Comparisons of some learning models for response bias in signal detection

MORTON P. FRIEDMAN, EDWARD C. CARTERETTE, LLOYD NAKATANI AND AL AHUMADA[2]
UNIVERSITY OF CALIFORNIA, LOS ANGELES

The effects of variations in signal probability and varying degrees of correct feedback on response bias were studied in a yes-no auditory signal detection task. The main finding was that the bias towards saying yes was an increasing function of the frequency of signal feedback events, but did not depend on the correctness of the feedback. Several learning models coupled with a simple psychophysical and decision model yielded predictions about overall biases and certain sequential statistics. Only one model, which can be described as an "informational" model, gave a good account of both observed overall biases and sequential statistics. This model assumes the observer's response bias is strengthened for the feedback-reinforced response when the observer's sensory information is ambiguous or is contradicted by the feedback information.

This report presents an empirical and theoretical analysis of the effects of informational feedback variables and signal-presentation probabilities on response bias in a yes-no auditory detection task. The general approach is similar to that employed by Atkinson and Kinchla (1965) for forced-choice tasks, though we use a broader range of informational variables and theoretical models. Our primary interest is in the effects of varying degrees of correct feedback on response bias.

The main theoretical analysis compares a number of learning models for response bias. The learning models assume that the feedback and signal probability conditions lead to trial-by-trial bias adjustments which at asymptote yield statistically stable response biases. The various models differ as to which psychophysical and trial outcome events are presumed to occasion adjustments in the bias.

Because the main interest is response bias, only a very primitive psychophysical model is employed. It is assumed that the stimulus event on each trial (signal plus noise or noise alone) can lead to one of three possible sensory events, or states, in the subject. The intuitive meaning of the sensory states is that, for purposes of deciding whether to respond yes or no on a given trial, the sensory mechanism is set to give only one of three outputs: State D, which roughly corresponds to certainty that the signal was presented; State U, uncertainty whether signal or noise was presented; or State $\underline{D}$, certainty that the signal was not presented. The three states, D, U, and $\underline{D}$, are a simple way of representing the continuum of sensory events which arise from fluctuations in internal and external noise for fixed physical signal and noise parameters. Exactly which physical aspects of the stimulus might give rise to the continuum of sensory states is still an open question and an area of active research (see, for example, Green & Swets (1966) and the recent work of Ahumada (1967).)

We relate the psychophysical and the bias models by combining them with a simple "threshold" decision model. The model is ubiquitous and has been suggested by a number of workers (e.g., Atkinson, 1963; Blackwell, 1963; Parducci & Sandusky, 1965) as a way of representing the effects of bias. It evolved from the ordinary "correction for guessing" threshold model. The particular version we use is that given by Atkinson and Kinchla (1955) and Atkinson, Bower, and Crothers (1965). Figure 1 outlines the assumptions of the model and their relationship to the experimental procedure.
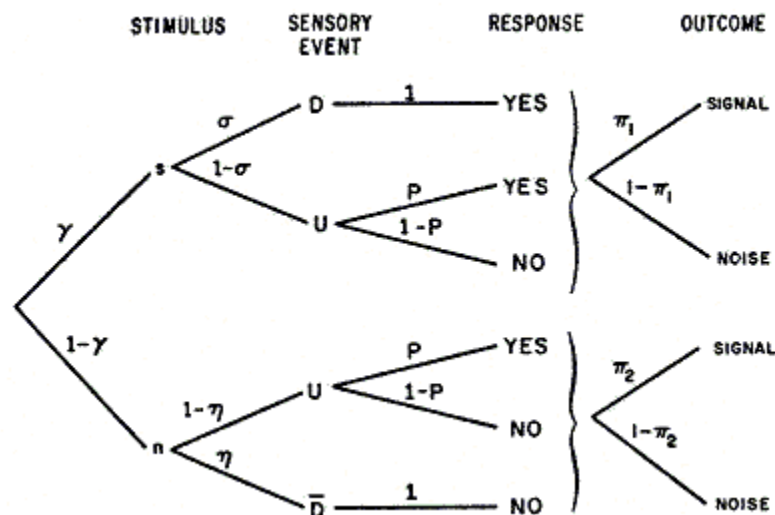

Fig. 1. An outline of the simple decision model and how it relates to the experimental procedure.

