ORGANIZATIONAL BEHAVIOR AND HUMAN PERFORMANCE 24, 111–125 (1979)

Detection of Change in Nonstationary, Random Sequences

DONALD M. BARRY

Rutgers University

AND

GORDON F. PITZ

Southern Illinois University at Carbondale

In each of two experiments, subjects were presented with randomly generated sequences of binary information. At some point during each sequence, a change in the probabilities of the events occurred; the task was to detect this change Late detections resulted in a loss proportional to the number of events between the occurrence of the change and detection, while false alarms were penalized a constant amount. Experiment I used two levels of task difficulty and two payoff conditions; Experiment II employed a single level of task difficulty and five payoff conditions. Subjects' detections were affected by payoffs in the appropriate direction, but not to as great a degree as was the optimal Bayesian model. Two parameter-free descriptive models of behavior were examined: a fixed sample size model, which postulated that subjects observe a predetermined number of items, was rejected; a critical odds model, assuming that subjects use as a criterion the probability that change has occurred, received some support, except for a slight tendency for subjects to relax their criterion as sample size increased. A simple heuristic model of performance in the task was examined, and implications of alternative utility assumptions were explored.

Experiments in risky decision making can be classified according to whether stationary or nonstationary probabilities govern various experimental events. In the stationary task, events are generated by probabilities that remain fixed over time; the familiar bookbag and poker chip experiment, in which subjects make inferences about populations with fixed characteristics, is an example. On the other hand, the nonstationary task involves probabilities that may change over time; it closely parallels many situations that confront decision makers in real-world settings, but has received relatively little attention. The paradigm of the present study could, in theory, be extended to model the task of, say, a clinician attempting to detect personality changes on the basis of periodic observa-

This research was supported by The National Science Foundation through Grants GS-2616 and GB-28708X to the second author.

Requests for reprints should be sent to Donald M. Barry, School of Criminal Justice, Rutgers University, 15 Washington Street, Newark, NJ 07102.

0030-5073/79/040111-15$02.00/0

tions of behavior. Since the human perceptual system appears to be constructed to detect changes more easily than identify steady states, it is worth studying the detection of change at the cognitive level also.

Prior studies have yielded mixed results concerning performance in nonstationary tasks. Some have concluded that subjects deal quite effectively with nonstationarity (Rapoport, 1964; Robinson, 1964; Pitz, 1966; Chinnis & Peterson, 1970). However, Theios, Brelsford, and Ryan (1971) concluded that their subjects performed poorly when required to estimate transition points of binary sequences in which the two generating probabilities differed only slightly. And, Brown and Bane (1975), in a nonstationary probability estimation task, found systematic overestimation of increasing probabilities and underestimation of decreasing ones. Burkheimer and Rapoport (1971) have presented a thorough treatment of detection of change from a normative point of view, and their paper contains a detailed discussion of the important parameters in such a task.

In the present study, a decision maker is presented with a sequence of binary information, whose elements we arbitrarily term R (red) and B (blue). Initially, the probability of an R item is fixed at some value P_0; sometime during the course of the sequence, this probability changes to a new value P_1. The decision maker's task is to detect the shift on the basis of the sequence of items he has seen; the sequence continues until a detection response occurs. As an example of the detection of change task, consider a quality-control inspector who must examine the output of a manufacturing device. He knows that the machine's normal output results in, say, a 20% rate of defective items ($P_0 = .20$). At some point in time, the defective rate shifts to 60% ($P_1 = .60$), remaining there until a stopping decision is made. Assuming that the inspector must base his decisions on samples of the machine's output taken at regular intervals, under what conditions should he halt production so that the machine can be checked?

In order to answer this question, a payoff structure must be imposed on the task. We note that two types of errors are possible: first, the change may not yet have taken place when stopping occurs (false alarm); in this case, the decision maker is penalized a constant amount F. On the other hand, he may stop after the change has occurred (late detection); here, the penalty is proportional to the number of observations between the change and the decision to stop; i.e., the penalty for a late detection is $w(T - t)$, where T is the number of observations taken before stopping, t is the number of observations prior to the occurrence of the change, and w is a payoff parameter representing the penalty per post-change observation.

It is also necessary to specify a probability distribution that governs the time of occurrence of the change. A geometric distribution was chosen for this purpose; it has the advantage of providing a conceptually random distribution of change times, in that the probability of the change occur-

ring before the next observation, if it has not ye
The geometric distribution is characterized by a
resenting this probability. Hence, the probabilit
immediately before observation n is given by:

$$p_n = \alpha (1 - \alpha)^{n-1},$$

and the probability of the change occurring at sc
tion n is:

$$P_n = 1 - (1 - \alpha)^n.$$

Finally, the probabilities of an R item both befo
and P_1) must be specified; the respective proba
$(1 - P_0)$ and $(1 - P_1)$.

Once the five parameters above (F, w, α, P
find the optimal strategy, i.e., the rule prescrib
decision maker must stop if he wishes to min
The strategy employs Bayes' Theorem to calc
tion, a revised probability that the change ha
probability depends only on the parameters
quence of R and B items that have been ob
probability is independent of any payoff consi

Pollock (1967) has shown that there exists
having the following property: whenever the B
ity favoring change equals or exceeds q, it is
decision; as long as the posterior probability is
tions should be taken. The value of q depends,
all five parameters noted above. When there
total number of observations that can be taken,
set of parameters, i.e., does not vary as a
observations. Pollock (1967) has also provid
mining the value of q, employing the technique
(see Bellman, 1957).

