JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1969, 12, 475-480 NUMBER 3 (MAY)

SIGNAL DETECTION THEORY AND OPERANT BEHAVIOR
A Review of David M. Green and John A. Swets' Signal Detection Theory and Psychophysics.[1]

JOHN A. NEVIN[2]

COLUMBIA UNIVERSITY

A review of *Signal detection theory and psychophysics* in the *Journal of the Experimental Analysis of Behavior* is justified by the growing interest in signal detection among many psychologists concerned with operant behavior. The interactions of wavelength discrimination and reinforcement (Boneau, Holland, and Baker, 1965) have been examined within a detection-theory framework (Nevin, 1965), and a full-blown decision theory of animal discrimination performance has been published (Boneau and Cole, 1967). Blough (1967) has used signal-detection analysis to present wavelength generalization gradients in a novel way, and Rilling and McDiarmid (1965) and Stubbs (1968) have extended this analysis to the study of response-produced and temporal stimuli. Most recently, Blough (1969) has considered signal detection theory as an alternative to traditional formulations of stimulus generalization. A knowledge of signal detection theory is therefore important for its relevance to operant research, as well as for its central position in sensory psychology.

A number of psychologists have noted some interesting similarities between research in signal detection and operant conditioning. As Goldiamond (1962) pointed out, workers in both areas emphasize rigorous control of the experimental environment, with automatic scheduling of stimuli and recording of responses. Typical experiments involve very few subjects, whose performance is studied intensively for many sessions. For each subject, measures of asymptotic performance are obtained after prolonged training on several conditions. In laboratory terminology, both areas employ steady-state, single-organism methodology.

The distinctive feature of signal-detection research, as opposed to classical psychophysics, is the emphasis on immediate feedback for each response in the form of rewards and penalties. The standard "yes-no" signal-detection experiment involves two stimulus classes: signal superimposed on noise, and noise alone; and two response classes: "yes" (there was a signal) and "no" (there was no signal). There are four stimulus-response events to be considered here: hits and correct rejections lead to rewards, while false alarms and misses lead to penalties. The explicit use of these reinforcement and punishment contingencies provide direct contact with the analysis of operant behavior.

In the language of the analysis of behavior, the yes-no experiment is a multiple schedule. In the presence of signals, positive reinforcement for one response (yes) is scheduled concurrently with punishment for the alternative response (no). These contingencies alternate irregularly with the reverse—punishment for yes and reinforcement for no—on noise trials. One variable of the yes-no experiment is the *a priori* probability of a signal, which determines the relative frequency of exposure to these contingencies. Another is the amount of reinforcement or punishment, which may be scheduled independently for each contingency. A third variable, presumably of most interest to psychophysics, is signal strength. From the standpoint of behavior analysis, this variable determines whether the schedule is multiple, as when the stimuli are readily discriminable, or mixed, as in the case of zero signal strength. When signal strength is zero, the yes-no detection experiment reduces to the so-called "binary guessing" or "probability learning" experiment, which has been discussed a number of times within this journal in relation to concurrent schedules of reinforcement (*e.g.*, Herrnstein, 1964; Shimp, 1966).

[1]New York: Wiley, 1966. Pp. xi +455, $13.95

[2]Reprints may be obtained from J. A. Nevin, Dept. of Psychology, Schermerhorn Hall, Columbia University, New York, N.Y. 10027.

