A LEARNING MODEL FOR FORCED-CHOICE DETECTION EXPERIMENTS

by

R. C. Atkinson
Stanford University

R. A. Kinchla
New York University

TECHNICAL REPORT NO. 70

February 10, 1965

PSYCHOLOGY SERIES

INSTITUTE FOR MATHEMATICAL STUDIES IN THE SOCIAL SCIENCES

STANFORD UNIVERSITY

STANFORD, CALIFORNIA

A Learning Model for Forced-Choice Detection Experiments[1]

by

R. C. Atkinson
Stanford University

R. A. Kinchla
New York University

Abstract

Several signal detection experiments are analyzed in terms of a model that incorporates two distinct processes: an activation process and a decision process. The activation process specifies the relation between external signal events and hypothetized sensory states of the subject. The decision process specifies the relation between the sensory states and the observable responses of the subject. The activation process is assumed to be fixed throughout an experiment, whereas the decision process is viewed as varying from trial to trial as a function of the particular sequence of preceding events. The changes in the decision process are governed by a simple stochastic learning mechanism, and the experimental studies reported here are designed specifically to test the adequacy of this and related representations.

---

[1]Support for this research was provided by the National Aeronautics and Space Administration.

I. Introduction

This paper examines a model for choice behavior in a two-alternative forced-choice detection task. The model is restricted to experimental situations where the subject is given feedback on every trial regarding the correctness of his response, and to situations with a simple outcome structure. Thus, the model has a restricted range of applicability, but for appropriately contrived experiments it appears to provide an accurate account of the gross aspects of the data and certain sequential effects. The model represents a special case of a more general theory proposed by Luce (1963); it is also very similar in most details to a model of forced-choice behavior proposed by Atkinson (1963). The relation of the model considered here to these other theories will be discussed later.

The model postulates that the observable relations between stimulus events and responses are a product of two processes: an activation process and a decision process. The activation process specifies the relation between the external stimulus event and hypothetical sensory states of the subject. The decision process specifies the subject's response in terms of his current sensory state and information that he has acquired during the course of a given experiment. Roughly speaking, the stimulus is fed into the activation process which converts the pattern of external energy changes into sensory information (sensory events); the decision process then operates on the sensory information to determine a response.

In the literature on signal detection, some theories have assumed a continuum of sensory states (Green, 1960; Swets, 1961; Tanner and Swets, 1954), whereas others have argued for a finite representation (Atkinson, Carterette and Kinchla, 1962; Fechner, 1860; Luce, 1963; Norman, 1964). Further, some theories have assumed that the activation process is static

over trials, whereas others have proposed that it varies within certain limits from trial to trial as a function of immediately preceding events (Atkinson, 1963). One point of agreement among all theories is that the decision process is dynamic, and may undergo change when the experimenter manipulates the presentation schedule or outcome structure. However, for a given experimental schedule, some theories treat the decision process as fixed (trial-wise independent), whereas others represent it as changing from trial to trial as a function of the particular sequence of preceding events. This latter way of representing the decision process is an important feature of the model considered in this paper. The subject is viewed as adopting a characteristic pattern of decision making in each experimental situation by means of a simple stochastic learning mechanism. The learning mechanism that will be examined is similar to learning models proposed by Bush and Mosteller (1955).

As noted above, we shall only consider detection experiments involving the two-alternative forced-choice design. Thus, there are two possible stimulus presentations, one of which occurs on each trial: signal plus noise in the first temporal interval and noise alone in the second interval, or noise alone in the first interval and signal plus noise in the second interval. After the stimulus presentation the subject responds "1" or "2" to indicate which interval he believes is more likely to have contained the signal.

To make matters concrete we shall describe the experimental procedure used in an acoustic experiment involving the two-interval forced-choice design; data from this experiment will be presented later. Band-limited Gaussian noise was presented binaurally in the subject's headphones throughout a test session and the signal was a 1000 cycles per second sinusoidal tone; the tone was presented for 100 milliseconds, including equal fall and rise times of 20 milliseconds. The subject was seated before a display board. On each

trial three lights flashed on briefly in succession: a red light, an amber light, and another amber light. Each light was on for 100 milliseconds with a 500 millisecond delay between each successive on period. The red light was simply a warning light, while the amber lights defined two observation intervals. The onset of the signal occurred simultaneously with the onset of one of the amber lights. After the second amber light went off, the subject had 2.5 seconds to indicate his response by pressing a push button located under the appropriate amber light. At the conclusion of the response period, a green light flashed on for 700 milliseconds above the correct response button. There was a 1.5 second intertrial period, thus each trial lasted for 6 seconds. A typical experimental session ran from 300 to 500 trials (thirty to fifty minutes).

With this experimental procedure in mind, let us introduce some notation. The presentation of a signal plus noise in the first interval and noise along in the second interval on trial n will be denoted as $S_{1,n}$, and the presentation of noise in the first observation interval followed by signal plus noise in the second observation interval as $S_{2,n}$. Further, the subject's responses will be denoted $A_{1,n}$ and $A_{2,n}$ to indicate which interval he reported contained the signal on trial n. Finally, $E_{1,n}$ and $E_{2,n}$ will denote the occurrence of an event at the end of trial n informing the subject that stimulus S_1 or S_2, respectively, was presented. Thus

$S_{i,n}$ = the presentation of stimulus S_i on trial n,

$A_{j,n}$ = the occurrence of response A_j on trial n,

$E_{k,n}$ = information event at the end of trial n indicating that stimulus S_k was presented.

Each of the indices i, j, and k can take on the values 1 or 2.

In experiments of the type described above, the following variables can

be manipulated: (a) physical parameters of the situation; (b) presentation schedule of signals; (c) information feedback; and (d) outcome structure. The presentation schedule refers to the scheme used to generate the sequence of stimuli. In general, experimenters have adopted a simple probabilistic schedule for presenting stimuli; namely, to let the events S_1 and S_2 form a binomial sequence with parameter γ. However, more complex schedules have been used; e.g., the stimulus on trial n can depend on the stimulus on trial n-k, the response on trial n-k', or both (Friedman and Carterette, 1964). Generally, an analysis of the simpler schedule is sufficient for most purposes. Manipulation of the physical parameters refers to any change in the physical aspects of the experimental situation; in particular to changes in the level of the background noise and/or the level of the signal. The nature of the information feedback may also be manipulated. In the experiment described above, the subject was always given information regarding the correctness of his response, but one can omit such information, or even give false information. The last variable deals with the outcome structure of the experiment. In general, the outcome structure of a psychophysical experiment is specified by giving the subject a payoff function; i.e., a list of rewards and penalties that he receives depending on what he does under the various stimulus conditions.

In this paper we shall refer to experimental manipulations that involve all four of these variables, but by and large our analyses will be concerned with a very special case of the two-alternative forced-choice design. The presentation schedule will be a binomial sequence with a parameter γ and the outcome structure will involve no explicit payoffs. The subjects will simply be instructed to make a correct response as often as possible. Further, on each trial the subject will be given information regarding the correctness

of his response.

The basic dependent variable is the probability of an A_j response on trial n, given that stimulus S_i occurred. The four outcomes are represented by the matrix

$$\underset{\sim}{P}_n = \begin{array}{c} \\ S_{1,n} \\ S_{2,n} \end{array} \begin{array}{c} \begin{array}{cc} A_{1,n} & A_{2,n} \end{array} \\ \begin{bmatrix} \Pr(A_{1,n}|S_{1,n}) & \Pr(A_{2,n}|S_{1,n}) \\ \Pr(A_{1,n}|S_{2,n}) & \Pr(A_{2,n}|S_{2,n}) \end{bmatrix} \end{array} . \tag{1}$$

This matrix will be called the <u>performance matrix</u>. Note that the rows of the matrix sum to one, for on every trial the subject makes either an A_1 or A_2 response. Thus, if one entry in a row is known, so also is the other. Typically the performance matrix is specified by giving the entries in the first column; namely $\Pr(A_{1,n}|S_{1,n})$ and $\Pr(A_{1,n}|S_{2,n})$. In the literature the occurrence of an A_1 response to an S_1 stimulus is called a <u>hit</u>, and the occurrence of A_1 response to an S_2 stimulus is called a <u>false alarm</u>. We shall use this terminology, denoting them as H_n and F_n; i.e.,

$$\Pr(H_n) = \Pr(A_{1,n}|S_{1,n})$$
$$\Pr(F_n) = \Pr(A_{1,n}|S_{2,n}) \quad .$$

Fixing $\Pr(H_n)$ and $\Pr(F_n)$, then, completely specifies the performance matrix.

Other quantities of interest can be defined in terms of the hits and false-alarms. Frequently we want to know the probability of an A_1 response on trial n independent of the stimulus event; namely,

$$\Pr(A_{1,n}) = \Pr(H_n)\Pr(S_{1,n}) + \Pr(F_n)\Pr(S_{2,n}) \quad . \tag{2}$$

Also of interest is the probability of a correct response on trial n (which we denote as C_n):

$$\Pr(C_n) = \Pr(H_n)\Pr(S_{1,n}) + [1 - \Pr(F_n)]\Pr(S_{2,n}) \quad . \tag{3}$$

II. Assumptions and Rules of Identification

Activation and Decision Processes

The model assumes that one and only one sensory state can occur on each trial of the experiment. The sensory states will be denoted as s_0, s_1, s_2, s_3, ... We do not suppose that the same sensory state necessarily results whenever a particular stimulus is presented, but rather that the state is determined by a random process. The activation process on trial n of an experiment can be represented by the stochastic matrix

$$\underset{\sim}{A}_n = \begin{array}{c} \\ S_1 \\ S_2 \end{array} \begin{array}{c} \begin{array}{cccc} s_0 & s_1 & s_2 & \cdots \quad s_x \end{array} \\ \begin{bmatrix} a_{10}^{(n)} & a_{11}^{(n)} & a_{12}^{(n)} & \cdots \quad a_{1x}^{(n)} \\ a_{20}^{(n)} & a_{21}^{(n)} & a_{22}^{(n)} & \cdots \quad a_{2x}^{(n)} \end{bmatrix} \end{array},$$

where $a_{ij}^{(n)}$ denotes the probability of eliciting sensory state s_j on trial n given stimulus S_i on that trial. Similarly, the decision process can be represented by the matrix

$$\underset{\sim}{D}_n = \begin{array}{c} \\ s_0 \\ s_1 \\ s_2 \\ \cdot \\ \cdot \\ \cdot \\ s_x \end{array} \begin{array}{c} \begin{array}{cc} A_1 & A_2 \end{array} \\ \begin{bmatrix} d_{01}^{(n)} & d_{02}^{(n)} \\ d_{11}^{(n)} & d_{12}^{(n)} \\ d_{21}^{(n)} & d_{22}^{(n)} \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ d_{x1}^{(n)} & d_{x2}^{(n)} \end{bmatrix} \end{array},$$

where $d_{ij}^{(n)}$ is the probability of eliciting response A_j on trial n given sensory state s_i on that trial. Then, the performance matrix specified by Eq. 1 is obtained by taking the product of the activation matrix and the

decision matrix; i.e.,

$$\underset{\sim}{P}_n = \underset{\sim}{A}_n \underset{\sim}{D}_n \quad .$$

The model that we shall examine postulates three sensory states for the two-alternative forced-choice task:

s_0 = no detection

s_1 = detection in observation interval 1

s_2 = detection in observation interval 2.