Given an optimal strategy for the detection o
present themselves. What is the relationship b
the optimal model and the decision behavior
the critical parameters affect subjects in the sa
model, and to the same degree? If not, is it po
that can describe behavior? The experiments r
to examine the effects of two parameters upo
models of decision strategies for the detectio

Several studies of decision making in sta
noted a failure of subjects to respond appropr
payoffs. When payoffs have been manipulate

e the human perceptual system appears to be con-
nges more easily than identify steady states, it is
ection of change at the cognitive level also.
ielded mixed results concerning performance in
me have concluded that subjects deal quite effec-
rity (Rapoport, 1964; Robinson, 1964; Pitz, 1966;
70). However, Theios, Brelsford, and Ryan (1971)
jects performed poorly when required to estimate
ary sequences in which the two generating prob-
slightly. And, Brown and Bane (1975), in a
ty estimation task, found systematic overestima-
abilities and underestimation of decreasing ones.
ort (1971) have presented a thorough treatment of
om a normative point of view, and their paper
ssion of the important parameters in such a task.
a decision maker is presented with a sequence of
ose elements we arbitrarily term R (red) and B
bability of an R item is fixed at some value P_0;
rse of the sequence, this probability changes to a
ion maker's task is to detect the shift on the basis
ns he has seen; the sequence continues until a
rs. As an example of the detection of change task,
ol inspector who must examine the output of a
e knows that the machine's normal output results
efective items (P_0 = .20). At some point in time,
o 60% (P_1 = .60), remaining there until a stopping
ing that the inspector must base his decisions on
s output taken at regular intervals, under what
production so that the machine can be checked?
question, a payoff structure must be imposed on
wo types of errors are possible: first, the change
place when stopping occurs (false alarm); in this
is penalized a constant amount F. On the other
he change has occurred (late detection); here, the
the number of observations between the change
i.e., the penalty for a late detection is $w(T - t)$,
of observations taken before stopping, t is the
prior to the occurrence of the change, and w is a
nting the penalty per post-change observation.
pecify a probability distribution that governs the
change. A geometric distribution was chosen for
advantage of providing a conceptually random
nes, in that the probability of the change occur-

ring before the next observation, if it has not yet occurred, is a constant. The geometric distribution is characterized by a single parameter, α, representing this probability. Hence, the probability of the change occurring immediately before observation n is given by:

$$p_n = \alpha (1 - \alpha)^{n-1}, \tag{1}$$

and the probability of the change occurring at some time prior to observation n is:

$$P_n = 1 - (1 - \alpha)^n. \tag{2}$$

Finally, the probabilities of an R item both before and after the change (P_0 and P_1) must be specified; the respective probabilities of a B item are then $(1 - P_0)$ and $(1 - P_1)$.

Once the five parameters above (F, w, α, P_0, P_1) are known, one can find the optimal strategy, i.e., the rule prescribing that point at which the decision maker must stop if he wishes to minimize his expected losses. The strategy employs Bayes' Theorem to calculate, after each observation, a revised probability that the change has occurred. This posterior probability depends only on the parameters α, P_0, and P_1, and the sequence of R and B items that have been observed; i.e., the posterior probability is independent of any payoff considerations.

Pollock (1967) has shown that there exists a critical probability, q, having the following property: whenever the Bayesian posterior probability favoring change equals or exceeds q, it is optimal to make a stopping decision; as long as the posterior probability is less than q, more observations should be taken. The value of q depends, of course, on the values of all five parameters noted above. When there is no constraint upon the total number of observations that can be taken, q is constant for any fixed set of parameters, i.e., does not vary as a function of the number of observations. Pollock (1967) has also provided an algorithm for determining the value of q, employing the techniques of dynamic programming (see Bellman, 1957).

Given an optimal strategy for the detection of change, several questions present themselves. What is the relationship between the prescriptions of the optimal model and the decision behavior exhibited by subjects? Do the critical parameters affect subjects in the same way that they affect the model, and to the same degree? If not, is it possible to construct a model that can describe behavior? The experiments reported here were designed to examine the effects of two parameters upon performance, and to test models of decision strategies for the detection of change task.

Several studies of decision making in stationary environments have noted a failure of subjects to respond appropriately to modification of the payoffs. When payoffs have been manipulated so as to favor particular decision alternatives, strategies have shifted in the correct direction, but not to the extent required by an optimal model (Pitz & Downing, 1967; Pitz & Reinhold, 1968; Ulehla, 1966). In other words, performance has been more optimal when payoffs induced no bias, and responses could be made on the basis of the posterior probabilities alone. For the detection of change task, it is not clear exactly what might constitute an unbiased set of payoffs, since there is no symmetry between the two kinds of costs incurred.

General descriptive models of subjects' strategies in a dynamic decision task may be tested. The models and the tests of their adequacy are based on two descriptive models for a deferred decision making task, proposed by Pitz, Reinhold, and Geller (1969), and are concerned with the general form of strategies that subjects may adopt. The critical odds model states that the decision to stop occurs when the odds favoring change (equivalently, the probability favoring change) reaches some predetermined cutoff level, which may itself vary from trial to trial. The optimal strategy is a special case of the critical odds model, in which the critical odds is determined by the optimal stopping probability q. Thus, a subject could perform in accordance with the critical odds model without necessarily performing optimally.

The second model tested was a fixed sample size model, which specifies that subjects base their decisions to stop solely on the number of observations that have been taken. While Pitz *et al.* (1969) found that a modification of the fixed sample size model approximated behavior in a stationary information-purchase task, we did not view the fixed sample size model as a serious candidate for a descriptive model in the present nonstationary task. Rather, it was included in order to demonstrate that an implausible model can indeed be rejected via the analytic techniques we employed.

Next, a simple heuristic model requiring two parameter estimates per subject was examined, and finally, the effects of (1) a risk-averse utility assumption and (2) a utility for being correct assumption, were explored.

EXPERIMENT I

Method

Design and subjects. Two determinants of the optimal strategy, task difficulty and penalty for a false alarm, were varied in a between-subjects factorial design. Task difficulty is inversely related to the difference between the pre- and postchange probabilities, P_0 and P_1. Under the easier condition, these values were 0.2 and 0.6, respectively; under the more difficult condition they were 0.4 and 0.6. The penalties for a false alarm (F) were \$7 and \$14, in token money. For all groups, the probability of change conditional on its having not yet occurred (α) was 0.1, and the ategies have shifted in the correct direction, but d by an optimal model (Pitz & Downing, 1967; Jlehla, 1966). In other words, performance has payoffs induced no bias, and responses could be posterior probabilities alone. For the detection of r exactly what might constitute an unbiased set no symmetry between the two kinds of costs

dels of subjects' strategies in a dynamic decision nodels and the tests of their adequacy are based s for a deferred decision making task, proposed eller (1969), and are concerned with the general ojects may adopt. The critical odds model states occurs when the odds favoring change (equiv- 'avoring change) reaches some predetermined self vary from trial to trial. The optimal strategy itical odds model, in which the critical odds is l stopping probability q. Thus, a subject could ith the critical odds model without necessarily

sted was a fixed sample size model, which e their decisions to stop solely on the number of en taken. While Pitz *et al.* (1969) found that a sample size model approximated behavior in a rchase task, we did not view the fixed sample andidate for a descriptive model in the present r, it was included in order to demonstrate that indeed be rejected via the analytic techniques

model requiring two parameter estimates per d finally, the effects of (1) a risk-averse utility for being correct assumption, were explored.

