ESCAPE, AVOIDANCE, PUNISHMENT: WHERE DO WE STAND?

JAMES A. DINSMOOR¹

INDIANA UNIVERSITY

This paper reviews progress since the author's previous writings in three areas. In escape training, the preparatory-response explanation of bar holding still appears to be valid. In avoidance, the newer safety-signal version of two-factor theory has much to recommend it and readily incorporates Anger's conditioned aversive temporal stimuli formulation. Shock-density reduction is rejected as a substitute for two-factor theory. Finally, criticisms of the avoidance interpretation of punishment are answered and recent empirical data are cited in its support.

One of my major interests over the years has been in the general topic of aversive control—in particular escape training, avoidance, and punishment. Inevitably, as time has passed, my published views on these matters have become somewhat out of date. New arguments have been presented and new findings have been obtained. The purpose of the present paper is to offer you an updated version, in which I attempt to point out what I consider to be the advances in this area and to criticize some of the missteps along the way. I will also make use of the opportunity to reply to some of my critics.

I will begin with escape training, a topic on which I have published several empirical studies, most of them summarized in my presentation at the Miami Symposium in 1967 (Dinsmoor, 1968). Escape has remained fairly quiescent as an area of investigation, and most of the interest in the intervening years appears to have centered on a problem that many of you may find a bit esoteric—an analysis of the holding response. To summarize the basic phenomenon, it has repeatedly been found that rats exposed to escape training, on a continuous reinforcement schedule at least, typically

spend a large part of the experimental session with their paws on the bar, holding it in a depressed position. When the shock arrives, the animal lets the bar up again momentarily, only to return to the resting position. The release of the bar and the subsequent press are sufficient, under most arrangements, to terminate the shock. The animal then remains on the bar until the next shock appears. Holding the bar is of interest not only because it is an experimental nuisance, making it difficult to compare escape behavior with that observed when food or water is used as the reinforcer, but also because it is a puzzle, a pattern of behavior not readily explained in terms of standard conditioning principles.

We were, of course, aware from the beginning of the possibility that unlearned behavior might be involved. This was the obvious explanation. Indeed, Mowrer (1940) had implied as much in describing the behavior of his rats during an earlier study of escape from shock: "On the second trial, the animal usually hit the pedal within 2 to 4 minutes and very commonly 'froze' as soon as it did so, in precisely the position that it happened to be in at the instant when the shock went off (reminding one of the game of 'Statue' played by children)" (p. 64). But in those days it was not the fashion, as it is today, to take innate behavior as the target of investigation in a learning experiment. Matters of this kind were treated merely as difficulties to be circumvented in order to study more general laws, those governing learned behavior. In the years since our initial work, renewed emphasis has been

¹Revision of an Invited Address delivered at the Second Annual Meeting of the Midwestern Association of Behavior Analysis, Chicago, May 1, 1976. I would like to acknowledge the influence of earlier, unpublished critiques by Douglas Anger of the shock-density-reduction interpretation of avoidance learning. Reprints may be obtained from the author, Department of Psychology, Indiana University, Bloomington, Indiana 47401.

placed on the role of innate responses in escape from shock (*e.g.*, Bolles and McGillis, 1968; Davis and Burton, 1974; Davis, Hirschorn, and Hurwitz, 1973), but it is not clear to what type of generalization such investigations are likely to lead.

Furthermore, there were indications that freezing was not the *only* factor involved. For one thing, Keller (1941), using light as the aversive stimulus, had also noted prolonged holding of the bar. It seemed unlikely that holding was an innate response to light as well as to shock. For another thing, freezing did not inevitably lead to continued pressure on the manipulandum. For example, in the previously cited study, Mowrer's "pedal" was actually a panel with a vertical surface mounted in one wall of the experimental chamber. Although the animal was said to freeze following termination of the shock, what it specifically did during most of the time between trials was to "sit (often on its haunches facing the pedal) in a position that was expeditious for making the next pedal reaction" (Mowrer, 1940, p. 64). This description suggested that more was going on than was implied by the simple label "freezing".

Seeking procedural modifications that might reduce or eliminate holding, we first checked a plausible hypothesis suggested by Fred S. Keller (1941). When food or water was used as the reinforcer, Keller had pointed out, it was necessary for the animal to release the lever to obtain the reinforcing substance. But when termination of shock was used, the last act in the sequence was to press the bar down. No release was necessary. Therefore, the animal remained on the bar. Unfortunately, when we added release of the bar as a requirement for termination of shock, the results were not very conclusive. Although we noted some reduction, a good deal of holding continued to occur (Dinsmoor, Hughes, and Matsuoka, 1958).

Some data collected by the late Sam L. Campbell, as part of his doctoral dissertation, provided the necessary lead to a more fruitful hypothesis, the hypothesis that holding down the bar served as a preparatory response, permitting more rapid termination of the shock. Campbell (1962) compared the speed (latency) with which the shock was terminated on trials on which the rat was holding at the moment of shock onset and on trials when it was not holding. His records indicated that the animal ter-

minated the shock more promptly when on the bar than when off the bar. While this finding suggested that holding the bar down permitted more rapid termination of the shock and that more rapid termination of the shock selectively reinforced the behavior of holding the bar down, we thought of this more as a hypothesis than as a conclusion. Campbell's data were correlational, and correlational data are never as persuasive as those obtained with a design in which one variable is freed or isolated from its covariation with other factors by direct intervention of the experimenter.