Further, the activation process and the decision process are defined by the following matrices:

$$\underset{\sim}{A}_n = \begin{array}{c} \\ S_1 \\ S_2 \end{array} \begin{array}{c} \begin{array}{ccc} s_0 & s_1 & s_2 \end{array} \\ \begin{bmatrix} 1-\sigma & \sigma & 0 \\ 1-\sigma & 0 & \sigma \end{bmatrix} \end{array} \tag{4}$$

$$\underset{\sim}{D}_n = \begin{array}{c} \\ s_0 \\ s_1 \\ s_2 \end{array} \begin{array}{c} \begin{array}{cc} A_1 & A_2 \end{array} \\ \begin{bmatrix} p_n & 1-p_n \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \end{array} . \tag{5}$$

There are several points to note about these matrices. First, the entries in $\underset{\sim}{A}_n$ are constants independent of the trial number; thus the sensory process is assumed to be fixed over all trials of the experiment. In contrast, the decision process may vary as a function of the trial number, and this dependence is indicated by affixing the trial index n to p. Also, note that s_1 can occur only if S_1 is presented, and s_2 can occur only if S_2 is presented. Thus, these sensory states have an unambiguous relation to the stimulus, since the signal event can be inferred with probability 1 when s_1 or s_2 occur. However, sensory state s_0 is ambiguously related to the stimulus, for it can occur following either signal event. The parameter σ

characterizes this stimulus ambiguity in the output of the sensory system. Both loss of stimulus information due to external noise and loss due to limitations on the resolving power of the sensory system are summarized in the parameter σ. Thus σ may be interpreted as a measure both of the physical stimulus and of the subject's sensitivity; we shall refer to σ as the sensitivity parameter.

The decision matrix D_n reflects the relative ambiguity of the sensory states. If the subject's instructions are to make an A_i response given an S_i stimulus, then the correct response is completely determined when an s_1 or s_2 sensory state occurs. However, the subject faces a dilemma if he must make a response on the basis of s_0; either stimulus could have evoked s_0, so the subject needs some strategy by which he can resolve the ambiguity and select a response. The quantity p_n is a measure of the subject's tendency to resolve the ambiguity by making an A_1 response rather than an A_2; p_n will be referred to as the response bias on trial n.

As indicated earlier, the parameter σ represents the subject's sensitivity to the signal and p_n is a response bias more or less under the control of the subject. Of the experimental variables discussed earlier, we assume that the presentation schedule, information feedback, and the outcome structure influence p_n, but do not affect the value of the sensitivity parameter. Also, we assume that the sensitivity parameter for a given subject is determined solely by the physical aspects of the experimental situation. It is, of course, necessary to show experimentally that these interpretations are correct, and to examine how the parameters σ and p_n are related to the physical characteristics of a given experimental situation.

In order to see how the sensitivity parameter and the bias parameter interact let us examine the relation between hits and false alarms as one or

the other of these parameters is manipulated. Taking the product of the matrices in Eqs. 4 and 5 yields the performance matrix $\underset{=}{P}_n$ for this model. The entries in the first column of $\underset{=}{P}_n$ are as follows:

$$\Pr(H_n) = (1 - \sigma)p_n + \sigma \qquad (6a)$$

$$\Pr(F_n) = (1 - \sigma)p_n \quad . \qquad (6b)$$

If the sensitivity parameter σ is held constant and p_n is manipulated, an exchange relation is established between $\Pr(H_n)$ and $\Pr(F_n)$; if one probability is changed the other is also, and in a predictable way. To find the equation of this relation, we eliminate p_n from Eq. 6 yielding

$$\Pr(H_n) = \sigma + \Pr(F_n) \quad . \qquad (7)$$

Thus, if σ is held constant (fixed signal and noise levels) and p_n is forced to vary (manipulations in the presentation schedule, outcome structure, etc.), the relation between hits and false alarms should be characterized by a linear function with slope 1. Plots of the relation between $\Pr(H_n)$ and $\Pr(F_n)$ under experimental conditions where the signal to noise ratio is held fixed and other variables are allowed to vary are often referred to as receiver-operating-characteristic curves, or more simply as ROC curves. Generally $\Pr(F_n)$ is plotted along the abscissa and $\Pr(H_n)$ along the ordinate. When this is the case, the theoretical ROC curve intersects the ordinate at a point whose value is σ; as p_n goes from zero to one, a straight line is traced from the point $(0, \sigma)$ to the point $(1 - \sigma, 1)$.

If p_n is held constant and the sensitivity parameter changed, there is a well-defined relation between hits and false alarms. Eliminating σ from Eq. 6 yields

$$\Pr(H_n) = 1 - \Pr(F_n)\left[\frac{1 - p_n}{p_n}\right] \quad . \qquad (8)$$

Plots of the relation between $Pr(H_n)$ and $Pr(F_n)$ when p_n is constant and σ is varied are called iso-bias curves. As σ goes from one to zero we trace out an iso-bias curve that goes from the point $(0, 1)$ to the point (p_n, p_n).

Learning Process

As indicated earlier, an important feature of the present analysis is to represent changes in the bias probability in terms of a learning process of the type proposed by Bush and Mosteller (1955). We assume that the bias on trial $n+1$ is a linear function of its value on trial n. Specifically, if s_0 occurs and is followed by E_1 (i.e., the experimenter informs the subject that the signal was in the first interval) then p_n will increase. If s_0 occurs and is followed by information event E_2, then p_n will decrease. For all other contingencies no change will occur in p_n. These statements can be summarized as follows:

$$p_{n+1} = \begin{cases} (1-\theta)p_n + \theta, & \text{if } s_{0,n} \,\&\, E_{1,n} \\ (1-\theta')p_n, & \text{if } s_{0,n} \,\&\, E_{2,n} \\ p_n, & \text{otherwise,} \end{cases} \qquad (9)$$

where $0 < \theta, \theta' \leq 1$. Justification for this equation is postponed until later.

We now want to derive an expression for the expected value of p_n as a function of the presentation schedule and the sensitivity parameter. Recall that γ is the probability of an S_1 signal event and $1 - \sigma$ is the probability of activating sensory state s_0 given either S_1 or S_2. Hence

$$\begin{aligned} Pr(s_{0,n} \,\&\, E_{1,n}) &= \gamma(1-\sigma) \\ Pr(s_{0,n} \,\&\, E_{2,n}) &= (1-\gamma)(1-\sigma) \\ Pr(\text{otherwise}) &= \sigma . \end{aligned}$$

To compute the expected value of the bias probability on trial $n+1$, we simply weight each of the possible outcomes listed in Eq. 9 by its probability of

occurrence given above. That is, the expected value on trial $n+1$ given a fixed value p_n on trial n is

$$E(p_{n+1}) = \gamma(1-\sigma)[(1-\sigma)p_n + \theta] + (1-\gamma)(1-\sigma)(1-\theta')p_n + \sigma p_n$$
$$= \{1 - (1-\sigma)[\theta\gamma + \theta'(1-\gamma)]\}p_n + \theta\gamma(1-\sigma) .$$

Without going into the mathematical details, it can be shown that p_n in the above equation can be replaced by its expected value. Consequently, we have a linear first-order difference equation in $E(p_n)$ whose solution is

$$E(p_n) = p_\infty - (p_\infty - p_1)G^{n-1}$$

where

$$p_\infty = \frac{\gamma}{\gamma + (1-\gamma)\varphi} \quad , \qquad G = 1 - (1-\sigma)[\theta\gamma + \theta'(1-\gamma)] \tag{10}$$

and $\varphi = \frac{\theta'}{\theta}$. Note that p_∞, which is $\lim_{n\to\infty} E(p_n)$, does not depend on the absolute values of θ and θ' but only on their ratio.

Combining the results in Eqs. 6 and 10 yields

$$Pr(H_n) = \sigma + (1-\sigma)[p_\infty - (p_\infty - p_1)G^{n-1}] \tag{11a}$$

$$Pr(F_n) = (1-\sigma)[p_\infty - (p_\infty - p_1)G^{n-1}] . \tag{11b}$$

From these equations it is clear that hits and false alarms will depend on p_1 at the start of an experimental session; however, over trials the subject's performance will change at a rate controlled by the quantity G and at asymptote will be determined by σ and p_∞. The change in performance predicted by Eq. 11 is a well-known experimental phenomenon. Generally, however, most research workers have tended to ignore the changes that occur at the beginning of an experimental session, and instead have concentrated on an analysis of data after performance has settled down to a stable level. For the experiments analyzed in this paper we shall adopt this policy; to do so makes matters simpler because fewer parameters need to be estimated.