EXPERIMENT I

wo determinants of the optimal strategy, task false alarm, were varied in a between-subjects culty is inversely related to the difference be- ange probabilities, P_0 and P_1. Under the easier ere 0.2 and 0.6, respectively; under the more re 0.4 and 0.6. The penalties for a false alarm ken money. For all groups, the probability of having not yet occurred (α) was 0.1, and the

penalty per observation for a late detection (w) was \$1. Forty male students enrolled in an introductory psychology course served as subjects in partial fulfillment of a course requirement. Each subject was randomly assigned to one of the four conditions, and served in a 50-min session. Subjects were told that their payment would depend upon how well they performed in terms of the token money; actual payments averaged approximately \$1.70.

Apparatus. The stimuli consisted of two lamps, one red and one blue, mounted on a control panel in front of the subject. He began each sequence by pressing a start button, whereupon the stimuli were presented at the rate of approximately 40 per min; when he wished to indicate that the change had occurred, he pressed a stop button. There was no limit on the number of observations that might be taken. At the end of a sequence one of two state indicators lighted, informing the subject whether the change had in fact taken place. The amount lost for that decision was displayed on a counter, and another counter displayed the total amount lost thus far. A third meter displayed the total number of stopping decisions that had been made. A card mounted on the panel listed the values of P_0, P_1, F, and w. Stimuli and occurrences of the change were generated randomly, according to the specified parameters, by a solid state probability unit. All events were automatically recorded on punched paper tape.

Procedure. At the beginning of the experimental session, the subject was given a booklet of instructions, which served to acquaint him with the concept of probabilistic information and its relevance to the detection of change task. He was instructed to imagine himself in the position of a quality control inspector whose job was to decide whether or not a manufacturing device was functioning properly, with emphasis on the fact that his job was to minimize the amount of token money lost for each stopping decision made. The various parameters involved were explained in the context of this situation, and the functions of the control panel were explained. Any questions he had about the task were answered, and several practice trials were run before the experiment proper began. Each subject was then run for 25 min; the number of sequences averaged about 80. Subjects were told how long the sessions would last, and that their payment would depend only upon the average token money loss per stopping decision; this was done to eliminate inappropriate strategies such as performing slowly in order to minimize total amount lost, or performing quickly in order to leave the session early.

Data analysis. The probability favoring change at the time each decision was made and the number of information items observed for each decision were obtained. For purposes of comparison, the same measures were obtained for the performance of the optimal model. When actual behavior is compared with optimal, a problem arises if the subject decides to stop earlier than the optimal model would; since the sequence is terminated, the outcome of the model's decision is unknown. To evaluate the model's behavior in this situation, a Monte Carlo technique was employed; each time the subject stopped before the model, the sequence was later continued via computer simulation until the model specified that stopping should occur. The values of the probability favoring change when stopping occurred were transformed into log odds (log Ω) measures according to:

$$\log \Omega = \log \frac{P}{(1 - P)},$$

where P is the probability favoring change.

Results

The mean log Ω for the subjects' decisions was compared with mean log Ω for the optimal model's decisions, for the same sequences of information as seen by the subjects. A repeated measures comparison of the difference between observed and optimal mean log stopping odds was significant, $F(1,39) = 12.82$; the likelihood ratio, L, in favor of the null hypothesis (Jeffreys, 1961) was .0062, indicating a clear discrepancy between observed and optimal stopping points.[1] Table 1 shows group means for the four conditions; generally, subjects tended to wait too long before stopping, although in one condition ($P_0 = .2$, $F = \$14$) their decisions tended to come too early. An analysis of variance, restricted to the left half of Table 1, indicated a significant effect due to Task Difficulty, $F(1,36) = 6.95$, $L = .343$. The effect due to False Alarm Penalty was not significant, nor was the Difficulty by Penalty interaction.

Next, the fixed sample size and critical odds models were examined (see Pitz *et al.*, 1969, for a discussion of a similar technique for evaluating the two models in an information seeking task). According to the fixed sample size model, subjects base their decisions to stop only upon the number of items seen since the beginning of a sequence. If the model were correct, the actual value of the criterion (i.e., number of items before stopping occurs) would vary from subject to subject, and might fluctuate within individuals as well; a test of the model must allow for such variability. If subjects do employ a fixed sample size strategy, then the terminal value of log Ω becomes a random variable, with expected value and variance that depend only on n.

Expected values of log Ω conditional on $n[E(\log \Omega|n)]$ were estimated

[1] The likelihood ratio is the ratio of the probability of the data given the null hypothesis to the probability of the data given all alternatives to the null. Values less than 1.0 indicate evidence against the null hypothesis.

TABLE 1

OBSERVED AND OPTIMAL VALUES OF MEAN LOG Ω AS A FUNCT AND PENALTY FOR A FALSE ALARM

	Observed		
	$P_0 = .2$	$P_0 = .4$	
$F = \$7$	.572	.414	$F = \$7$
$F = \$14$	.575	.459	$F = \$14$

by 2000 Monte Carlo simulations for each value of These are compared with observed means, poolec across payoff conditions, in Fig. 1. The dotted lin 95% confidence intervals around $E(\log \Omega|n)$ and w basis of the number of decisions made for the given if the model is correct, all values of log Ω are indep even within the data for a single subject. Means values of n at which 10 or more individual decisio based on less than 10 decisions were considered t sion. It can be seen that the trend of observed mea

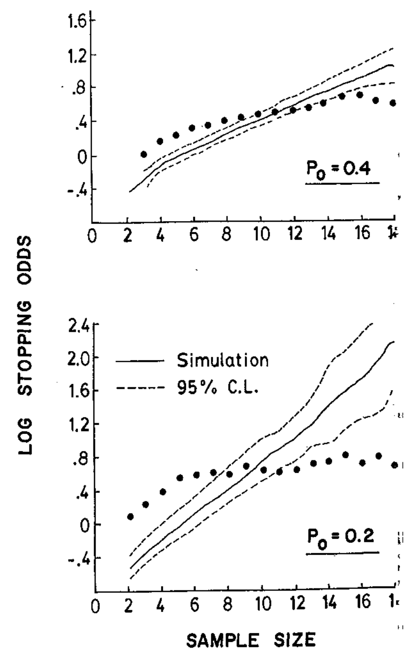

FIG. 1. Comparison between expected mean log stopping od size model (solid line), and actual mean log stopping odds e Broken lines indicate 95% confidence limits about the means c

optimal, a problem arises if the subject decides
imal model would; since the sequence is termi-
model's decision is unknown. To evaluate the
s situation, a Monte Carlo technique was
ubject stopped before the model, the sequence
nputer simulation until the model specified that
he values of the probability favoring change
ere transformed into log odds (log Ω) measures

$$g\ \Omega = \log \frac{P}{(1-P)},$$

favoring change.

