

*"AVOIDANCE" IN BEHAVIOR THEORY*<sup>1,2</sup>

W. N. SCHOENFELD

QUEENS COLLEGE, THE CITY UNIVERSITY OF NEW YORK  
AND CORNELL UNIVERSITY MEDICAL COLLEGE

I. Some years ago, I tried to formulate an account of the acquisition and maintenance of operant avoidance behavior. In that paper, I could draw upon all that was known at the time about avoidance, since the number of available studies was small. The literature has grown enormously since then, and it would no longer be practicable today to review it all in any meaningful way in a single paper. Actually, I do not think that is what is most needed in this area. In any case, I became dissatisfied with my own theoretical treatment of avoidance, and now it seems to me that the outline of a more adequate formulation has become clearer in the light of the remarkable experimental advances made since my first writing. The relevant information is in the current literature, and I am here only recognizing its implications for a new theoretical approach to the avoidance problem. This new approach is essentially a restatement leading not to the "solution" of the problem, but to its dissolution; framed in proper terms, the phenomenon of avoidance simply disappears as a separate category of behavior, or at least becomes a derivable spin-off from a more general treatment of behavior. It seems to me that because "aversive" responding is not different in either principle or process from the acquisition and maintenance of other types of responding, it is involved in the broader question of "reinforcement" in general.

II. Behavioral science has drawn many words from the layman's vocabulary, even though such words are seldom good ways to

describe the response material of conditioning theory. They reflect social attitudes and beliefs about behavior; they are defined by social criteria; they group acts by their social outcomes; and, they are almost always mistaken in what they accept as their behavioral referents. More often than not, they only obstruct a proper analysis of behavior by clouding issues, generating profitless debates, and side-tracking research. These words range from seemingly simple responses to broad action patterns: memory, lying and stealing, thinking, aggression-submission, cooperation, courtship, motivation, stress. Even as seemingly specific a word as "swim" can mislead: if it is asked how a rat can know how to swim the first time it is put in water, the reply might be that it is not swimming at all, just doing what it knows, walking, but that "swimming" is how the rat walks in a medium like water, and that a continuous series of modifications could be traced in its walking through media and on surfaces of different densities and firmnesses. Some words like "pain" have been extended quite far in anthropomorphic application to animals. In fact, of course, "pain" is a socially taught verbal response reporting the occurrence of certain stimuli and preceding responses; animals do not have "pain" even when responding under those circumstances that we generalize with what we would ourselves report socially as causing "pain" to us. In other cultures, and in varying situations in our own culture, we might report those same stimuli and responses quite differently, whereupon animals would have those other feelings imputed to them.

Other words from the lay lexicon refer to a restricted response range that is conventionally set off for special naming within a broader continuum of behavior effects. Aversive behavior words like "avoidance" and "escape" are of this sort. As a lay word, "avoid-

<sup>1</sup>Dedicated to B. F. Skinner in his sixty-fifth year.

<sup>2</sup>This work was supported in part by Public Health Service Research Grant MH 12964 from the National Institute of Mental Health. Reprints may be obtained from the author, Dept. of Psychology, Queens College of the City University of New York, Flushing, N.Y. 11367.

ance" has a kind of commonsense meaning of doing something to prevent a hurtful experience. In a general way, "avoidance" like this does occur, but it is careless to permit the lay sense to define the scientific problem. To import the term into a science of behavior as defining a problem for study is to import along with it the same problems the layman eventually encounters in understanding and using it. An example may be found in the *t*-system wherein a recycling time interval, *T*, has some part,  $t^D$ , in which a response may have an experimentally specified environmental consequence. If *T* is very short, and if  $t^D$  is made coextensive with a stimulus like strong electric shock and is given an appropriate value, then the shock itself will appear continuous to a subject, and an "avoidance" response, if allowed to produce a timeout, is identical to an "escape" response. Avoidance and escape are thus products of the parameter of *T* length, and each term will be used by the lay community for that part of the *T* range where it seems acceptable to tradition and valid on anthropomorphic-empathic grounds.