Accordingly, the work reported by Dinsmoor, Matsuoka, and Winograd (1958) included two experimental manipulations, each designed to show that altering the consequences would affect the amount of holding that occurred. In the first of these two experiments, the sequence of actions necessary to terminate the shock was varied. Both groups of rats faced the same requirement if they were off the bar when the shock arrived. They had to press the bar down and then let it up again. But if the rat was on the bar when the shock arrived—that is, holding—the requirements differed for the two groups. In the first case, it was necessary to let the bar up, press it, and then let it up a second time. Three successive actions were required, compared with two actions when not holding. In the second case, it was necessary only to let the bar up if the rat was already on it. Only one action was required, compared with two actions when not holding. For the first group, holding made the sequence longer; for the second group, holding made the sequence shorter, while the terminal action was in both cases the same. The result was that the group allowed to profit from holding was soon spending about 50% more time on the bar than the group that was, perhaps, penalized.

In the second experiment, the benefit accruing to holding was reduced or removed by the imposition of an arbitrary delay. One group of rats was trained in the usual way, as a basis of comparison. But members of a second group were subject to a delay of 1.5 sec if they were on the bar when the shock began. During this delay, manipulation of the bar was ineffective. After the delay was over, a subsequent release and press of the bar were still required before the shock was terminated. Members of the third group were subjected to a similar contin-

gency, but the delay lasted for 3 sec. The result was that the zero-delay group showed an increase in holding during training from about 3 sec per trial, at the beginning, to about 8 sec at the end—a value very close to that for the corresponding group in the previous experiment. But the group subject to a 1.5-sec delay showed a much smaller rise, from about 4 sec to 6, on the average. And the group subject to a 3-sec delay showed little, if any, increase in holding during the 200 training trials. In short, when the benefit to be gained from remaining on the bar between presentations of the shock was eliminated by the imposition of a delay, little or no learning to hold occurred. While holding was substantial, it remained at its initial level.

To summarize the situation, Campbell's (1962) data had shown that holding the bar down before onset of the shock led to more prompt termination, once it arrived, and our data showed that prompt termination of the shock affected the amount of holding that occurred. The cycle was complete. I cite these studies to show you that we had fairly solid grounds for concluding that one of the reasons, at least, why the rat learns to hold the bar down is because this leads to faster termination of the shock, once it appears.

The only serious challenge that I have seen to this conclusion comes from an experiment conducted by Keehn (1967). Under Keehn's procedure, the onset of the shock was preceded by 5 sec of light and tone. The performance required of the animal was a chain in which it had to press and release bar A and then press bar B. If this occurred during the 5 sec of light and tone preceding the shock, the shock was avoided; if it occurred after the onset of the shock, the shock was terminated. Only one of the four animals actually avoided the shock with any frequency, so that basically the procedure turned out to be escape training. Keehn's main finding was that the rat held down the bar that terminated the shock, bar B, rather than the bar with which the chain of behavior was initiated, bar A. That is, the behavior appeared to be perseverative rather than preparatory.

Keehn's procedure was not, of course, one that offered much advantage to preparatory holding. Even if the rat was on bar A when the shock began, it was required to complete a chain that occupied an appreciable period of

time before it could turn off the shock. When the delay in termination is relatively long either way, the difference produced by holding is smaller in proportion to total time than when the delay is relatively short. To put the matter another way, picture the classic gradient describing the effectiveness of the reinforcer as a function of the length of time by which it follows the response (e.g., Chung, 1965). At zero delay, the curve drops quite rapidly. But at longer delays the gradient becomes relatively shallow: a difference of given magnitude along the horizontal axis is accompanied by a much smaller change along the vertical axis.

Furthermore, Keehn's procedure minimized the subject's exposure to the preparatory contingency. Under normal circumstances, any tendency to freeze as an unlearned response to the experience of electric shock leaves the animal with its paws on the crossbar, exposed to the selective reinforcement that results from being in this position. But under Keehn's procedure, freezing would occur on the bar that terminated the shock, and this behavior would not contact the contingency by means of which the preparatory response is learned. Therefore, Keehn's data represent a special case and do not apply to more conventional escape training procedures.

A subsequent experiment added further support to the preparatory response interpretation. After training their rats initially on the usual single-lever procedure and obtaining substantial amounts of holding, Marriott and Russell (1969) shifted them to a procedure involving two levers, side by side. On each trial, a lamp was lit over one lever or the other, and the rat was required to operate the lever under the lamp that remained dark. Under this procedure, which greatly reduced the effectiveness of holding as a preparatory response (since the rat had no advance indication of which lever to hold), the percentage of session time spent on the lever was reduced to a relatively low value. It should be noted, perhaps, that the final action required by Marriott and Russell for termination of the shock was release of the bar, while the final action required by Keehn had been a press.

The data obtained by Davis, Hirschorn, and Hurwitz (1973) indicate that some attention must also be paid to the issue of what response topographies are natural or convenient to the

subject. (See also Flint, 1969.) These authors changed the requirement for termination of the shock from pushing the bar down to pushing it up. With this arrangement, the rats typically spent the time between shocks underneath the bar rather than on top of it. From this vantage point the rat could quickly raise the crossbar by thrusting its back upward against it when the shock came on. Davis *et al.* felt that their rats' behavior was consistent either with a preparatory account or with one based on freezing. But they noted that the rat did not actually maintain contact with the bar after the shock terminated, and suggested that this observation raised a problem for Keehn's characterization of holding as perseverative. It seems to me that it also raised equivalent problems for freezing and preparatory accounts. If the rat freezes before it releases the bar under normal arrangements, it should freeze in contact with the bar under those employed by Davis *et al.* And presumably the rat could terminate the shock more quickly if it were actually in contact with the bar, rather than merely positioned beneath it. All three accounts, it would appear, must make some concession to what is easy and natural for the rat.

