JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1968, 11, 777-786 NUMBER 6 (NOVEMBER)

INDIFFERENCE BETWEEN PUNISHMENT AND FREE SHOCK: EVIDENCE FOR THE NEGATIVE LAW OF EFFECT[1]

RICHARD SCHUSTER AND HOWARD RACHLIN

HARVARD UNIVERSITY

Pigeons were trained to respond under two conditions with two identical variable-interval schedules of positive reinforcement. While the schedules operated for separate response keys, they were not available concurrently. During one condition, each response was punished with electric shock. During the other condition, shocks were delivered independently of responding. The punishment suppressed responding but the free shocks did not. However, when allowed to choose, the pigeons preferred the condition associated with the lowest rate of shock regardless of whether or not the shock was dependent on responding. In general, shocks exerted their greatest effect on whichever response had the greatest influence on shocks. In this respect, punishment is instrumental in suppressing behavior and the properties of punishment are symmetrical to those of reinforcement. This empirical symmetry dictates a corresponding conceptual symmetry in terms of a positive law of effect accounting for response increments and a negative law accounting for response decrements.

When Thorndike (1911) first proposed the Law of Effect, positive and negative stimuli were given symmetrical roles. "Pleasure" was identified with increased responding and "pain" with decreased responding. Current theories of reinforcement do not accept this symmetry. Their cornerstone is still the positive version: that certain stimuli called positive reinforcers can act directly on a preceding response to increase its frequency (Skinner, 1938, 1953; Hull, 1943, 1951). But the negative law, which would say that aversive stimuli can directly weaken responding, is no longer widely held despite recent evidence that brief electric shocks can effectively suppress responding (Azrin, 1956; Rachlin, 1966, 1967). Instead, the suppressive effects of punishment are viewed as the by-product of two processes by which some other behavior increases in frequency (Skinner, 1953; Dinsmoor, 1954; Mowrer, 1960). According to such theories, in the first stage an arbitrary stimulus becomes a conditioned aversive stimulus (CAS) owing to pairings with a primary aversive stimulus such as electric shock. Subsequently, in the second stage, responding can be reinforced if it removes the CAS. When extended to punishment, the increased response is a non-specific event defined as not-responding. In short, two-factor theories of punishment recognize only the positive Law of Effect; punishment, which produces a decremental effect, is interpreted as the outcome of negative reinforcement, an incremental process.

Despite the different processes thought to govern reinforcement and punishment, many effects of aversive stimuli seem to mirror the effects of reinforcers, as Thorndike originally proposed. For instance, when an arbitrary stimulus is terminated with an unavoidable shock, responding during the stimulus is suppressed (Estes and Skinner, 1941). This is sometimes called a conditioned emotional response. A symmetrical "joy" effect can be produced with positive reinforcement, when an arbitrary stimulus is terminated with a free reinforcement. In this case, response rates during the stimulus can be enhanced (Morse and Herrnstein, 1957). Mirror-image generalization gradients can be produced with reinforcement and punishment. When responding is reinforced during one stimulus, lower rates of responding occur in the presence of neighboring stimuli on the same continuum with the reinforced stimulus; when responding is punished during one stimulus, higher rates occur in the presence of neighboring stimuli (Honig, 1966).

[1]This research was supported by National Science Foundation Research Grant GB 7334. Reprints may be obtained from Howard Rachlin, Department of Psychology, Harvard University, William James Hall, 33 Kirkland Street, Cambridge, Massachusetts 02138.

A further point of symmetry is the fact that behavioral interactions between components of multiple schedules can be produced by changing the frequencies of reinforcement (Reynolds, 1961) or punishment (Rachlin, 1966) during one component. Also, both reinforcement and punishment produce effects appropriate to the schedules according to which they occur (Ferster and Skinner, 1957; Azrin and Holz, 1966; Morse and Kelleher, 1966).

Perhaps the most fundamental symmetry between reinforcement and punishment is that both kinds of stimuli have their greatest effect on whichever response has the greatest effect on the stimuli. Reinforcement increases responding more when responding produces reinforcement than when reinforcements are delivered randomly (Ferster and Skinner, 1957), and brief electric shocks suppress responding more when responding produces shock than when shocks are independent of responding (Azrin, 1956; Rachlin, 1967). This instrumentality of punishment is revealed still more vividly by comparing the effects of reinforcers and punishers on rate of responding and choice. For example, in a typical variable-interval schedule of positive reinforcement, response rate has no influence on parameters of reinforcement such as delay and amount, and little influence on frequency of reinforcement. Correspondingly, response rate is found to be relatively insensitive to these parameters (Ferster, 1953; Catania, 1963; Neuringer, 1967; Catania and Reynolds, 1968). On the other hand, when differential responding can affect which value of a reinforcer will occur, choices are sensitive to parameters such as frequency (Herrnstein, 1964; Autor, 1960), delay (Chung and Herrnstein, 1967), and amount (Catania, 1963; Neuringer, 1967). Moreover, in the case of frequency, choices between various schedules of reinforcement vary with the frequencies of reinforcement, and are independent of response rates, whether reinforcements are independent of responding (Autor, 1960), dependent on not responding (Autor, 1960), or dependent on different schedules of responding which generate different response rates (Herrnstein, 1964). In other words, these studies showed that the effect of reinforcement on choice seems to depend on the parameters of reinforcement being chosen, and not on the correlation between a reinforcer and another response. It is possible that the same is true of aversive stimuli since Rachlin (1967) demonstrated that choices between conditions of shock are sensitive to the frequency and intensity of shock even though, in separate experiments, shocks were independent and dependent on responding.

