JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR VOLUME 9, NUMBER 4 JULY, 1966

NEGATIVE REINFORCEMENT AS SHOCK-FREQUENCY REDUCTION[1]

R. J. HERRNSTEIN AND PHILLIP N. HINELINE[2]

HARVARD UNIVERSITY

Is a conditioned aversive stimulus necessary in avoidance conditioning? Or is a reduction in the rate of aversive stimulation alone sufficient to generate and maintain an avoidance response? Rats were subjected to an avoidance procedure in which shocks occurred randomly in time, but a response could reduce the overall rate of shock. Fifteen acquisition curves, obtained from 16 animals, showed both immediate and delayed, rapid and gradual increases in response rate; there was no representative acquisition curve. Response rates were directly related to the amount by which the response reduced shock frequency. In extinction, when shock rates were not affected by responding, the response total was inversely related to the amount by which the response had reduced shock frequency during prior conditioning, with as many as 20,000 extinction responses when the shock frequency reduction had been relatively small. Responding on this procedure shows that avoidance conditioning can occur without benefit of either classical exteroceptive stimuli or covert stimuli inferred from the temporal constancies of a procedure. It also shows that reduction in shock rate is alone sufficient to maintain avoidance.

In Sidman's (1953a, b) procedure for avoidance conditioning, brief, intense and inescapable electric shocks are presented at regular time intervals to an animal—most often a rat—confined in a chamber containing a response lever. Each time the rat depresses the lever, the train of shocks is interrupted for a given period of time. Shocks can be indefinitely postponed by a sufficiently high rate of responding. The procedure is fully specified by two time parameters: the interval between shocks in the absence of responding (the S-S interval) and the interval between a response and the next shock if there are no further responses (the R-S interval). As an explanation of the behavior generated by his method, Sidman initially favored a version of Schoenfeld's (1950) theory of avoidance. Dinsmoor (1954) subsequently expanded this view. The rat, it was said, depresses the lever because everything else it has done has been paired with shock. In its efforts to escape from its punished behavior, the rat is left with the one unpunished response: lever-pressing. Aside from the virtual impossibility of a direct experimental test, this account suffers from sheer implausibility. Rats often learn to avoid very quickly, long before even a minor fraction of their potential behavior can have been shocked.

Two recent accounts seem more plausible. One, proposed by Anger (1963), says that the avoidance response is reinforced because it is followed by covert stimuli which are less aversive than those immediately preceding it. Anger arrived at this conclusion by applying well-established facts and concepts to Sidman's procedure. For example, it has often been shown that rats can react to the passage of time (Anger, 1956). It has also been shown that originally neutral stimuli may become aversive when paired with aversive stimuli (Miller, 1948). In Sidman's procedure, shocks are paired with a given period of time since a response or since a shock. It follows that the passage of these time periods may affect the rat as a growing level of aversiveness, dropping suddenly to a minimum after either a shock or an avoidance response, which is thus reinforced. In short, the psychological literature contains ample evidence that a rat should be responsive to the safe period that follows lever-pressing in Sidman's procedure. Anger's ac-

---

[1] This research was supported by a grant from the National Science Foundation to Harvard University and was conducted with the helpful assistance of Mrs. Antoinette Papp and Mr. Wallace Brown. Reprints may be obtained from either author at the Psychological Laboratories, William James Hall, Harvard University, Cambridge, Mass. 02138.

[2] Work begun under support of a National Institutes of Health Predoctoral Fellowship, and continued under the Graduate Psychology Student Program of the United States Army.

count discussed other sources of reinforcement, but ranked them as secondary in comparison with the removal of the aversive temporal stimulus.

The other recent formulation, proposed by Sidman (1962), attributes avoidance simply to the reduction in shock frequency which it produces. He drew support for this view from: (a) the rapidity of learning when the S-S interval is much shorter than the R-S interval, so that each depression of the lever produces a large decrement in shock frequency; (b) the difficulty of maintaining behavior with a short R-S interval relative to the S-S interval, so that the depression of the lever at times increases shock frequency; (c) the results of an experiment in which an independent R-S interval was associated with each of two levers, and in which the rats tended to distribute their behavior to minimize the total shock frequency.

Applied to Sidman's procedure, the two accounts are essentially identical, making the same empirical predictions with equivalent precision (or lack thereof). Both theories strive to specify a plausible reinforcer, and both succeed. Anger's reinforcer is the reduction of conditioned aversive stimuli; Sidman's is the reduction of shock rate. The advantage of familiarity enjoyed by Anger's reinforcer seems to be fully balanced by the advantage of direct measurability enjoyed by Sidman's. In both theories, the hypothesized reinforcers are maximized by spaced responding. For Anger, the rat reduces aversiveness by greater decrements if it responds just before a shock is due. For Sidman, such spaced responding earns a greater decrement in shock rate than responses distributed more randomly in time. And rats do not respond randomly in time in Sidman's procedure, just as either theory implies.

The difficulty in distinguishing the two theories may be a peculiarity of Sidman's procedure, and not of the theories themselves, for shock rate and conditioned aversiveness are not linked by logical necessity. Bolles and Popp (1964) attempted to separate the two variables in an experiment in which R-S intervals initiated during S-S intervals took effect only after the shock at the end of that S-S interval. Responses during the subsequent R-S interval then postponed shock in the usual manner. On this procedure, in which the response-produced delay of shock was preceded by a shock, only one of 14 animals learned to avoid, and that only with a very large difference between the S-S and R-S intervals (5 *vs* 45 sec). Bolles and Popp suggested that the shortcoming of this avoidance procedure was the elimination of low aversiveness for short post-response times, thereby implying support for Anger's hypothesis. However, since the procedure delays any reinforcement due to the reduction of shock rate, Sidman's hypothesis also predicts impaired performance.