ıbjects' decisions was compared with mean log
decisions, for the same sequences of informa-
cts. A repeated measures comparison of the
ed and optimal mean log stopping odds was
2; the likelihood ratio, L, in favor of the null
was .0062, indicating a clear discrepancy be-
l stopping points.[1] Table 1 shows group means
erally, subjects tended to wait too long before
condition ($P_0 = .2$, $F = \$14$) their decisions
An analysis of variance, restricted to the left
a significant effect due to Task Difficulty,
'he effect due to False Alarm Penalty was not
ficulty by Penalty interaction.
ize and critical odds models were examined
iscussion of a similar technique for evaluating
mation seeking task). According to the fixed
s base their decisions to stop only upon the
the beginning of a sequence. If the model were
f the criterion (i.e., number of items before
y from subject to subject, and might fluctuate
a test of the model must allow for such vari-
y a fixed sample size strategy, then the termi-
a random variable, with expected value and
on n.
conditional on $n[E(\log \Omega | n)]$ were estimated

of the probability of the data given the null hypothesis to
ll alternatives to the null. Values less than 1.0 indicate
ıs.

TABLE 1

Observed and Optimal Values of Mean Log Ω as a Function of Task Difficulty and Penalty for a False Alarm

	Observed			Optimal	
	$P_0 = .2$	$P_0 = .4$		$P_0 = .2$	$P_0 = .4$
$F = \$7$	.572	.414	$F = \$7$	.351	.002
$F = \$14$	.575	.459	$F = \$14$	.694	.340

by 2000 Monte Carlo simulations for each value of n from 1 through 25. These are compared with observed means, pooled across subjects and across payoff conditions, in Fig. 1. The dotted lines in Fig. 1 represent 95% confidence intervals around $E(\log \Omega | n)$ and were calculated on the basis of the number of decisions made for the given value of n. Note that, if the model is correct, all values of log Ω are independent of each other, even within the data for a single subject. Means are plotted for those values of n at which 10 or more individual decisions were made; means based on less than 10 decisions were considered too unstable for inclusion. It can be seen that the trend of observed mean log stopping odds is

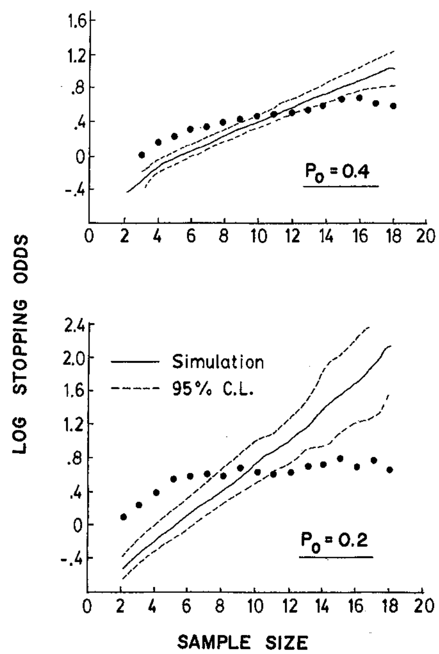

Fig. 1. Comparison between expected mean log stopping odds assuming a fixed sample size model (solid line), and actual mean log stopping odds exhibited by subjects (dots). Broken lines indicate 95% confidence limits about the means obtained via simulation.

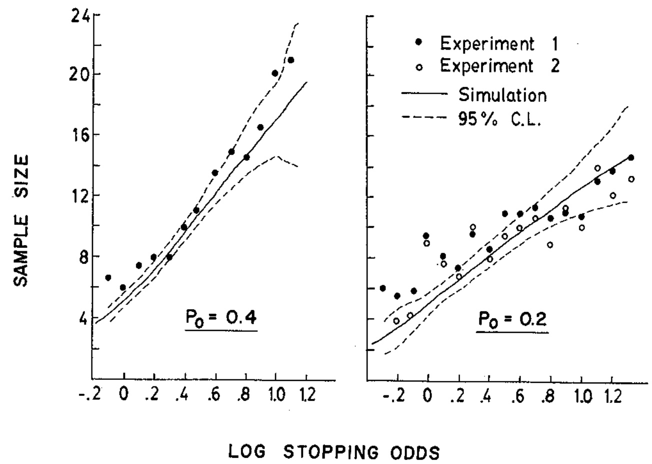

FIG. 2. Comparison between expected mean sample sizes assuming a critical odds model (solid line), and actual mean sample sizes exhibited by subjects (dots). Broken lines indicate 95% confidence limits about the means obtained via simulation.

not in accordance with the predictions of the fixed sample size model. The conclusion is that subjects were not solely influenced (if they were influenced at all) by the number of items appearing in a sequence.

A similar test was carried out for the critical odds model, which predicts that subjects decide to stop when log Ω (or, equivalently, probability favoring change) reaches or exceeds some predetermined level. If this model is correct, then n becomes a random variable, with expected value and variance that depend only upon the value of log Ω. Thus, for each value of log Ω, the critical odds model predicts a value for mean terminal n. Figure 2 shows the predictions of the critical odds model. Observed values of log Ω were grouped into intervals in order to simplify analysis; the abscissa values in Fig. 2 are the midpoints of those intervals, which ranged from −0.35 to 1.85 in steps of 0.1 log Ω units. The dotted lines again represent 95% confidence intervals. As before, the data were pooled across subjects and across payoffs, and means for intervals containing less than 10 individual decisions were omitted.

It can be seen that the critical odds model can serve at least as a reasonable first approximation to observed behavior. Subjects, on the average, did base their decisions on the probability that change had occurred. Distortions that did occur consist mainly of increased sample sizes at smaller values of log Ω.