This brief review of experimental procedures suggests ample opportunity for methodological integration of signal detection and operant conditioning. I would now like to examine some detection data in a fashion suggested by recent work on operant behavior. This exercise is intended more to indicate possible directions for future work than to provide a definitive integration of these research areas. Consider the data from a yes-no experiment described by Green and Swets (pp. 88-90) which was designed to trace out an ROC curve, the well-known receiver-operating-characteristic function relating the probability of hits to the probability of false alarms, with signal strength constant. The *a priori* probability of signal, and the payoffs and costs consequent upon responding, were varied separately for a single subject. The subject's unconditional probability of saying "yes", that is, the proportion of trials on which he said "yes" regardless of the stimulus presentation, was related to the relative frequency or amount of reinforcement scheduled for saying "yes." For example, if the payoffs were equal, and the *a priori* probability of signal was 0.70, reinforcement was available for "yes" on 70 trials out of 100, and for "no" on 30. The scheduled relative frequency of reinforcement for "yes" was therefore 0.70. If the *a priori* probability of signal was 0.50, but correct detections were reinforced with two points while correct rejections produced only one point, the relative amount of reinforcement for "yes" was $2/2 + 1 = 0.67$. (The punishment contingencies were ignored.) As shown in the upper panel of Fig. 1, the proportion of trials on which "yes" occurred was about equal to the relative frequency or amount of reinforcement scheduled for "yes". Similar matching functions have appeared repeatedly in studies of concurrent schedules (see, for example, Catania, 1963*a*; Herrnstein, 1961; Reynolds, 1963), and the equivalence of relative frequency and amount of reinforcement in concurrent schedules is well established (Catania, 1963*b*; Rachlin and Baum, 1969; Shimp, 1968).

The matching relation shown here must depend on signal strength. If signals are intense, the subject will say "yes" whenever the signal is presented, and never otherwise, regardless of the relative payoffs for hits and correct rejections. If signal intensity is zero, the subject will nearly always make the response having the larger probability or amount of reinforcement. In like fashion, the matching relation obtained with concurrent reinforcement schedules must depend on the contingencies of reinforcement (see discussion by Catania, 1966). Speculatively, then, it may be fruitful to regard signal intensity and the contingencies of reinforcement as functionally similar determinants of choice behavior.

A further similarity is suggested by another way of examining the detection data. As the reinforcement for "yes" decreased relative to total available reinforcement, the proportion of correct "yes" responses increased. This relation is shown for Green and Swets' data in the lower panel of Fig. 1. A similar relation has been reported by Pliskoff, Shull, and Gollub (1968) in a study involving multiple and concurrent variable-interval schedules. They arranged a multiple schedule on one key, with 40 reinforcements per hour when the key was red, alternating with green which was correlated with 10 reinforcements per hour. The relative rate of responding in the presence of red increased as the frequency of concurrent reinforcement on a second key increased. If the multiple key is equated with the "yes" response, red with signal, and green with noise, the function obtained by Pliskoff *et al.* may also be plotted in the lower panel of Fig. 1. On the face of it, the identification of the component stimuli of a multiple schedule with signal and noise in a detection experiment is implausible. The component stimuli control different response rates by virtue of their correlated reinforcement schedules rather than their physical difference. However, the similarities in results encourage speculation on the functional similarity between signal intensity in detection research and reinforcement scheduling in research on operant behavior.

The remainder of this review will attempt to make explicit some features of detection theory which may be worth pondering in any future theorizing about operant behavior.

A persistent problem in sensory psychology is the dependence of measured thresholds on procedural variables. A variety of direct measures of sensory sensitivity are easy to derive from the response data, but it has never been obvious how they should be treated to test for