The present experiment tested this possibility directly by shocking pigeons under both conditions in the same experiment. In one condition, the punisher was delivered after every response; in another condition, the punisher was delivered at various rates independently of responding. Response rates during exposure to these two conditions were compared. In addition, the subjects were offered a choice between the opportunity to respond for a time under one or the other condition. If the shock itself, and not the nature of the correlation between the shock and responding, is important, then choice behavior should vary only with the relative number of shocks delivered under the two conditions. If the correlation between responding and shock is important, then choice behavior should depend also on the relation of shock to responding.

METHOD

Subjects

Four adult male White Carneaux pigeons were maintained at about 80% of free-feeding weights. All were experimentally naive at the start of the experiment. The birds were implanted around the two pubis bones with two gold wires through which brief electric shocks were delivered (Azrin, 1959).

Apparatus

The apparatus was a modified chamber designed for pigeons (Ferster and Skinner, 1957). The positive reinforcer was 3-sec access to grain. The box was equipped with two 1-in. diameter keys mounted above the grain hopper, side by side, and separated by about 4 in. The keys required at least 15 g of force to operate. The left key could be transilluminated by either a white or red lamp; the right key by either a white or orange lamp.

The shock was a 35-msec pulse of 60 cps through an 11 K ohm fixed resistance plus a variable resistance adjusted to deliver 7 ma (Azrin, 1959).

Procedure

The procedure throughout the experiment was a concurrent chain schedule (Herrnstein, 1964), diagrammed in Fig. 1, which offered a choice between response-dependent and response-independent shock superimposed on identical schedules of positive reinforcement. These conditions were arranged as terminal links of two two-link chain schedules; the initial links of these two schedules were separate, concurrent variable-interval (VI) schedules.

Although the concurrent chain procedure appears complicated, it is to some extent analogous to a simple T-maze with only three possible states, shown as separate boxes in Fig. 1: the concurrent initial links correspond to the choice point in a T-maze; the terminal links correspond to the goal boxes. The principal departure from the analogy with a T-maze procedure is what occurs at the choice point. In a T-maze, every choice usually produces one or the other alternative, so that a subject tends always to choose a single alternative, even when it is only slightly preferable. During the concurrent initial links of the present procedure, on the other hand, the subject was encouraged to sample continually both alternatives because, by this strategy, reinforcement frequency can be maximized. In other words, distributed responding on concurrent VI 2-min schedules insures that one of the terminal links will occur on the average of once per minute, whereas choice of only one alternative produces that link on the average of once per 2 min. Most subjects in schedules of this kind distribute their choices over all the available alternatives. In both the T-maze and the present procedure, choice is measured by the "behavior ratio", responses to one side (or key) divided by total responses. However, in the case of the T-maze, the behavior ratio is directly proportional to the relative number of entries into the goal boxes; in the present procedure the relative number of terminal links will remain around 50% for almost any behavior ratio. Another property of the concurrent chain schedule is