On each trial, either a signal, a 1000 Hz tone imbedded in band-limited Gaussian noise is presented, or the noise alone is presented, with probabilities γ and $1 - \gamma$, respectively. These stimulus events are assumed to lead to one of the three sensory states in the subject. If the signal is presented, then State D is elicited with probability σ and State U with probability $1 - \sigma$. If noise alone is presented, then Sensory State $\underline{D}$ is elicited with probability η and State U with probability $1 - \eta$.

The judgment is determined by the sensory state according to a very simple response rule. If State D is elicited, the subject says "yes"; if State $\underline{D}$ is elicited, the subject says "no"; and if State U is elicited, the subject biases his response and says "yes" or "no" with probabilities P and 1- P. The trial ends with the presentation of one of two trial outcomes: (1) the subject is told that a signal was In fact presented, or (2) the subject is told that noise alone occurred. The probability of a signal outcome is π_1 on signal trials and π_2 on noise trials. The subject thus receives correct feedback on all trials when $\pi_1=1$ and $\pi_2=0$. It is assumed that the value of the guessing or bias parameter P depends on the outcome and stimulus presentation probabilities π and γ but the values of the stimulus parameters σ and η do not.

The theoretical expressions for the "hit rate," the proportion of correct detections of the signal P(H), and the "false alarm rate," the proportion of incorrect yes responses on noise alone trials P(F) , are simply:

$$P(H) = \sigma + (1- \sigma)P \qquad (1)$$
$$P(F) = (1- \eta) P \qquad (2)$$

The ROC (iso-sensitivity) function relating the hits and false alarms is:

$$P(H) = \sigma + ((1 - \sigma)/(1 - \eta)) P(F) \qquad (3)$$

Thus the theory predicts a linear ROC function whose slope and intercept depend only on the stimulus parameters , σ and η.

We should point out that there are a variety of other conceptions of the detection process which yield curvilinear and segmented linear ROC functions. Some of these are discussed by Luce (1963). Recent results (Norman, 1964; Atkinson & Kinchla, 1965; Markowitz & Swets, 1967) make it clear that the shape of the ROC function and the relationship of parameters of the ROC function to parameters of the physical stimulus is an open question. For our purpose of studying response bias, the simple model outlined here is convenient. And, as we shall see, it is reasonably consistent with the data.

<div align="center">METHOD</div>

Design

A yes-no auditory signal detection task was employed. The S's task was to report on each trial whether a 1000 Hz sinusoidal signal was or was not presented against a continuous background of white noise during a brief observation interval.

Three variables were manipulated in the experiment: (1) signal intensity , (2) probability of occurrence of the signal, and (3) feedback. The levels of the variables were combined to yield a modified three-way, independent groups, factorial design.

There were two levels of signal intensity .These will be referred to as the hard (low signal intensity) and easy (high signal intensity) conditions. There were three probabilities of occurrence of the signal: .25, .50, and .75. Finally, there were five different feedback conditions: No Feedback on any trial; 100% Correct Feedback on all trials; and three misinformation conditions. The three misinformation conditions were: (1) Random Feedback on all trials; (2) Random Signal-random feedback on all signal trials but correct feedback on noise trials; and (3) Random Noise-random feedback on all noise trials but correct feedback on signal trials.

When Random Feedback was given on all trials, 50% of the trials were followed by signal outcome events and 50% were followed by noise outcome events independent of the stimulus actually presented and the S's response. In the Random Signal condition. 50% of the trials on which the signal was actually presented were called signal trials, and the other 50% were called noise trials independent of the S's response; correct feedback was given on all noise trials. The Random Noise condition was similarly defined.
The conditions of No Feedback. 100% Correct Feedback. and Random Feedback were run under both signal intensity levels and all three levels of signal probability. The Random Noise and Random signal feedback conditions were run under both signal intensity levels, but only with the .50 signal probability. It should be noted that, ignoring the correctness of the feedback in the Random Noise condition, 75% of the

trials were followed by signal outcomes and the other 25% by noise outcome events; and in the Random Signal condition, 25% of the trials were followed by signal outcomes and 75% by noise outcomes. Altogether, these combinations of the three variables yielded 22 independent groups. There were 15 Ss in each group.