The present experiment was designed to test shock rate as a controlling factor in avoidance. To an approximation, shocks occurred randomly in time, but a response could reduce the overall rate of shock.

METHOD

Subjects

Fourteen brown female rats of the Lashley strain and four hooded male rats served. D-1, D-2, and D-3, brown females, were 200 days old at the beginning of the experiment; they had previously avoided on a variable S-S—R-S Sidman avoidance schedule. The remaining brown rats, F-5, F-6, M-1, M-2, M-3, N-4, N-5, N-6, P-4, P-5, and P-9, were experimentally naive, and ranged from 90 to 150 days old when the experiment began. The four hooded males, G-1, G-2, L-1, and L-2, were all experimentally naive and 200 days old at the beginning of training. Animals designated with the same letter are littermates. All subjects were kept in individual home cages with food and water freely available.

Apparatus

The experimental chamber was a standard rat box with a floor made of stainless steel grid bars 2.5 mm in diameter and spaced ½ in. apart center-to-center. The response lever, a modified telegraph key requiring about 30 g of force to operate, was 3 in. above the grid floor and protruded 1⅛ in. into the box. A false, transparent ceiling, sloping towards the lever, prevented animals from climbing onto the lever. Shocks of 0.3 sec duration and approximately 0.8 ma intensity were delivered to the walls, lever and grid bars with polarities scrambled randomly.

A Trans-Lux Transmitter (punched-tape reader) and conventional switching circuitry were used to program the contingencies and record the data. White noise was supplied during all sessions.

Procedure

In general outline, the procedure comprised two separate and independent programs of shocks delivered randomly in time. One program was in force as long as the rat failed to depress the lever; the other was in force after lever depressions. During conditioning, the program associated with responding produced the smaller frequency of shocks; during extinction the two programs produced shocks equally often.

Throughout each session a multi-channel paper tape was advanced one step every 2 sec without exception. Only two channels were used, each corresponding to one of the two random sequences of shock, and at any given time only one of the channels could program a shock. If the channel in control turned up with a punched hole, the brief shock was delivered; if there was no hole, there would be no shock. The frequency and pattern of holes in each channel were taken from a random-number table; hence the holes turned up irregularly but with a specified probability. One channel was in control until a response transferred it to the other channel. Control reverted to the original channel with the next shock. The two channels are referred to as the post-shock distribution and the post-response distribution, in accordance with the events that put them in control. A single response transferred control to the post-response distribution until the next shock, and additional responses in the interim were totally without effect. The procedure is schematized in Fig. 1, in which the various operations are plotted in time as a function of a hypothetical sequence of responses and shocks. The parameters of the procedure are specified by a pair of numbers, the first of which is the probability of shock in the post-shock distribution and the second is the probability of shock in the post-response distribution. Thus .3-.1 means that a response changes the probability of shock per 2-sec period from .3 to .1.

Daily sessions lasted 100 min, during which the house light was on continuously, and each lever press produced an audible click, except during the shocks.

RESULTS

Acquisition. Of the 18 animals submitted to this procedure, only one failed to develop stable avoidance. This animal, F-6, remained crouched in the corner, a typical response for animals which fail to avoid on other procedures for avoidance conditioning.

Figures 2 and 3 show the course of acquisition for three naive rats, littermates M-1, M-2, and M-3, with parameters .3-.1. Other rats did not appear to differ consistently from these. Figure 2, showing total responses per minute as a function of successive sessions, reveals a variety of acquisition curves. M-1 pressed the lever 130 times during the first 100-min session, but then dropped to a low rate which increased only gradually, leveling off at about the 38th session; M-2 emitted 60 responses

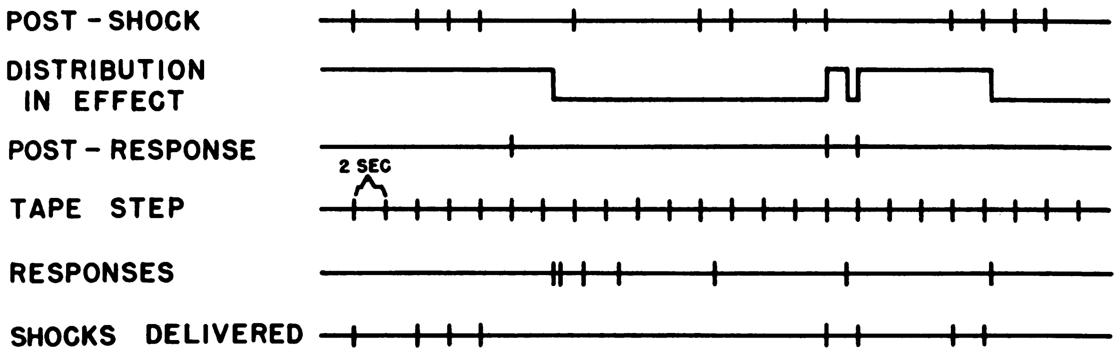

Fig. 1. Schema of the experimental procedure. The tape advances at regular 2-sec intervals. Deflections on the lines marked "post-shock" and "post-response" indicate holes in these respective channels of the punched tape. The "distribution in effect" line shows which of these channels controls the delivery of shock. Responses and shocks are shown as deflections on the lines indicated. The delivery of shock coincides with the occurrence of a hole in the tape channel currently in control. Control is changed from one channel to the other by a shock if a response has occurred since the last shock, and by a response if a shock has occurred since the last response.