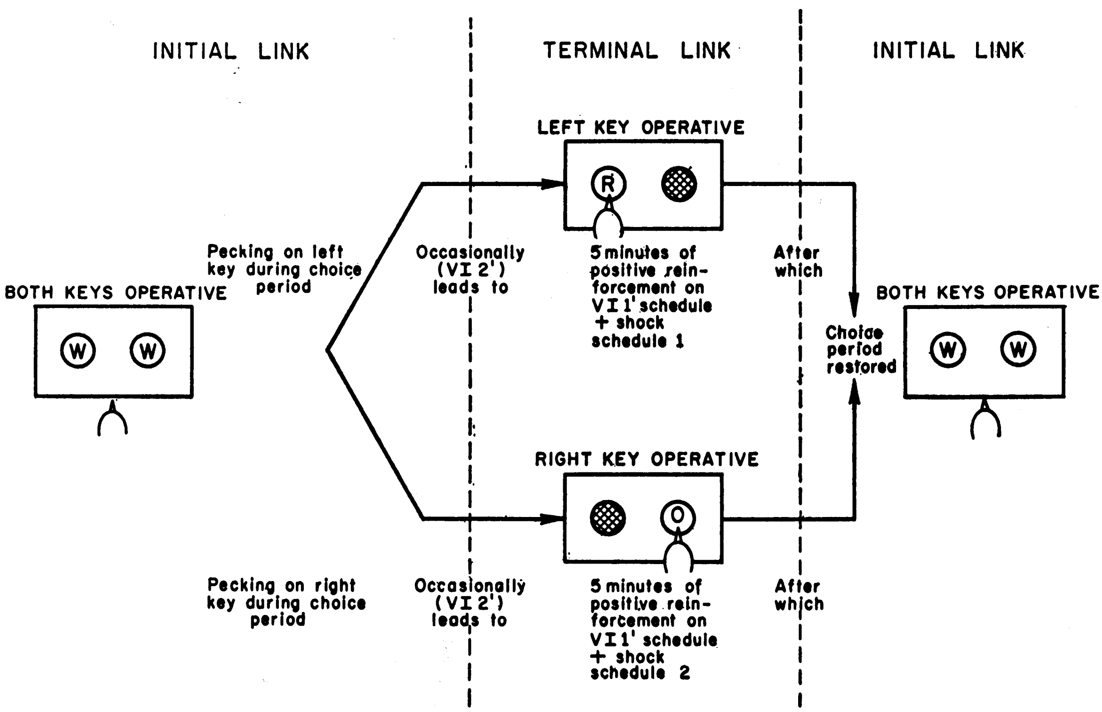

Fig. 1. Diagram of the concurrent chain procedure. During the initial link both keys are illuminated with white light and operative. Pecking on either key leads on a VI 2-min schedule to a terminal link during which the key just pecked is illuminated with red or orange light and the other is dark and inoperative. Pecking on the red or orange key during the terminal link is reinforced with food on a VI 1-min schedule. Shock schedule 1 in this experiment consisted of a brief shock for each peck. Shock schedule 2 consisted of free shocks, at various rates, independent of pecking.

that it offers a choice between extended "stays in the goal box" during which responding is examined under the alternative conditions of reinforcement and punishment. The schedule thus measures the preference between the two terminal links, and the behavior during these terminal links, within the same experimental setting.

The details of the procedure were as follows: during the concurrent initial links, both keys were operative and transilluminated with white light. Responding on the left key occasionally (on a VI 2-min schedule) changed key color to red indicating that the terminal link with response-dependent shock (punishment) was in effect; similarly, responding on the right key occasionally (on another VI 2-min schedule) changed key color to orange, correlated with the terminal link containing response-independent shock (free shock). Terminal links lasted 5 min, during which schedules of reinforcement and punishment, described below, were arranged. Both VI programmers of the initial component stopped during either terminal component. The session ended after ten 5-min terminal links were produced which, combined with initial link durations, resulted in a session of about 65 min.

During both 5-min terminal links, positive reinforcement was delivered on a VI 1-min schedule. Initially, there was no shock during either terminal link. This was continued as a baseline condition until choices and terminal link response rates were stable over at least 10 sessions. Then, the different shock conditions were introduced during the two terminal links, providing the only asymmetry between the terminal-link alternatives. On the right key, every response was punished with a single shock pulse throughout the remainder of the experiment. On the left key, single shock pulses were delivered independently of responding at constant intervals except during presentation of food. The frequency of free shock, the only independent variable in the experiment, was varied between 0 and 120 pulses per min in the following sequence: 0, 6, 12, 30, 60, 120, 60, 30, 12, 6, and 0. There were 10 sessions at each point in the sequence.

RESULTS

All data shown are medians from the last three days at each condition. Where there were two determinations for a point, the average of the two determinations is shown.

Response Rate

The effect of free shock on terminal link response rate is shown in Fig. 2 as a function of the shock frequency. The dependent variable is suppression, defined here as one minus the ratio of the rate of response during shock to the rate before shock was introduced. A ratio of 0.0 indicates no suppression of responding while a ratio of 1.0 indicates complete suppression. Overall, the effect of free shock was small; three subjects (469, 324, and 249) suppressed only at the high shock frequencies and one subject had an increased rate.

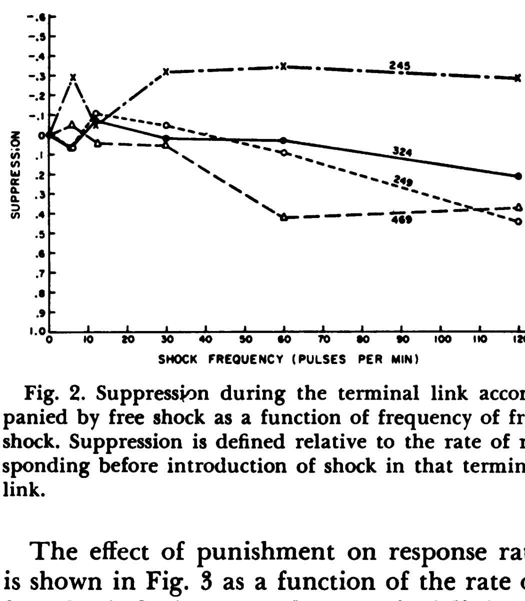

Fig. 2. Suppression during the terminal link accompanied by free shock as a function of frequency of free shock. Suppression is defined relative to the rate of responding before introduction of shock in that terminal link.