Apparatus

There were four subject booths at the narrow end of a 15 x 21 ft experimental room. The warning, observation-interval, and feedback lamps were located on the wall opposite the booths and 12 ft from the Ss. The warning lamp was placed above the observation-interval lamp at about the Ss' eye level. The two feedback lamps were located slightly below and symmetrically about the observation-interval lamp. The Ss' response switches were on large armrests on the chairs in the booths. The switches were double-pole, center-off, spring-return, silent knife switches. Ss indicated their responses by pressing the switch momentarily to the left or right. Programming and recording equipment were in an adjoining control room.

The white background noise was band-pass filtered (at a rate of 24 dB per octave) between 100 and 5000 cps. On signal trials, an electronic switch gated the 1000 cps sinusoid during the entire .5 sec observation interval with equal rise and decay times of 25 msec. The calculated E/No values were 4.97 and 6.30 for the hard and easy signal intensity conditions, respectively. Ss listened binaurally over PDR-600 earphones.

Procedure

Insofar as possible, Ss were run in squads of four. In the instructions, all groups receiving feedback were led to believe that the feedback was correct. Equal emphasis was placed on correctly detecting signal and noise trials. Ss were given examples of the signal alone, the signal in noise, and 10 practice trials (six signal and four noise trials in haphazard order, in which the signal started out 10 dB above its experimental level and was gradually decreased over presentations). Then the 350 experimental trials were run without any break. Previous research has shown that this is more than enough training to reach stable levels of responding when examples of signal and noise trials are explicitly demonstrated to Ss (Carterette & Wyman, 1962; Gundy, 1963; Swets & Sewell, 1963). On each trial, the warn1ng light came on for .5 sec followed by the .5 sec observation interval. The S was allowed 2.75 sec to respond. Then one of the feedback lamps came on for .75 sec in the feedback conditions. The left feedback lamp indicated a signal trial, and the right feedback lamp indicated a noise trial. In the No Feedback condition, both lamps came on during the feedback interval. The intertrial interval was 1.0 sec. The experimental sessions lasted about 45 min.
Different signal and event presentation sequences were used for each squad of Ss in each group. The 350-trial sequences were random with the restriction that obtained proportions of trials of the various types matched the expected proportions.

To increase the homogeneity within the two signal intensity conditions, Ss were eliminated who either did not respond on a high proportion of the trials or who showed low sensitivity to the signal. The criteria for eliminating Ss were: (1) 10 or more trials on which there were no responses, or (2) a d' sensitivity measure of less than .33 in the hard condition or .54 in the easy condition. There were 32 Ss eliminated for these reasons in the easy condition and 41 Ss eliminated in the hard condition. The experimental plan called for 15 Ss in each condition. Because squads of varying sizes were run and because varying numbers of Ss were eliminated in each condition, there were actually from 15 to 19 "good" Ss run in each of the 22 conditions. To simplify some of the calculations, Ss were eliminated randomly to make equal the sizes of the groups at 15 Ss per group.

Subjects

The 330 Ss were students recruited from introductory psychology courses at the University of California, Los Angeles. Participation in a certain number of experiments was a course requirement.

RESULTS AND DISCUSSION

This report will discuss only the results for the feedback conditions. A more complete discussion of the experiment including an analysis of the No Feedback results is given elsewhere (Friedman, Carterette, Nakatani, & Ahumada, 1967).

Preliminary Analysis

A rough check on the adequacy of the simple threshold decision model may be had by considering the ROC functions for the hard and easy signal intensity conditions. Figure 2 shows the ROC functions, plots of the hit rates (proportion of yesses on signal trials) versus the false-alarm rates (proportion of yesses on noise trials) for all experimental conditions. The left panel contains points for groups run under the hard