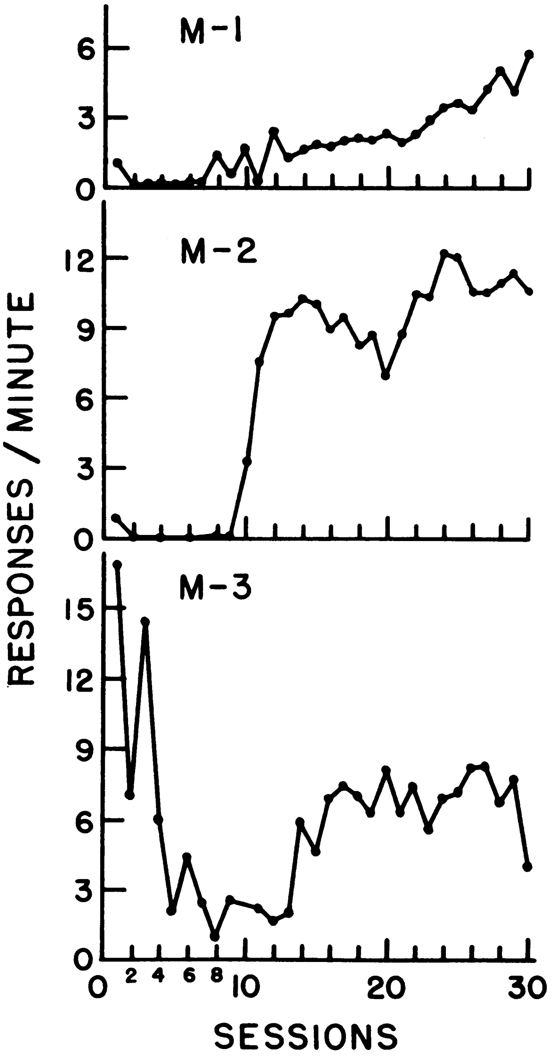

Fig. 2. Responses per minute for three rats, M-1, M-2, and M-3, during the first 30 sessions of conditioning with shock probabilities equal to .3-.1. Each point represents a 100-min session.

during the first 100 min, dropped to negligible rates until session 10, and then climbed to its final level within three sessions; M-3 began responding at a high rate within 15 min after the first session began, decreased until the eighth session, after which the rate slowly increased to a stable level. Figure 3 shows acquisition measured as a reduction in the logarithm of the latency of post-shock responses, the only functional responses in the procedure.

Since the procedure favors the response after a shock over all other responses, it seems appropriate to compare the two classes of responding, henceforth called effective and in-

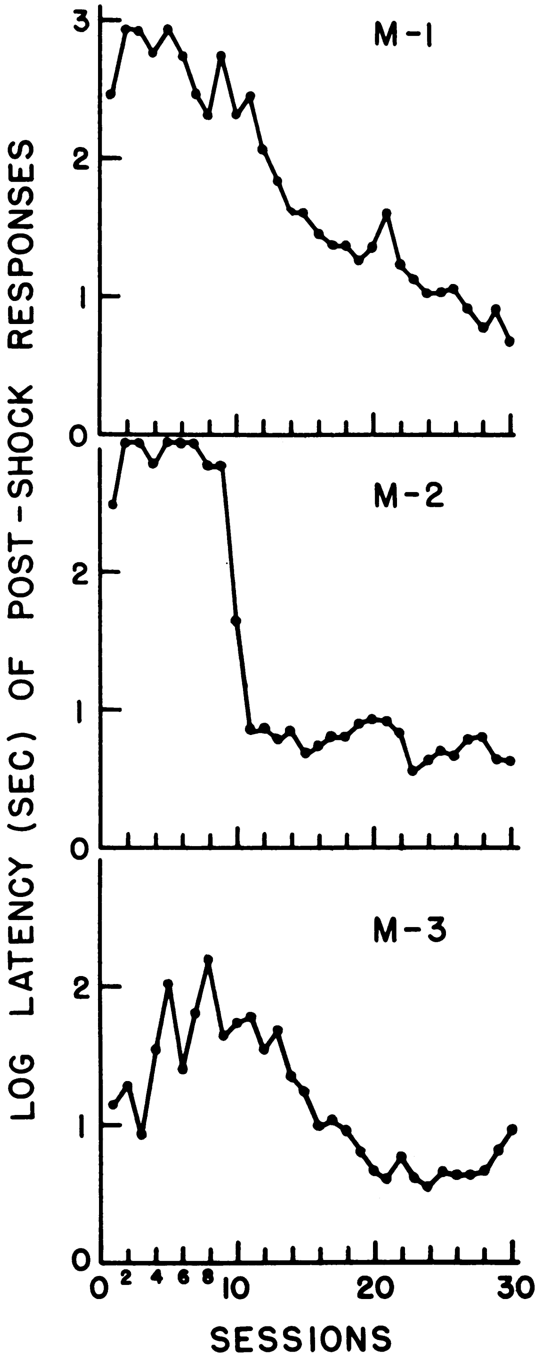

Fig. 3. Common logarithm of the latency (in seconds) between shocks and responses for three rats, M-1, M-2, and M-3, during the first 30 sessions of conditioning with shock probabilities equal to .3-.1. Each latency was measured from the shock which followed the previous response. Each point is a mean taken over a 100-min session.