The effect of punishment on response rate is shown in Fig. 3 as a function of the rate of free shock during the other terminal link. In general, the rate of responding under the punished condition (equal to the rate of punishment) hardly varied.

When response rates during the terminal links are compared, the punishment had a much greater suppressive effect. This is shown in the dashed line of Fig. 4 as the relative response rate during free shock, *i.e.*, the response rate during the terminal link containing free shock divided by the sum of the rates during both terminal links. All subjects responded faster during free shock, even at the shock rate of 120 pulses per min (about three to four times the shock rate during punishment). All subjects, though, gradually increased response rates during punishment, and gradually dropped response rates during free shock, as the frequency of free shocks was increased. This is shown in Table 1, which gives the response rates during the two terminal links at each frequency of free shock.

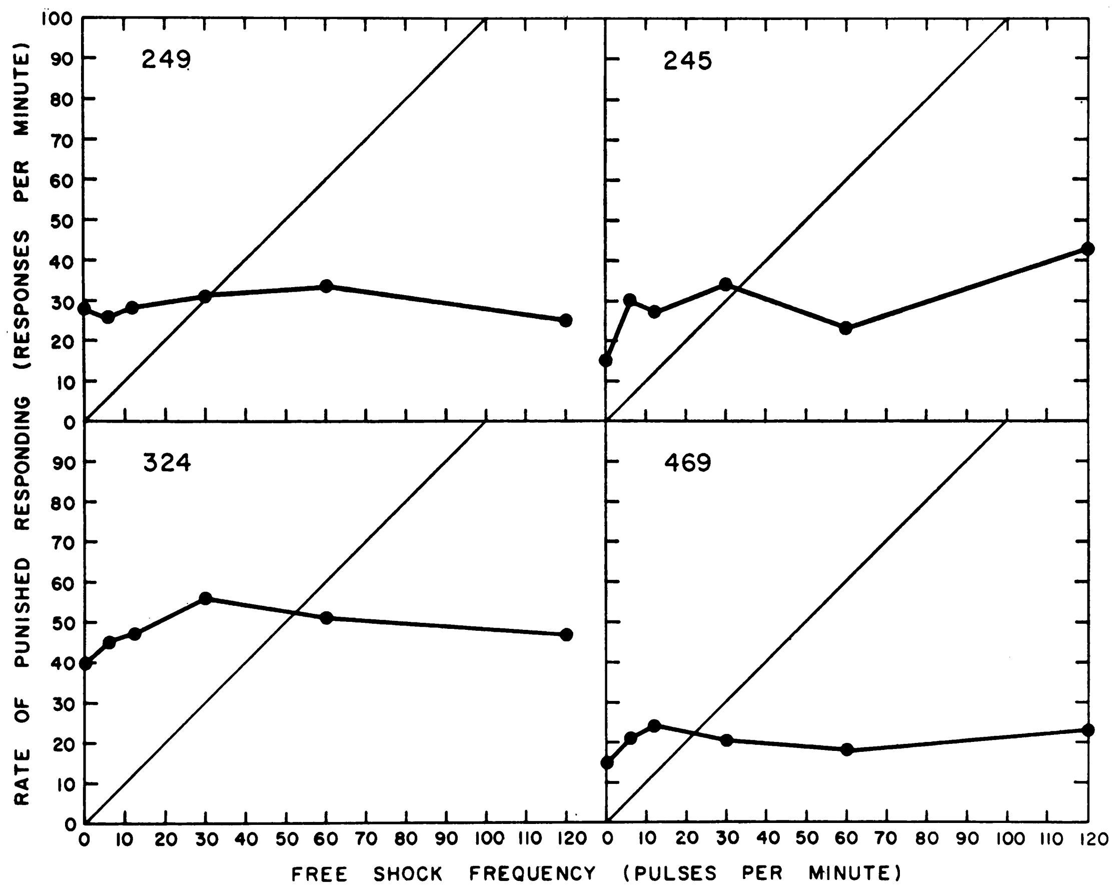

Fig. 3. Rate of pecking (equal to rate of shock) during the terminal link where pecks were shocked as a function of the rate of free shocks during the other terminal link. The diagonal lines are the loci of equal overall rates of shock during the two terminal links.