The same issue arises in a subsequent study by Davis and Burton (1974). These authors kept a continuous record of the pressure or force exerted by their animals as they rested on the bar. Typically, the level was higher than that required to be registered as a response, and they pointed out that holding with a force below the criterion would have permitted quicker termination of the shock. Under such circumstances, the initial downward thrust against the bar as the animal reacted reflexively to the onset of the shock would be effective, rather than subsequent restoration of pressure after it had "lurched" off the bar. But probably the animal rested its weight on the bar, and this determined the amount of pressure it exerted. Holding at a lower force may have been awkward or effortful: the level appeared to be insensitive to other variables, such as changes in the criterion set for recording as a response. And holding with a force below the criterion level apparently did not occur with sufficient frequency to provide very effective contact with its consequences.

Although the position taken in a brief article by Bolles and McGillis (1968) has sometimes been cited as an *alternative* to the preparatory

response hypothesis, it appears to be better characterized as an *extension* of that hypothesis to special circumstances. Using a lever that "could not be perceptibly moved" (p. 261) and a procedure that required a release of the lever, following onset of the shock, and a subsequent press, these authors found that their rats often terminated the shock within 0.05 sec. They maintained that responses which occurred that quickly could not be operants but must be unlearned responses elicited by the shock. They also noted that these short-latency responses occurred when the rat was holding. What their rats had learned, they suggested, was a response of "freezing" in such a position that the subsequent onset of shock led to "a reflexive lurch or perhaps a direct current-induced muscular contraction that results in the bar being depressed" (p. 261). Freezing began as an innate response to the shock but persisted only if it took a form that permitted these unlearned responses to terminate the shock. "Freezing on the bar", they concluded, "is an operant" (p. 261).

The difference between this analysis and the preparatory response hypothesis lies in the nature of the final behavior that holding is said to make possible. According to the preparatory response hypothesis, as originally stated, the final behavior is a more efficient form of the escape response. According to Bolles and McGillis, it is an unlearned, reflexive response to the shock. And in their apparatus such a response may indeed have played a role. With their lever, very little movement was required, and small fluctuations in pressure were recorded as responses. (See also Davis and Burton, 1974.) Like any other behavior that led to especially prompt termination of the shock, these responses presumably produced selective reinforcement of holding. But it is doubtful that the Bolles and McGillis analysis applies to other situations, in which larger movements are required. In most of our escape work, including our studies of holding, a much larger excursion (about 1 cm) was required, and muscular twitches of the sort they described were not likely to have been recorded. The temporal distribution of bar depressions following termination of shock was quite different (Dinsmoor, Hughes, and Matsuoka, 1958) from that noted by Bolles and McGillis. Nevertheless, the amount of holding increased as training continued. A comprehensive explanation of

holding should make note of both types of behavior: elicited responses when these are effective in terminating the shock and, more commonly, operant presses reinforced by termination of the shock. The only requirement is that these responses produce earlier termination of the shock when the animal is holding the bar down at the moment of onset than when it is somewhere else, off the bar.

Bolles and McGillis also referred to the "non-operant nature of the bar-press escape response", and it is easy to overgeneralize their message. A close reading of the text indicates that they were not claiming that all bar presses to shock were reflexive but only some of the responses, under certain conditions. I would like to make it clear that the responses studied in my laboratory were legitimate escape responses. Their latencies were of the order of seconds rather than fractions of a second, and early in training, at least, the majority of them were initiated from positions other than on the bar at the moment of shock onset. Even when we shifted, in our later studies, to a lever with a smaller excursion, the operant nature was clear. In the studies I conducted with Marilyn Clayton (1963, 1966), for example, we required an average of three presses to terminate the shock from the very beginning of training. Although a reflexive response to the onset of shock might have facilitated acquisition, it certainly could not have done the whole job, and it was clear to casual observation that a different form of behavior was involved. Note also that escape-reinforced behavior has been placed on various schedules, producing appropriate patterns of responding, and that it has been placed under discriminative control, shaped, and chained. If we make some allowance for special features of the situation, we can reproduce most of the phenomena observed in operant work with food or water as the reinforcer (Dinsmoor, 1968). This, to my mind, is a much more important point than the possibility that reflexive responses may sometimes occur.

Before leaving the topic of escape training, I would like to respond to LoLordo's criticism (1969) of some work conducted by Dinsmoor and Clayton (1963, 1966). The purpose of the Dinsmoor and Clayton studies was to show that a stimulus that preceded the termination of shock was capable of serving as a conditioned reinforcer. I will skip over the prelimi-

nary training and describe the final maintenance procedure. First, the shock was turned on. This was applied to the rat's feet through a grid floor. Once the shock was on, pressing the lever produced a white noise—the intended reinforcer—on a variable-ratio schedule requiring an average of three responses. If no further presses ensued, the white noise and the shock terminated 30 sec later. On the other hand, if the animal continued to press the bar, the timer was reset by each successive response, and the noise and the shock did not terminate until 30 sec after the last press had occurred.

Under this procedure, our rats pressed the bar until the noise was produced and then quit. A similar performance was also obtained with a group of rats for which we used the same procedure, except that no contingency of any kind was applied to pressing after the noise came on. Since responding was maintained at a substantial level until the noise appeared, but not thereafter, we concluded that the onset of the noise was the reinforcing event for this response (Dinsmoor and Clayton, 1966).

The major control procedure used to buttress this conclusion was to present the noise on an arbitrary schedule for a number of sessions, rather than contingent upon the bar-pressing response. Under this procedure, the animals' rates of pressing dropped to relatively low levels; when the contingency was restored, so was the rate of pressing.