Since asymptotic performance will be stressed in subsequent discussions, the following notation will be useful:

$$\lim_{n\to\infty} \Pr(H_n) = \Pr(H)$$

$$\lim_{n\to\infty} \Pr(F_n) = \Pr(F) \ .$$

That is, asymptotic expressions will be indicated by simply deleting the trial subscript. Making the appropriate substitutions in Eq. 11 yields

$$\Pr(H) = \sigma + \frac{(1-\sigma)\gamma}{\gamma + (1-\gamma)\varphi} \tag{12a}$$

$$\Pr(F) = \frac{(1-\sigma)\gamma}{\gamma + (1-\gamma)\varphi} \ . \tag{12b}$$

Similarly, for the asymptotic proportion of correct responses we obtain (see Eq. 3)

$$\Pr(C) = \sigma + (1-\gamma)(1-\sigma) + \frac{(1-\sigma)\gamma(2\gamma-1)}{\gamma + (1-\gamma)\varphi} \ . \tag{13}$$

And, for the asymptotic proportion of A_1 responses (see Eq. 2)

$$\Pr(A_1) = \gamma\sigma + \frac{\gamma(1-\sigma)}{\gamma + (1-\gamma)\varphi} \ . \tag{14}$$

Related Models

The model described in this section is very similar to one proposed by Luce (1963). His article presents a theory of signal detection that is applicable to a wide range of experimental procedures including both yes-no and the forced-choice designs. When the theory is applied to the two-alternative forced-choice experiment a model is obtained that has four sensory states. There are four such states, because it is assumed that in each observation interval a hypothetical event D or $\bar{D}$ will occur; hence, for a two-interval problem, the sensory states are the ordered pairs $\langle DD\rangle$, $\langle D\bar{D}\rangle$, $\langle \bar{D}D\rangle$, and $\langle \bar{D}\bar{D}\rangle$. The D event will occur with probability q when a signal is presented, and with probability q' when the signal is not presented. Further, the subject always makes the A_1 response when $\langle D\bar{D}\rangle$ occurs, never when

$\langle\bar{D}D\rangle$ occurs, with probability v_n when $\langle DD\rangle$ occurs, and another probability w_n when $\langle\bar{D}\bar{D}\rangle$ occurs. These assumptions can be represented in matrix form as follows:

$$\underset{\sim}{A} = \begin{array}{c} \\ s_1 \\ s_2 \end{array} \begin{array}{c} \begin{array}{cccc} \langle DD\rangle & \langle D\bar{D}\rangle & \langle\bar{D}D\rangle & \langle\bar{D}\bar{D}\rangle \end{array} \\ \begin{bmatrix} qq' & q(1-q') & (1-q)q' & (1-q)(1-q') \\ q'q & q'(1-q) & (1-q')q & (1-q')(1-q) \end{bmatrix} \end{array}$$

$$\underset{\sim}{D}_n = \begin{array}{c} \\ \langle DD\rangle \\ \langle D\bar{D}\rangle \\ \langle\bar{D}D\rangle \\ \langle\bar{D}\bar{D}\rangle \end{array} \begin{array}{c} \begin{array}{cc} A_1 & A_2 \end{array} \\ \begin{bmatrix} v_n & 1-v_n \\ 1 & 0 \\ 0 & 1 \\ w_n & 1-w_n \end{bmatrix} \end{array} .$$

Luce also postulates that the bias parameters v_n and w_n undergo trial-by-trial changes of the form specified by Eq. 9.

When $q' = 0$ the event D will never occur in the absence of a signal, and then the above matrices reduce to those presented in Eqs. 4 and 5 with $p_n = w_n$ and $\sigma = q$. Under these conditions Luce's model is precisely the same as the one presented here. Although we will not present the analyses, it can be shown that the fit of the model cannot be significantly improved by letting q' be nonzero for the data treated in this paper.

The model proposed here also is very similar to a special case of a theory proposed by Atkinson (1963) and Atkinson, Carterette and Kinchla (1962). The difference is that their bias process was formulated in terms of the multi-element pattern model of stimulus sampling theory (Atkinson and Estes, 1963). However, the two models make identical predictions for all of the statistics analyzed in this paper, and differ only on certain predictions such as

sequential statistics that depend on previous responses. It should be pointed out that although both Luce's theory and Atkinson's reduce to essentially the same model in the two-alternative forced-choice case, they make markedly different predictions for yes-no experiments and for forced-choice experiments with more than two intervals.

III. Data Analysis

We now examine data collected from eight subjects in the forced-choice acoustic experiment described earlier. In this study the signal and noise levels were held constant throughout the experiment and the subject was always given information at the end of each trial regarding the correctness of his response. The only experimental manipulation involved the use of three different presentation schedules. The probability, γ, of an S_1 event took on the following values:

Schedule A: $\gamma = .25$

Schedule B: $\gamma = .50$

Schedule C: $\gamma = .75$.

Test sessions of 350 trials each were run on consecutive days. Each day a subject ran on one of the three schedules for the entire session. In successive 3-day blocks a subject ran one day on each of the three schedules; within each 3-day block the order was randomly determined. The experiment involved 15 experimental sessions and therefore each schedule was run on five separate days.

Table 1 presents the proportion of A_1 responses on both S_1 and S_2 trials over the last 250 trials of replications two through five of each presentation schedule; thus each estimate is based on 250 x 4 = 1000 trials. The first replication of each presentation schedule has been deleted, because we view the subject as adapting to the detection task on early days of the

Table 1

Predicted and Observed Proportions of Pr(H), Pr(F), Pr(C), and $Pr(A_1)$

(The observed proportions are in parentheses.)

Subject	Schedule A				Schedule B				Schedule C			
	Pr(H)	Pr(F)	Pr(C)	$Pr(A_1)$	Pr(H)	Pr(F)	Pr(C)	$Pr(A_1)$	Pr(H)	Pr(F)	Pr(C)	$Pr(A_1)$
1	.601 (.622)	.154 (1.63)	.785 (.783)	.266 (.278)	.744 (.714)	.297 (.260)	.724 (7.27)	.521 (.487)	.877 (.890)	.430 (.462)	.800 (.802)	.765 (.783)
2	.543 (.529)	.125 (.136)	.792 (.780)	.229 (.234)	.680 (.654)	.262 (.249)	.709 (.702)	.471 (.451)	.832 (.854)	.414 (.397)	.771 (.791)	.727 (.740)
3	.597 (.626)	.106 (.107)	.820 (.826)	.229 (.237)	.716 (.707)	.225 (.210)	.746 (.748)	.470 (.459)	.849 (.842)	.358 (.384)	.797 (.786)	.726 (.728)
4	.529 (.517)	.127 (.122)	.787 (.788)	.227 (.221)	.669 (.649)	.267 (.242)	.701 (.703)	.468 (.446)	.825 (.857)	.424 (.454)	.763 (.779)	.725 (.756)
5	.520 (.546)	.120 (.142)	.790 (.780)	.220 (.243)	.658 (.650)	.258 (.240)	.700 (.705)	.458 (.445)	.816 (.799)	.416 (.413)	.758 (.746)	.716 (.703)
6	.542 (.547)	.141 (.139	.780 (.783)	.241 (.241)	.689 (.680)	.287 (.279)	.701 (.701)	.488 (.479)	.841 (.847)	.440 (.451)	.771 (.772)	.741 (.748)
7	.618 (.627)	.125 (.136)	.810 (.805)	.249 (.259)	.744 (.742)	.252 (.251)	.746 (.746)	.498 (.496)	.872 (.864)	.379 (.369)	.809 (.806)	.749 (.740)
8	.570 (.552)	.125 (.108)	.799 (.807)	.236 (.219)	.704 (.687)	.258 (.244)	.723 (.722)	.481 (.465)	.847 (.887)	.401 (.438)	.785 (.806)	.735 (.775)
Average	.565 (.571)	.128 (.132)	.795 (.794)	.237 (.241)	.700 (.685)	.263 (.247)	.719 (.719)	.482 (.466)	.845 (.855)	.408 (.421)	.782 (.786)	.735 (.746)

experiment and prefer to treat his data only after he clearly understands the experimental routine and is well experienced. Also, the first 100 trials of each of the subsequent experimental sessions were deleted because, as noted earlier, we shall confine the analysis to asymptotic performance.

In this experiment the signal and noise levels were constant over all sessions and only the presentation schedule varied. Therefore, σ should be fixed throughout the experiment, but p_∞ should vary with changes in γ. It has already been shown that in theory hits and false alarms should fall on the straight line

$$Pr(H) = \sigma + Pr(F) \ .$$

We now wish to fit this equation to the three data points corresponding to presentation schedules A, B, and C. Figure 1 presents plots of $Pr(H)$ and $Pr(F)$ for individual subjects. In order to fit the above equation to the three points for a given subject, we use the method of least squares; i.e., we select the value of σ so that it minimizes the sum of squared deviations between observed values and those predicted by the above equation. Applying the least squares method yields the estimates of σ that are given in Figure 1; these estimates were used to generate the ROC curves displayed in the figure. By inspection of the figures we see that there is good agreement between the observed data points and the predicted ROC curves. Recall that the signal and noise levels were set at the same values for all subjects and consequently variations in σ represent inter-subject variations in sensitivity level. The maximum sensitivity level is displayed by Subject 7, with $\sigma = .492$; whereas the minimum sensitivity level is displayed by Subject 5, with $\sigma = .400$.

We now evaluate the bias process with regard to the data presented in Table 1. First, however, note that if γ and σ are fixed in Eq. 12 and