Discussion

The fact that subjects in general tended to take too many observations might be accounted for in part by the nature of the task itself. A stopping

decision required an active response, namely, pres
decision to continue, on the other hand, was a
overt response was required. The tendency might
if both "stop" and "continue" decisions require
The results indicate that subjects' decisions w
difficulty; however, as can be seen in Table 1,
affect behavior to as great a degree as it affected t
it is somewhat surprising that, within each of the t
there was essentially no effect due to false alarm
The fixed sample size model did not provide a
of behavior in the present experiment; subjects v
decisions on something other than the number o
This result was expected; the fixed sample size
actual values (R or B) of information items into a
of the critical odds model were generally in ac
havior; however, some discrepancies did occur. T
Fig. 2 indicate a slight tendency among subject
criterion (i.e., reduce the critical probability) a
This result is indicated by subjects stopping too
log Ω following longer sequences, with the resul
large for these smaller stopping odds.

EXPERIMENT II

Method

The apparatus and procedure were the same
male students enrolled in introductory psychol
partial fulfillment of a course requirement. Ea
assigned to one of five payoff conditions; none o
pated in Experiment I. Again, subjects were told
depend upon how well they performed, and
around $1.70.

Five levels of F, the penalty for a false ala
between-subjects design: $4, $6, $10, $18, and
five groups was the same as in the easier condit
.2, P_1 = .6. The values of α and w were, resp
Experiment I.

Results

A comparison of observed and optimal mea
cated, as before, that the optimal model itself
performance, $F(1,49)$ = 21.45, L = .0014. Fig
parison, and also suggests the existence of son
not apparent in Experiment I. First of all, it
tended to stop at higher log Ω values for higher

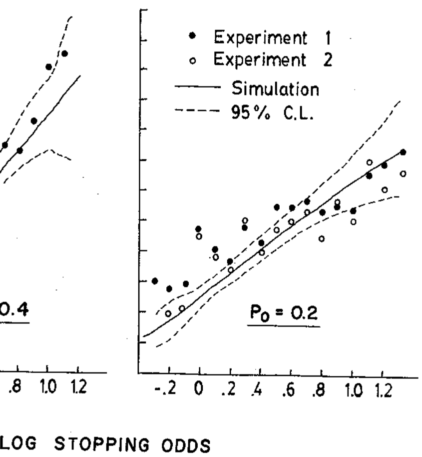

pected mean sample sizes assuming a critical odds model
e sizes exhibited by subjects (dots). Broken lines indicate
neans obtained via simulation.

redictions of the fixed sample size model. The
were not solely influenced (if they were influ-
of items appearing in a sequence.
out for the critical odds model, which predicts
p when log Ω (or, equivalently, probability
r exceeds some predetermined level. If this
omes a random variable, with expected value
nly upon the value of log Ω. Thus, for each
dds model predicts a value for mean terminal
ictions of the critical odds model. Observed
ed into intervals in order to simplify analysis;
2 are the midpoints of those intervals, which
in steps of 0.1 log Ω units. The dotted lines
nce intervals. As before, the data were pooled
payoffs, and means for intervals containing
ions were omitted.
critical odds model can serve at least as a
ion to observed behavior. Subjects, on the
sions on the probability that change had oc-
occur consist mainly of increased sample
g Ω.

eneral tended to take too many observations
rt by the nature of the task itself. A stopping

decision required an active response, namely, pressing the stop button. A decision to continue, on the other hand, was a passive one, in that no overt response was required. The tendency might not have been observed if both "stop" and "continue" decisions required an active response.

The results indicate that subjects' decisions were influenced by task difficulty; however, as can be seen in Table 1, this parameter did not affect behavior to as great a degree as it affected the optimal model. Also, it is somewhat surprising that, within each of the two difficulty conditions, there was essentially no effect due to false alarm penalty.

The fixed sample size model did not provide a satisfactory description of behavior in the present experiment; subjects were clearly basing their decisions on something other than the number of items in the sequence. This result was expected; the fixed sample size model does not take the actual values (R or B) of information items into account. The predictions of the critical odds model were generally in accord with observed behavior; however, some discrepancies did occur. The discrepancies seen in Fig. 2 indicate a slight tendency among subjects to relax their stopping criterion (i.e., reduce the critical probability) as sample size increased. This result is indicated by subjects stopping too often at small values of log Ω following longer sequences, with the result that average n was too large for these smaller stopping odds.

EXPERIMENT II

Method

The apparatus and procedure were the same as in Experiment I. Fifty male students enrolled in introductory psychology served as subjects in partial fulfillment of a course requirement. Each subject was randomly assigned to one of five payoff conditions; none of the subjects had participated in Experiment I. Again, subjects were told that their payment would depend upon how well they performed, and were paid an average of around $1.70.

Five levels of F, the penalty for a false alarm, were employed in a between-subjects design: $4, $6, $10, $18, and $28. Task difficulty in all five groups was the same as in the easier condition of Experiment I; $P_0 = .2$, $P_1 = .6$. The values of α and w were, respectively, .1 and $1, as in Experiment I.

Results

A comparison of observed and optimal mean log stopping odds indicated, as before, that the optimal model itself did not describe subjects' performance, $F(1,49) = 21.45$, $L = .0014$. Figure 3 illustrates this comparison, and also suggests the existence of some payoff effects that were not apparent in Experiment I. First of all, it can be seen that subjects tended to stop at higher log Ω values for higher values of F, as is appropri-

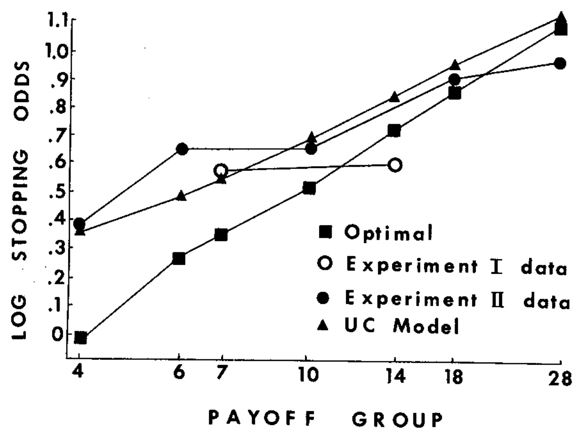

FIG. 3. Mean log stopping odds for each payoff group in Experiment I and II and predictions of a *UC* (utility for being correct) model.

ate. Figure 3 also illustrates, however, that payoffs did not affect subjects' decisions as much as they affected those of the optimal model. That is, the observed log Ω values did not rise as quickly as the optimal stopping values. A test of the simple main effects of payoffs, restricted to subjects' decisions, was also significant, $F(4,45) = 13.54$, indicating that payoffs in Experiment II did indeed affect subjects' stopping decisions. And finally, the Subjects vs Optimal by Payoffs interaction was also significant, $F(4,90) = 13.71$, supporting the interpretation that payoffs affected subjects differently from their effect on the optimal model.