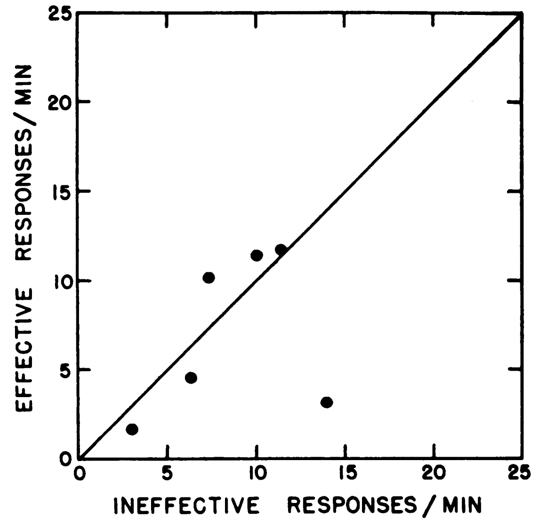

Fig. 4. Showing the correlation between effective (post-shock) response rates and ineffective (post-response) response rates during avoidance acquisition. Each point is the mean response rate for one rat, taken over five sessions beginning with the first session in which avoidance responding became apparent.

effective responding, respectively. Figure 4 is a scatter diagram of the two types of responding, taken from the initial stages of acquisition for six representative rats. The rate of effective responding is merely the reciprocal of the latencies whose logarithms are shown in Fig. 3. The rate of ineffective responding was calculated by dividing the number of such responses by the total time during which they could have occurred, that is, the time during which the post-response distribution was in effect. The points hover close to the line of equality, except for one rat, M-3, which was emitting its responses in bursts and thereby had a relatively high rate of ineffective responses. In Fig. 5, however, the scatter diagram is markedly different. This figure presents data from the same six rats under the same procedure, but at a later stage in training, when behavior had reached asymptotic levels. Now the points are consistently above the line of equality, showing that the rate of effective responding had increased relative to the rate of ineffective responding, presumably because the rats were reacting to the difference in effectiveness. There remains, however, a clear positive correlation between the two classes of lever pressing.

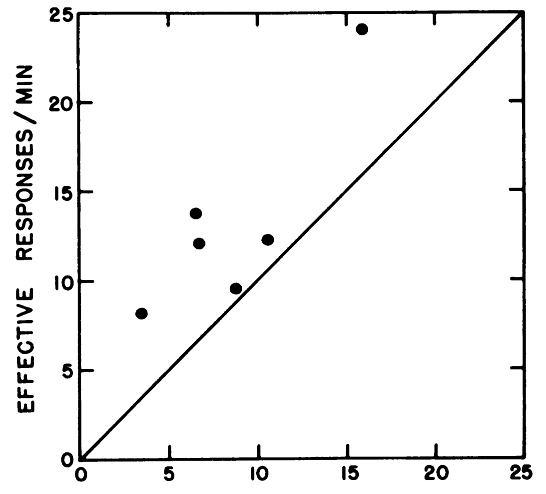

Fig. 5. Showing the correlation between effective and ineffective response rates for the same rats as in Fig. 4 for steady-state avoidance on the same shock probabilities as during acquisition, with no differing intervening conditions. Each point is a mean of seven sessions late in conditioning for one rat.

Steady-state performance. The slight, but persistent, tendency for responses to occur in bursts is documented in Fig. 6, which shows the conditional probability of a response as a function of interresponse time (IRT) in 2-sec class intervals, for four typical animals. For each animal these data were pooled over a minimum of seven sessions. The conditional probability for a given class interval is obtained by dividing the number of IRTs falling within the interval by all IRTs within or greater than the interval. For random distributions of responses in time, this measure, IRTs/Op, gives a horizontal line (*c.f.* Anger, 1956, 1963 for discussion). The elevated first point for each rat is the consequence of bursts of responses. What is more unusual about these curves, however, is that they are consistently horizontal beyond the first point, showing that, aside from the bursts, behavior is randomly distributed in time. This is strikingly different from performance under Sidman's procedure, where a temporal discrimination is the norm, revealed by higher conditional probabilities of response at longer IRTs.

Figure 6 shows how behavior is ordered with respect to behavior. Figure 7 shows how behavior is ordered with respect to the programming circuitry. For the purposes of Fig. 7, the 2-sec cycle that advanced the programcal animals. All three show a gradually decreasing function, with a maximum in the ming tape was divided into 10, .2-sec compart-