signal intensity condition, and the right panel contains points for groups run under the easy signal intensity condition. Each data point represents the mean of the responses of the 15 Ss taken over the final 150 trials of the experiment. A preliminary analysis of the data showed that the mean hit and false-alarm rates were stable over it least that part of the data for all conditions. The straight lines in Fig. 2 are linear ROC functions with slope = 1. They were not fit by the best methods but are shown to illustrate predictions of the theory.
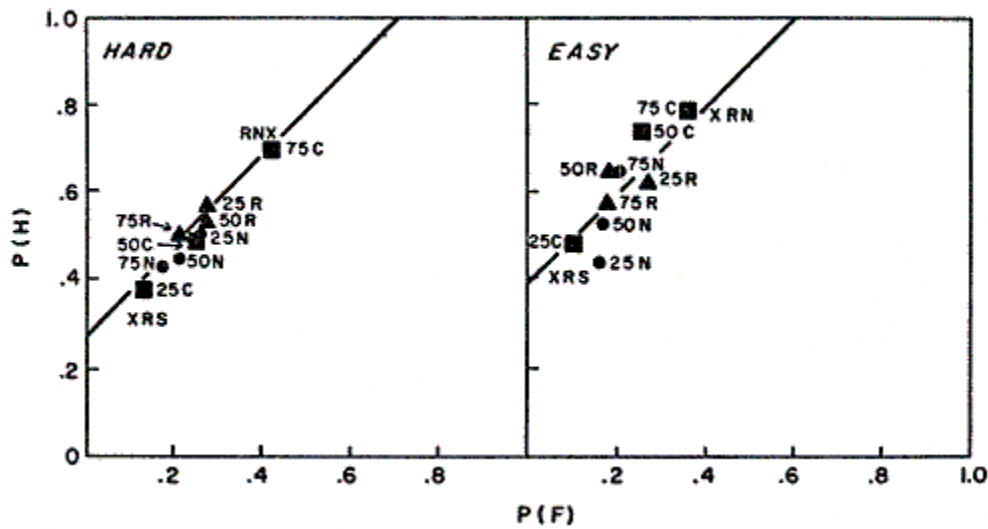

Fig. 2. ROC (Receiver Operating Characteristic) plots of the hit rates P(H) (proportion of yeses on Signal trials) versus the false alarm rates P(F) (proportion of yeses on Noise trials) for all experimental conditions over the last 150 trials. The left panel shows the results for groups run under the Hard condition, and the right panel shows the results for groups run under the Easy condition. The conditions are labeled according to their signal probability (25, 50, or 75) and feedback condition (C = Correct, N = No Feedback, R = Random, RS = Random Signal, RN = Random Noise). The RS and RN conditions were run only with a signal probability of .50. The lines are linear theoretical ROC functions which illustrate the special case of Eq. 3 with $\sigma = \eta$.

We shall not argue here that these data support the three-state model over the many other alternatives. However, we do say that a linear function is a reasonable approximation for the obtained range of hits and false alarms and that the model appears to be a reasonable framework in which to study the effects of our experimental variables on response bias, which we now consider.

Analysis of Bias

To simplify our analysis of bias we consider the special case of the model in which $\sigma = \eta$. (This is roughly analogous to the "equal variance" assumption of the theory of signal detectability (Green & Swets, 1966).) This special case of the theory generates the ROC function with slope = 1 shown in Fig. 2. But more to the point, small changes in the slope of the linear functions (or, for that matter, reasonable curvilinear ROC functions) will not change the qualitative relationships between the experimental conditions and the relative degree of bias, and will make only small quantitative changes.
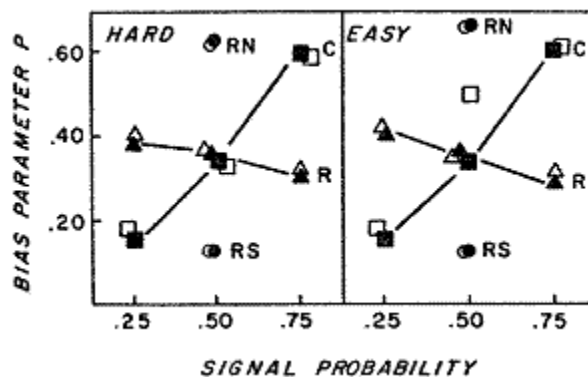

Fig. 3. Obtained (open symbols) and predicted (closed symbols) values of the bias parameter P for the Hard condition (left panel) and the Easy condition (right panel). C = Correct Feedback, R = Random Feedback, RN = Random Noise, RS = Random Signal.