LoLordo was not impressed with this evidence. Since the shock terminated 30 sec after the onset of the noise, he pointed out, eliminating the connection between the response and the noise also eliminated the connection between the response and the shock. That is, termination of the shock no longer depended on the occurrence of a press 30 sec before; it depended wholly on the passage of time. We were well aware of this issue and had discussed it at some length. Once again it is necessary to consider the shape of the delay-of-reinforcement gradient, *i.e.*, of the functional relationship between the effectiveness of the reinforcer and the length of the interval by which it follows the response. (For a review of the literature, see Renner, 1964.) The relationship is curvilinear. The decline is steep at first but soon becomes more gradual. By 30 sec, the function has necessarily become almost horizontal in slope. Therefore, at this point on the

function the addition of another few seconds between the occurrence of the last previous press and the termination of the shock should not make much difference. It does not appear to provide an adequate explanation for the collapse of a bar-pressing performance that had been well maintained up to that time. On the other hand, the white noise had always followed some depression of the bar with a zero delay. At this point, the gradient is usually quite steep. Addition of even a few seconds to the average delay might well have a profound effect on the rate of pressing.

AVOIDANCE

In a paper published in 1954, I presented an outline of what is sometimes called the two-factor analysis of avoidance behavior, as a prelude to extending that analysis to account for the suppressive action of punishment. Although relatively little theoretical material on punishment has appeared since that time, a great deal has been published on avoidance, and I shall have to be fairly selective in what I choose to review. I will restrict my comments to three major developments that have occurred since my 1954 paper: (a) the safety-signal interpretation of avoidance; (b) Douglas Anger's CATS approach; and (c) shock-density reduction as an alternative formulation.

The safety-signal interpretation of avoidance might be characterized as an inside-out or upside-down version of the conventional or traditional two-factor theory. The two processes alluded to in the conventional theory are as follows: first, as a consequence of being paired with the shock, a warning signal becomes aversive to the subject; second, termination of the warning signal following the avoidance response provides the reinforcement for that response. But if there is a stimulus positively correlated with the shock, as is the case with the warning signal, the absence of that signal must be negatively correlated with the shock. Thus, for every warning signal there is a safety signal; and conversely, of course, for every safety signal there must be a warning signal. Therefore, the two processes of two-factor theory can be characterized in another set of terms: first, as a consequence of being inversely correlated with the receipt of shock, a safety signal becomes reinforcing to the subject; second, production of the safety signal maintains the

avoidance response. The main problem with this interpretation is that it is all but impossible to distinguish the two versions of two-factor theory from each other. In one version, reinforcement is provided by termination of the warning signal; in the other version, it is provided by production of the safety signal. Unfortunately, termination of the warning signal and production of the safety signal are one and the same event.

Nevertheless, Dinsmoor and Sears (1973) suggested that the two theories may indeed be separated. In this experiment, we trained our pigeons to step on a treadle to postpone unsignalled shock. Each such response was followed by 5 sec of tone at 1000 Hz. Thus, the tone was roughly equivalent in its temporal position and relationships to the natural feedback from the response. Since the response that produced the tone also postponed the shock, the tone was inversely correlated with the shock. It met the procedural specification for a safety signal. Following this training, we subjected our pigeons to a test designed to measure the effectiveness of the tone as a conditioned reinforcer. First we gave them a half-hour warm-up, in which the standard procedure was employed. Then, we disconnected the shock stimulator but continued to provide the tone. On different test sessions, interspersed with additional training sessions, we conducted this test not only with the original 1000-Hz tone but also with tones of one half and one quarter and with tones of twice and four times that frequency, so that we could plot a gradient of generalization. This strategy was reasonably successful. From each bird a gradient was obtained that peaked at the original training value of 1000 Hz and tailed off fairly consistently for higher and lower values. Note that these gradients of generalization reflected variation along a dimension that was specific to the safety signal. The warning signal in this experiment was simply the absence of the tone. Since the frequency dimension is orthogonal to the absence of tone, the gradient could not be explained in terms of degree of similarity to the warning signal. The effectiveness of each frequency as a conditioned reinforcer could be attributed only to its similarity to the safety signal, the tone inversely associated with the receipt of the shock.

Two notes should be added. First, our demonstration that the safety signal provides posi-

tive reinforcement does not preclude the possibility that termination of the warning signal also makes a separate contribution. Indeed, it seems likely that both types of reinforcement are operating in most avoidance procedures. Second, the Dinsmoor and Sears experiment not only provides evidence for the reinforcing role of the safety signal; in addition, it suggests that whatever natural feedback is received by the animal when it performs the avoidance response should also serve as an effective safety signal.

With the safety-signal theory in hand, it is relatively easy to integrate Douglas Anger's Conditioned Aversive Temporal Stimuli or CATS approach. Noting the existence of temporal discrimination in free operant (unsignalled) avoidance, Anger (1963) suggested that there must be some type of stimulation correlated with the dimension "time since last response". Stimulation associated with low values along this dimension should be relatively low in aversiveness, and stimulation associated with values approaching the length of the response-shock interval should be relatively high. Therefore, when a response occurs and resets the clock, so to speak, back to zero, a reduction occurs in the level of aversive stimulation. In the context of a safety-signal approach, Anger's theory can be reworded without change in the functional relationships it subsumes. The feedback stimulation produced by the response is assumed to fade with time, following the model of a stimulus trace or, if you wish, a short-term memory. During this process, the feedback gradually loses its value as a safety signal. Whenever the designated response occurs again, however, the safety signal is restored to its full value. In other words, the CATS approach is virtually a corollary to the response-feedback safety-signal theory, and the arguments set forth by Anger (1963) for the CATS approach become arguments for the safety-signal theory.

