JOURNAL OF MATHEMATICAL PSYCHOLOGY 32, 91–134 (1988)

Psychological Models of Deferred Decision Making

JEROME R. BUSEMEYER

Purdue University

AND

AMNON RAPOPORT

University of North Carolina at Chapel Hill

In a two-state deferred decision making task one of two mutually exclusive states of nature is responsible for generating a sequence of independent, identically distributed, and costly observations. After purchasing each observation, the decision maker must either (a) stop purchasing costly observations and make a terminal choice favoring one of the two states, or (b) continue purchasing at least one more observation. We describe a new method, called pattern analysis, for distinguishing alternative models of deferred decision making. Seven different psychological models are evaluated including the optimal stopping rule, fixed sampling, random walk, fixed forgetting, horse race or accumulator, runs, and hybrid stopping rules. Violations of basic properties implied by each of these seven models are reported. The most promising psychological model was a myopic stopping rule, which prescribes purchasing observations until the expected loss of making a terminal decision after purchasing n observations is less than or equal to the sum of the costs of purchasing $n+1$ observations.

PSYCHOLOGICAL MODELS OF DEFERRED DECISION MAKING

How do physicians decide when to stop conducting medical tests and make a final diagnosis? How do school psychologists decide when to stop administering IQ tests and categorize a child as mentally gifted, normal, or retarded? How do scientists decide when to stop performing experiments on a new phenomenon and publish their conclusions? How do military commanders decide when to stop collecting intelligence reports and take action? These decision situations are all examples of deferred decision problems. The present article investigates rules that individuals may actually use to decide when to stop purchasing costly observations and commit to a terminal decision in a two-state deferred decision task.

General Task Description

The two-state deferred decision task is characterized by a data generating process, choice alternatives, and a payoff structure.

0022-2496/88 $3.00

Data generator. The decision maker is informed that one of two stochastic processes, labeled S_A or S_B, is generating a sequence of costly observations $[Z(1), Z(2), ..., Z(n)]$. For example, in a medical decision context S_A may represent a patient with "disease present," S_B may represent a patient with "disease absent," and $Z(n)$ may represent the nth observation in a sequence of costly laboratory tests. The models evaluated in this article generally assume that the sequences are identically and independently distributed random variables, and that the decision maker is provided information about the probability laws and the prior probability of each process. Of course, many deferred decision problems encountered in real life involve sequences that are not identically and independently distributed. Although the models can be adapted to the dependent observation case, it may be wise to begin with the simplest case that assumes independence, and subsequently extend the models to the more complex dependent case.

Choice alternatives. After purchasing n observations, the decision maker has a choice among three alternatives: stop the sequence and choose R_A (e.g., diagnose the patient as having disease S_A), stop the sequence and choose R_B (e.g., diagnose the patient as having disease S_B), or continue the sequence by purchasing another observation (e.g., conduct another laboratory test). The random variable N represents the total number of observations purchased before making a terminal decision. In some cases only a limited number of observations (denoted M) can be purchased.

Payoff structure. If a terminal decision R_j $(j = A, B)$ is made when, in fact, the sequence was generated by process S_i $(i = A, B)$, then a terminal monetary payoff, v_{ij}, is incurred depending on the event (S_i, R_j). If the decision maker purchases another observation, then a monetary fee, $c(n)$, must be paid, which may depend on the number of observations already sampled. The loss incurred at the end of each sequence is therefore defined as $L = [v_{ii} - v_{ij}] + \sum c(n)$, where S_i is the process that generated the sequence, R_j is the terminal response, and summation extends from $n = 1$ to N.

Fixed vs sequential sampling tasks. Sequential sampling tasks differ from fixed sampling tasks in the following way. With fixed sampling, the decision maker must select the number of samples before observing any of the sample values, wheras with sequential sampling, he or she must decide whether to stop sampling after observing each sample value. Whereas the random variable N is independent of the obtained sample values for fixed sampling, it is dependent on the obtained sample values for sequential sampling.