For the $\sigma = \eta$ case of the theory, Equations (1) and (2) for P(H) and P(F) may be solved for the bias parameter P to yield:

$$P = P(F) /(1 - P(H) + P(F)) \qquad (4)$$

This relationship defines the bias associated with a given point (P(H), P(F)) with respect to the linear ROC function with slope of unity through that point. Equation (4) says that the bias associated with a given point is equal to the distance of that point from the x-intercept along the ROC function relative to the total length of the ROC function.

Using Equation (4), estimates of P were calculated for each S from his hit and false-alarm rates over trials 201-350. Figure 3 shows the mean obtained bias scores (open symbols) for all feedback conditions as a function of signal probability along with some predicted values (closed symbols) which will be discussed shortly. The left panel shows the scores for the hard signal intensity condition, and the right panel shows the scores for the easy condition. The main result of note here is that the bias is an increasing function of the probability of signal outcomes, independently of the correctness of the information. According to a Newman-Keuls analysis run at the .05 level, all conditions in which signal outcome probabilities were equal were reliably different from conditions run under other signal outcome probabilities, but did not differ among themselves. Thus, the biases obtained in the Random Noise conditions (labeled RN) in which the true signal probability was .5 but the overall outcome signal outcome probability was. 75 were not significantly different from the correct feedback conditions run with signal probabilities of .75. Similar results hold for the Random Signal conditions (in which the signal out- come probabilities were .25) and the Correct Feedback conditions run with .25 signal probability. There are two other consistent (though not significant by the Newman-Keuls test) relationships in Fig. 3 that we wish to point out: First, for both signal intensity conditions, the Random Signal and Random Noise conditions generated the most extreme biases. Second, the effects of signal probability on bias in the Random Feedback conditions (in which the signal outcome probability was always .5) were small and opposite to the effects of signal probability in the Correct Feedback conditions. For both intensity conditions, there was a tendency for the bias to decrease as the signal probability increased.

Learning Models for the Bias Parameter P

We now consider a class of alternative learning models for the bias parameter P. The learning models may be understood in the context of an isomorphism between the probability learning experiment and the simple psychophysical judgment experiment. The fundamental notion is that the probability learning experiment may be regarded as a limiting case of psychophysical judgment in which there is no stimulus information and the S must guess on every trial (Friedman & Carterette, 1964). For the present threshold model, the limiting case is one in which the psychophysical parameters σ and η are both zero and the uncertain state, U, is evoked on every trial so that the S must guess. All of the learning models employ simple linear operators which assume that the value of the bias parameter P on trial n+1, P_{n+1} is a linear function of the value of the bias parameter on trial n, P_n. Whether the bias increases or decreases on trial n + 1 and the amount of change depend on the sensory state and trial outcome on trial n. The only restriction placed on the operators when feedback is employed is that the bias can change only in the direction of the response specified by the trial outcome. Six of the models which were considered for the feedback conditions are outlined in Table 1. All of the models require two learning parameters.

Table I. Some Alternative Models for the Bias Parameter P ($\alpha = 1 - \theta$).

Trial n		The Value of the Bias Parameter on Trial n+1, P_{n+1}					
Sensory State	Trial Outcome	Information Models		Reinforcement Models		Confirmation Models	
		A	B	A	B	A	B
D	S	P_n	P_n	$\alpha_1 P_n + \theta_1$	$\alpha_1 P_n + \theta_1$	$\alpha_1 P_n + \theta_1$	$\alpha_1 P_n + \theta_1$
D	N	$\alpha_2 P_n$	P_n	$\alpha_2 P_n$	P_n	$\alpha_2 P_n$	P_n
U	S	$\alpha_1 P_n$	$\alpha_1 P_n + \theta_1$	$\alpha_1 P_n + \theta_1$	$\alpha_1 P_n + \theta_1$	P_n	P_n
U	N	$\alpha_2 P_n$	$\alpha_2 P_n$	$\alpha_2 P_n$	$\alpha_2 P_n$	P_n	P_n
D	S	$\alpha_1 P_n + \theta_1$	P_n	$\alpha_1 P_n + \theta_1$	P_n	$\alpha_1 P_n + \theta_1$	P_n
D	N	P_n	P_n	$\alpha_2 P_n$	$\alpha_2 P_n$	$\alpha_2 P_n$	$\alpha_2 P_n$

The various models arise from three alternative conceptions of bow the events on each trial will change the S's biases. The three classes of models are Reinforcement Models. Information Models, and Confirmation Models.