Of course, scientific study historically begins with words from the lay vocabulary, even though they must be refined and perhaps rejected as knowledge accumulates, and as discoveries are made that permit sharper definitions of phenomena, of their determining conditions, and so on. This has happened with the word "avoidance". We can often obtain responding, or produce response changes, of various sorts by applying known experimental operations to the stimuli and responses composing the behavior stream. If we wish, we can also by our procedures and apparatuses permit certain environmental outcomes to occur as consequences of an ordained response, and if the resulting behavioral changes match a lay term in our vocabulary, such as "avoidance", we apply it. But then, as research into the matter develops, new or "strange" observations crop up that puzzle us, such as the effects of "free" stimulations, or the parametric kinships among conditioning procedures once thought of as unrelated. At present, a set of behavioral phenomena are familiar to us, as they are to laymen, under the guise of "avoidance". It is within this rubric that we at one time formulated the "problem of avoidance", and were led to ask the question of what the reinforcement for avoidance responding might be. But the question in this form was an en-

trapment. We are beginning to see that set of phenomena as part of a larger but continuous range along which appear related behavioral phenomena that impress us as strange just yet. Once these are linked to the familiar, once they are seen as parts of the same parametric continua of experimental operations and variables, the strangeness disappears. And with the strangeness, the problem also disappears. We have been carrying the avoidance problem down this road for many years, traveling what seemed a long distance toward this goal, but the conceptual distance, as science has often found, is quite short.

III. The problem of operant avoidance behavior has seemed to be identifying the reinforcement for it. The "avoidance paradox" arose from the definition of avoidance as behavior leading to non-occurrence of certain stimuli, with the consequent need to rationalize stimulus non-occurrence as reinforcing. This difficulty originally suggested to me that the very existence of avoidance as a behavior phenomenon might be doubted, but—amusingly enough in view of my present belief—I dismissed those misgivings and concentrated on a possible solution. It was organized along a few main lines. The emphasis was placed on the cued avoidance case, although the place of non-cued avoidance was plain. Avoidance was reached as the last in a progression of paradigms starting with experimental anxiety (more currently called CER or "conditioned suppression"), followed by escape, and then avoidance. The paradigms were based upon a sequence of two stimuli  $S_1 \rightarrow S_2$  (where  $S_1$  is the traditional "neutral" stimulus or cue,  $S^N$ , and  $S_2$  the "aversive" stimulus or negative reinforcer,  $S^{-R}$ ), and upon the experimentally defined role of the response (*R*) to these two stimuli. The progressive paradigms were treated as operationally discontinuous, but the behavioral products were taken as continuous, and this continuity provided the basis for identifying the reinforcement for avoidance. (My present approach to avoidance not only removes the operational discontinuity but reaffirms the behavioral continuity and extends it to include all behavior, not just anxiety and escape.) The reinforcing function was assigned to proprioceptive-kinesthetic stimuli, and the avoidance response was the behavior that survived after all other responses were suppressed by the "punishing" stimulus.

At least two weaknesses in that explanation were caused by approaching avoidance through the question of what its reinforcement is, and they did not escape the notice of later theorists. They were tolerated at the time because no alternative seemed available while the question was in that form. One was that the actual role of proprio-kinesthetic stimuli is not ascertainable; although the class of stimulus called proprioceptive is known to exist, such stimuli are not yet accessible for measurement or experimental manipulation. The second was that the idea of suppressing all behavior except the avoidance response left open the problem of specifying "all behavior" without invoking infinity. Even if the behavior stream were broken into a finite number of blocks of responses within which all responses were affected in some measure by reinforcement or extinction of any one response member, it still is not prudent to rely on elimination of all blocks save the one containing the avoidance response, partly because no experimental information is available on the sizes and boundaries of the "other response" classes, and partly because the speed with which avoidance is often learned has implications for the speed with which weakening of all the other blocks would have to be accomplished.

But now there are more persuasive reasons for relinquishing that early explanatory approach. In the intervening years, we have learned so much more about avoidance behavior and its properties that it is no longer necessary to try to reach avoidance responding from supposedly simpler paradigms like anxiety and escape. Indeed, it might be more useful to work from avoidance to the other forms of "aversive" behavior, since some things we have learned about acquisition and reinforcement from the avoidance case can be transferred to the other cases. For example, avoidance responding without reduction in frequency of aversive stimulus falls into the purview of experimental anxiety, and broadens the latter category both theoretically and in experimental prediction of behavioral possibilities extending beyond the simple response suppression view of the CER paradigm.