The Optimal Model

The major purpose of the present article is to empirically evaluate psychological models of deferred decision making. These models are reviewed in the discussion section after presenting the empirical results. However, the optimal model is briefly described at this point because it is useful for suggesting important task properties.

The optimal decision rule is defined as the decision rule that minimizes the expected loss, $E[L]$ (cf., DeGroot, 1970). This rule prescribes making a terminal decision after purchasing n observations if the expected loss of making a terminal decision on the basis of the n obtained sample values is less than or equal to the expected loss of deferring the decision and purchasing additional observations. It is formulated in terms of two components—an evidence accumulator and a stopping rule:

Evidence accumulator. Define $s(n)$ as the evidence engendered by the observation $Z(n)$, where $s(n) > 0$ represents evidence favoring the presence of process S_A, and $s(n) < 0$ represents evidence favoring the presence of process S_B. According to the optimal model, $s(n)$ is set equal to the log odds corresponding to the observation $Z(n) = z$ (i.e., $s(n) = \ln[f_A(z)/f_B(z)]$, where f_i is the density function for process S_i). The log posterior odds after purchasing n observations equals the sum of the log odds, $s(0) + s(1) + s(2) + \cdots + s(n)$, where $s(0)$ represents the log of the prior odds. Previous research on probabilistic inference using tasks similar to the present one supports the additive rule for combining evidence (Shanteau, 1970; Wallsten and Sapp, 1977).

Stopping rule. If the log posterior odds is greater than or equal to some upper bound, $\alpha(n)$, then R_A is selected, whereas if the log posterior odds is less than or equal to some lower bound, $-\beta(n)$, then R_B is chosen; otherwise sampling continues.

In general, the criterion bounds $\alpha(n)$ and $-\beta(n)$ vary within a sequence of observations depending on the number of remaining observations. For a given pair of stochastic processes, the rules for assigning criterion bounds after each observation are influenced by the cost of observation, $c(n)$, the terminal payoffs, v_{ij}, and the limit, M, on the number of observations that can be purchased. (For real time decisions, the interval of time required to reach a decision may be a critical factor.) The exact values of the upper and lower criterion bounds can be calculated by dynamic programming (see Rapoport and Burkheimer, 1971, for details).

Rapoport and Burkheimer's (1971) theoretical investigation of the optimal policy revealed several interesting properties. Assume that a total of M observations can be purchased, the prior probabilities are equal, and the payoffs are symmetric (i.e., $v_{AA} = v_{BB}$, $v_{AB} = v_{BA}$). When the observation costs are constant ($c(n) = c$), the criterion bounds remain constant as long as the number of observations remaining to be purchased is large. When the observation costs increase ($c(n+1) > c(n)$), the criterion bounds decrease in magnitude after each observation converging to a common meeting point. Finally, when the observation costs decrease ($c(n+1) < c(n)$), the criterion bounds initially increase in magnitude after each observation but eventually converge toward a meeting point as the number of observations purchased approaches the limit, M.

Situations where the observation costs either increase or decrease are abundant. For an example of increasing costs, consider the use of X-rays in a medical diagnosis problem—the health hazards produced by another X-ray test are much more severe after several tests are administered. As an example of decreasing costs, consider the costs associated with the initiation of a new line of research. The initial set-up costs are very high, but replications of experiments become cheaper because the same resources are reused each time.

PREVIOUS RESEARCH

Much of the previous research employed a simple paradigm similar to the following example. Subjects are asked to decide which of two possible diseases, S_A and S_B, are present based on a sequence of independent laboratory tests. Each test yields a binary outcome, say $Z(n) = +1$ or -1. The conditional probability of a positive test given disease S_A equals the conditional probability of a negative test given S_B, and both conditional probabilities exceed 0.5. The prior probabilities are equal, the terminal payoffs are symmetric, the observation costs are constant, and there is no limit on the number of tests.