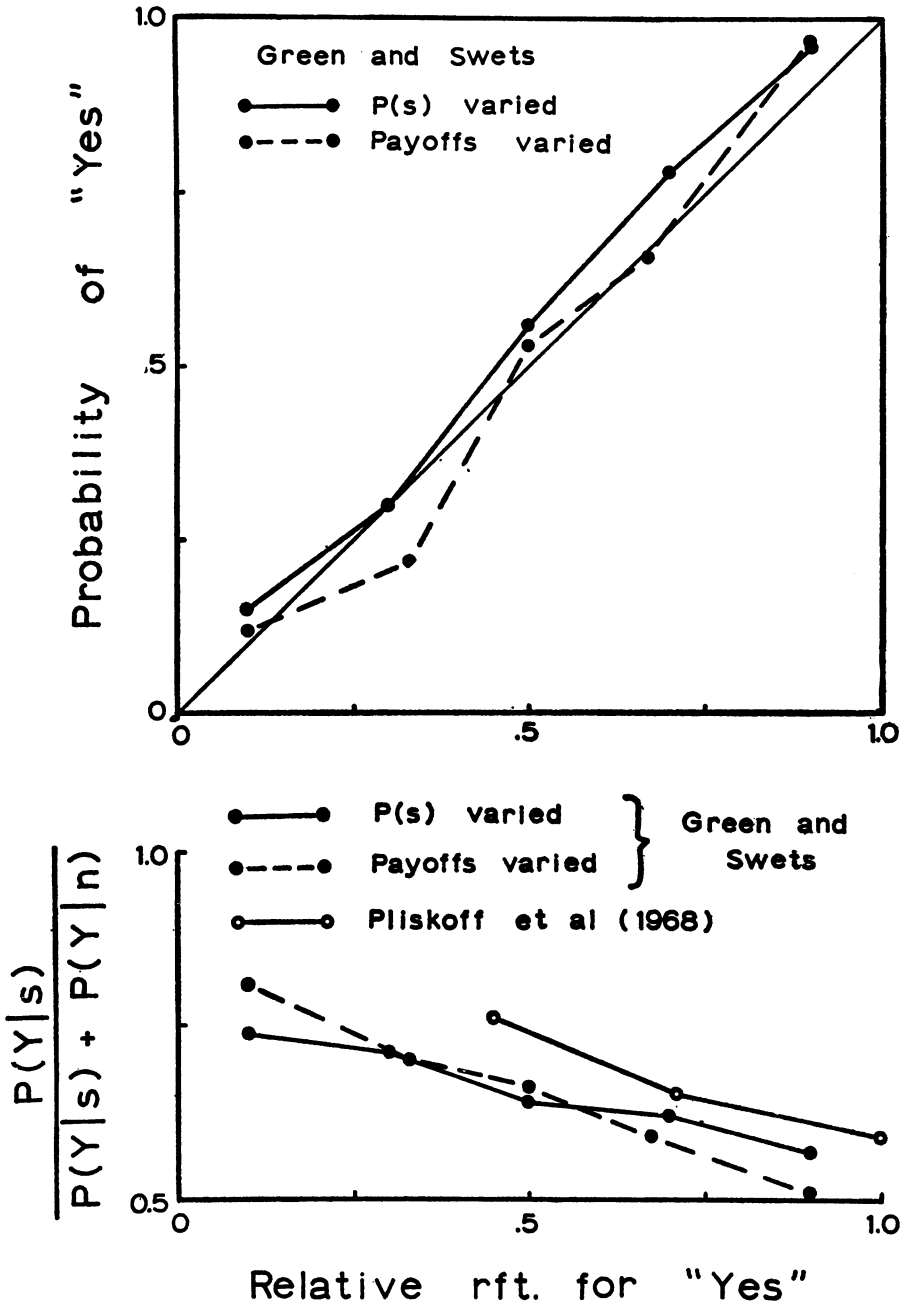

Fig. 1. Upper panel: The probability of responding "yes" (there was a signal), regardless of whether signal or noise was presented, as a function of the relative reinforcement for "yes" when signal probability and payoffs were varied. The data are for Green and Swets' Observer 1. Lower panel: The relative probability of saying "yes" given a signal as a function of relative reinforcement for "yes" for Green and Swets' Observer 1. Also presented is the average rate of responding by pigeons during the VI 40-rft/hr component of a multiple schedule, relative to total responding on the multiple schedule during the VI 40-rft/hr component and the VI 10-rft/hr component as a function of the frequency of reinforcement for pecks on the multiple-schedule key relative to the total frequency of reinforcement on both multiple-schedule and concurrent-schedule keys (adapted from Pliskoff *et al.*, 1968, second sequence).