The intervening years have also provided the opportunity to learn more about reinforcement schedules in general, including those involving "positive" reinforcers ( $S^+R$ ), and thereby to observe and conceptualize more

broadly—as in the  $t-r$  systems—the kinships and overlaps among schedules once thought to be different in kind. We have before us today attempts from several researchers to systematize schedules along new theoretical lines, and to deal with the stimulus and response contents of the behavior stream in new ways.

All these factors are working toward a re-assessment of the avoidance problem. New experimental information exerts pressure on older formulations, and, conjoined with the latter's weaknesses, creates the need for new theoretical efforts.

IV. There are a number of considerations, both specific and general, that inform a recasting of the avoidance problem. The following seem just now to be the more important ones.

(1) Preliminary considerations alone show the difficulty of even defining avoidance behavior, and suggest that avoidance need not be regarded as a separate behavioral problem. The individual response itself is not an issue, since any response may be selected from the behavior stream to do the "avoiding" and be measured. The strategy that was adopted historically was to define avoidance in terms of outcome; that is, in terms of what the response accomplishes. For the free operant non-cued avoidance case ("Sidman avoidance"), this outcome is usually defined as a reduction in  $S_2$  frequency; in a trial-by-trial procedure, where each trial is an opportunity for  $S_2$  to occur, the outcome is an increase in inter- $S_2$  time. There are two experimental findings, however, that undermine any such reliance on outcome as a defining criterion of avoidance. Indeed, these findings are crucial for any theory of avoidance, and we will call upon them several times. The first, already well validated by experiment, is that behavior acquired under "avoidance" procedures can be maintained indefinitely by "free"  $S_2$ s delivered non-contingently upon any stated response and without reduction in frequency. In fact, there are schedules under which responding is maintained even though shock frequency is increased. The second, as yet reported only by isolated investigators but probably soon to be firmly established, is that a pre-selected response, which would be the recognized "avoidance" response if acquired in an experiment where reduction of shock frequency were its "contingent" outcome, can also be acquired *ab initio* under certain conditions by the sim-

ple procedure of intruding "free"  $S_2$ s into the behavior stream. Moreover, any avoidance schedule can involve, as a parameter, proportions of "free" to "contingent"  $S_2$ s varying from zero to unity. Facing these facts, it cannot be correct to define "avoidance" in terms of reduction in number of  $S_2$ s.

This conclusion is supported by several other thoughts. To look to reduction in  $S_2$  frequency as the cause of avoidance responding is once again to appeal to the non-occurrence of a stimulus as reinforcing. The absurdity of this appeal is evident because at any moment an infinite number of stimuli are not occurring, and we would therefore have to assume that all behavior is avoidant. There is also the unacceptable implication that avoidance behavior is unique because no parallel case can be drawn from schedules involving "positive reinforcement"; it is not the *reduction* in frequency of food delivery on partial schedules which achieves a relatively greater durability of a response as compared to the 100% or regular reinforcement schedule. Furthermore, avoidance regarded as an outcome would have "food-procurement" as a parallel in positive reinforcement (rather than, say, bar pressing) as the "response" being manipulated. Even if one wished to deal with outcomes in this way, there is the problem of how outcome itself is to be defined. For example, if the "timeout from  $S_2$ " in the free operant case involves the cancellation of many scheduled  $S_2$ s, shall an "avoidance" be tallied for each forestalled  $S_2$ , or is each timeout or train of timeouts to be taken as a single "avoidance"? Or, if the outcome schedule depends upon a number of responses ("ratio" avoidance responding) rather than a single response, how shall "avoidances" be tallied? Still other questions flow from the measurement of the avoidance response. A traditional one for the free operant non-cued case is that of response rate, while a traditional one for trial procedures is that of latency. We have long realized, of course, that the proper gauge of conditioning is not response rate alone, but rather the degree of control over the response exercised by the experimental variable. Thus, DRL procedures that lower response rate demonstrate control as well as procedures which raise rate, and both DRL escape and DRL avoidance responding may be achieved experimentally. Avoidance is not response-specific; it needs to be analyzed in terms of the

actual stimulus-response events of which it is the product; that is, in terms of what the organism is *doing* rather than of the outcomes of his behavior.

(2) Avoidance schedules are conventionally classed as "aversive", leading us to ask first what aversive schedules are and then how they might resemble positive reinforcement schedules.