Table 1

Terminal link response rates during free shock and during punishment, for various rates of free shock (resp./min).

Subject	*Terminal Link Shock Condition*	*No Shock*	*Punishment Only*	*Free-Shock Frequency (Shocks per Min)* 6	12	30	60	120	60	30	12	6	*Punishment Only*
245	Free	49	49	105	79	127	119	117	128	113	112	124	132
	Punishment	59	0	12	15	19	27	43	19	48	39	47	30
324	Free	58	49	57	66	59	62	45	48	52	56	48	64
	Punishment	67	30	56	44	56	52	47	50	56	49	33	50
249	Free	51	79	80	84	98	86	48	80	93	118	90	102
	Punishment	56	17	23	21	27	35	25	32	34	34	28	38
469	Free	50	42	45	40	46	24	27	26	36	41	45	43
	Punishment	42	13	20	21	19	16	23	20	22	26	22	18

Choice

Choice is here defined as the number of responses in a session during the initial link on the key leading to free shock divided by the total number of responses during the initial link on both keys. In other words, choice is the relative rate of responding on the free-shock key during the initial link. The effect of the two shock conditions on choice is shown in the solid line of Fig. 4 as a function of the frequency of free shocks. In all subjects, the effect of this shock frequency on choice in the initial link was substantially greater than its effect on relative response rates in the terminal links; free shock was preferred at low shock frequencies but, as frequency increased, punishment was eventually preferred. In all subjects, it was common for the relative rate of responding during the terminal links to be above 50% at the same time that choices were below 50%, showing that subjects were responding faster during free shock even though they preferred punishment.

The effect of the two shock conditions on preference is shown in Fig. 5 as a function of the relative frequency of free shock. In all subjects, this function passes close to the 50% point on both ordinate and abscissa, suggesting that choices varied with the number of shocks actually delivered during the two conditions, regardless of the difference in correlations between responding and shock delivery. This is confirmed when the choices are mea-