For this simple case, the log posterior odds after n tests is proportional to the difference, denoted $d(n)$, between the number of positive and negative tests results observed after n tests. The optimal decision rule is to (a) decide R_A as soon as $d(n) \geqslant K$, (b) decide R_B as soon as $d(n) \leqslant -K$, and (c) continue sampling otherwise (cf. Edwards, 1965). Note that for this particular task, the optimal model predicts that subjects use a constant critical difference, K, as the stopping rule.

One of the most interesting conclusions from this research—which contradicts the optimal model—is that subjects tend to require less evidence to make a terminal decision as the number of observations purchased increases (Pitz, Reinhold, and Geller, 1969; Sanders and Ter Linden, 1967; Viviani, 1979; Wallsten, 1968). Apparently, the critical difference K decreases as the number of observations purchased increases. For example, Sanders and Ter Linden (1967) and later Viviani (1979) found that the tendency to stop after observing a subsequence of strong evidence was much greater when the subsequence was preceded by a non-diagnostic subsequence. Even more interesting is the finding by Pitz *et al.* (1969) that subjects frequently terminate information purchasing with $d(n) = 0$ (no evidence) late within a sequence of tests, which is impossible if subjects use a constant critical difference stopping rule.

Several alternative models may account for the results. Most of these will be considered in the discussion section after presenting the new results of the present experiments. However, two explanations can be ruled out by prior research. One is a "noisy counter" hypothesis which states that decisions are based on an estimate of $d(n)$ that is perturbed by error. This hypothesis may be plausible when a large number of test results are presented very rapidly as in the Sanders and Ter Linden (1967) study. However, it is less plausible when the studies by Pitz *et al.* (1969) are considered. Pitz *et al.* compared the performances of two groups—one group was required to keep track of the difference mentally because physical records of the test results were neither available nor allowed; the second group was provided with physical counters so that perfectly accurate estimates were always available. For both groups, presentation of the sample observations was self paced. The results of the study indicated that the presence or absence of physical counters produced no effect; both groups displayed an equal tendency to stop on a difference of zero after long sequences of non-diagnostic information.

The second hypothesis, called the "sample size" hypothesis, states that when the log posterior odds are held constant, long sequences have greater impact on subjective probabilities than short sequences. Pitz and Barrett (1969) tested the sample size hypothesis by initially presenting subjects with free samples that were equal in terms of log posterior odds, but varied according to sample size. They found that the size of the initial free sample had little effect on the subsequent number of observations purchased, and concluded that this hypothesis was inadequate.

The purpose of the present article is to investigate rules that individuals use to decide when to stop purchasing information and make a terminal decision. Three experiments are reported. In the first two, the observation costs varied across samples within a sequence and the test results were binomially distributed. A new method called pattern analysis is proposed to evaluate alternative stopping rules. In the third experiment the number of observations was limited, the test results were normally distributed, and the terminal payoffs were manipulated. This last study provides a replication of some of the findings from the first two studies using a continuous rather than a discrete distribution of costly observations.

Experiments 1 and 2

Method

Subjects

Three male and three female subjects from Purdue University volunteered to participate in the first experiment. Five of these subjects were undergraduates enrolled in a statistics course. The remaining subject, labeled S5, was a graduate student in quantitative psychology.

Four subjects volunteered to take part in Experiment 2. All four were male graduate students in psychology at Purdue. Two were experimental psychologists who had completed a graduate psychology course in statistics. The other two were quantitative students with substantial training in statistics (one of these was S5 from Experiment 1).

Procedure

In both Experiments 1 and 2, subjects were to imagine that they were physicians specializing in the diagnosis of cancer. Each sequence of observations was said to be generated by a randomly selected patient who had one of two diseases. At any stage within a sequence, the subject had a choice between terminating the decision by making a diagnosis favoring one of the two diseases or sampling another test result. Each test produced a positive (evidence for disease S_A) or a negative (evidence for disease S_B) result. Subjects were told that they could sample as many tests as they wished, including none at all.

