to that for the initial rise in the random treatment: initially AX may approach its asymptote more rapidly than A approaches its final level. Consequently, after AX has ceased to grow, A continues to increase and during AX trials the compound V is redistributed among the components at the expense of

X. The magnitude of this overshooting in V_X will depend upon the relative rates of conditioning of A and AX. Thus, the same parameters will affect the rise in the case of positive correlations as were important in the case of no correlation.

Negative CS–US correlations. One of the more interesting results from experiments exploring correlations between CSs and USs is the finding that negative correlations lead to CSs which are conditioned inhibitors. Furthermore, the magnitude of the conditioned inhibition is a function of the degree of negative correlation (Rescorla, 1969). The present model is in general agreement with these findings; however, conditions which asymptotically generate negative Vs may, according to the model, initially generate positive Vs. For instance, the .4–.8 condition in Figure 7, although asymptoting at a level of −.33, attains considerable positive associative strength early in conditioning. Again, this initial rise is controlled by parameters affecting relative rates of conditioning for A and AX; it is only when V_A is sufficiently large that V_{AX} exceeds its asymptote that X can begin to acquire negative associative value. Thus prior to the setting up of X as a conditioned inhibitor, A must first be established as a conditioned excitor. Furthermore, notice that if a negatively correlated treatment is terminated early in conditioning, conditioned excitation may be observed despite the fact that the negative correlation was in force from the outset.

In summary, the present model seems consistent with the major asymptotic results of arranging various correlations between a CS and US. Furthermore, it specifies the set of experimental manipulations which might be expected to influence these results. In addition, it makes a number of interesting predictions about preasymptotic consequences of arranging correlations between CSs and USs.

RELATION TO ATTENTIONAL THEORY

The model we have presented was designed to account for instances in which identical stimuli, although associated with equal reinforcement schedules, nevertheless acquire different associative strengths, as a result of the stimulus context in which they are imbedded. The correspondence between the model and data, as described in the preceding sections, would appear to offer encouragement to the line of theorizing which has been developed. There is, however, another plausible, more conventional theoretical approach to this same general problem, and one which has otherwise received some measure of support (e.g., Mackintosh, 1965). It thus becomes pertinent to ask what relative advantage, if any, is enjoyed by the present theory.

The alternative which must be considered is an "attentional" or

"stimulus selection" interpretation (e.g., Sutherland, 1964). The familiar notion is that an organism can learn only about those cues to which it is attending, and that it has a limited attentional capacity. Attending to one cue is thus presumed to decrease the likelihood of attending to, and hence learning about, other available cues. Such theory has no apparent difficulty, for example, in accounting for the fact that the reinforcement of one stimulus element will have less incremental effects upon the learning to that element, if there is concurrently present a stronger (better attended-to) cue.

The kind of "in principle" arguments in favor of attentional theory are very seductive. Certainly the organism does not have an unlimited capacity to process sensory information. Thus, it must be expected that under some circumstances an environmental stimulus will not be reacted to, or may be less reacted to, as a result of the processing of concurrent stimuli. We would hardly quarrel with this. The proper question, however, is whether the organism's capacity is *so* limited that it is necessary to assume that several highly distinctive stimuli, as employed in compound-stimulus Pavlovian training, *cannot* generally be simultaneously attended to. And even if it were advisable to make such assumptions, would an attentional theory still account for the range of data with which we are presently concerned?

In many instances an attentional theory appears quite adequate. If we pretrain associative strength to a stimulus (A) and then reinforce the same stimulus in conjunction with a novel stimulus (X), the resultant associative strength of X will be reduced compared with that of a group not pretrained on A. It seems natural to assume that pretraining on A leads the animal to attend to A to the deteriment of X, and thus to show little conditioning to X. A similar account seems applicable to the failure of a common cue to acquire considerable associative strength in a discriminative conditioning situation. The animal may be assumed to attend to the dimension defining the primary discriminanda and thereby fail to attend to stimuli less well correlated with the US. And the same reasoning may appear to apply to the failure to observe substantial evidence of learning in the "truly random" conditioning procedure. The CS is no more informative than are the contextual cues concerning the occurrence of the US, is thus not especially attended to, and is not learned about.

A major difficulty with this approach is that we are not provided with a specification of the trial-by-trial events which control attention in Pavlovian conditioning. It is not sufficient to argue that stimuli uncorrelated with reinforcement are not attended to; one needs to know how the trial-by-trial events which compose the uncorrelated treatment are processed by the animal so as to generate failure to attend to the

CS. In the absence of such mechanisms the attentional account is little more than a redescription of the data.