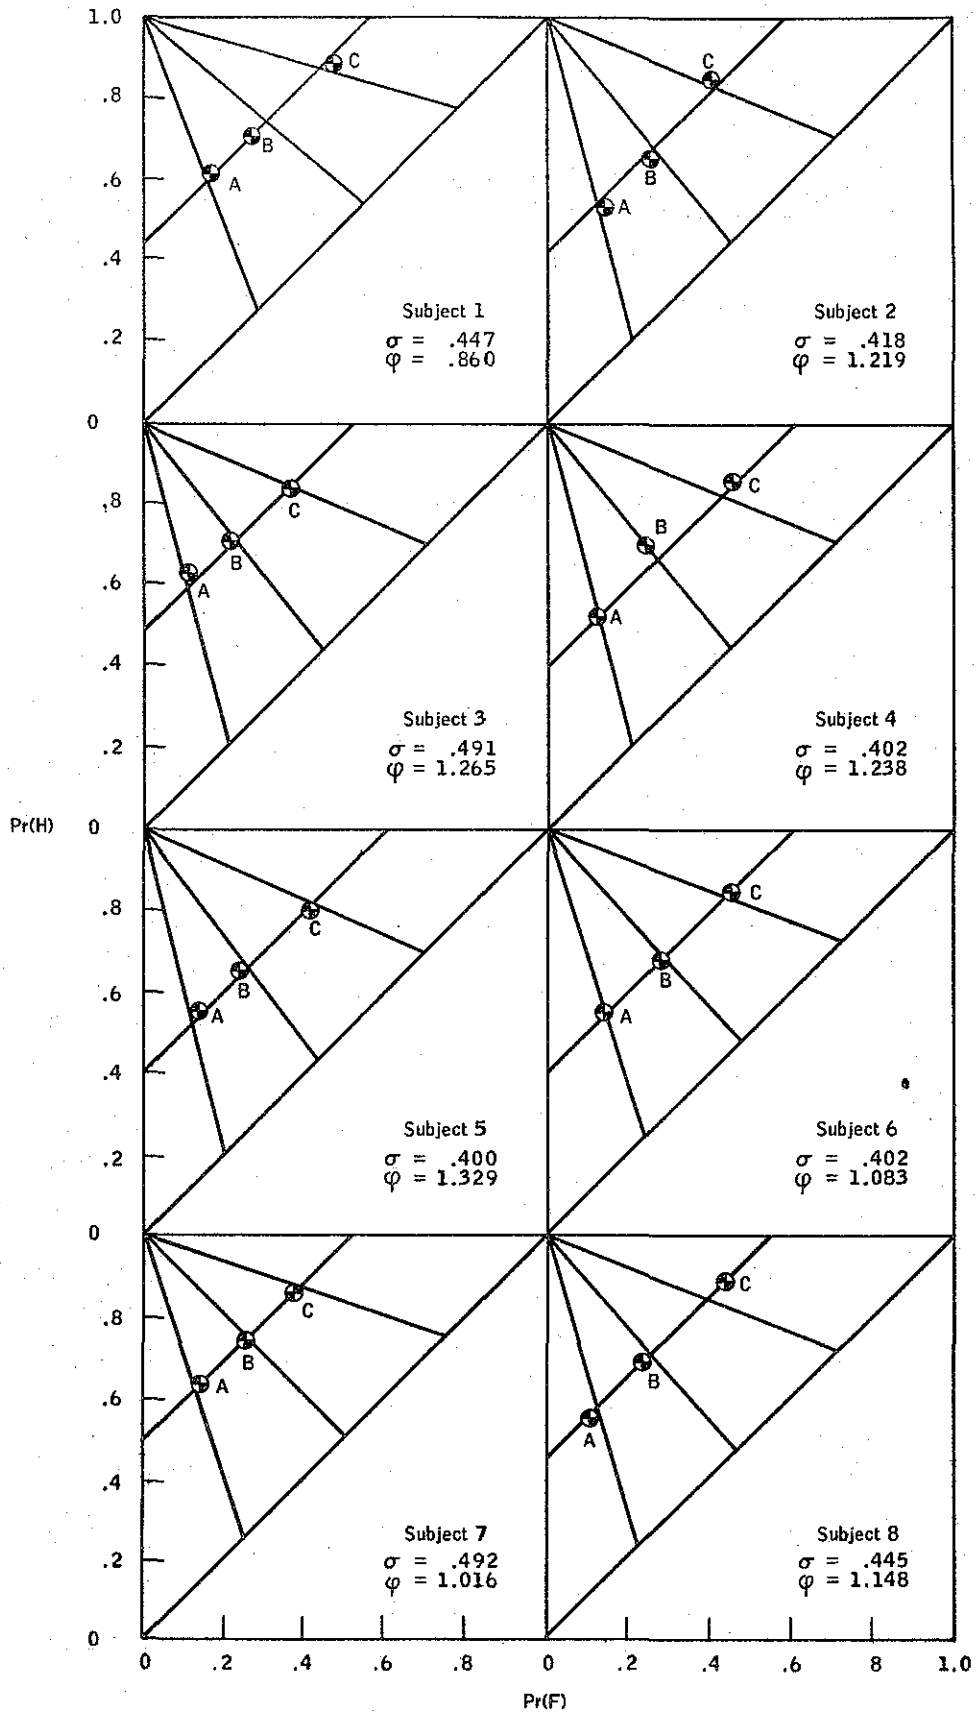

Figure 1. Observed and predicted values for Pr(H) and Pr(F)

φ is varied from 0 to ∞, then the point [Pr(F), Pr(H)] moves along an ROC curve and approaches the lower-left point $(0, \sigma)$ as $\varphi \to \infty$, and the upper-right point $(1-\sigma, 1)$ as $\varphi \to 0$. Stated differently, no matter where the point may fall on the ROC curve (for fixed values of γ and σ), there exists a corresponding value of φ. Hence, if the three observed points [Pr(F), Pr(H)] in our experiment fall on a straight line with slope 1, then perfect fits of the data can be obtained by estimating separate values of φ for each presentation schedule.

Performing separate estimates of φ, however, would violate our original intentions. In formulating Eq. 9 it was assumed that the parameters θ and θ' characterize a subject's trial-to-trial adjustments to stimulus and information events, and do not depend on the overall presentation schedule. The values of θ and θ' may vary from subject to subject reflecting individual differences; however, for a given subject θ and θ' are assumed to be fixed and invariant with regard to the presentation schedule and the signal intensity. Earlier we required that σ be independent of the presentation schedule, and now the same constraint is placed on φ. Thus for each subject we want a single estimate of φ which then can be used to make predictions for all three presentation schedules.

To obtain an estimate of φ we use the observed proportion of A_1 responses given in Table 1. Eq. 14 gives the theoretical expression for $Pr(A_1)$; solving for φ yields

$$\varphi = \frac{\gamma(1-\sigma)}{[Pr(A_1) - \sigma\gamma](1-\gamma)} - \frac{\gamma}{1 - \gamma} \quad .$$

For each presentation schedule we have substituted the estimated value of σ and the observed value of $Pr(A_1)$ in the above equation to obtain an estimate of φ. For example, for Subject 1, $\sigma = .447$, $\hat{Pr}(A_1) = .278$, and $\gamma = .25$

on schedule A; hence substituting in the above equation yields $\hat{\varphi}_A = .777$. Similarly $\hat{\varphi}_B$ and $\hat{\varphi}_C$ can be computed using the appropriate values of γ and $\hat{Pr}(A_1)$. To obtain an overall estimate of φ for each subject we have taken the average of the three estimates; namely

$$\hat{\varphi} = \tfrac{1}{3}[\hat{\varphi}_A + \hat{\varphi}_B + \hat{\varphi}_C] .$$

The various estimates of $\hat{\varphi}$ are presented in Table 2. Note that the value of $\hat{\varphi}$ averaged over subjects is somewhat greater than one, indicating that $\theta' > \theta$. The interpretation of this result is that the E_2 event has a slightly greater effect on increasing the probability of an A_2 response than the E_1 event has on increasing the probability of an A_1 response.

Using the estimates of σ and φ, predictions can be computed for $Pr(H)$, $Pr(F)$, $Pr(C)$, and $Pr(A_1)$ from Eqs. 12-14. These predicted values and the corresponding observed quantities are presented in Table 1. Also in Fig. 1 the predicted and observed values of $Pr(H)$ and $Pr(F)$ are plotted in the ROC space. In this figure the predicted point for each presentation schedule is at the intersection of the computed iso-bias curve and the ROC curve.

By and large, the correspondence between predicted and observed values is very good. Only Subject 8 shows a systematic discrepancy between predicted and observed quantities. For this subject $\hat{\varphi} = 1.148$ and hence $Pr(A_1)$ should be about .236 for schedule A and about .735 for schedule C. However, what we observe is that $Pr(A_1)$ overshot its predicted value for schedule C and went below its predicted value for schedule A. To a degree, this subject's performance deviated from the theoretical values in the direction of optimizing the probability of a correct response. Specifically, consider the function for the probability of a correct response; namely $Pr(C) = \sigma + (1-\sigma)[\gamma p + (1-\gamma)(1-p)]$. For a fixed σ, to maximize this

Table 2

Estimates of φ

Subject	$\hat{\varphi}$	$\hat{\varphi}_A$	$\hat{\varphi}_B$	$\hat{\varphi}_C$
1	.860	.777	1.099	.705
2	1.219	1.162	1.400	1.096
3	1.265	1.155	1.390	1.251
4	1.238	1.324	1.446	.945
5	1.329	1.065	1.449	1.472
6	1.083	1.085	1.147	1.018
7	1.016	.914	1.028	1.105
8	1.148	1.384	1.284	.775
Average	1.145	1.108	1.280	1.046

function the subject should set the bias parameter as follows:

$$p = \begin{cases} 1, & \text{if } \gamma > \frac{1}{2} \\ 0, & \text{if } \gamma < \frac{1}{2} \end{cases}.$$

When $\gamma = \frac{1}{2}$ any value of p yields a maximum. If a subject adopted the strategy specified by the above equation, then the ROC curve would reduce to three points; one at $(0, \sigma)$ for $\gamma < \frac{1}{2}$, another at $(1-\sigma, 0)$ for $\gamma > \frac{1}{2}$, and a third point for the presentation schedule where $\gamma = \frac{1}{2}$. The behavior of Subject 8 tends to move away from the theoretical predictions in the direction of maximization, but of course to nowhere near the extent indicated by the above equation. It is our contention that if monetary payoffs for correct responses and penalties for incorrect responses were introduced into the experimental situation, then more subjects would tend to deviate from the theoretical values, the deviation being in the direction of optimization. Thus, under conditions of monetary payoff the model would have to be generalized to account for such effects. We shall return to a discussion of this point later.

Time-order Effect

In the forced-choice detection task the terminology time-order effect is used to refer to the fact that subjects generally are more accurate in detecting signals embedded in the second observation interval than in the first interval. For example, on schedule B (which has S_1 and S_2 events occurring equally often), every subject had a higher probability of being correct when the signal was in the second interval than in the first interval; i.e., $Pr(A_1|S_1) < Pr(A_2|S_2)$. In terms of the present analysis there are two possible explanations for this time-order effect. One is that the bias parameter tends to favor the A_2 response. Hence when sensory state s_0

is activated, the subject makes the A_2 response more frequently, which insures that he will have a higher probability of being correct on S_2 than on S_1. Another possibility is that the time-order effect occurs because the sensitivity level changes from one observation interval to the next; specifically, that there are two sensitivity parameters σ_1 and σ_2 associated with the two intervals and that $\sigma_2 > \sigma_1$.

Thus a time-order effect can be accounted for by postulating a bias process that tends to favor the A_2 response, or by postulating a sensory mechanism that is more sensitive to stimuli presented in the second observation interval. The first argument explains the time-order effect in terms of the decision process, whereas the second accounts for the effect in terms of the activation process. Both of these explanations are tenable and one would like to have some rationale for selecting between them. Fortunately the model makes quite different predictions depending on which process is used to account for the time-order effect. If the explanation is in terms of the bias function (as was the case in our analysis of these data) then the ROC curve has slope 1 and the time-order effect is simply due to the fact that $\varphi > 1$. If, however, the effect is explained in terms of different sensitivity levels, then

$$\Pr(H) = \sigma_1 + (1-\sigma_1)p$$
$$\Pr(F) = (1-\sigma_2)p \ .$$

Under these conditions the ROC curve is

$$\Pr(H) = \frac{1 - \sigma_1}{1 - \sigma_2} \Pr(F) + \sigma_1 \ .$$

If $\sigma_2 > \sigma_1$, then the slope of the ROC curve is greater than one. Thus to decide whether the time-order effect is due to the bias process alone, or whether it also may be due to differential sensitivity levels, we must ask

whether the ROC curve has slope greater than one. One can see by inspection of Fig. 1 that for our subjects there is no evidence (except possibly for Subject 2) to suggest that the observed points would be better fit by a line with slope greater than one. Therefore, for this experiment, the conclusion is that the time-order effect is due to the bias process, and cannot be explained by changes in sensitivity over the two observation intervals.