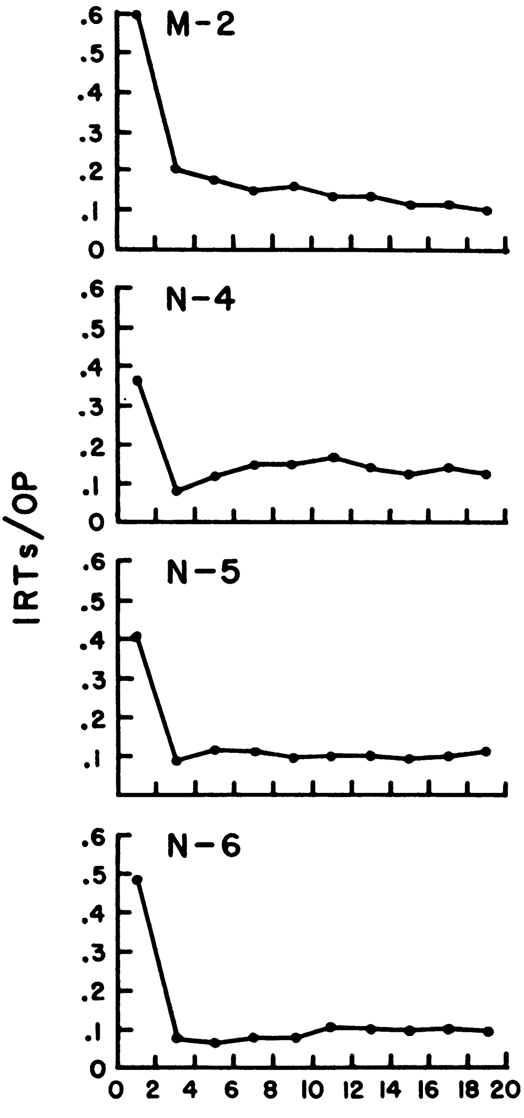

Fig. 6. Conditional probability of response as a function of interresponse time for rat M-2 with shock probabilities of .3-.1, and for rats N-4, N-5, and N-6 with shock probabilities of .3-.2. IRT shows interresponse time in 10 2-sec class intervals. IRTs/Op shows the number of any IRT divided by the number of IRTs at least that large. A final, 11th class interval is not plotted since it necessarily has the trivial value of 1.0.

ments and a tally of responses in each was kept. Figure 7 shows the mean count per session, taken over seven sessions, for three typi- second (.2-.4 sec) compartment. The elevation of the early portion of the curve is a consequence of a tendency for responses to follow closely upon shocks, which can occur only just before the first compartment (or, equivalently, just after the 10th). If responses were distributed randomly with respect to the programming cycle, this function would be flat, a condition not far from the observed results.

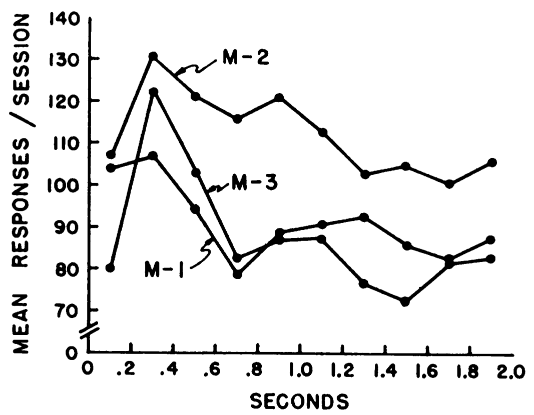

Fig. 7. Distribution of responses in successive 10ths of the 2-sec tape-advancing program for three rats, M-1, M-2, and M-3, with shock probabilities of .3-.1. The ordinate gives the mean number of responses per 100-min session in each .2-sec class interval. The end of the 10th class interval is contiguous in time with the beginning of the first class interval, when the probability distributions were sampled by the programming apparatus.

Although the present experiment was concerned primarily with the mere possibility of obtaining avoidance with a new procedure, a sufficiently varied group of parameters was explored to permit some tentative conclusions about steady-state performance. Figure 8 shows rate of responding as a function of variations in the two parameters. The open circles are for times when the post-shock probability was held constant at .3 and the post-response probability varied, in steps of .1, from .1 to .3. The filled circles are for times when the post-response probability was held constant at .1, and the post-shock probability was .1, .2, .3, and .5. All points are means of the mean performance—taken over seven sessions—of at least two rats, and most points include data from four or more rats. The different probability values were applied in an irregular order and were maintained until at least 15 sessions of stable responding had occurred. The small number of data points barely allows any more precise characterization than that the greater the reduction in shock probability, the higher the rate of responding. The two functions in Fig. 8, however, appear to be mutually symmetrical and suggest downward concavity,

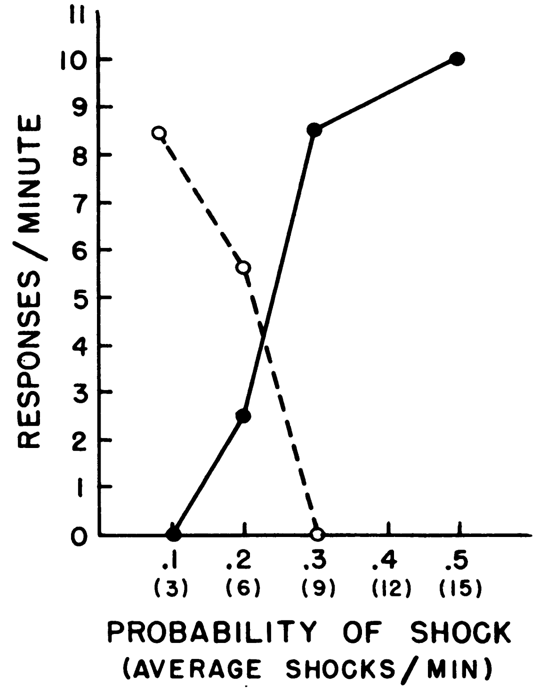

Fig. 8. Response rate as a function of shock probabilities (and equivalent shock rates). The open circles show rates obtained when the post-shock probability was constant at .3, and the post-response probabilities varied as indicated on the abscissa. The filled circles represent response rates obtained when the post-response probability was constant at .1, and the post-shock probability varied as indicated on the abscissa. Each point is a mean of the performances of at least two rats, and, in most cases, of four or five rats. The mean rates for individual rats were taken over the final seven days of stable performance on each set of probabilities.

which may prove to be based on a logarithmic relation between rate of responding and magnitude of reduction in shock probability.

Extinction. At .1-.1 and .3-.3 in Fig. 8 may be found the extinction performance, since, with the parameters equal, level pressing is ineffective. As is to be expected, the terminal rate of responding at these values is virtually zero. More interesting, however, is the resistance to extinction of this sort of avoidance behavior.