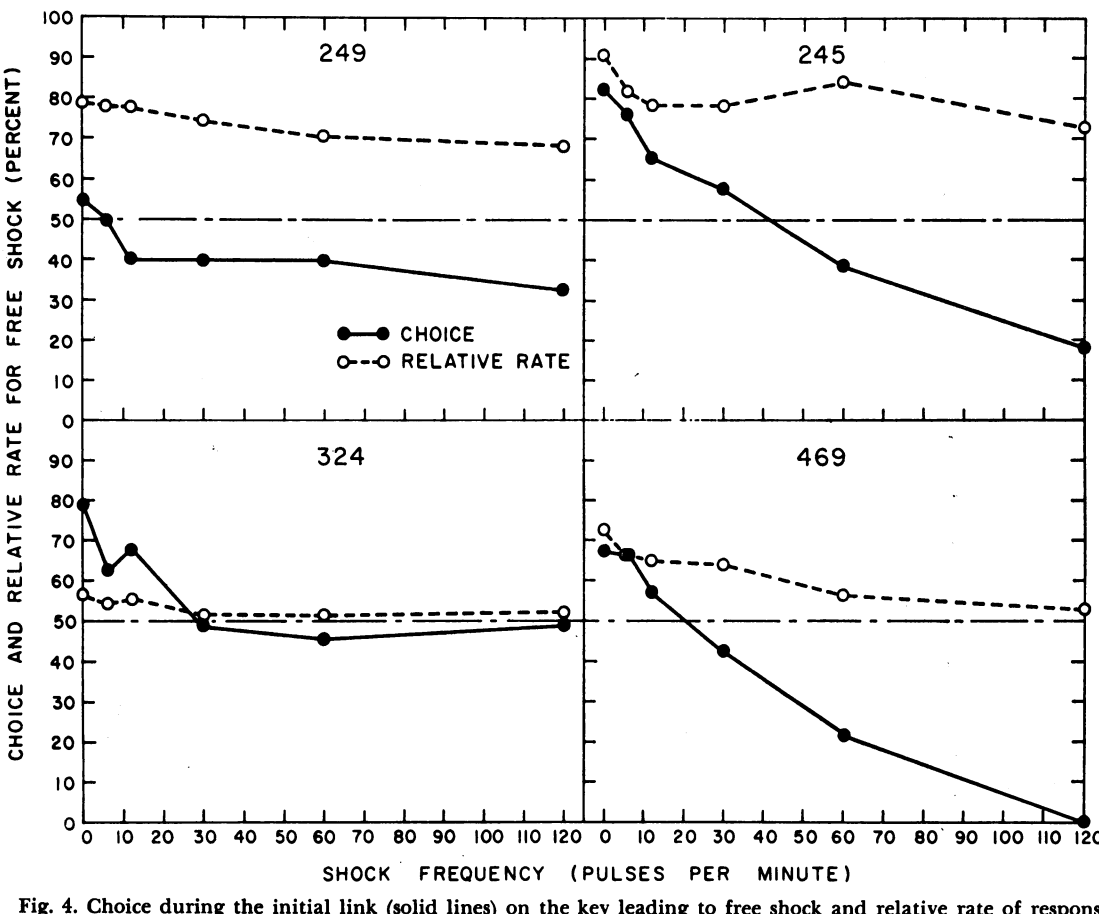

Fig. 4. Choice during the initial link (solid lines) on the key leading to free shock and relative rate of response during the terminal link (dotted lines) on the free-shock key as a function of the frequency of free shock. Both choice and relative rate are calculated as the rate of responding on the free-shock key divided by the sum of the rates of responding on both keys expressed as a percent. Choice is also equal to the relative number of responses on the two keys. Both choice and relative rate equal 50% when rates of responding are equal on the two keys. They both would equal 100% if all responding were on the free-shock key and 0% if all responding were on the key where responses were punished.

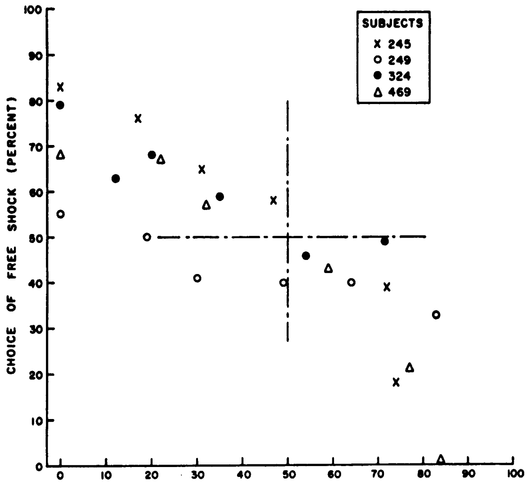

Fig. 5. Choice of free shock as a function of the relative rate of shock on the free-shock key. The ordinate values were determined as for Fig. 4. The abscissa values were determined by dividing the rate of free shock (the shock pulse frequency) by the sum of the rate of free shock and the rate of punishment (necessarily equal to the rate of responding) during the other terminal link.

Table 2

Subject	*Rate of Free Shock Equal to Rate of Punishment (Shocks per Minute)*	*Choice of Free Shock at This Point (from Fig. 4) (Percent)*
249	31	40
245	33	55
324	52	49
469	22	50

sured at the point where the rates of free shock and punishment are equal. This point of equality is shown for each subject in Fig. 3 which plots the rate of punished responding against the rate of free shock. The diagonal lines cross the individual functions at the point where the rates of shock during the two terminal links were equal. At this rate of shock, choice of free shock can be read from Fig. 4. These values, shown in Table 2, average to 48.5%, which is close to indifference (50%). Two subjects preferred punishment slightly, and one subject preferred the free shock slightly.

DISCUSSION

This experiment offers additional support for the position that punishment is an effective means for controlling behavior, contrary to the view that it is not (Thorndike, 1932; Estes, 1944). During the terminal links, punishment had a greater effect on response rate when dependent on responding than when independent of responding, as in previous studies by Azrin (1956, 1958) and Rachlin (1967). In the present study, in fact, punishment caused greater suppression than free shock, even when the rate of free shock was much higher.

The effects of shock on choice shown here also agree with previous findings in one respect, that choices for free shock bear an orderly relation to the frequency and intensity of shock pulses (Rachlin, 1967). But the present results also extend previous findings in one respect: choices for shock bear an orderly relation to the frequency of shocks, and *do not depend on the rate of responding and the correlation between shock and responding.*

In this respect, the effects of shocks are completely analogous to the effects of positive reinforcement. Autor (1960) used a concurrent chain schedule to measure choices between two conventional response-dependent schedules, and also between two schedules of free reinforcement. Choices were found to vary with the relative frequencies of reinforcement in both cases. Herrnstein (1964) supported this finding with choices between different types of reinforcement schedules, interval and ratio. Even though rates of responding for given frequencies of reinforcement were different, due to the difference between interval and ratio requirements, choices were found to vary with the relative frequencies of reinforcement.

Despite this apparent symmetry between the effects of reinforcement and punishment, it was noted in the Introduction that this symmetry is not recognized by current learning theories. Although they handle reinforcement effects by the positive Law of Effect, they explain the effects of punishment with two factors, neither of which is the negative Law of Effect. The first factor is classical, or Pavlovian, conditioning according to which a stimulus becomes a CAS owing to pairings with a primary aversive event such as electric shock. The second factor is instrumental conditioning, that a response will increase in frequency if it is followed by negative reinforcement, specifically, removal of the CAS. As previously pointed out in Rachlin and Herrnstein (1967), this theory views both avoidance and punishment as incremental processes.