IV. Blank Trials and False Information

In this section we consider two modifications of the typical forced-choice detection task. One involves the introduction of blank trials, and the other the use of false-information feedback. By blank trials we mean that on occasion a trial will occur on which the signal has been omitted entirely; the subject is not told that blank trials are being introduced and (because of the forced-choice nature of the task) continues to make A_1 and A_2 responses. A blank trial will be denoted as S_0. By false-information feedback we mean that on some trials the subject will be told that a signal occurred in a particular observation interval when in fact it did not. The introduction of these two modifications in the detection task permits us to make some very sharp predictions that differentiate this model from several others with similar assumptions.

In the study to be analyzed, the subject was given the same instructions that were used in the other experiment; i.e., he was led to believe that a signal would occur on every trial and that the information events reliably indicated the interval in which the signal occurred. Actually, however, the presentation schedule involved S_1, S_2 and S_0 type trials; on S_1 trials an E_1 always occurred, on S_2 trials an E_2 always occurred, and on S_0 trials sometimes E_1 occurred and sometimes E_2. The presentation schedule used in this study can be characterized by the parameters γ, π,

and x as follows: (a) with probability $x\gamma$ a signal was presented in the first interval and, after the response, E_1 occurred, (b) with probability $x(1-\gamma)$ a signal was presented in the second interval and followed by E_2, and (c) with probability $1-x$ a blank trial was presented and an E_1 occurred with probability π and an E_2 event with probability $1-\pi$. Thus, the probability of presenting a signal in the first interval was $x\gamma$; but the probability of telling the subject that the signal occurred in the first interval was

$$\Pr(E_{1,n}) = x\gamma + (1-x)\pi .$$

Similarly, the probability of presenting the signal in the second interval was $x(1-\gamma)$; however, the probability that the subject was told that the signal occurred in the second interval was

$$\Pr(E_{2,n}) = x(1-\gamma) + (1-x)(1-\pi) .$$

The model generalizes directly to this experiment. No new assumptions are necessary; we need only apply the axioms and carry out the appropriate derivations. First of all, consider the activation matrix for this experiment. In terms of the assumptions

$$\underset{\sim}{A}^* = \begin{array}{c} \\ S_1 \\ S_2 \\ S_0 \end{array} \begin{array}{c} \begin{array}{ccc} s_0 & s_1 & s_2 \end{array} \\ \begin{bmatrix} 1-\sigma & \sigma & 0 \\ 1-\sigma & 0 & \sigma \\ 1 & 0 & 0 \end{bmatrix} \end{array} .$$

Using the matrix $\underset{\sim}{A}^*$ and the decision matrix $\underset{\sim}{D}_n$ specified by Eq. 5, we can derive a performance matrix $\underset{\sim}{P}^*_n$ whose rows are the events S_1, S_2, and S_0 and whose columns are the responses A_1 and A_2. The entries in the first column of the matrix $\underset{\sim}{P}^*_n$ are

$$\Pr(H_n) = \Pr(A_{1,n}|S_{1,n}) = \sigma + (1-\sigma)p_n \tag{15a}$$

$$\Pr(F_n) = \Pr(A_{1,n}|S_{2,n}) = (1-\sigma)p_n \tag{15b}$$

$$\Pr(A_{1,n}|S_{0,n}) = p_n \ . \tag{15c}$$

By inspection of Eqs. 15a and 15b, it is clear that the ROC curve is the same as one given in Eq. 7 for the no-blank trial case. Also, from Eqs. 15a and 15c it follows that the function relating $\Pr(H_n)$ and $\Pr(A_{1,n}|S_{0,n})$ is a straight line with slope $1-\sigma$ and intercept σ; namely,

$$\Pr(H_n) = \sigma + (1-\sigma)\Pr(A_{1,n}|S_{0,n}) \ . \tag{16}$$

Now let us derive an expression for the response bias. Eq. 9 presents the axioms describing possible changes in p_n. These axioms are directly applicable to the experiment involving blank trials and false-information feedback. Given Eq. 9, we need only to compute the probability of the events $(s_{0,n} \ \& \ E_{1,n})$ and $(s_{0,n} \ \& \ E_{2,n})$.

The tree in Fig. 2 describes the possible events that can occur in a given trial. From the figure we obtain

$$\Pr(s_{0,n} \ \& \ E_{1,n}) = x\gamma(1-\sigma) + (1-x)\pi$$

and

$$\Pr(s_{0,n} \ \& \ E_{2,n}) = x(1-\gamma)(1-\sigma) + (1-x)(1-\pi) \ .$$

Given these results we can now derive $E(p_n)$. We shall not carry out the derivation, for it involves precisely the same arguments that were employed in developing Eq. 10. Invoking these arguments yields the following equation:

$$E(p_n) = p_\infty - [p_\infty - p_1]G^{n-1} \ .$$

Here

$$G = 1 - \theta[x\gamma(1-\sigma) + (1-x)\pi] - \theta'[x(1-\gamma)(1-\sigma) + (1-x)(1-\pi)] \ ,$$

and

$$p_\infty = \frac{x\gamma(1-\sigma) + (1-x)\pi}{[x\gamma(1-\sigma) + (1-x)\pi] + [x(1-\gamma)(1-\sigma) + (1-x)(1-\pi)]\varphi} \ . \tag{17}$$

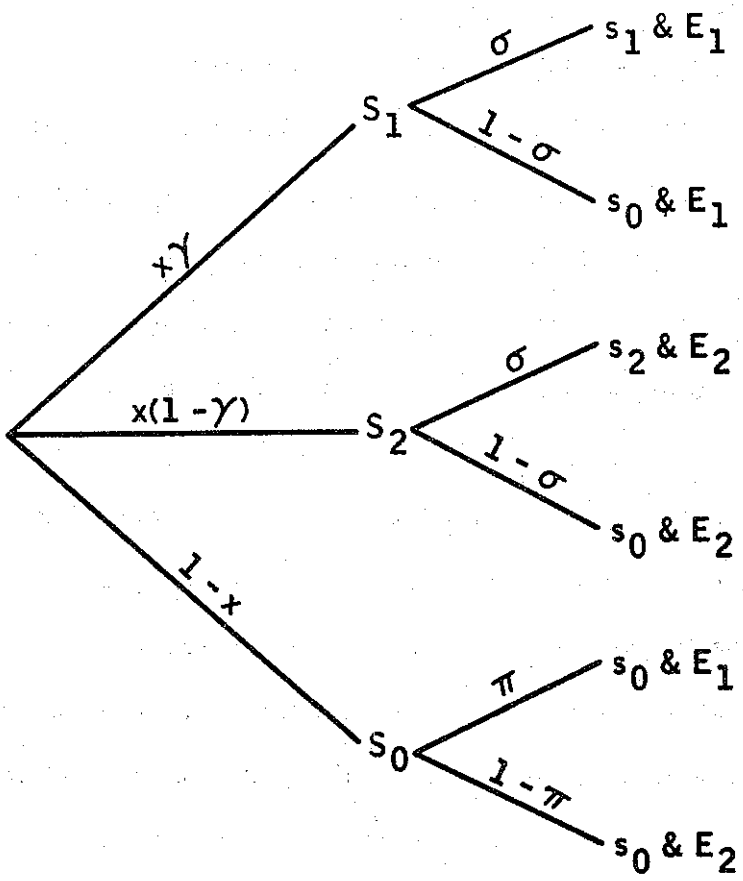

Figure 2. A tree describing possible events and their related probabilities for the blank-trial experiment

where $\varphi = \frac{\theta'}{\theta}$.

Empirical Analysis

We now examine some data from a forced-choice detection experiment that was run using presentation schedules involving blank trials and false-information feedback. The same experimental procedures were employed in this study as in the earlier one except for the pretraining phase. Pretraining lasted three days and involved running each subject on the schedule B routine used in the first experiment (during pretraining, a signal always occurred in one of the two observation intervals, information feedback was reliable, and $\gamma = .5$). The signal intensity was held fixed throughout the experiment, but the experimenter manipulated the noise level during pretraining in an attempt to establish a signal-to-noise ratio for each subject that yielded a correct response percentage of approximately 79; the theoretical rationale for selecting this particular value will be given later. The manipulation of the noise level was done strictly by trial and error, but the procedure proved to be quite successful for by the end of pretraining a level had been established for each subject that yielded a correct response probability fairly close to the desired value of .79. During the remainder of the experiment the noise level was fixed for each subject at the value determined during pretraining. Also, during pretraining any subject who tended to strongly favor one response over the other was eliminated from the experiment. Only subjects whose overall proportion of A_1 responses was between .40 and .60 for the second and third days of pretraining were included in the main experiment. Four subjects from a group of 18 were eliminated on this basis. Since $\gamma = \frac{1}{2}$ during pretraining, this selection procedure guaranteed that φ would be in the neighborhood of one for all subjects.

Pretraining, therefore, involved two special features: (a) noise levels were determined individually for each subject, and (b) subjects were eliminated from the experiment who showed a strong preference for one of the response alternatives. The first requirement guaranteed that the sensitivity parameter σ was approximately the same for all subjects. The second insured that φ was fairly close to 1 for all subjects. Thus, in a rough sense, a homogeneous group of subjects was formed by using this pretraining procedure; homogeneous in the sense that all subjects were characterized by approximately the same values of σ and φ.

In the experiment proper, four presentation schedules were used. The probability x of a signal trial was .50 for all schedules, but the schedules differed in the values of γ and π as follows:

Schedule A': $\gamma = .25, \ \pi = .25$

Schedule B': $\gamma = .75, \ \pi = .25$

Schedule C': $\gamma = .25, \ \pi = .75$

Schedule D': $\gamma = .75, \ \pi = .75$.

Test sessions of 400 trials were run on consecutive days. Each day a subject ran on one of the above presentation schedules for the entire session. In successive 4-day blocks a subject completed one day on each of the four schedules; within each 4-day block the order of schedules was randomly determined. The experiment involved 20 test sessions and therefore each schedule was repeated on five separate days.