The main alternative to a two-factor theory, within an S-R framework at least, is the shock-density-reduction theory of avoidance. This theory was first proposed by Sidman (1962) and later argued at greater length by Herrnstein and Hineline (1966) and by Herrnstein (1969). Although they correspond in vocabulary, the Sidman and the Herrnstein versions appear to differ in content. Sidman's account appeals to the effect of a given pattern of performance on

the overall number of shocks, Herrnstein's to localized reductions in density immediately following individual instances of the designated response. Sidman's account was based primarily on some data he obtained using a procedure in which shocks were concurrently arranged on independent avoidance schedules controlled by two separate levers (Sidman, 1962). Some of his animals responded almost exclusively on one lever, and the lever these animals chose was not the one that postponed its shock for the greater length of time but the one that produced the greater decrement in the number of shocks per session (Sidman, 1966, p. 479 ff.). However, data collected by Bolles and Popp (1964, Exp. IV) called Sidman's version of the theory into question. Bolles and Popp employed a procedure in which pressing the lever produced an interruption in the sequence of shocks, but only after the next scheduled shock had been delivered. Their procedure was equivalent to the avoidance procedure normally used by Sidman, except that the interval between the response and the very next shock was not altered. The overall density of shock was reduced to the same degree, but the rats did not learn to press the bar. The opposite case, in which the overall density of shock was *not* reduced, but the response was nevertheless established, is represented by an experiment by Hineline (1970, Experiment I) which will be described later in this paper.

Herrnstein and Hineline's (1966) comment that the Bolles and Popp procedure "delays any reinforcement due to the reduction of shock rate" (p. 422) made it clear that their version of the theory differed from that originally proposed by Sidman. Overall correlations between the rate of responding and the rate of shock delivery were not deemed sufficient. Instead, reliance was placed on reinforcement by a temporary decrement in the density of the shock, localizable in time, following individual instances of the response. Although there is some ambiguity in his remarks, Herrnstein's later comment that "the one-factor form of avoidance theory . . . is but the law of effect itself" (1969, p. 59) seemed to imply the same thing. This version of the shock-density-reduction hypothesis sounded more promising than the earlier version, since it appeared to provide for selective contact between a reinforcing event and a specific response. Indeed, Herrn-

stein and Hineline were able to translate this version directly into an experimental procedure. In their experiment, no ostensible change in stimulation occurred as a result of the rat pressing the lever. All that happened was that control over the delivery of shocks was transferred from a random schedule of relatively high density to one of somewhat lower density. This procedure did work: 17 of the 18 rats employed by these authors learned to perform the response. The rate of pressing was a positive function of the probability of shock prior to the response and an inverse function of the probability of shock during the period immediately following the response. Since the procedure was designed on the basis of the shock-density hypothesis, its success in maintaining behavior seemed to attest to the validity of that hypothesis. But the results also fit quite readily into the safety-signal version of two-factor theory. Only the first shock following the response was arranged on the low-density schedule; at that point, the program reverted to the high-density schedule. Therefore, all that Herrnstein and Hineline actually programmed was a CS-US interval for the stimulation accompanying the bar-press response that was longer, on the average, than the CS-US interval for any other stimulus occurring in the experimental situation. Clearly, the natural feedback from the bar-pressing response should function as a safety signal and should therefore serve to reinforce that response.

In any case, it is difficult, logically, to see how a reduction in shock density could, by itself, serve as a reinforcing event. It has no locus in time. Let us assume for purposes of argument that the organism has all the logic, the speed, and the accuracy of a modern high-speed computer. Moment by moment it calculates, for some sample of time that has just passed, the frequency with which shock is being delivered. Now, the *increases* in estimated rate of shock delivery are sudden. They are caused by a discrete event, the delivery of a shock, and they can be localized in time. But the *decline* in estimated shock density is gradual, since it depends on the passage of time. It continues throughout all parts of the session, except when interrupted by the delivery of a shock, for all passage of time without shock is effective. There is no moment ever, however, when a sudden decrement occurs, and therefore there is no way of locating such a decre-

ment in time. If the reduction in estimated shock density is a continuous process, not localized in time, how can it become correlated with any specific item of behavior? It follows all behavior, but none more closely or more consistently than any other.

To preserve the shock-density formulation, it might be argued that the subject can average the rate or density of shock over several successive samples of postresponse time and can then compare the resulting value with the density of shock at other times during the experimental session. However, this raises a new problem. Average shock densities are computed by tallying the number of shocks over certain selected periods of time, summing the durations of those same periods of time, and dividing the first term by the second. To perform this calculation, the animal must in some fashion determine which shocks are to enter into the numerator of the fraction and which are not. The only way in which it can allocate a shock to one category or the other, so far as I can see, is on the basis of what stimuli (or stimulus traces) are present when that shock arrives. In similar fashion, the animal must sort out samples of time, assigning some to the denominator of the fraction and rejecting others. It must also determine when each sample begins and when it ends. The only basis on which one sample of time can be distinguished from another, so far as I can see, is again in terms of the stimuli that are present. And the only way in which the boundaries of a given sample can be marked is in terms of a change in stimulation. Therefore, if shock-density-reduction theory takes the route of pooling successive samples of postresponse time in order to compute relative density, it will be forced to recognize the role of the stimuli by means of which successive samples are identified. And this, of course, is precisely the outcome that one-factor, shock-density-reduction theory has attempted to avoid.