Avoidance behavior in general enjoys a certain notoriety for its persistence (Solomon, Kamin, and Wynne, 1953; Sidman, 1955), which has weighed heavily in the argument that psychopathology may bear some relation to avoidance conditioning. From this point of view, Fig. 9 is significant, for it presents what may be some sort of record for persistence in avoidance behavior. It presents in cumulative form the experimental history of a rat originally trained with parameters .3-.1 and extinguished with parameters .1-.1. Total, cumulated presses are plotted against cumulated time under experimentation. After a gradual acquisition period lasting about 3500 min (*i.e.*, 35 sessions), the rat responded at a fairly uniform rate for another 3000 min before being submitted to extinction. During about 17,000 min (170 sessions) of extinction, the rat emitted about 20,000 responses at a decreasing rate. Eventually lever pressing virtually stopped, but was readily restored by the original parameters, .3-.1.

Altogether, nine extinction curves, after various parameter values, were obtained. The unmistakable conclusion to be drawn from these is that extinction is slower the closer the parameter values are to one another, which is to say that the resistance to extinction is inversely related to the pre-extinction rate of responding (given in Fig. 8) for this procedure. Figure 10 shows the average number of presses in extinction as a function of the difference between the post-shock and post-response probabilities during the preceding condition. The n's in the figure refer to the number of extinctions, which is equal to the number of rats except for the .4 value where the number of rats was only two.

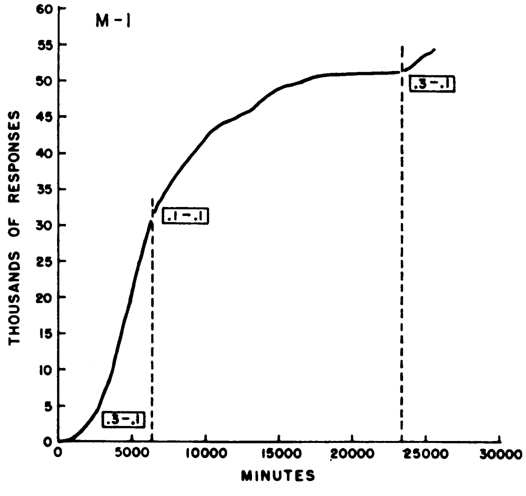

Fig. 9. Cumulative response curve for rat M-1 showing conditioning, extinction and reconditioning of avoidance. The vertical dotted lines mark transitions from one pair of shock probabilities (post-shock and post-response) to another, the new set of values being given to the right of each dotted line.

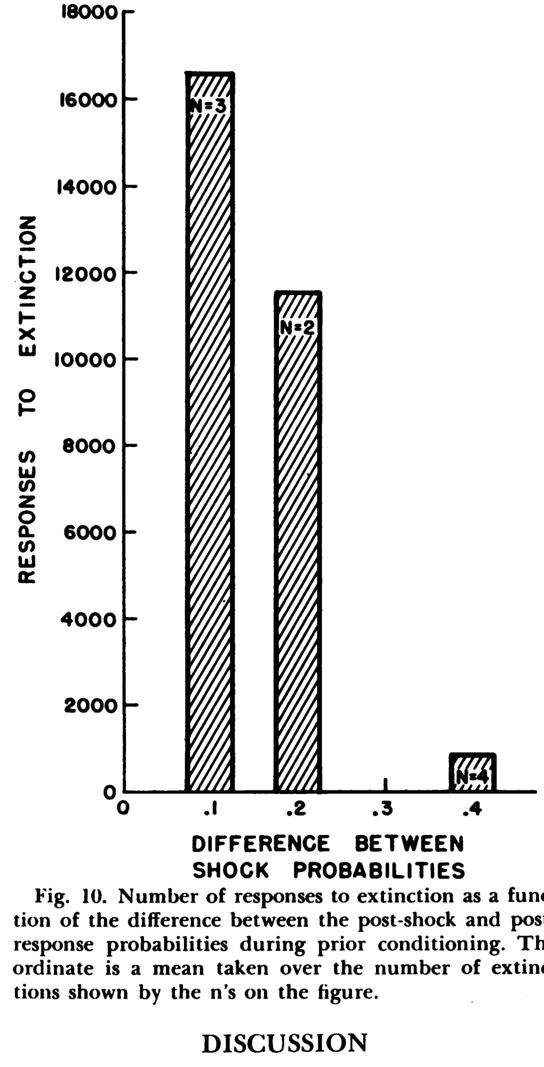

Fig. 10. Number of responses to extinction as a function of the difference between the post-shock and post-response probabilities during prior conditioning. The ordinate is a mean taken over the number of extinctions shown by the n's on the figure.

DISCUSSION

For almost three decades, investigators failed to make a clear distinction between instrumental avoidance conditioning and defensive classical conditioning (for a review, see Solomon and Brush, 1956). It was Schlosberg (1934) who first designed an experiment specifically to see if sheer contiguity between a neutral and an aversive stimulus is adequate to produce a conditioned response or whether the conditioned response need be instrumental in some way. Although his pioneering study produced no definitive answer, it set off a chain of experiments ultimately demonstrating (*e.g.*, Brogden, Lipman, and Culler, 1938) that instrumental avoidance conditioning is a more effective procedure for producing motor behavior than the classical, contiguity procedure, a conclusion lurking unrecognized in the experiments done under Bekhterev's guidance in 1907 and afterwards (see Razran, 1956).