In the case of avoidance, an incremental process is demonstrably occurring; a specific response is indeed increased in frequency when an aversive stimulus can be avoided. For example, responses such as shuttling or bar-pressing both may be made to avoid the primary aversive stimulus and terminate an arbitrary stimulus which precedes the primary stimulus (Sidman, 1953; Mowrer, 1947). In the case of punishment, however, the only observable change is a decrease in responding. The incremental process is only inferred, as the increased response is said to be anything but the punished response, that is, anything but what is measured.

A similar problem arises on the stimulus side of two-factor theories. In avoidance situations, the CAS is typically, but not necessarily, an external "warning" stimulus. But punishment situations usually do not include an exteroceptive stimulus to announce punishment. In order to extend the two-factor apparatus to punishment, the CAS is said to be proprioceptive feedback from the punished response (Dinsmoor, 1954; Mowrer, 1960). In other words, the CAS for punishment is also unobserved.

Because both the reinforcer and the response are unobserved and unobservable with conventional techniques, the two-factor theory of punishment poses a serious problem for the experimenter who wishes to test it: how can it be disproved? All the critical events are assumed to occur within the organism being punished. This is not always necessary. for, according to the proprioceptive escape mechanism outlined above, the animal must begin to engage in movements which lead to punishment in order to escape from conditioned aversive stimuli. Moreover, the frequency of these incipient movements must be assumed to remain constant, regardless of the amount of punishment received in the past. While such incipient responses are occasionally observed in punishment situations, they are not common enough to satisfy the theory's requirements. To retain the two-factor approach, theorists have been forced to move these events inside the organism where they cannot be measured.

Nevertheless, the present experiment offers suggestive evidence against the proprioceptive-CAS mechanism. This mechanism means that the ability of a stimulus to suppress responding depends on the specific correlation between the stimulus and the response which precedes the stimulus, for it is the response-produced stimuli which are supposed to provide the only reinforcement. One prediction from this theory is that, in the present experiment, subjects should not match their choices to the relative number of shocks during the two terminal links. Instead, when equal numbers of shocks occur under the two shock conditions, the response-dependent shock condition should be less preferred. One reason is that the theory is implying that the effectiveness of shock depends on response-produced stimuli (although, as we shall see shortly, it is still possible to predict the results shown here from two-factor theory). A second reason is that the response-dependent condition should include a high frequency of CASs, aversive feedback from every response, which should be added to the pool of aversiveness created by the shocks themselves. A third reason is that, in the concurrent chain choice procedure used here, the response which produced shock and the choice for that shock condition had the same topography—both were pecks on the same key. If proprioceptive feedback became aversive, then both of these responses should have decreased together. In fact, this result was not obtained; the response-dependent and free shocks had equal effects on choice, and initial link choices and terminal link response rates varied independently.

It must be emphasized that this result is not impossible to reconcile with two-factor theories. But this reconciliation is highly complex. To explain fully the effects of punishment shown here, the two-factor apparatus must be enlarged. During all three states of the concurrent chain shown in Fig. 1, which really consists of four states—two links on each key—key-pecking presumably generated the same proprioceptive feedback. This should have become aversive since, during one terminal link, every response was immediately followed by shock. Yet responding was suppressed only during this one link. Further assumptions must therefore be made about the inside of the organism. We would have to propose an interaction between the proprioceptive feedback and the external stimuli on both keys, such that the external stimuli signal differing levels of aversiveness for the same proprioceptive feedback. This would result in four separate interactions, two for the two states on each key, resulting in four different states of conditioned inhibition, each reinforcing the learned response of not-responding. In this way, the results shown here can be fitted to the two-factor approach (which as noted, is difficult to disprove), but not without considerable effort.

This complexity offers one benefit for learning theory: it attempts to explain the effects of positive and negative stimuli with the single concept of response increments. But, even this attempt is not without a major flaw. For reasons that remain obscure, this unification has never been completed except by Mowrer (1960) who, having first developed the two-factor approach for avoidance (Mowrer, 1947), now extends the theory to reinforcement (Mowrer, 1960). Other theorists persist in treating reinforcement and punishment differently, even though response increments are claimed in both cases. In the case of punishment, conditioned aversive stimuli are used to explain the apparent effectiveness of primary aversive stimuli, because these stimuli are not thought to have any genuine instrumental effect upon behavior. In the case of positive stimuli, however, the positive Law of Effect says that primary reinforcers such as food do have a direct effect. Conditioned positive stimuli are recognized, though. These are called conditioned (Kelleher, 1966; Skinner, 1953) or secondary (Hull, 1943; Mowrer, 1960) reinforcers, and are thought to exercise an independent reinforcing function like primary reinforcers; they are not used to explain how primary reinforcers can increase responding. This inconsistency is totally unexplained. If negative stimuli need a two-factor theory, why has it not been more widely extended to positive stimuli? Or, to reverse the question, if two-factor theory is not needed for positive stimuli, why has it been applied to negative stimuli? It might also be asked why a theory of response decrements, *i.e.*, punishment, has not been applied to the case of avoidance, in which it could be argued that the avoidance response increases in frequency because every other response is punished. And, given this, another kind of parsimony, analogous to two-factor theories, could be proposed in which all learning, both positive and negative, is explained by response decrements. In short, the parsimony implicit in two-factor theories is illusory. It seems, rather, that there have been arbitrary decisions, perhaps, as suggested here, the result of historical accident, to accept a positive Law of Effect and to deny a negative Law of Effect.