Subjects were instructed that a computer program began each sequence by randomly selecting a patient from a population afflicted with either disease using a procedure similar to flipping a coin so that each disease was equally likely to be present before testing began. Subjects were further told that if a patient had disease S_A then the probability of a positive test equaled .65, and if a patient had disease S_B then the probability of a negative test equaled .65. The computer algorithm was programmed exactly in this manner.

In Experiment 1, each subject received initially \$10.00 per session. For each sequence that terminated with an incorrect decision the subject lost 25 cents. Subjects received nothing for correct decisions ($v_{AA} = v_{BB} = 0$, $v_{AB} = v_{BA} = -25$). Three observation cost conditions were employed: for the decreasing condition $c(n) = 2^{3-n}$ cents; for the constant condition $c(n) = .25$ cents; for the increasing cost condition, $c(n) = 2^{n-4}$ cents. There was no limit on the number of tests that could be purchased.

In Experiment 2, each subject started out with \$12.00 per session. For each sequence that terminated with an incorrect decision the subject lost 20 cents. Subjects received nothing for correct decisions ($v_{AA} = v_{BB} = 0$, $v_{AB} = v_{BA} = -20$). Three observation cost conditions were employed again: for the slow increasing cost condition $c(n) = (.1)(1.4)^{n-1}$ cents; for the constant cost condition $c(n) = \frac{1}{3}$ cents; for the fast increasing cost condition $c(n) = (.1)(1.8)^{n-1}$ cents. There was no limit on the number of tests that could be purchased.

Each subject was individually tested during 15 sessions. With few exceptions, daily sessions were scheduled with at most one or two sessions per day. Sessions scheduled on the same day were separated by several hours. During each session, a total of 75 sequences were presented with blocks of 25 sequences under each observation cost condition. The order of observation cost conditions was counterbalanced across subjects and sessions. Altogether $75 \cdot 15 = 1125$ sequences were presented to each subject.

The experiment was conducted on an IBM-PC microcomputer; stimuli were presented on a color monitor and responses were recorded by depressing a key. Each block of 25 sequences began with a message indicating the observation cost condition. Each sequence began with a random sequence of tones (to alert the subject that a new patient was selected) and a random patient number. Each choice opportunity began with the following information displayed on the monitor: (a) the cost of the next test, (b) the loss produced by an incorrect decision, (c) the total amount of money spent buying previous sample tests on the current patient, and (d) the three choice alternatives.

If a sample test was requested, the result was graphically displayed on the monitor using a striking combination of colors, tones, and movements to distinguish positive and negative tests. The subject was required to press one of two keys after observing the test depending on the test result, and the program would not continue until the correct key was selected.

If a correct terminal diagnosis was made, then a brief tune was played followed by a message on the monitor indicating that the decision was correct. If the incorrect terminal diagnosis was selected, then a different tune was played, and a message that the decision was incorrect and the amount lost was printed on the monitor.

After completing each block, subjects were shown the total amount of money lost during that block (summed across payoffs and observation costs). At the end of each session, subjects were shown their percentage correct for each disease and observation cost condition, and the average number of tests purchased for each observation cost condition. During the entire experiment, a chart illustrating (a) the prior and conditional probabilities, and (b) the cost of each observation depending on the number of tests was always visible. Subjects were asked not to record test results with a paper and pencil because this would greatly increase the amount of time required to complete a session. A session usually lasted about 45 min.

There was one major difference between the stimulus displays of Experiments 1 and 2. In Experiment 1, subjects had to remember the number of positive and negative test results. In Experiment 2, the exact number of positive and negative test results obtained from a patient was displayed on the video monitor prior to each choice.