But even though this distinction was clarified and substantiated, the typical avoidance procedure (*e.g.*, Hunter, 1935; Mowrer and Lamoreaux, 1942; Solomon and Wynne, 1953) long consisted of a pairing of neutral and aversive stimuli, with the neutral preceding the aversive by a few seconds. A successful avoidance response—one that occurs in the interval between the onset of the two stimuli—terminates the neutral stimulus and causes the omission of the aversive stimulus for that trial. Should the designated form of behavior occur after this interval ends, it will then terminate both stimuli simultaneously, serving thus as an escape response. Virtually every feature of this procedure mimics the Pavlovian paradigm, especially the paired neutral and effective stimuli, with the former slightly preceding, and a response appropriate to the effective stimulus whose transference to the neutral stimulus is the mark of successful conditioning. Even the vocabulary of avoidance conditioning retained the Pavlovian stamp, with CS-US intervals, CRs and URs, delay and trace conditioning, *etc.* The similarity in procedure and language persisted long after the existence and importance of instrumental conditioning, as distinct from classical conditioning, were widely appreciated.

The typical avoidance procedure described above has been supplemented, if not replaced, by others less similar to the classical paradigm. Mowrer and Lamoreaux (1946) showed successful avoidance conditioning when the escape response was morphologically different from the avoidance response. Sidman (1953a) disposed of the escape contingency altogether by using an electric shock so brief that virtually no instrumental responses could occur in its presence. His rats, too, learned the avoidance response. Sidman's procedure differed from the classical paradigm in another respect; it omitted any exteroceptive conditioned stimulus. The electric shock was not preceded by an originally neutral stimulus. It was, however, preceded by a certain time since the prior response or shock, and so the absence of an exteroceptive stimulus was perhaps vitiated by the presence of temporal regularity. The present procedure, with its random presentation of shock, continues the evolution away from the classical paradigm. Here the avoidance response occurs without benefit of a conditioned stimulus or an escape contingency of any sort.

These procedural changes have tended to accompany revisions in theories of avoidance behavior. In the earliest theories, formulated before the classical-instrumental dichotomy and of which Hull's (1929) is a good example, the avoidance response was merely the escape response "moved forward" in time. The reliance on Pavlovian theory was here complete and obvious. The escape response, which was the response appropriate to the unconditioned stimulus, was analogous to the classical UR, which became the CR when elicited by the conditioned stimulus. The procedural change that displaced this theory was the shift from the strict Pavlovian paradigm, in which the response had no effect on the presentation of the stimuli, to the instrumental paradigm, in which the response avoided or terminated the unconditioned stimulus. The superiority of the instrumental procedure required a theory built on more than simple contiguity. Since termination of the conditioned stimulus was a convention in all avoidance procedures, the newer theories took this as the point of departure. The various two-factor theories are the result, one factor to motivate CS-termination, the other to handle the consequences of CS-termination. Solomon's (1964) presidential address to the Eastern Psychological Association and Anger's (1963) monograph on avoidance are each recent examples.

If avoidance behavior is not escape behavior "moved forward," must it be escape from the CS? Kamin (1956, 1957) considered the possibility of a negative answer, but rejected it. He modified the usual avoidance procedure so that one group of rats terminated the CS without avoiding the US, while another group avoided the US without terminating the CS. Both of these groups learned to make the required response, but neither learned as well as a group trained on the standard procedure. Kamin explored the notion that US-avoidance can, in and of itself, maintain avoidance behavior, but settled on an admittedly tortuous rendering of two-factor theory.

The other theorist who has looked beyond two-factor theory, specifically to US-avoidance, is Sidman (1962), as noted above (Introduction). Two-factor theories are rooted in the standard CS-US pairings of classical conditioning, a procedural constraint from which Sidman almost escaped with his novel method. But as already described, Anger (1963) pointed out the conditioned stimulus implicit in the temporal regularities of Sidman's procedure.

Although it is possible to think of avoidance procedures that leave out an escape contingency or CS-US pairings, it is virtually impossible to conceive of one that omits US-reduction. A response-dependent change in the amount of subsequent aversive stimulation appears to be the *sine qua non* of avoidance conditioning. The present finding is, simply, that behavior can be maintained by reducing frequency of shock, in much the way that it can be maintained by reducing intensity (Campbell, 1956). From the point of view of reinforcement theory, the only difficulty with such a conclusion is that a change in the rate of occurrence of a stimulus cannot easily be depicted as a momentary contingency between a response and a consequence. We are familiar with theories that say a response is influenced by its having terminated an electric shock or by having terminated a conditioned stimulus associated with an electric shock. But we are unaccustomed to the notion that a response can be influenced by changing the rate of a stimulus, a change that itself can be manifested only over some period of time. This difficulty in conceptualization may, however, ultimately be submerged by the weight of experimental evidence supporting such a conclusion. In the field of positive reinforcement, a similar conclusion has been drawn regarding the importance of the rate of occurrence of the reinforcer, above and beyond the simple contingency between response and consequence (Anger, 1956; Herrnstein, 1964).