consistency across different experiments. Signal detection theory has dealt with this problem by proposing an indirectly derived measure of sensitivity, known as d'. This measure is defined as the separation, in standard-deviation units, between a pair of hypothesized normal density functions representing the internally observed effects of signal plus noise, and noise alone. Other assumptions about the distributions representing signal and noise lead to different definitions of the sensitivity parameter, and a distribution-free measure of sensitivity is also available. These matters are treated quite clearly in Green and Swets' Chapters 2 and 3. Under any particular assumptions, the sensitivity parameter may be estimated from ROC plots of the covariation of correct detections and false alarms in yes-no experiments; from experiments permitting the subject to rate his degree of confidence that a signal was presented; and from forced-choice procedures in which the observer must state which of two (or more) presentations contained the signal. These basic experiments, together with theoretical assumptions and calculation procedures, are very clearly set forth in Chapter 4. The most striking achievement of the theory is that if signal strength is constant, d' remains the same within variations of any given experiment; for example, the payoff conditions in a yes-no study or the number of alternatives in a forced-choice procedure. More importantly, it remains constant for individual subjects across different classes of experiment. d' has also been shown to vary directly with signal strength (several studies of this relation are reviewed in Chapter 7). Therefore, d' is an exceedingly valuable measure of sensitivity.

One classical problem in sensory psychology is the existence and measurement of a true threshold—a limit on sensitivity below which detection is impossible and different stimuli are indiscriminable. The history of this problem is reviewed in Chapter 5, and the threshold question is shown to reduce to the question of whether the function relating d' (on the ordinate) to signal intensity (on the abscissa) intersects the abscissa to the right of the origin. The advantages of a relatively pure measure of sensitivity for this analysis are obvious.

At the same time, signal detection theory provides an independent parameter, β, which quantifies the response criterion or bias of the observer. Detection theory explicitly assumes that the observer's response is based on nonsensory information, as well as on the physical stimuli presented on each trial. The subject's pre-experimental biases, his expectations based on the instructions and the *a priori* probability of signal, and the effects of the consequences of responding, are all subsumed under the parameter β. The subject is assumed to transform his observations into a likelihood ratio, which is the ratio of the probability density of an observation if a signal is present to the probability density of that observation in the absence of a signal. He is assumed, further, to partition the likelihood-ratio continuum so that one response occurs if the likelihood ratio exceeds β, and the other if it is less than β. The estimation of β from the response data is also described in Chapter 4. β may be related to non-stimulus variables just as d' may be related to signal strength. In principle, these two independent parameters suffice to make detailed predictions of responding in any detection experiment.

Operant conditioning is in much the same state as psychophysics before the advent of detection theory. Numerous parametric relations between response rates and reinforcement variables have been reported, but as the analysis of behavior has proceeded, it has become increasingly clear that there is little correlation between response rates within different experimental conditions, and that rate itself is a conditionable property of responding. Therefore, the parametric findings may be specific to the situations examined, and it is exceedingly difficult to achieve a systematic unification of our knowledge. To accomplish the desired unification, it may be necessary to derive a theoretical measure of the effects of reinforcement schedules which is conceptually independent of the deprivation and reinforcement parameters of any particular experiment. Conceivably, a second theoretical variable might summarize the latter conditions. The intended analogy to the stimulus and non-stimulus parameters of detection theory is, I hope, obvious.

One particularly interesting feature of signal detection theory is that it is possible to define, under various assumptions, an ideal observer which makes the best possible use of

both sensory and nonsensory information. The ideal observer has, in a sense, two components corresponding to the processing of sensory information and the establishment of a decision criterion. The sensory side of the ideal observer is discussed in detail, with respect to auditory detection, in Chapters 6, 7, and 8. This material is invaluable for students of hearing, but is too specialized for detailed review here.