The same pressure to recognize the role of the stimuli associated with the different densities of shock emerged in Herrnstein's (1969) review article. In that article, Herrnstein summarized an experiment conducted in his laboratory by Hineline (1970, Experiment I). Each 20-sec segment of the experimental session included one shock, no matter what the rat did. The 20-sec segment was divided into two parts, which were of equal length if the animal failed

to respond. During the first part, the lever was in the chamber, ready to be pressed; during the second, the lever was retracted. For some animals a buzzer accompanied the lever's presence, for others the lever's absence. If no press occurred during the first part, the shock was delivered after 8 sec and the lever retracted 2 sec later. But if the animal pressed the lever, the first stimulus condition was immediately replaced by the second one, and the shock normally delivered in the first condition was postponed until late in the second condition. This procedure was effective in establishing and maintaining pressing, despite the fact that such behavior produced no reduction in the total number of shocks delivered. This finding appears to refute the early version of shock-density-reduction, at least, as an all-encompassing explanation of avoidance behavior. But Herrnstein suggested that the success of Hineline's procedure could be attributed to the reduction in shock rate associated with the change from one stimulus to the other. "To the rat, then, each cycle presents a choice between a pair of stimuli in whose presence the potential shock rates are different" (1969, p. 65). Although Herrnstein called the two stimuli discriminative, rather than conditioned stimuli, there was nothing else in his account to distinguish it from a conventional two-factor formulation. "Two-factor theory says, in effect, that an avoidance procedure must contain a potential difference in shock frequencies which is under the control of the animal's behavior and which is signalled by an initially neutral stimulus" (p. 59).

However, regardless of whether it should be characterized as a one-factor or as a two-factor account, Herrnstein's analysis turned out to be incorrect. His estimation of relative shock densities was based on only those cases in which a shock did, in fact, occur in the presence of the particular stimulus under consideration; it did not take into account the question of whether the single shock on a given trial was delivered in the presence of the first or in the presence of the second stimulus. As the frequency of responding increased, so did the resulting frequency with which the shock appeared in the presence of the second stimulus. By the time performance approached asymptote, for two of Hineline's three rats the density of shock was higher in the presence of the second stimulus than it was in the presence of the first (Hine-

line, 1970, Figure 5). That is, the rats were terminating a lower-density stimulus and substituting one associated with a higher density of shock.

A reformulation is required. As in many another situation, the key to success of Hineline's procedure appears to be the mean length of the interval between the response and the shock. If the animal pressed the bar promptly when the first stimulus came on, shock did not appear for nearly 18 sec. If it responded at any time during the first 8 sec, shock did not appear for at least 10 sec. Since 10 sec was the average time to shock at any randomly chosen moment, under this procedure, it is clear that for the stimuli associated with bar pressing the CS-US interval was longer than for the stimuli associated with any other item of behavior. In other words, the feedback from the response should have served as a safety signal.

It may also be necessary to comment on the results of Hineline's second experiment. The procedure was identical to that used in the first experiment, except that the shock appeared precisely 8 sec after the bar was pressed, and a new cycle was initiated with the insertion of the bar again 2 sec after that. Under this procedure, pressing the bar reduced the length of the cycle and increased the overall density of the shock. It was to this feature—the increase in shock density—that Hineline attributed the failure of conditioning under this procedure (1970, p. 265). But it should be noted that when the animal did *not* press the bar, shocks came every 20 sec, and the mean interval between the stimulation accompanying random behavior and the shock was 10 sec. On the other hand, when the animal *did* press the bar, the interval between the accompanying stimulus pattern and the shock was only 8 sec. The CS-US interval for bar pressing was shorter than the CS-US interval for other behavior. Thus, the sensory feedback from bar pressing seems likely to have become a conditioned punisher, rather than a safety signal. No appeal to the increase in overall shock density seems to be necessary to account for the failure of the animals to acquire the bar-pressing response under this procedure.

PUNISHMENT

Finally, let me comment on one more issue. The purpose of the articles in the *Psychologi-*

cal Review (Dinsmoor, 1954, 1955) was to present my arguments for a theoretical interpretation of punishment. In those articles, I suggested that the suppressive action of punishment could be attributed to the establishment of avoidance reactions that were opposite in their physical consequences to the sequence of behavior being punished and which therefore interfered with its occurrence. Time has dealt kindly with this hypothesis, perhaps because it has been difficult to test. At what might be called the philosophical or logic-of-science level, the avoidance hypothesis was criticized by Azrin and Holz (1966) in their widely cited and highly influential chapter on punishment. "It appears more parsimonious", they commented, "to consider punishment at the . . . descriptive level . . . without appeal to changes in behavior that thus far have been inferred rather than directly measured" (1966, pp. 435-436). To be purely descriptive may be a very safe road to travel, scientifically, but I would not categorize it as parsimonious. To me, the objective of science is to account for the greatest possible number of facts with the smallest possible number of rules. When the issue is stated in these terms, it seems parsimonious, *i.e.*, economical, to account for the suppressive action of certain response-contingent events on the basis of rules that have already been formulated to deal with some other body of data. That way, no new concepts are needed. On the other hand, it is not parsimonious or economical to account for punishment on an *ad hoc* basis, by postulating a negative law of effect. Such a practice could lead to a horrendous proliferation of theoretical principles, each independent of the other.

Indeed, in another influential critique of the avoidance theory of punishment, Rachlin and Herrnstein (1969) conceded this point. It is conceptually simpler, they granted, to derive both increments and decrements in behavior from a single law of effect than to write separate laws for each. Rachlin and Herrnstein offered other criticisms, however. For one thing, they pointed out that the avoidance hypothesis does not account for the suppressive action of shock that is not contingent upon the recorded response, as in conditioned suppression. Actually, I had attempted such an extension in the 1955 paper (pp. 102-103). But the main point here is that while an explanation of conditioned suppression would add to the economy

of the system, it is not therefore incumbent upon the avoidance theory of punishment to offer such an interpretation. Conditioned suppression could be an entirely independent phenomenon involving, for example, an innate response of crouching or freezing. The avoidance hypothesis is not required to take on the burden of explaining some other effect.