Table 3 presents the observed average proportion of A_1 responses conditionalized on the various trial types; these averages are based on 14 subjects. Proportions were computed for each subject based on the last 350 trials of replications two through five of a given presentation schedule; thus the estimates for each subject are based on a sequence of $4 \times 350 = 1400$

Table 3

Observed and predicted values

for the blank-trial study

	Schedule A'		Schedule B'		Schedule C'		Schedule D'	
	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.
Pr(H)	.641	.672	.755	.734	.820	.820	.903	.886
Pr(F)	.086	.100	.174	.162	.227	.248	.344	.314
$Pr(A_1 \mid S_0)$	.213	.234	.401	.378	.553	.578	.765	.733
$Pr(A_1)$	.219	.238	.505	.485	.464	.484	.764	.738

trials. The averages of these individual subject proportions are the quantities presented in the table. Although data were analyzed on an individual subject basis in the first experiment, there are at least two justifications for presenting group averages now. One reason is that it greatly simplifies the analysis, and the second is that there is a theoretical rationale for treating group data in the present experiment. The rationale is based on the pretraining procedure, which was designed to insure that both σ and φ would be approximately the same for all subjects. By inspection of Eqs. 15 and 17 we see that the asymptotic expressions for $Pr(H)$, $Pr(F)$ and $Pr(A_1|S_0)$ depend on only σ and φ. If σ and φ are identical for all subjects, then the theory makes the same predictions for the group average as for individual subjects.

Figure 3 presents plots of the observed values of $Pr(H)$ and $Pr(F)$ as given in Table 3. The theory predicts that these points should fall on a linear curve with slope 1 and intercept σ. We estimated σ from these four data points by using the method of least squares and obtained

$$\hat{\sigma} = .572.$$

This estimate was used to generate the straight-line ROC curve displayed in Fig. 3. The four observed points (one from each schedule) fall quite close to the predicted line.

Figure 4 presents a plot of $Pr(A_1|S_0)$ versus $Pr(H)$. As indicated in Eq. 16, these points should be related by a linear function with slope $1-\sigma$ and intercept σ. Using our estimated σ, we generated the straight line in Fig. 4. Once again the linear relation seems to be reasonably well supported.

In order to make numerical predictions for $Pr(A_1|S_i)(i = 0, 1, 2)$ we need an estimate of φ. Estimation of this parameter is attained using the same method employed earlier. The overall probability of an A_1 response is

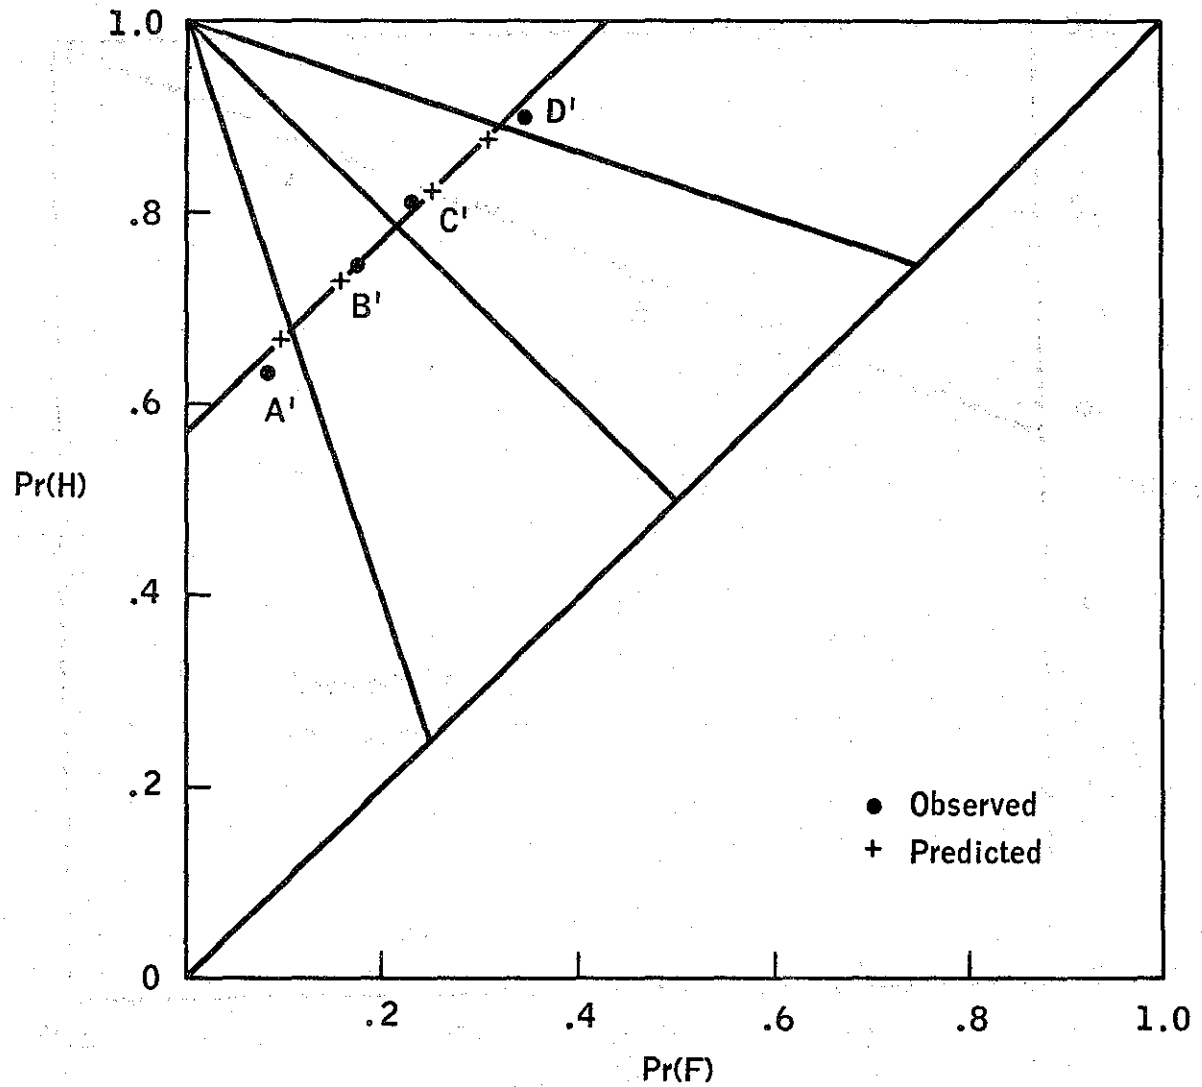

Figure 3. Observed and predicted values for Pr(H) and Pr(F)

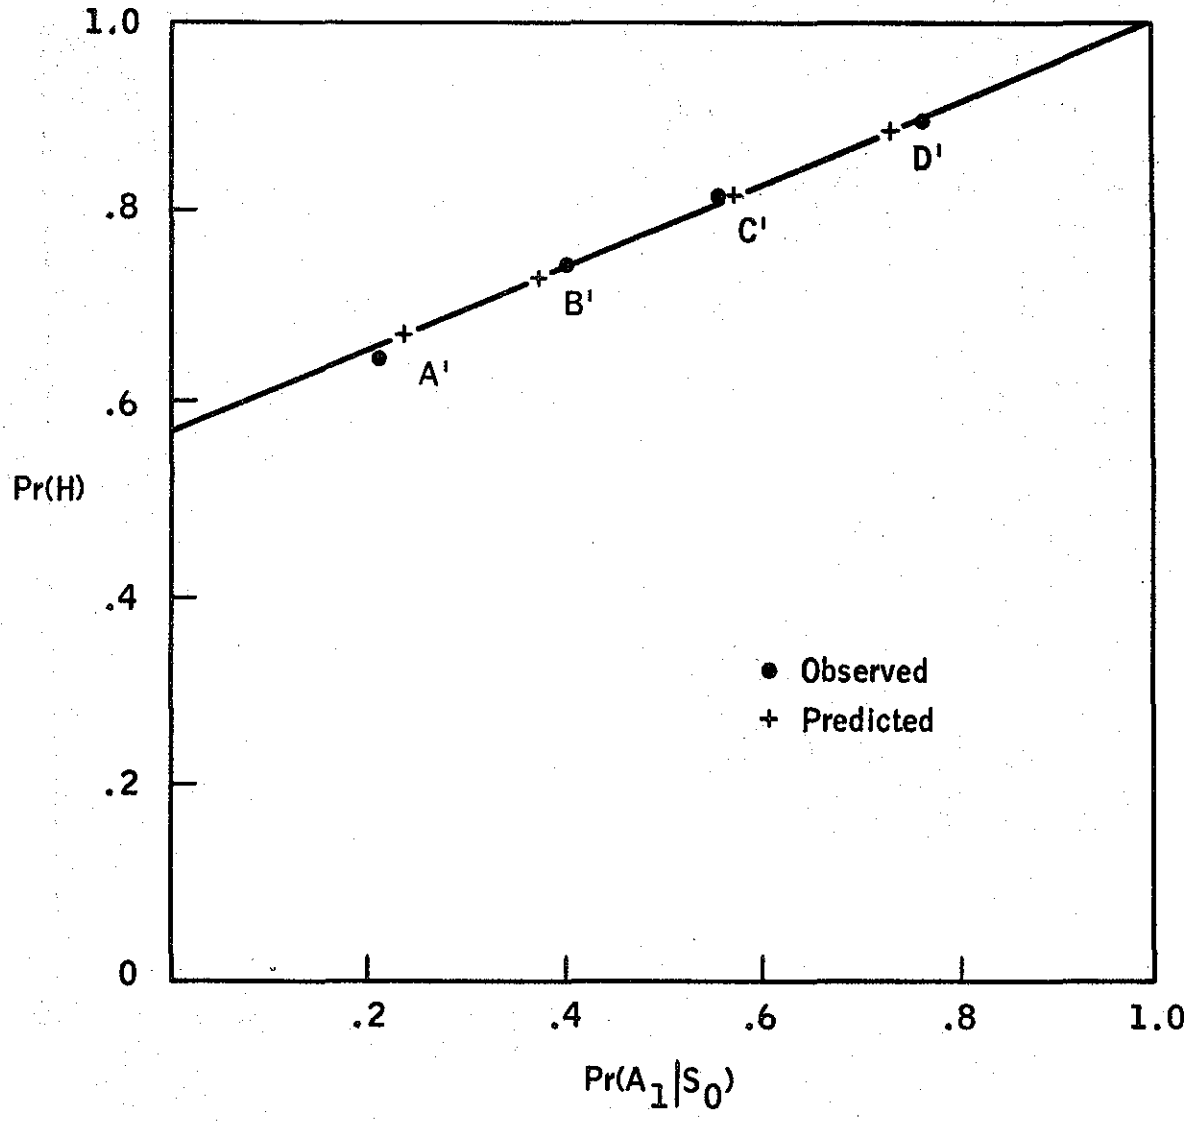

Figure 4. Observed and predicted values for Pr(H) and $Pr(A_1|S_0)$

$$\begin{aligned} \Pr(A_1) &= x\gamma\Pr(A_1|S_1) + x(1-\gamma)\Pr(A_1|S_2) + (1-x)\Pr(A_1|S_0) \\ &= \sigma x\gamma + (1-\sigma x)p_\infty \ . \end{aligned} \tag{18}$$

Substituting in the expression for p_∞ given in Eq. 17 yields an expression in φ. For each presentation schedule we have substituted the estimated value of σ and the observed value of $\Pr(A_1)$ in the above equation and solved for φ. For example, for schedule A' the observed value of $\Pr(A_1)$ is .219; letting $\hat{\sigma} = .572$, $\gamma = .25$, $\pi = .25$, and $\Pr(A_1) = .219$ in the above equation yields $\hat{\varphi}_{A'} = 1.281$. Similarly, for the other three schedules we obtain $\hat{\varphi}_{B'} = .969$, $\hat{\varphi}_{C'} = 1.229$, and $\hat{\varphi}_{D'} = .897$. It is interesting to note that $\hat{\varphi}$ seems to be correlated more with γ than with π. Schedules A' and C' ($\gamma = .25$) both yield $\hat{\varphi} > 1$, whereas schedules B' and D' ($\gamma = .75$) yield $\hat{\varphi} < 1$. Recalling that $\varphi = \frac{\theta'}{\theta}$ and that γ is the probability of a signal in the first interval (if there is a signal), these estimates suggest that $\theta' > \theta$ if the probability of the signal being in the second interval exceeds $\frac{1}{2}$. Hence the change in the bias parameter p_n seems to be somewhat dominated by the interval with the higher probability of bracketing the signal. Despite this suggestion of a departure from independence of the parameters φ and γ, very little damage is done to the accuracy of the model's predictions, as will be seen shortly.