But, to our view, the most significant fact is that while there are obvious symmetries in the results we have discussed, the attentional account seems only to apply to portions of the data. For instance, just as prior reinforcement of A will reduce the amount learned about X on subsequent reinforced AX trials, so prior inhibitory training of A will augment the amount learned about X on reinforced AX trials. It is not clear what modification in attention to A could be produced by inhibitory training which would enhance the amount learned about X on AX trials. It does not seem plausible to argue that inhibitory training of A makes the animal especially *fail* to attend to that cue since Rescorla (1969) has shown that such training gives A decremental control over responding. In addition, an attentional account of the effects of prior reinforcement of A upon the subsequent nonreinforcement of AX does not seem satisfactory. Why should training an animal to attend to A make nonreinforced AX trials especially potent in conditioning inhibition to X?

Even in the case of certain phenomena which attentional theory has been thought to handle well, the apparent adequacy of the theory may be illusionary. Consider Kamin's (1968) blocking experiment in which prior conditioning to A makes the subsequent reinforcement of AX practically ineffective in conditioning X. According to attentional theory, pretreatment of A causes the animal to attend to A on AX trials, so that all of the reinforcement effects that occur go to A. However, since A is already well conditioned, the influence of this reinforcement is difficult to detect. According to the present model, the prior conditioning of A results in an AX with a high associative value which in turn devalues the reinforcer; consequently, neither A nor X should receive additional conditioning.

An unpublished experiment from Rescorla's laboratory was designed to evaluate these alternative interpretations. The strategy was to produce "blocking" using a high V_{AX} arrived at through a low level of pretraining to both A and X rather than considerable training of A alone. According to the present model, any way of arriving at the same V_{AX} should interfere equally with the effectiveness of the reinforcer and preclude further conditioning to *either* stimulus on the AX trials. According to an attentional notion, however, the reinforcer remains effective. Since both A and X separately, should have low associative values, any selection of one of the stimuli to which to attend should result in further conditioning of that stimulus. Consequently, some conditioning should occur to at least one of the stimuli.

Four groups of rats were bar-press trained on a VI schedule of food

reinforcement. They then all received 6 conditioning trials while bar-pressing, with a 2-min. CS and a 0.5-sec., 1-ma. shock; on three trials the CS was a flashings light, on three a 1200 hz tone. This training resulted in less than asymptotic suppression to both CSs, such that subsequent TL compound trials could be shown to produce more complete suppression. Group TL then received 10 conditioning trials on which the tone-light compound terminated in shock. Groups T and L each received 10 reinforced trials with the tone or light alone, respectively. Group N received no further conditioning with either stimulus. All animals were then tested with 2 nonreinforced presentations of each component stimulus over each of 5 test days. Finally, each animal received 3 test sessions during which the TL compound was presented 4 times.

Figure 8 shows the mean suppression ratios for the light and tone separately for each of the 4 groups over the initial 5 test days. Looking first at the suppression observed to the tone CS, as represented in the left panel of Figure 8, it is clear that further conditioning to the tone alone (Group T) resulted in more suppression than did failure to give additional training (Group N). Group L, which had received only further conditioning to the light, showed suppression to the tone similar to that of Group N. The most interesting result, however, was that the suppression in Group TL was not different from that of Group N, but was considerably less than that of Group T. Reinforcing the tone in the presence of the light evidently produced no additional acquisition of fear to the tone.

The suppression observed to the light CS, as shown in the right panel of Figure 8, indicates that the findings for the tone in Group TL were not due simply to all animals attending to the light during compound training. The overall suppression to the light was greater than that to the tone, but the pattern of results was similar. Groups N, T, and TL did not differ in responding to the light, but all suppressed less than did Group L. Thus, additional training to the light alone, but not to the light in compound with the tone, yielded further conditioning to the light.

The results of the subsequent compound test trials were consonant with the above results obtained with the components. Over the three days of testing the TL compound, Groups TL and N gave a mean suppression ratio of .23 and .28; the combined T and L groups gave a ratio of .13. Thus, additional training to either T or L alone yielded more subsequent suppression to the TL compound than did additional training to the compound itself.

These data clearly demonstrate that by giving a small amount of prior conditioning to each of two stimuli, it is possible to interfere severely with further conditioning to both stimuli when their compound presentation is reinforced. This finding is difficult to reconcile with an


Figure 8. Mean suppression ratio for each of four experimental groups over daily test sessions. Plotted separately in the left panel is the suppression to a light element, and in the right panel is the suppression to a tone element.

attentional theory, but is an obvious deduction from the present model.