Optimal Decision Rule

For the first two experiments, the log posterior odds is proportional to $d(n) = Z(1) + \cdots + Z(n)$, or in other words, the difference between the number of positive and negative test results. In this case, the optimal terminal decision is to choose R_A if $d(N)$ is positive, and to choose R_B if $d(N)$ is negative. Note that the optimal terminal decision is not necessarily correct. The correct decision is defined as the event $(S_i, R_j, i=j)$.

The optimal stopping rules were calculated using the observed monetary payoffs and observation costs[1]. Table 1 shows the upper bound, $\alpha(n)$, ordered according to the number of tests already purchased, for each cost condition of Experiments 1 and 2. For example, with the constant cost condition of Experiment 1, the optimal model prescribes that a terminal decision should be made as soon as the difference equals four in magnitude. With the increasing cost condition of Experiment 1, the optimal model prescribes that a terminal decision should be made after purchasing the second observation if the difference equals two in magnitude.

[1] The dynamic programming algorithm used to compute the upper criterion bound $\alpha(n)$ for the optimal rule is based on the assumption that there is a finite limit on the number of observations that can be purchased, denoted M. Although an unlimited number of observations could be purchased in Experiments 1 and 2, the criteria in Table 1 were computed by setting $M = 50$. Thus $\alpha(n)$ represents the upper cirterion bound when n observations have already been purchased and there are $(M-n) = (50-n)$ observations remaining to be purchased. The solutions are the same for sufficiently large M and small n.

TABLE 1

Upper Criterion Bound for the Optimal Stopping Rule

	Expt 1			Expt 2		
n	*D*	*C*	*I*	*S*	*C*	*F*
0	1	1	1	1	1	1
1	2	2	2	2	2	2
2	3	3	2	3	3	2
3	4	4	1	2	3	2
4	5	4	1	2	3	1
5	6	4	0	2	3	1
6	7	4	0	1	3	0
7	8	4	0	1	3	0
8	9	4	0	1	3	0
9	10	4	0	1	3	0
10	11	4	0	0	3	0
11	12	4	0	0	3	0
12	13	4	0	0	3	0
13	14	4	0	0	3	0
14	15	4	0	0	3	0
15	16	4	0	0	3	0
16	17	4	0	0	3	0
17	18	4	0	0	3	0
18	18	4	0	0	3	0
19	19	4	0	0	3	0
20	19	4	0	0	3	0
21	20	4	0	0	3	0
22	20	4	0	0	3	0
23	20	4	0	0	3	0
24	20	4	0	0	3	0
25	20	4	0	0	3	0

Note. n = number of observations purchased, D = decreasing cost, C = constant cost, I = increasing cost, S = slow increasing cost, F = fast increasing cost.

Note that for the decreasing cost condition, the criterion magnitude is greater than the number of observations purchased for $n < 18$. In this case, the optimal model prescribes purchasing another test. The extremely large criterion produced by the decreasing cost condition results from the fact that the observation costs are rapidly approaching zero. A more realistic model which includes a subjective cost for waiting would not produce such an extremely large criterion.

Results of Experiments 1 *and* 2

The following results were obtained by pooling sequences across sessions 2 through 15. The first session was not included because subjects were unfamiliar with the task.

Marginal Statistics

Relative frequency distribution. Table 2 presents the cumulative percentages of the number of tests purchased, N. The percentages within each column are based on $25 \times 14 \times 6 = 2100$ observations in Experiment 1, and $25 \times 14 \times 4 = 1400$ observations in Experiment 2. The first column indicates the possible values for N. The next six columns show the cumulative percentages pooled across subjects for all sequences within each cost condition. For example, consider the constant cost condition of Experiment 1. The percentage of observed sequences with $N \leqslant 7$ was 53.1%.

For the increasing cost condition of Experiment 1 and the fast increasing cost condition of Experiment 2, the cost of the 10th test exceeded the loss produced by an incorrect decision. Consequently, almost all sequences terminated by $N = 10$. For the slow increasing cost condition of Experiment 2, the cost of the 17th test exceeded the loss produced by an error. Consequently, almost all sequences terminated by $N = 17$.