First consider the Information models. The psychological notion behind these models is that the bias changes only when: (1) a trial outcome reduces the observer's uncertainty concerning which stimulus occurred on a trial, or (2) a trial outcome provides information which conflicts with that of the sensory state.

For the Information Model A, whenever the sensory state of uncertainty U occurs, the bias changes in the direction of the trial outcome, because information is provided about the stimulus occurrence on that trial. Similarly, it is assumed that the bias changes in the direction of the trial outcome whenever a trial outcome conflicts with the certainty states, D and $\underline{D}$, and thus provides information to the S. (For the threshold model, this can only occur when the feedback is sometimes erroneous.) Whenever certainty states and trial outcome agree, it is assumed that the bias is unchanged, that is, $P_{n+1} = P_n$ since there was no uncertainty about the stimulus event and both sources of information, the sensory state and the trial outcome, agree. For the Information Model A, the learning parameters depend only on the trial outcomes.

Information Model B differs from Information Model A in assuming that whenever the sensory state and trial outcome conflict, the S ignores both sources of information and the bias is unchanged, i.e., $P_{n+1} = P_n$. Information Models A and B are equivalent for correct feedback conditions.

The basic psychological notion behind the Reinforcement Models is that trial outcomes and not sensory states are the crucial events which change the bias. For Reinforcement Model A, the bias changes on every trial and it changes in the direction of the trial outcome. The learning parameters depend only on the trial outcomes. Reinforcement Model B is like Information Model B in ignoring the trial outcome when it is inconsistent with the sensory state. Reinforcement Models A and B are equivalent for the correct feedback conditions.

According to Confirmation Model A, the bias changes only when one of the certain states D or $\underline{D}$ is evoked, and then only when the sensory information is confirmed or disconfirmed by the trial outcome information. Confirmation Model B is equivalent to Confirmation Model A for the correct feedback case, but assumes that only agreement between the sensory and outcome information can change the bias.

Evaluation of the Learning Models

The learning models were evaluated by comparing their predictions with the observed values of the bias scores and with the first-order sequential response statistics. A complete account of these comparisons and a discussion of several additional models is available elsewhere (Friedman, Carterette, Nakatani, & Ahumada, 1967), so we shall only summarize the main results here.

The goodness of fit of the various models with respect to the bias scores was evaluated by a procedure which allowed the treatment effects (signal probabilities, feedback correctness, task difficulty) and the parameters of the learning models to be separated out as additive effects. All of the models yield identical predictions for the correct feedback conditions, but differed as to their predictions for the partially correct and random feedback conditions. The main result of this analysis was that the predictions of the Information Model A were quite good. All of the models yielded reasonably good predictions for the correct feedback conditions, but only the predictions of Information Model A were consistent with the observed scores in the random and partially correct feedback conditions. The predicted values for Information Model A are shown in Fig. 3. Only Information Model A correctly predicted the obtained inverse relationship between the direction of bias and the signal probability in the random feedback condition, and also the large biases which occurred in the partially correct feedback conditions.

In a second set of analyses the predictions of the learning models were compared with the observed estimates of the first-order sequential statistics, i.e., the hits and false-alarm rates on trial n conditional on the stimulus, response, and outcome on the preceding trial. The main result of these analyses was that a number of models yielded adequate fits to the data when separate learning parameters were estimated for each condition, but Information Model A again yielded the best fits when only a single set of learning parameters was allowed for all conditions.

Our results for the feedback conditions are consistent with those of Atkinson and Kinchla (1965) for a forced-choice auditory detection experiment. Atkinson and Kinchla compared what we here call Information and Reinforcement models, and found that the Information model yielded the best fits for the feedback conditions.[3] However, there were no conditions in their experiment in which the stimulus information could contradict the feedback information, so a comparison of the Information A and B models cannot be had from their results.

All of the models considered here assume that the critical events which can change the bias are sensory events and trial outcomes, and that the response made by the observer on each trial is irrelevant. Other learning models have been suggested (for example, Sklansky, 1964) which assume that the adjustments in the bias depend on the correctness of the response. The "error-correcting" model developed by Sklansky assumes that the bias changes only following an error and remains the same following a correct response.