The fixed sample size and critical odds models were tested as in Experiment I. The fit of the critical odds model for Experiment II is shown in Fig. 2. The data for the fixed sample size model are not shown, being very similar to those in Fig. 1. Again, the critical odds model was superior to the fixed sample size model, and the deviations from the predictions of the former were similar to those in Experiment I.

The reasonably good fit of the critical odds model does not, of course, imply that subjects actually calculated posterior odds in the manner of a mathematical model. We next examined a simple heuristic[2] that seemed well within human information processing limits: "Stop as soon as k of the n most recent information items have been red." The model allows for individual differences in memory capacity (n) and stopping criterion (k); perhaps the model's strongest assumption is that for a given subject, k and n remain fixed throughout the experimental session. Given values of k and n, the model specifies the exact point at which stopping will occur. For

[2] We are grateful to a referee for suggesting the k/n model and for comments that led to our examination of a risk-averse utility explanation of the results, described later.

each of the 50 subjects in Experiment II, k and n we maximize the proportion of stopping decisions corre model. Table 2 summarizes the results.

For values of n = 1 through 8, there are $n(n + 1$ combinations. Table 2 lists only those k/n parame found to be best-fitting for three or more subjects listed were able to account for 39 of the 50 subjects, specification rate of 52.4%; i.e., approximately on decisions among these 39 subjects were specified ex ate k/n model.

Another approach to understanding behavior in t the effects of alternative assumptions about the util sign to payoffs. Specifically, might a risk-averse u the tendency for subjects to underreact to the false present case, the answer appears to be "no."

Intuitively, if subjects are risk-averse but otherwi to become progressively overcautious, relative to th increases. Such a tendency would be reflected in a ing from the optimal function in Fig. 3, but this was

TABLE 2

MEAN PROPORTIONS OF TOTAL STOPPING DECISIONS COR BEST-FITTING k/n MODELS AS A FUNCTION OF FALSE A EXPERIMENT II

	F			
k/n	$4	$6	$10	$18
2/3	.632 (2)	—	—	.703 (1)
2/4	.400 (1)	.446 (2)	—	—
3/4	—	.646 (1)	.378 (1)	.451 (2)
3/5	.657 (1)	—	.492 (2)	.311 (1)
3/6	—	.564 (3)	.557 (1)	.537 (5)
3/8	.575 (3)	.489 (4)	.718 (2)	—
4/8	—	—	—	—
All others	.640 (3)	—	.616 (4)	.707 (1)

Note. Parentheses contain the number of subjects for wh best-fitting.

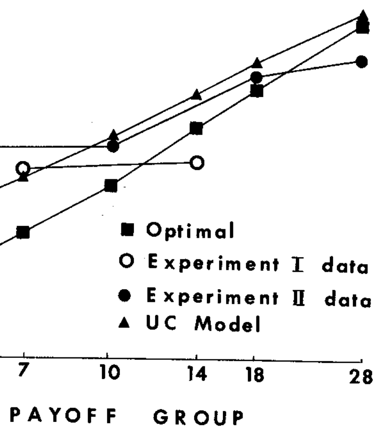

for each payoff group in Experiment I and II and predic-
rect) model.

however, that payoffs did not affect subjects'
fected those of the optimal model. That is, the
not rise as quickly as the optimal stopping
main effects of payoffs, restricted to subjects'
nt, $F(4,45) = 13.54$, indicating that payoffs in
fect subjects' stopping decisions. And finally,
y Payoffs interaction was also significant,
the interpretation that payoffs affected sub-
effect on the optimal model.
d critical odds models were tested as in Ex-
ical odds model for Experiment II is shown in
sample size model are not shown, being very
gain, the critical odds model was superior to
and the deviations from the predictions of the
in Experiment I.
f the critical odds model does not, of course,
calculated posterior odds in the manner of a
xt examined a simple heuristic[2] that seemed
on processing limits: "Stop as soon as k of the
tems have been red." The model allows for
mory capacity (n) and stopping criterion (k);
t assumption is that for a given subject, k and
e experimental session. Given values of k and
xact point at which stopping will occur. For

uggesting the k/n model and for comments that led to our
explanation of the results, described later.

each of the 50 subjects in Experiment II, k and n were estimated so as to maximize the proportion of stopping decisions correctly specified by the model. Table 2 summarizes the results.

For values of $n = 1$ through 8, there are $n(n + 1)/2 = 36$ possible k/n combinations. Table 2 lists only those k/n parameter pairs which were found to be best-fitting for three or more subjects. The seven models listed were able to account for 39 of the 50 subjects, at an overall correct specification rate of 52.4%; i.e., approximately one-half of all stopping decisions among these 39 subjects were specified exactly by the appropriate k/n model.

Another approach to understanding behavior in this task is to explore the effects of alternative assumptions about the utilities that subjects assign to payoffs. Specifically, might a risk-averse utility function explain the tendency for subjects to underreact to the false alarm penalty? In the present case, the answer appears to be "no."

Intuitively, if subjects are risk-averse but otherwise optimal, they ought to become progressively overcautious, relative to the optimal model, as F increases. Such a tendency would be reflected in a log Ω function diverging from the optimal function in Fig. 3, but this was clearly not the case.

TABLE 2

MEAN PROPORTIONS OF TOTAL STOPPING DECISIONS CORRECTLY SPECIFIED BY BEST-FITTING k/n MODELS AS A FUNCTION OF FALSE ALARM PENALTY (F), EXPERIMENT II

	F					
k/n	\$4	\$6	\$10	\$18	\$28	All groups
2/3	.632 (2)	—	—	.703 (1)	—	.656 (3)
2/4	.400 (1)	.446 (2)	—	—	—	.431 (3)
3/4	—	.646 (1)	.378 (1)	.451 (2)	—	.482 (4)
3/5	.657 (1)	—	.492 (2)	.311 (1)	.396 (2)	.457 (6)
3/6	—	.564 (3)	.557 (1)	.537 (5)	.466 (1)	.540 (10)
3/8	.575 (3)	.489 (4)	.718 (2)	—	—	.568 (9)
4/8	—	—	—	—	.502 (4)	.502 (4)
All others	.640 (3)	—	.616 (4)	.707 (1)	.460 (3)	.588 (11)

Note. Parentheses contain the number of subjects for whom the specified model was best-fitting.