Do the present results actually violate two-factor theory? The answer depends not so much on the data, but on the definition and specification of two-factor theory itself. Unfortunately, the literature on the theory of avoidance conditioning does not provide a sufficiently clear consensus for a crucial test. Given one set of definitions, it is easy to recast the present results in the terms of two-factor theory, postulating some covert stimulus whose properties vary in concert with the changes in shock rate, as follows. The rat presses the bar and thereby accomplishes two things: it reduces the objective rate of shock, and it changes the aversiveness of a postulated internal stimulus from a higher to a lower level, the variation in aversiveness in the latter being the result of pairings with the variations in shock rate in the former. Thus, two-factor theory can be saved, but in this distended form is it worth saving? Unlike more familiar avoidance procedures, the conditioned aversive stimulus is here operationally equivalent to the primary aversive event, separable only by a gratuitous distinction between interoceptive and exteroceptive events. The primary reinforcer is presumably the reduction in shock rate and the conditioned reinforcer would be, by two-factor theory, the change in aversiveness of the hypothetical internal stimulus, once again based on a superfluous distinction having no basis in fact. In the usual avoidance procedure, the experimenter does two things on each trial. He turns on the CS, and then he turns on the US. The subject, by responding, also does two things. It turns off the CS, and it eliminates the next programmed US. Here, the experimenter does but one thing. He presents the shock at a specified rate, and the subject, by responding, also does but one thing, which is to alter the rate. It is our judgment that to extend two-factor theory to the present results is to make it essentially irrefutable, for as long as a primary reinforcer occurs in any procedure it is always possible to postulate an internal event to correspond to it, and to endow the internal event with the properties of a conditioned stimulus. The present findings are thus best viewed as evidence against the universality of a two-factor mechanism in avoidance conditioning.

The present results bear directly only on the present procedure. Indirectly, however, they suggest that shock-rate reduction may contribute to the strength of the avoidance response in other procedures, as Sidman has argued that it does in his. The extent of its contribution only further experiments can tell.

REFERENCES

Anger, D. The dependence of interresponse times upon the relative reinforcement of different interresponse times. *J. exp. Psychol.*, 1956, **52**, 145-161.

Anger, D. The role of temporal discriminations in the reinforcement of Sidman avoidance behavior. *J. exp. Anal. Behav.*, 1963, **6**, 477-506.

Bolles, R. C. and Popp, R. J., Jr. Parameters affecting the acquisition of Sidman avoidance. *J. exp. Anal. Behav.*, 1964, **7**, 315-321.

Brogden, W. J., Lipman, E. A., and Culler, E. The role of incentive in conditioning and extinction. *Amer. J. Psychol.*, 1938, **51**, 109-117.

Campbell, B. A. The reinforcement difference limen (RDL) function for shock reduction. *J. exp. Psychol.*, 1956, **52**, 258-262.

Dinsmoor, J. A. Punishment: I. The avoidance hypothesis. *Psychol. Rev.*, 1954, **61**, 34-46.

Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *J. exp. Anal. Behav.*, 1964, **7**, 27-36.

Hull, C. L. A functional interpretation of the conditioned reflex. *Psychol. Rev.*, 1929, **36**, 498-511.

Hunter, W. S. Conditioning and extinction in the rat. *Brit. J. Psychol.*, 1935, **26**, 135-148.

Kamin, L. J. The effects of termination of the CS and avoidance of the US on avoidance learning. *J. comp. physiol. Psychol.*, 1956, **49**, 420-424.

Kamin, L. J. The effects of termination of the CS and avoidance of the US on avoidance learning: an extension. *Canad. J. Psychol.*, 1957, **11**, 48-56.

Miller, N. E. Studies of fear as an acquirable drive: I. Fear as motivation and fear-reduction as reinforcement in the learning of new responses. *J. exp. Psychol.*, 1948, **38**, 89-101.

Mowrer, O. H. and Lamoreaux, R. R. Avoidance conditioning and signal duration—a study of secondary motivation and reward. *Psychol. Monogr.*, 1942, **54**, No. 5 (Whole No. 247).

Mowrer, O. H. and Lamoreaux, R. R. Fear as an intervening variable in avoidance conditioning. *J. comp. Psychol.*, 1946, **39**, 29-50.

Razran, G. Avoidant *vs* unavoidant conditioning and partial reinforcement in Russian laboratories. *Amer. J. Psychol.*, 1956, **69**, 127-129.

Schlosberg, H. Conditioned responses in the white rat. *J. genet. Psychol.*, 1934, **45**, 303-335.

Schoenfeld, W. N. An experimental approach to anxiety, escape, and avoidance behavior. In P. H. Hoch and J. Zubin (Eds.), *Anxiety*. New York: Grune and Stratton, 1950. Pp. 70-99. Reprinted by New York: Hafner, 1964.

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

Sidman, M. Two temporal parameters of the maintenance of avoidance behavior by the white rat. *J. comp. physiol. Psychol.*, 1953b, **46**, 253-261.

Sidman, M. On the persistence of avoidance behavior. *J. abnorm. soc. Psychol.*, 1955, **50**, 217-220.

Sidman, M. Reduction of shock frequency as reinforcement for avoidance behavior. *J. exp. Anal. Behav.*, 1962, **5**, 247-257.

Solomon, R. L. Punishment. *Amer. Psychol.*, 1964, **19**, 239-253.

Solomon, R. L. and Brush, E. S. Experimentally derived conceptions of anxiety and aversion. In *Nebraska Symposium on Motivation*. Lincoln, Nebr.: Univ. Nebr. Press, 1956. Pp. 212-305.

Solomon, R. L., Kamin, L. J., and Wynne, L. C. Traumatic avoidance learning: the outcomes of several extinction procedures with dogs. *J. abnorm. soc. Psychol.*, 1953, **48**, 291-302.

Solomon, R. L. and Wynne, L. C. Traumatic avoidance learning: acquisition in normal dogs. *Psychol. Monogr.*, 1953, **67**, No. 4 (Whole No. 354).

Received August 31, 1965