A simpler alternative, and one more directly linked to observable events, is to reassert Thorndike's negative Law of Effect. This is simpler because, when linked with the widely accepted positive Law of Effect, we would have parsimony of a different sort: the effects of primary stimuli, both positive and negative, would depend on the parameters of the primary event itself, and not on the nature of the correlation between the event and behavior. It could also be added that the maximal effect of those events on behavior will occur when the behavior interacts in some way with the occurrence of these events. This union is also defensible because the present experiment, coupled with previous work, shows that positive and negative stimuli do, indeed, have effects entirely consistent with a symmetrical Law of Effect based on reinforcing and aversive events.

REFERENCES

Autor, S. M. The strength of conditioned reinforcers as a function of frequency and probability of reinforcement. Unpublished doctoral dissertation, Harvard University, 1960.

Azrin, N. H. Some effects of two intermittent schedules of immediate and non-immediate punishment. *Journal of Psychology*, 1956, **42**, 3-21.

Azrin, N. H. A technique for delivering shock to pigeons. *Journal of the Experimental Analysis of Behavior*, 1959, **2**, 161-163.

Azrin, N. H. and Holz, W. C. Punishment. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 380-447.

Catania, A. C. Concurrent performances: A baseline for the study of reinforcement magnitude. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 299-300.

Catania, A. C. and Reynolds, G. S. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 327-383.

Chung, S. H. and Herrnstein, R. J. Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 67-75.

Dinsmoor, J. A. Punishment I. The avoidance hypothesis. *Psychological Review*, 1954, **61**, 34-46.

Estes, W. K. An experimental study of punishment.

Psychological Monographs, 1944, **53**(3), (Whole No. 263).

Estes, W. K. and Skinner, B. F. Some quantitative properties of anxiety. *Journal of Experimental Psychology*, 1941, **29**, 390-400.

Ferster, C. B. Sustained behavior under delayed reinforcement. *Journal of Experimental Psychology*, 1953, **45**, 218-224.

Ferster, C. B. and Skinner, B. F. *Schedules of reinforcement*. New York: Appleton-Century, 1957.

Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 27-35.

Honig, W. K. The role of discrimination training in the generalization of punishment. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 377-389.

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

Hull, C. L. *Essentials of behavior*. New Haven: Yale University Press, 1951.

Kelleher, R. F. Chaining and conditioned reinforcement. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 160-212.

Morse, W. H. and Herrnstein, R. J. Some effects of response-independent positive reinforcement on maintained operant behavior. *Journal of Comparative and Physiological Psychology*, 1957, **50**, 461-467.

Morse, W. H. and Kelleher, R. T. Schedules using noxious stimuli. I. Multiple fixed-ratio and fixed-interval termination of schedule complexes. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 267-293.

Mowrer, O. H. On the dual nature of learning—a reinterpretation of "conditioning" and "problem-solving". *Harvard Educational Review*, 1947, **17**, 102-148.

Mowrer, O. H. *Learning theory and behavior*. New York: Wiley, 1960.

Neuringer, A. J. Effects of reinforcement magnitude on choice and rate of responding. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 417-425.

Rachlin, H. C. Recovery of responses during mild punishment. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 251-263.

Rachlin, H. C. The effect of shock intensity on concurrent and single-key responding in concurrent-chain schedules. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 87-93.

Rachlin, H. C. and Herrnstein, R. J. Hedonism revisited: on the Negative Law of Effect. Paper read at Conference on Punishment, Princeton, N.J., 1967.

Reynolds, G. S. Behavioral contrast, *Journal of the Experimental Analysis of Behavior*, 1961, **4**, 57-72.

Sidman, M. Avoidance conditioning with brief shock and no exteroceptive warning signal. *Science*, 1953, **118**, 157-158.

Skinner, B. F. *The behavior of organisms*. New York: Appleton-Century, 1938.

Skinner, B. F. *Science and human behavior*. New York: Macmillan, 1953.

Thorndike, E. L. *Animal intelligence: experimental studies*. New York: Macmillan, 1911.

Thorndike, E. L. Reward and punishment in animal learning. *Comparative Psychological Monographs*, 1932, **8**, No. 39.

Received 15 April 1968.