More precisely, suppose that subjects adopted the following risk-averse disutility function:

$$d(F) = F^Y, \; Y > 1, \tag{3}$$

which implies that, for example, a loss of \$2F is more than twice as serious as a loss of \$$F$. The reason for the d function's inability to explain the present results can most easily be seen by referring to Fig. 3. Since abscissa values are logarithmically spaced, it can be seen that optimal log Ω is very nearly linear in log $F(r = .997)$; in fact, the relationship remains linear through the highest F value we examined for this purpose ($F = 2000$). Thus, at least for $4 \leq F \leq 2000$, we have a close approximation to the optimal model's stopping behavior as a function of log F:

$$\log \Omega = a \log F + b \quad \text{for some } a, b.$$

If subjects were behaving optimally but minimizing disutility rather than dollar loss in accordance with Eq. (3), then the functional relationship between log F and log Ω would be given by:

$$\begin{aligned} \log \Omega &= a \log [d(F)] \\ &= a \log F^Y \\ &= aY \log F. \end{aligned}$$

But since $a > 0$ and $Y > 1$, it follows that subjects acting in accordance with Eq. (3) would yield a log Ω function with slope greater than that of the optimal function in Fig. 3; this was clearly not the case.

Alternatively, one can postulate a utility for being correct (*UC*), suggested by Pitz and Downing (1967). The *UC* hypothesis received mild support from Pitz and Reinhold (1968) and Ullrich (1969) in decision making and information seeking situations, respectively. In the detection of change situation, the *UC* hypothesis would say, in effect, that subjects place some value on demonstrating an ability to detect the change when it occurs, independent of the monetary costs associated with their decisions; the effects of monetary rewards would thus be diminished.

Predictions of a *UC* model are shown in Fig. 3, and were obtained by adding the Constant 3 (estimated graphically) to each of the F values, and then determining optimal values of log Ω for the resulting increased penalties. The (post hoc) predictions of the UC model are clearly better approximations of behavior than those of the optimal model. However, the *UC* model predicts that log Ω will never be less than optimal, while for two conditions (F = \$14 in Experiment I and F = \$28 in Experiment II), this appears to be untrue.

GENERAL DISCUSSION

In view of the results of Experiment II, the comment made earlier regarding subjects' tendencies to wait too long before stopping requires

some modification. In the \$4, \$6, and \$10 groups,
again observed, but in the \$28 group, the reverse w
generally stopped too early; observed and optim
agreed rather closely in the \$18 condition. These res
"active vs passive response" hypothesis less tenabl
that the excess observations taken in Experiment I
by an under-reaction to the payoff (false alarm pe
biased payoff effects have been found in decision m
ing, 1967; Pitz & Reinhold, 1968) and perceptual ju
1966). In each of these studies, behavior was affecte
to the extent that an optimal model would prescrib
cates that an analogous effect is present in the detec

Both experiments demonstrated the superiority
model with respect to the fixed sample size model.
study of information seeking reported a marked dec
as a function of sample size; the critical odds mode
constant critical odds that is independent of sample
Such an effect apparently did occur in the present
much smaller degree. This difference presumably re
ences between the two tasks. When no possibility o
in the Pitz *et al.* information seeking study, all in
sented to subjects have equal status, regardless of th
are presented. For example, the sequence RRRBB
sequence BBBRRR in terms of the optimal decisio
tion of change task, however, recent items carry mo
ones; of the two sequences above, the latter is mor
than the former if P_0 is less than P_1. Subjects presu
ness of the relatively greater importance of recent i
ignore entirely those items seen early in a long se
case, there is less reason to expect a decreasing c
non. Regardless of the number of items seen, a d
only upon the most recent items. In contrast to th
case, subjects in a nonstationary situation might th
a sequence, so that no effect of the number of item
found.

The k/n model is attractive for at least two reas
for any given information sequence, the exact point
occur, thus permitting precise measures of the mod
its two parameters, reflecting a stopping criterion (
ity (n), are easily interpreted; the latter, in fact, ca
the detection of change context, thus providing a
cognition. An apparent shortcoming of the k/n mo
any discernable relationship between payoff conditi
parameters (see Table 2).

some modification. In the \$4, \$6, and \$10 groups, such a tendency was again observed, but in the \$28 group, the reverse was true, i.e., subjects generally stopped too early; observed and optimal log stopping odds agreed rather closely in the \$18 condition. These results tend to make the "active vs passive response" hypothesis less tenable, and suggest instead that the excess observations taken in Experiment I can be accounted for by an under-reaction to the payoff (false alarm penalty) factor. Similar biased payoff effects have been found in decision making (Pitz & Downing, 1967; Pitz & Reinhold, 1968) and perceptual judgment tasks (Ulehla, 1966). In each of these studies, behavior was affected by payoffs, but not to the extent that an optimal model would prescribe. Experiment II indicates that an analogous effect is present in the detection of change task.

Both experiments demonstrated the superiority of the critical odds model with respect to the fixed sample size model. The Pitz *et al.* (1969) study of information seeking reported a marked decrease in stopping odds as a function of sample size; the critical odds model, in that it predicts a constant critical odds that is independent of sample size, was inadequate. Such an effect apparently did occur in the present experiments, but to a much smaller degree. This difference presumably reflects inherent differences between the two tasks. When no possibility of change is present, as in the Pitz *et al.* information seeking study, all information items presented to subjects have equal status, regardless of the order in which they are presented. For example, the sequence RRRBBB is equivalent to the sequence BBBRRR in terms of the optimal decision rule. For the detection of change task, however, recent items carry more weight than earlier ones; of the two sequences above, the latter is more indicative of change than the former if P_0 is less than P_1. Subjects presumably have an awareness of the relatively greater importance of recent items, and in fact may ignore entirely those items seen early in a long sequence. If this is the case, there is less reason to expect a decreasing critical odds phenomenon. Regardless of the number of items seen, a decision may be based only upon the most recent items. In contrast to the information seeking case, subjects in a nonstationary situation might thus ignore the length of a sequence, so that no effect of the number of items observed would be found.