Note that the cumulative percentages for the decreasing and constant cost conditions of Experiment 1 crossover at $N = 6$. Prior to this point, the cumulative percentages for the decreasing condition dominate, but after this point the cumulative percentage for the constant cost dominate. This distribution crossover was consistent across all subjects of Experiment 1. One explanation for this result is the fact that tests were more expensive uner the decreasing cost condition for $N < 5$, but tests were more expensive under the constant cost condition for $N > 5$.

Proportion of correct decisions. Table 3 shows the proportion of correct terminal decisions for each disease state as well as the mean, median, and standard deviation of the number of tests purchased for each subject and cost condition of Experiments 1 and 2. The first column indicates the condition, the second indicates the subject number, the third and fourth columns show the proportion of correct decisions for states S_A and S_B, respectively, and the last three columns show in turn the mean, median, and standard deviation of the number of tests purchased.

The proportion of correct decisions increases almost linearly as the mean number of tests purchased increases. (A simple linear regression yields a square correlation of .85.) On the average, approximately a 2% gain in accuracy is obtained with each additional test. (This is restricted to the range of 4 to 13 tests.) As the cost of an error equaled 25 cents in Experiment 1, each additional test was worth approximately one-half cent on the average. (This cost analysis ignores the evidence provided by previous test results, which strongly influences the expected value of each new test.)

The proportions of correct decisions are nearly equal for each state, which suggests that the terminal decision was unbiased. The means and medians differ by a small amount, indicating a slight positive skew in the distribution of N. The standard deviations are farily large for the decreasing and constant conditions of Experiment 1, and tend to be an increasing function of the means.

TABLE 2

Cumulative Relative Frequency Distributions

	Experiment 1			Experiment 2		
n	*D*	*C*	*I*	*S*	*C*	*F*
.0	1.1	.1	.3	.1	.0	.0
1.0	2.5	.9	1.4	.1	.3	.3
2.0	4.8	1.9	4.4	4.1	5.9	7.8
3.0	15.1	10.7	25.1	18.7	22.0	26.4
4.0	27.1	25.6	48.2	36.1	39.8	52.6
5.0	37.4	35.9	75.5	53.7	56.2	72.1
6.0	45.5	45.8	89.2	69.1	68.1	82.9
7.0	52.4	53.1	99.4	79.0	76.4	91.7
8.0	57.6	59.8	99.9	83.9	83.4	97.2
9.0	63.1	64.4	99.9	88.1	86.5	99.6
10.0	67.6	70.3	100.0	91.9	90.0	100.0
11.0	70.9	74.3	100.0	94.8	91.9	100.0
12.0	74.6	79.4	100.0	97.3	93.7	100.0
13.0	77.1	82.3	100.0	98.8	95.4	100.0
14.0	79.6	84.6	100.0	99.6	96.1	100.0
15.0	82.8	87.0	100.0	99.8	96.7	100.0
16.0	85.1	89.5	100.0	99.9	97.3	100.0
17.0	86.8	91.1	100.0	100.0	97.7	100.0
18.0	88.2	92.5	100.0	100.0	98.1	100.0
19.0	89.7	93.6	100.0	100.0	98.3	100.0
20.0	91.1	94.9	100.0	100.0	98.6	100.0
21.0	92.5	95.5	100.0	100.0	98.6	100.0
22.0	93.5	96.3	100.0	100.0	98.9	100.0
23.0	94.1	96.6	100.0	100.0	98.9	100.0
24.0	94.8	97.3	100.0	100.0	99.0	100.0
25.0	95.4	97.5	100.0	100.0	99.1	100.0
26.0	96.2	97.9	100.0	100.0	99.2	100.0
27.0	97.0	98.1	100.0	100.0	99.3	100.0
28.0	97.4	98.8	100.0	100.0	99.3	100.0
29.0	97.9	99.0	100.0	100.0	99.3	100.0
30.0	98.2	99.1	100.0	100.0	99.5	100.0
40.0	99.4	99.9	100.0	100.0	99.9	100.0
50.0	99.8	100.0	100.0	100.0	99.9	100.0
65.0	100.0	100.0	100.0	100.0	100.0	100.0

Note. n = number of observations purchased, D = decreasing cost, C = constant cost, I = increasing cost, S = slow increasing cost, F = fast increasing cost.