Another objection raised by Rachlin and Herrnstein is that the two classes of behavior, responding and nonresponding, or the recorded response *versus* other behavior, are not symmetrical. It is notoriously difficult to establish a reliable bar-pressing performance, for example, let alone key pecking, under an avoidance procedure, but it is relatively easy to suppress it with a punishment procedure. The members of the class "other behavior" seem to be much easier to condition than the pressing or the pecking. In the preceding chapter in the same book, Estes (1969, p. 63) raised much the same point.

In answer to this criticism, first let me note that it is not all that clear that behavior is, in fact, rapidly suppressed by punishment. In conventional laboratory studies, the effects of punishment, defined in terms of response-contingent delivery of the stimulus, are often obscured by other effects of the shock. Brief shocks suppress responding even when they are delivered independently of any particular response. In early studies of punishment, no attention was ordinarily paid to this issue. But in more recent studies (*e.g.*, Bolles, Uhl, Wolf, and Chase, 1975; Church, Wooten, and Matthews, 1970; Matthews, McHugh, and Carr, 1974) attempts have been made to separate the two effects. These studies indicate that the suppressive effects specific to response-dependent shock take a long time to appear and require a substantial number of shocks. Note further that the same reactions (*e.g.*, freezing) that mimic the effects of punishment also tend to suppress the recorded response in studies of avoidance. In that context, however, they tend to slow down, rather than to speed up, the change in the subject's behavior. Finally, there is no reason why two classes of behavior as diverse as the recorded response and "other behavior" should be expected to condition at the same rate. They are quite different in operant level, for instance, as well as in topography and perhaps in terms of conditionability. Although the two classes of behavior may in-

terchange their roles as we shift from one type of training to the other, this does not imply that they must be quantitatively equivalent.

Rachlin and Herrnstein's (1969) final argument was that although later members are suppressed, the avoidance hypothesis requires that there be no change in the frequency with which the chain of responses leading to the punishing stimulus is initiated. That is, the first member of the chain should not be suppressed. In a brief parenthetical remark, they indicated that they believed this implication to be contrary to fact. I might comment, first, that it is not easy to specify where a chain of responses begins; it may very well originate with a nonrecorded response. Therefore, suppression of the first recorded response would not necessarily be evidence against the avoidance hypothesis. Second, at the empirical level, evidence bearing on precisely this point has subsequently been published by Millenson and Macmillan (1975). These authors trained their rats to perform a simple chain of responses, which was relatively easy to record. To receive a pellet of food, the rat was required to press the lever down, hold it down for at least 10 sec, and then release it. If the time requirement had been met, a pellet was delivered at the moment of release. After giving their rats 20 sessions of training on this procedure, to establish a baseline level of performance, Millenson and Macmillan added a punishment contingency, electric shock as a consequence of each completion of the chain. They found that the addition of the punishment reduced the mean time between each depression of the lever and its subsequent release and correspondingly reduced the proportion of chains that were completed.

In other words, punishment produced an increase in the frequency with which the rats interrupted the chain by releasing the lever before the 10-sec requirement had been met. Millenson and Macmillan have provided an empirical record of the avoidance behavior that intrudes within the punished sequence and cancels out the animal's progress toward its completion. Note also another finding: the mean time between the release of the lever and a subsequent depression did not increase. Therefore, the objection raised by Rachlin and Herrnstein (1969) proved to be unfounded in this instance, at least, since the first component of the chain occurred as promptly as ever;

the locus of the interference was found later in the chain.

The pattern of events described above is by no means an inevitable consequence of any and all reductions in the strength of the recorded behavior. Other patterns are possible, as indicated by the effects of satiation and of extinction. Extinction reduced both parts of the chain. It took the rat longer to make another press—initiating the chain, following a release of the lever—and the lever was not held down as long as before, once it had been pressed. Satiation, accomplished by prefeeding the animals with 12 g of laboratory chow, produced a pattern of events that was just the reverse of that produced by punishment. The mean time between the release of the lever and a subsequent press increased, indicating a reduction in the tendency to initiate the chain. But satiation also produced an increase in the proportion of times when the chain, once initiated, was successfully completed. Each of these procedures—punishment, satiation, and extinction—produced a reduction in the overall rate of occurrence of the chain, but each accomplished this result in a different manner.

Incidentally, the finding of a difference between the pattern for punishment and the pattern for extinction raises a problem for Leitenberg's contention (1965, pp. 433-435) that the two processes are basically the same.

It could be argued, perhaps, that the chain of behavior studied by Millenson and Macmillan (1975) was a special case, not typical of the ordinary bar-pressing experiment and therefore not representative in its results. In a subsequent experiment, however, Carman (*in press*) used a more typical sequence to study the selective effect of punishment on successive components of the chain. Under his procedure, no hold was required. Carman simply recorded touching the lever as his first component and pressing it as his second. In the beginning, pressing was reinforced with pellets of food on a variable-interval schedule. Later, a multiple schedule was used. During 3-min periods when a light was on, pressing the lever was followed by shock on a variable-interval 60-sec schedule with a limited hold of 15 sec. Both touches and presses were reduced in absolute frequency in the presence of the light. The number of touches increased in relative frequency, however, compared to the number of presses. That is, touching the bar was followed less often by

a subsequent press and more often by a retreat—another form of abortive responding. Carman's data, like those obtained by Millenson and Macmillan, showed frequent reversals in the animal's progress toward completion of the chain when the sequence was punished as well as reinforced. Thus, a certain amount of empirical data is now available supporting the avoidance theory of punishment, and after two decades it still appears quite viable.