Whatever the other virtues or liabilities of attentional theory, it simply does not fare well in relationship to the present model when applied to the range of Pavlovian conditioning arrangements under consideration.

CONCLUDING COMMENTS

We have attempted to point out some general principles governing the effectiveness of reinforcement and nonreinforcement in Pavlovian conditioning situations. Experiments from our own laboratories indicate that the incremental or decremental effects upon a component stimulus as a result of the reinforcement or nonreinforcement of a stimulus compound containing that component, depend upon the total associative strength of the compound, not simply upon the associative strength of the component. This general dependence incorporated within a more quantitative formulation of our earlier theoretical position (e.g., Wagner, 1969a, 1969b; Rescorla, 1969) provides a way of integrating a sizeable number of empirical findings. Several sample derivations made from the theory have been demonstrated to match well with available data. But, the greater value of the more specific theoretical formulation which has been proposed may be, as we have seen in several instances, in the identification of additional variables of importance to Pavlovian conditioning. It at least invites a fresh look at a data area in which the available theoretical alternatives have been meager.

References

Annau, Z., & Kamin, L. J. The conditioned emotional response as a function of intensity of the US. *Journal of Comparative and Physiological Psychology*, 1961, *54*, 428–432.

Bush, R. R., & Mosteller, F. *Stochastic models for learning*. New York: Wiley, 1955.

Egger, D. M., & Miller, N. E. Secondary reinforcement in rats as a function of information value and reliability of the stimulus. *Journal of Experimental Psychology*, 1962, *64*, 97–104.

Estes, W. K., & Burke, C. J. A theory of stimulus variability in learning. *Psychological Review*, 1953, *60*, 276–286.

Hull, C. L. *Principles of behavior*. New York: Appleton-Century-Crofts, 1943.

Kamin, L. J. Attention-like processes in classical conditioning. In M. R. Jones (Ed.), *Miami Symposium on the prediction of behavior: Aversive stimulation*. Miami: University of Miami Press, 1968.

Kamin, L. J. Predictability, surprise, attention, and conditioning. In R. Church and B. Campbell (Eds.), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969.

Konorski, J. *Conditioned reflexes and neuron organization*. Cambridge: The University Press, 1948.

Kremer, E. Pavlovian conditioning and the random control procedure. Unpublished doctoral dissertation, McMaster University, 1968.

Lovejoy, E. P. An analysis of the overlearning reversal effect. *Psychological Review*, 1966, *73*, 87–103.

Mackintosh, N. J. Selective attention in animal discrimination learning. *Psychological Bulletin*, 1965, *64*, 124–150.

Pavlov, I. P. *Conditioned reflexes*. London: Oxford University Press, 1927.

Rescorla, R. A. Pavlovian conditioning and its proper control procedures. *Psychological Review*, 1967, *74*, 71–80.

Rescorla, R. A. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, 1968, *66*, 1–5.

Rescorla, R. A. Conditioned inhibition of fear. In W. K. Honig and N. J. Mackintosh (Eds.), *Fundamental issues in associative learning*. Halifax: Dalhousie University Press, 1969.

Rescorla, R. A., & LoLordo, V. M. Inhibition of avoidance behavior. *Journal of Comparative and Physiological Psychology*, 1965, *59*, 406–412.

Restle, F. A theory of discrimination learning. *Psychological Review*, 1955, *62*, 11–19.

Sutherland, N. S. Visual discrimination in animals. *British Medical Bulletin*, 1964, *20*, 54–59.

Wagner, A. R. Stimulus validity and stimulus selection. In W. K. Honig and N. J. Mackintosh (Eds.), *Fundamental issues in associative learning*. Halifax: Dalhousie University Press, 1969, (a).

Wagner, A. R. Stimulus-selection and a "modified continuity theory." In G. H. Bower and J. T. Spence (Eds.), *The psychology of learning and motivation*. Vol. 3. New York: Academic Press, 1969, (b).

Wagner, A. R. Incidental stimuli and discrimination learning. In G. Gilbert and N. S. Sutherland (Eds.), *Discrimination learning*. London: Academic Press, 1968.

Wagner, A. R., Logan, F. A., Haberlandt, K., & Price, T. Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, 1968, *76*, 171–180.

Woodbury, C. B. The learning of stimulus patterns by dogs. *Journal of Comparative Psychology*, 1943, *35*, 29–40.