To obtain an overall estimate of φ we have taken the average of these four values:

$$\begin{aligned} \hat{\varphi} &= \tfrac{1}{4}(\hat{\varphi}_{A'} + \hat{\varphi}_{B'} + \hat{\varphi}_{C'} + \hat{\varphi}_{D'}) \\ &= 1.094 \ . \end{aligned}$$

Using the estimates of σ and φ, Eqs. 15 and 17 can now be used to generate predictions for $\Pr(H)$, $\Pr(F)$, $\Pr(A_1|S_0)$ and $\Pr(A_1)$. These predicted quantities are given in Table 3; they also are graphically displayed in Figs. 3 and 4 as cross marks on the appropriate line segments. It should

be pointed out that there are no constraints on the relations among the quantities $Pr(A_1|S_1)$, $Pr(A_1|S_2)$ and $Pr(A_1|S_0)$, and therefore twelve independent predictions are being made on the basis of two parameters. One need only inspect the array of observed and predicted quantities to realize that the correspondence between theoretical and observed values is quite satisfactory.

Recall that for both schedules B' and C' the E_1 and E_2 events occurred equally often; i.e., on both schedules the subject was being told (via the trial-to-trial feedback) that the signal was occurring equally often in the two observation intervals. However, actually, the signal occurred more frequently in the first interval for schedule B' than for schedule C'. These experimental manipulations are clearly reflected in the data. On an S_0 trial the probability of an A_1 response was greater for schedule C' than for schedule B' (.553 vs. .401), whereas over all trials the probability of an A_1 response was greater for schedule B' than for schedule C' (.505 vs. .464). Both of these relations are predicted by the model.

V. Sequential Effects

So far, our analysis has been restricted to fairly gross aspects of the data. However, the model provides a deeper analysis of the experiment than the foregoing results indicate. From the model we can predict not only hit and false alarm rates but also the sequential properties of response protocols. In terms of the axioms, sequential effects in the observable response events are produced by trial-to-trial fluctuations in p_n. Such fluctuations, of course, can take place on any trial and are not restricted to pre-asymptotic data. For example, even at asymptote the likelihood of making a correct response to an S_1 stimulus depends in a very definite way on whether an E_1 or an E_2 occurred on the preceding trial.

The sequential effects of particular interest deal with the influence of stimulus and response events on trial n as they affect the response on trial $n+1$; specifically

$$\Pr(A_{1,n+1} | S_{i,n+1} A_{j,n} S_{k,n}) .$$

However, we shall not examine the correspondence between these particular sequential effects and theoretical predictions, because there are 18 such independent quantities for each experimental conditions and the analysis would involve too much detail. Rather, we shall consider $\Pr(A_{1,n+1} | E_{1,n})$ and $\Pr(A_{1,n+1} | E_{2,n})$. Note that for these probabilities the stimulus events on trials n and $n+1$ are suppressed, and we only ask for the overall likelihood of an A_1 response conditionalized on the information event of the preceding trial. The A_1 could occur in response to any of the stimulus events S_1, S_2, or S_0 on trial $n+1$; similarly the information event E_1 on trial n could follow an S_1 or S_0 stimulus, and the E_2 an S_2 or S_0 stimulus. Asymptotic expressions for these quantities can be readily obtained (see Atkinson, Bower, and Crothers, 1965) and are as follows:

$$\begin{aligned} \lim_{n \to \infty} \Pr(A_{1,n+1} | E_{1,n}) &= \Pr(A_1) + (1-\sigma x)\theta(1-p_\infty) \frac{\pi(1-x) + x\gamma(1-\sigma)}{\pi(1-x) + x\gamma} \\ \lim_{n \to \infty} \Pr(A_{1,n+1} | E_{2,n}) &= \Pr(A_1) - (1-\sigma x)\theta' p_\infty \frac{(1-\pi)(1-x) + x(1-\gamma)(1-\sigma)}{(1-\pi)(1-x) + x(1-\gamma)} \end{aligned} \tag{19}$$

where p_∞ is given by Eq. 17 and $\Pr(A_1)$ by Eq. 18.

Table 4 presents the observed values for $\Pr(A_{1,n+1} | E_{1,n})$ and $\Pr(A_{1,n+1} | E_{2,n})$. Estimates of these quantities were obtained for individual subjects; the average of these estimates are the quantities presented in the table. These estimates are based on the same set of trials as the data presented in Table 3 and therefore will be regarded as asymptotic. We can now use the above equations to yield predictions for these observed values. By

inspection of these equations, we see that values are needed for σ, θ, and θ' in order to make numerical predictions. Since estimates of σ and $\varphi = \frac{\theta'}{\theta}$ have already been made, it is only necessary to estimate either θ' or θ.

Suppose we fix on some value of θ'; then θ is determined because $\frac{\theta'}{\theta}$ must equal the previous estimate of $\varphi = 1.094$. For a fixed θ' predictions can be calculated for the eight statistics displayed in Table 4; these calculations are made for each experimental schedule by substituting the appropriate values of γ and π in the above equations, along with $\sigma = .572$, $\varphi = 1.094$, and $\theta = \frac{\theta'}{1.094}$. Once numerical predictions have been generated for a particular value of θ', an evaluation of the goodness-of-fit can be made by computing the sum of squared deviations between predicted and observed values; i.e., for a specific value of θ' define the quantity

$$S(\theta') = \sum (\text{predicted-observed})^2 ,$$

where the sum is over the 8 entries in Table 4.

One method for estimating θ' is to select its value so as to minimize $S(\theta')$. To carry out this minimization analytically yields unwieldy expressions, and to avoid this complication we have simply calculated $S(\theta')$ for θ' ranging from .01 to 1.0 in successive increments of .01. Over this range of values the function $S(\theta')$ takes on its minimum when $\theta' = .08$. This value of θ' generates the predicted quantities in Table 4.

By and large the correspondence between predicted and observed sequential statistics is reasonably good. In evaluating the goodness-of-fit it should be kept in mind that all of the quantities in Table 4 are independent, and thus there are 8 degrees of freedom. The model requires that $\Pr(A_{1,n+1}|E_{1,n} > \Pr(A_1) > \Pr(A_{1,n+1}|E_{2,n})$, and this relation is supported by all four sets of data.

Table 4

Observed and predicted sequential quantities

for the blank-trial study

	Schedule A'		Schedule B'		Schedule C'		Schedule D'	
	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.	Obs.	Pred.
$\lim_{n\to\infty} \Pr(A_{1,n+1} \mid E_{1,n})$	.255	.267	.529	.503	.475	.503	.784	.748
$\lim_{n\to\infty} \Pr(A_{1,n+1} \mid E_{2,n})$	.207	.229	.482	.466	.453	.466	.716	.708

Also the model requires that

$$\Pr(A_{1,n+1}|S_{1,n+1}E_{1,n}) > \Pr(A_{1,n+1}|S_{1,n+1}E_{2,n})$$
$$\Pr(A_{1,n+1}|S_{2,n+1}E_{1,n}) > \Pr(A_{1,n+1}|S_{2,n+1}E_{2,n})$$
$$\Pr(A_{1,n+1}|S_{0,n+1}E_{1,n}) > \Pr(A_{1,n+1}|S_{0,n+1}E_{2,n}) \quad .$$

Although not presented here, a breakdown of the data into this form indicates that these inequalities hold over all four experimental conditions.

VI. Alternative Models

An alternative model for the bias process that initially appealed to us involved trial-to-trial changes in p_n that were determined solely by the information events E_1 and E_2. Formally stated, the idea was that

$$p_{n+1} = \begin{cases} (1-\theta)p_n + \theta \ , & \text{if } E_{1,n} \\ (1-\theta')p_n \ , & \text{if } E_{2,n} \ . \end{cases}$$

This formulation (which will be called Model II) is to be contrasted with Eq. 9 (Model I), where changes in p_n can occur only when sensory state s_0 is activated. In spite of the marked difference between these two sets of assumptions, the models yield identical predictions in the first experiment for the asymptotic probabilities of $\Pr(H)$, $\Pr(F)$, $\Pr(A_1)$ and $\Pr(C)$. Only by a detailed analysis of sequential statistics and pre-asymptotic data can it be shown that Model I is slightly better than Model II.

However, the two models make strikingly different predictions in the false-information study even for asymptotic hit and false alarm proportions. For example, applying Model II to the false-information study yields

$$p_\infty = \frac{x\gamma + (1-x)\pi}{[x\gamma + (1-x)\pi] + [x(1-\gamma) + (1-x)(1-\pi)]\varphi} \quad .$$

By inspection of this equation, we see that p_∞ is identical for both schedules B' and C' of the second experiment; whereas, using Model I, p_∞ is greater for schedule C' than for schedule B'. This relation, of course,

is reflected in Pr(H) and Pr(F). For Model II, Pr(H) and Pr(F) will be the same for both schedules B' and C'; whereas for Model I both Pr(H) and Pr(F) will be greater for schedule C' than for schedule B'. The ordering relation predicted by Model I for schedules B' and C' is borne out by the group averages presented in Table 3; it also is the case that the relation holds individually for all 14 subjects. Therefore Model I, but not Model II, appears substantiated by the data.