REFERENCES

- Anger, D. The role of temporal discriminations in the reinforcement of Sidman avoidance behavior. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 477-506.
- 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.
- Bolles, R. C. and McGillis, D. B. The non-operant nature of the bar-press escape response. *Psychonomic Science*, 1968, 11, 261-262.
- Bolles, R. C. and Popp, R. J., Jr. Parameters affecting the acquisition of Sidman avoidance. *Journal of the Experimental Analysis of Behavior*, 1964, 7, 315-321.
- Bolles, R. C., Uhl, C. N., Wolfe, M., and Chase, P. B. Stimulus learning versus response learning in a discriminated punishment situation. *Learning and Motivation*, 1975, 6, 439-447.
- Campbell, S. L. Lever holding and behavior sequences in shock-escape. *Journal of Comparative and Physiological Psychology*, 1962, 55, 1047-1053.
- Carman, J. B. Effect of discriminative punishment on a simple response chain. *Animal Learning and Behavior*, (in press).
- Chung, S. Effects of delayed reinforcement in a concurrent situation. *Journal of the Experimental Analysis of Behavior*, 1965, 8, 439-444.
- Church, R. M., Wooten, C. L., and Matthews, T. J. Contingency between response and an aversive event in the rat. *Journal of Comparative and Physiological Psychology*, 1970, 72, 476-485.
- Davis, H. and Burton, J. The measurement of response force during a lever-press shock-escape procedure in rats. *Journal of the Experimental Analysis of Behavior*, 1974, 22, 433-440.
- Davis, H., Hirschorn, P., and Hurwitz, H. M. B. Lever holding behavior during a leverlift shock escape procedure. *Animal Learning and Behavior*, 1973, 1, 215-218.
- Dinsmoor, J. A. Punishment: I. The avoidance hypothesis. *Psychological Review*, 1954, 61, 34-46.
- Dinsmoor, J. A. Punishment: II. An interpretation of empirical findings. *Psychological Review*, 1955, 62, 96-105.
- Dinsmoor, J. A. Escape from shock as a conditioning technique. In M. R. Jones (Ed), *Miami Symposium on the Prediction of Behavior*, 1967: *Aversive Stimulation*. Coral Gables, Florida: University of Miami Press, 1968. Pp. 33-75.
- Dinsmoor, J. A. and Clayton, M. H. Chaining and secondary reinforcement based on escape from shock. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 75-80.
- Dinsmoor, J. A. and Clayton, M. H. A conditioned reinforcer maintained by temporal association with the termination of shock. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 547-552.
- Dinsmoor, J. A., Hughes, L. H., and Matsuoka, Y. Escape-from-shock training in a free-response situation. *American Journal of Psychology*, 1958, 71, 325-337.
- Dinsmoor, J. A., Matsuoka, Y., and Winograd, E. Bar-holding as a preparatory response in escape-from-shock training. *Journal of Comparative and Physiological Psychology*, 1958, 51, 637-639.
- Dinsmoor, J. A. and Sears, G. W. Control of avoidance by a response-produced stimulus. *Learning and Motivation*, 1973, 4, 284-293.
- Estes, W. K. Outline of a theory of punishment. In B. A. Campbell and R. M. Church (Eds), *Punishment and aversive behavior*. New York: Appleton-Century-Crofts, 1969. Pp. 57-82.
- Flint, G. A. Bar orientation in operant escape training. *Behavior Research Methods and Instrumentation*, 1969, 1, 231-232.
- Herrnstein, R. J. Method and theory in the study of avoidance. *Psychological Review*, 1969, 76, 49-69.
- Herrnstein, R. J. and Hineline, P. N. Negative reinforcement as shock-frequency reduction. *Journal of the Experimental Analysis of Behavior*, 1966, 9, 421-430.
- Hineline, P. N. Negative reinforcement without shock reduction. *Journal of the Experimental Analysis of Behavior*, 1970, 14, 259-268.
- Keehn, J. D. Is bar-holding with negative reinforcement preparatory or perseverative? *Journal of the Experimental Analysis of Behavior*, 1967, 10, 461-465.
- Keller, F. S. Light aversion in the white rat. *Psychological Record*, 1941, 4, 235-250.
- Leitenberg, H. Is time-out from positive reinforcement an aversive event? A review of the experimental evidence. *Psychological Bulletin*, 1965, 64, 428-441.
- LoLordo, V. M. Positive conditioned reinforcement from aversive situations. *Psychological Bulletin*, 1969, 72, 193-203.
- Marriott, A. S. and Russell, I. S. Lever-holding and the operant nature of escape responding. *Psychonomic Science*, 1969, 15, 1-2.
- Matthews, T. J., McHugh, T. G., and Carr, L. C. Pavlovian and instrumental determinants of response suppression in the pigeon. *Journal of Comparative and Physiological Psychology*, 1974, 87, 500-506.
- Millenson, J. R. and Macmillan, A. St. C. Abortive responding during punishment of bar holding. *Learning and Motivation*, 1975, 6, 279-288.
- Mowrer, O. H. An experimental analogue of "regression", with incidental observations on "reaction formation". *Journal of Abnormal and Social Psychology*, 1940, 35, 56-87. (Republished: *Learning theory and personality dynamics*. New York: Ronald Press, 1950. Pp. 361-389.)
- Rachlin, H. and Herrnstein, R. J. Hedonism revisited: On the negative law of effect. In B. A. Campbell and R. M. Church (Eds), *Punishment and aversive*

- behavior*. New York: Appleton-Century-Crofts, 1969. Pp. 83-109.
- Renner, K. E. Delay of reinforcement: A historical review. *Psychological Bulletin*, 1964, **61**, 341-361.
- Sidman, M. Reduction of shock frequency as reinforcement for avoidance behavior. *Journal of the Experimental Analysis of Behavior*, 1962, **5**, 247-257.
- Sidman, M. Avoidance behavior. In W. K. Honig (Ed), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 448-498.

Received 3 June 1976.

(Final Acceptance 6 February 1977.)