The k/n model is attractive for at least two reasons. First, it specifies, for any given information sequence, the exact point at which stopping will occur, thus permitting precise measures of the model's accuracy. Second, its two parameters, reflecting a stopping criterion (k) and memory capacity (n), are easily interpreted; the latter, in fact, can be estimated outside the detection of change context, thus providing a link to other areas of cognition. An apparent shortcoming of the k/n model was the absence of any discernable relationship between payoff condition and best-fitting k/n parameters (see Table 2).

The results did not support a risk-averse utility interpretation. However, the utility for being correct (UC) hypothesis received mild support, and accounted in part for the diminishing effects of false alarm penalty. For example, in the $F = \$4$ condition, the optimal model requires a posterior probability of change of only .364 or greater before a stopping decision is indicated. While this strategy does minimize expected cost, it also yields a large proportion of false alarms (.507 as determined by Monte Carlo techniques, in contrast to an observed proportion of .299 for subjects). Subjects may consider false alarms undesirable, perhaps feeling somewhat foolish for consistently maintaining that the change has taken place when in fact it often has not. The optimal model, of course, has no such reservations.

The critical odds model, then, may serve as a general starting point for a descriptive model in the detection of change situation. That is, subjects do tend to behave as if they were stopping when the probability favoring change reaches some predetermined cutoff value. Exactly why their strategies approximate this appropriate way of responding in the nonstationary task, but not in a stationary task, needs to be explored more fully. Finally, implicit payoffs, such as utility for being correct, may also be important considerations in formulating a descriptive model for the detection of change.

REFERENCES

Bellman, R. *Dynamic programming.* Princeton, N.J.: Princeton Univ. Press, 1957.

Brown, E. R., & Bane, A. L. Probability estimation in a chance task with changing probabilities. *Journal of Experimental Psychology: Human Perception and Performance,* 1975, **104**, 183–187.

Burkheimer, G. J., & Rapoport, A. *Models for the detection of change.* Report No. 95, the L. L. Thurstone Psychometric Laboratory, University of North Carolina, June 1971.

Chinnis, J. O., & Peterson, C. R. Nonstationary processes and conservative inference. *Journal of Experimental Psychology,* 1970, **84**, 248–251.

Jeffreys, H. *Theory of probability.* London: Oxford Univ. Press (Clarendon), 1961.

Pitz, G. F. The sequential judgment of proportion. *Psychonomic Science,* 1966, **4**, 397–398.

Pitz, G. F. The influence of prior probabilities on information seeking and decision making. *Organizational Behavior and Human Performance,* 1969, **4**, 213–226.

Pitz, G. F., & Downing, L. Optimal behavior in a decision making task as a function of instructions and payoffs. *Journal of Experimental Psychology,* 1967, **73**, 549–555.

Pitz, G. F., & Reinhold, H. Payoff effects in sequential decision making. *Journal of Experimental Psychology,* 1968, **77**, 249–257.

Pitz, G. F., Reinhold, H., & Geller, E. S. Strategies of information seeking in deferred decision making. *Organizational Behavior and Human Performance,* 1969, **4**, 1–19.

Pollock, S. M. Minimum-cost checking using imperfect information. *Management Science,* 1967, **13**, 454–465.

Rapoport, A. Sequential decision making in a computer-controlled task. *Journal of Mathematical Psychology,* 1964, **1**, 351–374.

Robinson, G. H. Continuous estimation of a time-varying probability. *Ergonomics,* 1964, **7**, 7–21.

Theios, J., Brelsford, J. W., Jr., & Ryan, P. Detection of ch sequences. *Perception and Psychophysics,* 1971, **9**, 489–

Ulehla, Z. J. Optimality of perceptual decision criteria. *Journ ogy,* 1966, **71**, 564–569.

Ullrich, J. R. *Utilities in an information-seeking task when pay* doctoral dissertation, Southern Illinois University, 1969.

RECEIVED: October 4, 1976

ɔport a risk-averse utility interpretation. How-
correct (UC) hypothesis received mild support,
the diminishing effects of false alarm penalty.
4 condition, the optimal model requires a poste-
of only .364 or greater before a stopping deci-
is strategy does minimize expected cost, it also
of false alarms (.507 as determined by Monte
ast to an observed proportion of .299 for sub-
ider false alarms undesirable, perhaps feeling
istently maintaining that the change has taken
has not. The optimal model, of course, has no

then, may serve as a general starting point for a
tection of change situation. That is, subjects do
were stopping when the probability favoring
determined cutoff value. Exactly why their
his appropriate way of responding in the
in a stationary task, needs to be explored more
ffs, such as utility for being correct, may also
is in formulating a descriptive model for the

REFERENCES

ng. Princeton, N.J.: Princeton Univ. Press, 1957.
bability estimation in a chance task with changing prob-
ntal Psychology: Human Perception and Performance,

Models for the detection of change. Report No. 95, the
Laboratory, University of North Carolina, June 1971.
Nonstationary processes and conservative inference.
hology, 1970, **84**, 248–251.
London: Oxford Univ. Press (Clarendon), 1961.
of proportion. *Psychonomic Science,* 1966, **4**, 397–398.
robabilities on information seeking and decision making.
Human Performance, 1969, **4**, 213–226.
al behavior in a decision making task as a function of
al of Experimental Psychology, 1967, **73**, 549–555.
effects in sequential decision making. *Journal of Experi-*
249–257.
er, E. S. Strategies of information seeking in deferred
al Behavior and Human Performance, 1969, **4**, 1–19.
ing using imperfect information. *Management Science,*

making in a computer-controlled task. *Journal of*
4, **1**, 351–374.
tion of a time-varying probability. *Ergonomics,* 1964, **7**,

Theios, J., Brelsford, J. W., Jr., & Ryan, P. Detection of change in nonstationary binary sequences. *Perception and Psychophysics,* 1971, **9**, 489–492.

Ulehla, Z. J. Optimality of perceptual decision criteria. *Journal of Experimental Psychology,* 1966, **71**, 564–569.

Ullrich, J. R. *Utilities in an information-seeking task when payoffs are varied.* Unpublished doctoral dissertation, Southern Illinois University, 1969.

Received: October 4, 1976