To further illustrate the relations between Models I and II in the false-information study, we have presented iso-bias curves in Fig. 5; the curves were plotted for $\varphi \approx 1$. By inspection of the figure we see that the iso-bias curve for Model II is a straight line for all four presentation schedules, and the iso-bias curves for schedules B' and C' are identical. For Model I, the iso-bias curves for schedules A' and D' are the same as for Model II; however, under the assumptions of Model I schedules B' and C' generate different, nonlinear curves.

Adopting Model I, a distance function can be defined between corresponding points on the iso-bias curves for schedules B' and C'. The maximum of this function can be obtained by taking its derivative with respect to σ and setting the result equal to zero. Carrying out these operations yields

$$\sigma = 2 - \sqrt{2} \approx .59 .$$

Therefore, under the assumptions of Model I, the maximum difference between corresponding points on the iso-bias functions of schedules B' and C' will be observed when σ is approximately .59.

One of the principal reasons for running the false-information study was to determine whether such a difference would be observed. Therefore, to maximize the likelihood of discovering an effect if it existed, we wanted to set the noise level at a value corresponding to a σ of .59. Recall that

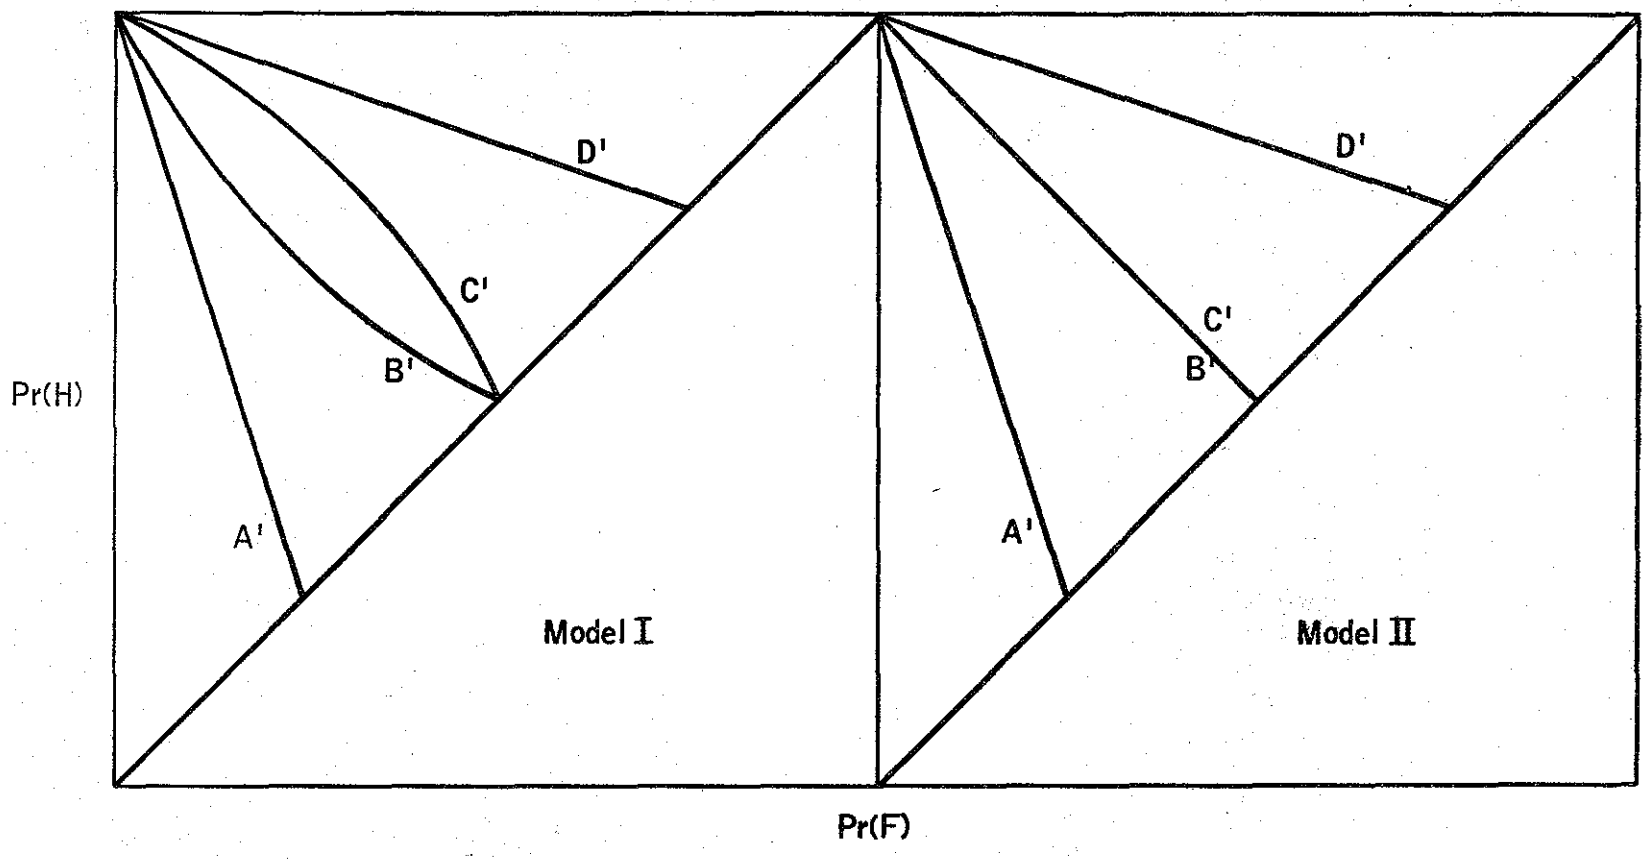

Figure 5. Iso-bias curves for models I and II

pretraining involved only S_1 and S_2 trials, and they were presented with equal likelihood; hence $Pr(C) = \sigma + (1-\sigma)\frac{1}{2}$. Consequently to fix σ at approximately .59 required adjusting the noise level during pretraining to yield a correct-response probability of approximately $.79 \cong .59 + (.41)\frac{1}{2}$. The pretraining procedure was fairly successful, inasmuch as the estimate of σ during the actual experiment was .572.

VII. Concluding Remarks

The applications of the model presented here have been confined to symmetric outcome structures involving no payoffs. If we were to generalize the model to situations involving manipulation of monetary payoffs then it would be necessary to offer a more complex theory of the decision process. Obviously there are outcome structures that will displace the subject's data point off the linear ROC curve specified by Eq. 7.[2] For example, consider the following payoff matrix:

	A_1	A_2
S_1	-1	+100
S_2	+100	-1

In this case the subject is heavily rewarded for incorrect detection responses and penalized for correct responses. Hence over time the subject would undoubtedly generate a point $[Pr(F), Pr(H)]$ that fell in the lower right-hand sector of the ROC space. That is, the probability of a false alarm would exceed the probability of a hit for this outcome structure. It is important to note that such effects cannot be predicted merely by generalizing the

[2] In fact, even for experiments discussed here, it is likely that the observed point $[Pr(F), Pr(H)]$ will fall below the predicted ROC curve when γ is close to zero or one (Atkinson, 1963).

assumptions governing p_n. No matter how p_n is permitted to vary, the model still requires that performance points fall on a linear curve with intercept σ.

Of course, several modifications of the theory seem able to account for experimental manipulations that generate performance points off the ROC curve. One approach is to develop a more elaborate conceptualization of the decision process. For example, one can replace the $\underset{=}{D}_n$ matrix of Eq. 5 with the matrix

$$\underset{=}{D}_n = \begin{array}{c} \\ s_0 \\ s_1 \\ s_2 \end{array} \begin{array}{c} \begin{array}{cc} A_1 & A_2 \end{array} \\ \begin{bmatrix} p_n & 1-p_n \\ d_n^{(1)} & 1-d_n^{(1)} \\ 1-d_n^{(2)} & d_n^{(2)} \end{bmatrix} \end{array} .$$

For this process experimental manipulations of the outcome structure might not only affect p_n but also the values of $d_n^{(i)}$. Thus, depending on the postulated relation of $d_n^{(i)}$ to the outcome structure, it would be possible to generate virtually any ROC curve. Of course, when this type of modification is introduced one obtains a model that is very close in structure to those proposed for discrimination learning (Atkinson and Estes, 1963, p. 238; Bush, Luce and Rose, 1964). Another possible modification of the detection model would be to develop a more general formulation of the sensory process. Pursuing this line, we might assume that the subject's sensitivity level could vary within certain limits as a function of the outcome structure and other variables.

Both of these alternatives represent potential lines of theoretical development for models of this type. They raise an important question: Can changes in performance induced by manipulation of the outcome structure be explained by elaborating the theory of the bias process, or do they also necessitate

postulating a more complex sensory mechanism? Developments of this sort are fairly complex and go beyond the scope of this paper.

References

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

Atkinson, R. C., Bower, G. H., and Crothers, E. J. (1965). An introduction to mathematical learning theory. New York: Wiley. (In press.)

Atkinson, R. C., Carterette, E. C., and Kinchla, R. A. (1962). Sequential phenomena in psychophysical judgments: a theoretical analysis. Inst. Radio Engineers Trans. on Information Theory, IT-8, S155-162.

Atkinson, R. C., and Estes, W. K. (1963). Stimulus sampling theory. In R. D. Luce, R. R. Bush, and E. Galanter (Eds.), Handbook of mathematical psychology, Vol. II. New York: Wiley. Pp. 121-268.

Bush, R. R., Luce, R. D., and Rose, R. M. (1964). Learning models for psychophysics. In R. C. Atkinson (Ed.), Studies in mathematical psychology. Stanford: Stanford Univer. Press. Pp. 201-217.

Bush, R. R., and Mosteller, F. (1955). Stochastic models for learning. New York: Wiley.

Fechner, G. T. (1860). Elemente der Psychophysik. Leip., Breitkoph and Härtel.

Friedman, M. P., and Carterette, E. C. (1964). Detection of Markovian sequences of signals. J. acoust. soc. Amer., 36, 2334-2339.

Green, D. M. (1960). Psychoacoustics and detection theory. J. Acoust. Soc. Amer., 32, 1189-1203. Reprinted in R. D. Luce, R. R. Bush and E. Galanter (Eds.), Readings in mathematical psychology, Vol. I. New York: Wiley, 1963. Pp. 41-66.

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

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

Swets, J. A. (1961). Detection theory and psychophysics: a review. Psychometrika, 26, 49-63.

Tanner, W. P., Jr., and Swets, J. A. (1954). A decision making theory of visual detection. Psychol. Rev., 61, 401-409.