An aversive schedule involves  $S^R$  and therefore generates "anxiety" and "escape" behavior as well as "avoidance". The difficulties of defining either anxiety or escape are like those encountered in defining avoidance, and need not be rehearsed. Similarly, to define an aversive schedule we must first define the terms "aversive stimulus" and "reinforcer". The effort to define  $S^R$  in terms of the behavior it can support does not advance our understanding; even if possible circularities are overcome, the behaviors themselves are suspect. The fact is that we are involved here with parametric aspects of certain experimental procedures, and with parameters of response measurement. This is a decisive point, and we shall return to it later.

A reinforcement schedule is an elaboration of the single and primitive operation of intruding a stimulus into the behavior stream. It is desirable to define a stimulus in physical terms; that is, as a physical event. The effects of any intruded stimulus on responding (whether "reinforcing" or "emotionalizing", *etc.*) depend on what behavior changes we are willing to accept as classes of effect, and what we are willing to include as part of any effect, as well as on a number of other variables such as the state of the organism, or the response composition of the behavior stream at the moment of  $S$  intrusion. Among these latter variables is one fundamental to any discussion of aversive schedules; namely, stimulus intensity.

For any stimulus, the intensity continuum ranges from zero to indefinitely large magnitudes. This continuum may be divided into sub-ranges, although these may have fluctuating and overlapping boundaries. Progressing from zero, they are: imperceptible intensities, intensities that make up the usual psychophysical range, "positively reinforcing" intensities, "aversive" intensities, and finally, fatal intensities. This is true of simple stimuli like tones, lights, or touches as well as complex stimuli or "objects" like food pellets. Thus,

an "intense" (heavy) pellet would be aversive if dropped onto a rat instead of to one side; on the other hand, an electric shock would *not* be aversive if delivered to one side instead of to the animal. (But if the side-delivered shock were eventually encountered, and each encounter were cued, the cues would secondarily function as or "be" the shock, just as the sight, smell, *etc.*, of the pellet "is" the pellet as "object".)

In short, no stimulus is aversive in itself without regard to the parameter of its intensity and the circumstances of its delivery. Joining this observation to the difficulty of defining aversive responding, we must conclude that "aversive schedules" are only arbitrarily, even conventionally and nominally, classified under that title. This conclusion holds equally for an aversive schedule which the animal seeks to escape in favor of some other "preferred" schedule. In the latter sense, of course, any circumstance affecting the animal may be labelled "aversive" since *some* other circumstance will always be preferred.

The thoroughly parametric nature of aversive conditioning operations (in common with all conditioning operations) cannot be exhaustively considered here, but a few of the many examples may be cited: (a) In free operant non-cued "Sidman avoidance", the  $S^R$  terminating an  $S-S$  interval might be different (in type, or intensity, *etc.*) from that terminating an  $R-S$  interval, and the probability of occurrence of each of those stimuli (the mean lengths of the  $S-S$  and  $R-S$  intervals) may be systematically varied either together or independently. (b) The "Sidman avoidance" paradigm is the same as that of the delayed punishment procedure which employs a reset of delay for intervening responses. While the delayed punishment procedure involves an initial response, "Sidman avoidance" does also, since with zero operant level no animal would ever learn that avoidance response (although the first shocks might re-make the operant level repertory to bring out at least one of the specified responses.) (c) Distinctions between cued and non-cued avoidance procedures are based on the temporal spacing of the stimuli in the  $S_1 \rightarrow S_2$  sequence. (d) In cued escape paradigms, if the interval before the next  $S_1 \rightarrow S_2$  trial is short, and if the probability of  $S_2$  occurrence is set between unity and zero, then the "escape"

response appears like an "avoidance" response. One may multiply such cases of parametric variation leading from one supposed type of aversive behavior to another.