Models of this sort appear reasonable and deserve to be tested and compared with the class of models presented here. However, it should be noted that the error-correction model is really quite similar to the present Information models, since response errors are highly correlated with what we are defining as Information. In fact, it would appear to be hard to distinguish between the two sorts of models.

Finally, it should be pointed out that all of the learning models which are used here in discrete sensory conceptions can, at the expense of a few more parameters, be applied to continuous sensory models. From our explorations with the continuous models, our impression is that comparisons of these models would yield a pattern of results close to that obtained here for the discrete models.

References

Ahumada, A. Detection of tones masked by noise: A comparison of human observers with digital-computer-simulated energy detectors of varying bandwidths. Unpublished Ph.D. dissertation, UCLA, 1967. (Available as Technical Report No.29, Human Communication Laboratory, Department of Psychology, UCLA.)

Atkinson, R. C. A variable sensitivity theory of signal detection. *Psychol. Rev.,* 1963,70, 91-106.

Atkinson, R. C., Bower, G. H., & Crothers, E. J. *An introduction to mathematical learning theory.* New York: Wiley, 1965.

Atkinson, R. C., & Kinchla, R. A. A learning model for forced-choice detection experiments. *Brit. J. math. & statist. Psychol.*, 1965, 18, 183-206.

Blackwel1, H. R. Neural theories of simple visual discriminations. *J. Opt. Soc. Amer.,* 1963, 53, 129-160.

Carterette, E. C., Friemnan, M. P., & Wyman, M. J. Feedback and psychophysical variables in signal detection. J. *Acoust. Soc. Amer.,* 1966, 39, 1051-1055.

Carterette, E. C.,& Wyman, M. J. Application of a Markov learning model to a simple detection situation involving social pressure. In J. Criswell, H. Solomon and P. Suppes (Eds.), *Mathematical methods in small group processes.* Stanford: Stanford University Press. 1962. Pp. 74-100.

Friedman, M. P.. & Carterette, E. C. Detection of Markovian sequences of signals. *J. Acoust. Soc. Amer.,* 1964, 36, 2334-2339.

Friedman, M. P., Carterette. E. C., Nakatani, L. n., & Ahumada. A. Feedback and frequency variables in signal detection. Technical Report No.30, Human Communication Laboratory, Department of Psychology, UCLA. 1968.

Green, D. M., & Swets, J. A. *Signal detection theory and psychophysics.* New York: Wiley, 1966.

Gundy, R. Auditory detection of an unspecified signal. *J. Acoust. Soc. Amer.,* 1961, 33, 1008-1012.

Kinchla, R. A. A learning factor in visual discrimination. In R. Atkinson (Ed.), *Studies in mathematical psvchology.* Stanford: Stanford University Press, 1963. Pp. 233-249.

Luce, R. D. A threshold theory for simple detection experiments. *Psychol. Rev.,* 1963, 70, 61-79.

Markowitz J., & Swets, J. A. Factors affecting the slope of empirical ROC curves: Comparison of binary and rating responses. *Percept. & Psychophys.,* 1967, 2: 91-100.

Norman, D. A. Sensory thresholds, response biases, and the neural quantum theory. *J. math. Psychol.,* 1964, 1, 88-120.

Parducci, A., & Sandusky, A. Distribution and sequence effects in judgment. *J. exp. Psychol.,* 1965, 89, 450-459.

Ronken. D. A. A note on some alternative models for response bias changes during forced-choice detection experiments. *Brit. J. math. & statist. Psvchol..* 1966, 19, 33-38.

Sklansky, J. Markov chain model of adaptive signal detection. *Proceedings of the* 1963 *Biomics Symposium.* Air Force Systems Command. Wright Patterson A. F. B., Ohio. March, 1963.

Swets. J. A., & Sewall, S. The invariance of signal detectability over stages of practice and levels of motivation. *J. exp. Psychol.,* 1963, 66, 120-126.

Notes

3. Ronken (1966) has suggested an alternative interpretation of Atkinson and Kinchla's results using a model which combines the assumptions of the Information and Reinforcement models.