The ideal decision-maker is considered in Chapter 1, which presents the assumptions of decision theory, and discusses several kinds of general decision goals. The observer may be viewed as attempting to maximize the difference between correct detections and false reports (or the per cent correct responses), to set an upper limit on the probability of false reports, or to maximize expected value. All of these goals may be expressed in terms of likelihood ratio. In Chapter 4, the performance which is plotted in Fig. 1 is compared with the performance expected of an ideal expected-value maximizor, and is found to depart systematically from the ideal. The authors suggest that the difference can be explained by the demand characteristics of sensory experiments, and by subjective transformations of the probabilities and payoff values. This is clearly not their major interest, and the analysis of sensory function can proceed without coming to grips with the problem be cause of the independence of d' and β.

It is interesting to observe the emergence of similar ideal decision-makers in theoretical treatments of schedules of reinforcement. For example, Morse (1966) has considered the possibility of explaining fixed-interval performance as resulting from the occurrence of responding when the probability of reinforcement reaches some threshold value, and Shimp (1966) has treated concurrent performances in terms of maximizing the momentary probability of reinforcement. It would be of considerable interest if a single decision mechanism could be worked out and applied to both free-operant schedule performances and discrete-trial choices in detection experiments. Once again, it is far from clear that the likelihood-ratio approach will be of value in this attempt, but those concerned with theoretical analyses of reinforcement-schedule effects should at least be familiar with the approach.

Finally, *Signal detection theory and psychophysics* contains an appendix of "lab lore" which is invaluable to an investigator attempting experiments in this area for the first time. Every practical aspect of detection research, including the setup and calibration of equipment, instruction of subjects and data-collection procedures, and data analysis and interpretation, is treated informally but comprehensively so that a novice in psychophysics can feel at home in a short time. For the researcher or theorist who is interested in combining the methods and ideas of modern psychophysics with operant conditioning, this book is essential.

REFERENCES

Blough, D. S. Stimulus generalization as signal detection in pigeons. *Science*, 1967, **158**, 940-941.

Blough, D. S. Generalization gradient shape and summation in steady-state tests. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 91-104.

Boneau, C. A. and Cole, J. L. Decision theory, the pigeon, and the psychophysical function. *Psychological Review*, 1967, **74**, 123-135.

Boneau, C. A., Holland, M. K., and Baker, W. M. Color-discrimination performance of pigeons: effects of reward. *Science*, 1965, **149**, 1113-1114.

Catania, A. C. Concurrent performances: reinforcement interaction and response independence. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 253-263. (*a*)

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

Catania, A. C. Concurrent operants. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 213-270.

Goldiamond, Israel. Perception. In A. J. Bachrach (Ed.), *Experimental foundations of clinical psychology*. New York: Basic Books, 1962. Pp. 280-340.

Herrnstein, R. J. Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 1961, **4**, 267-272.

Herrnstein, R. J. Secondary reinforcement and rate of primary reinforcement. *Journal of the Experimental Analysis of Behavior*, 1964, **7**, 27-36.

Morse, W. H. Intermittent reinforcement. In W. K. Honig (Ed.), *Operant behavior: areas of research and application*. New York: Appleton-Century-Crofts, 1966. Pp. 52-108.

Nevin, J. A. Decision theory in studies of discrimination in animals. *Science*, 1965, **150**, 1057.

Pliskoff, S. S., Shull, R. L., and Gollub, L. R. The relation between response rates and reinforcement rates in a multiple schedule. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 271-284.

Rachlin, H. and Baum, W. M. Response rate as a function of amount of reinforcement for a signalled

concurrent response. *Journal of the Experimental Analysis of Behavior*, 1969, **12**, 11-16.

Reynolds, G. S. On some determinants of choice in pigeons. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 53-59.

Rilling, M. and McDiarmid, C. Signal detection in fixed-ratio schedules. *Science*, 1965, **148**, 526-527.

Shimp, C. P. Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, 1966, **9**, 443-455.

Shimp, C. P. Magnitude and frequency of reinforcement and frequencies of interresponse times. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 525-535.

Stubbs, A. The discrimination of stimulus duration by pigeons. *Journal of the Experimental Analysis of Behavior*, 1968, **11**, 223-238.