Positive reinforcement is also of a thoroughly parametric nature and is on the same continua as aversive conditioning. At least some of the many similarities between the two types of schedule can be indicated, although an exhaustive list cannot be given, probably because at bottom the two schedules are not really different types at all. In both schedules an intruding stimulus is used, either as reinforcer or as cue, its presence alternating with its absence. In both schedules, all stimulus operations involve stimulus "presentation" and stimulus "removal", with both always simultaneously present as stimulus "change". In both schedules, the definition of the measured response class is arbitrary, and the exclusion from that class of closely related responses probably makes even the 100% reinforcement schedule effectively an intermittent schedule. In both schedules, because the behavioral stream has no "holes" in it or places empty of responses, behavioral change of any sort always means that some responses have gone up in strength, while others have gone down. This is, of course, the basis of the familiar theories of extinction known as "counter-conditioning" or "interference". Both schedules share characteristic problems of response acquisition and maintenance: in acquisition under positive reinforcement schedules, animals would be lost through failure to condition as they often are under aversive schedules if prior discriminative training were not given and prior behavior chains were not built up; the maintenance of aversive responding by "free" aversive stimulations is paralleled by maintenance of responding through "free positive" (often called "superstitious") reinforcements. In both schedules, either the same stimulus or an intensity variant of it may serve as the intruding stimulus or "reinforcing" stimulus event, and the same delivery schedule may be employed; a "positive reinforcer" may be an electric shock of appropriately weak intensity, and may be delivered on a "Sidman avoidance" schedule. (Schedules like these have actually been studied experimentally, although under the headings of "delayed reinforcement" and "reinforcement of not-R".)

Considerations of this sort inevitably weaken the assertion of a separate identity for

aversive schedules, including avoidance, since the differences between any two schedules can be arranged parametrically with respect to the details of independent variables and behavior measurements.

V. In treating the problem of avoidance conditioning, the first question to ask is not the traditional one of what the reinforcement for such responding is, but how avoidance responding is acquired—that is, what the determining conditions or variables are—and how it is maintained. We have tried to show that some familiar assumptions about conditioning an avoidance response are false. Given that a response has been arbitrarily chosen by the experimenter as the “avoiding” response, that response can be controlled by operations other than the frequency of aversive stimulation, or indeed of any stimulus contingent upon the response. (By “contingent” I mean a relationship whereby the temporal distribution of stimulus deliveries is determined by the temporal distribution of responses.) Avoidance conditioning commonly involves some arrangement by the experimenter of how, and under what schedule, certain stimuli are to be applied to an organism, and then simply projecting the organism into that arrangement. This rather insouciant casting of one’s behavioral bread upon scientific waters may occur in some measure in all schedules but is prominent in avoidance procedures including cued or non-cued avoidance training, free operant or trial-by-trial procedures, “Sidman avoidance”, or  $t-r$  system aversive schedules. This way of proceeding does not look at the experimental variables as they actually impact the animal, and it does not insure that the appropriate variables are brought to bear at any given instant in the training. It is not like the “shaping” whereby a pre-specified avoidance response is reached. As a result, of course, many animals fail to condition and are lost to the study. The failures are perhaps less instructive than the successes; why do some animals condition or shape up, and not others? Clearly, if we knew the conditions that produced the successful learning we could salvage every subject. Under these circumstances, we should not be surprised that a thoughtful researcher, by simply exposing an animal to a sequence of non-contingent aversive stimulations, has successfully produced acquisition of what under parametric variations of the procedure we

would call the “avoidance” response, and that he has therefore concluded that a response does not need a history of avoidance training for “free” aversive stimulations to be effective.

With regard to supporting an avoidance response, the familiarity of “free”  $S_2$ s as a maintaining schedule makes us readier to accept the possibility that with proper values of the relevant parameters the “free” procedure could maintain responding indefinitely. But it would be gratuitous for theory to endorse a distinction between the mechanisms of response acquisition and response maintenance, especially since we are beginning also to see that no such distinction need be made on the experimental level.

In addition to the parametric and classificatory aspects, our understanding of avoidance acquisition and maintenance must stand upon the temporal relations and the associated probability relations between the experimentally manipulated stimuli and those responses chosen from the behavior stream for observation. This conclusion is neither as onerous nor as vague as it may appear at first; in fact, it has always guided our research on behavior under “positive reinforcement”. Now it opens the way for operationally and theoretically linking aversive conditioning to the general theory and practice of reinforcement schedules. Efforts along this line have already been made by the  $t-r$  systems of schedule classification within which some currently popular aversive and positive schedules prove to be special cases. In addition to this linkage, a full treatment of aversive conditioning will surely involve new considerations about the behavior stream; there are already signs that this theme is re-emerging as a concern of behavior theory.

VI. The burden of this paper can now be summarized by two statements: (1) Any avoidance response, whether specified in advance or not, may be conditioned or controlled by an appropriate selection of values of the three parameters, among others, of operant level of the response, stimulus frequency, and stimulus intensity. (2) Since the same variables apply to “positive” and “negative” reinforcement schedules, separate or unique treatment is not required for either within a general theory of behavior.