TABLE 3

Percentage Correct for Each State, Mean, Median, and Standard Deviation of the Number of Observations Purchased

Condition	Subj.	$P[R_A \mid S_A]$	$P[R_B \mid S_B]$	Mean N	Median N	Std. N
D	1	.88	.89	10.47	9.00	7.09
C	1	.90	.87	9.12	7.00	5.89
I	1	.74	.74	5.19	5.00	1.53
D	2	.80	.83	5.36	4.00	3.07
C	2	.82	.82	5.18	5.00	2.23
I	2	.75	.73	4.45	4.00	1.46
D	3	.92	.94	12.19	10.00	8.60
C	3	.91	.88	11.32	8.00	7.97
I	3	.81	.70	4.43	4.00	1.24
D	4	.87	.89	9.84	8.00	7.75
C	4	.88	.85	8.87	8.00	5.53
I	4	.73	.75	4.37	5.00	1.22
D	5	.87	.84	10.95	9.00	8.75
C	5	.94	.84	11.04	10.00	5.50
I	5	.74	.75	4.38	5.00	1.24
D	6	.86	.81	8.62	7.00	6.26
C	6	.89	.81	7.86	6.00	5.37
I	6	.78	.78	4.59	4.00	1.60
S	5	.85	.82	5.19	5.00	1.21
C	5	.78	.79	5.21	5.00	1.68
F	5	.75	.81	4.09	4.00	0.88
S	7	.84	.88	8.40	8.00	3.03
C	7	.88	.88	9.54	8.00	6.76
F	7	.80	.73	5.98	6.00	1.76
S	8	.84	.76	5.11	5.00	2.37
C	8	.84	.80	5.10	5.00	2.55
F	8	.78	.74	4.55	4.00	1.91
S	9	.80	.81	4.69	4.00	2.13
C	9	.77	.82	4.97	4.00	2.64
F	9	.76	.77	4.15	4.00	1.56

Note. Subjects 1 to 6 participated in Experiment 1, subjects 5, 7, 8, and 9 participated in Experiment 2. D = decreasing cost, C = constant cost, I = increasing cost, S = slow increasing cost, F = fast increasing cost.

Proportion of optimal terminal decisions. The proportion of optimal terminal decisions was calculated separately for each subject as follows. The number of sequences terminating with $\{R_A$ and $d(N) > 0\}$ or $\{R_B$ and $d(N) < 0\}$ was divided by the number of sequences terminating with $\{d(N) < 0\}$ or $\{d(N) > 0\}$. The following proportions were obtained from subjects 1 through 6 of Experiment 1: .995, 1.00, .997, .991, .994, .987, with a mean equal to .994. In other words, the non-optimal decision was selected an average of 6 out of 1050 sequences. The corresponding proportions for subjects 5, 7, 8, and 9 of Experiment 2 were .996, .999, .999, .995, with a mean equal to .997. In other words, the non-optimal terminal decision was selected an average of 3 out of 1050 sequences.

For all practical purposes, *the terminal decision was determined by the final difference, $d(N)$, between the number of positive and negative tests.* This fact places a constraint on possible psychological models. For example, if subjects only paid attention to the $(n - m)$ most recent test results or if they randomly forgot m test results, then they would frequently fail to choose the optimal terminal decision. This results from the fact that the sign of the partial sum based on only $n - m$ test results will be imperfectly correlated with the sign of the total sum based on all n test results. The fact that non-optimal terminal decisions were almost never chosen implies that subjects were basing their terminal decision on a statistic that was nearly perfectly correlated with the sum of all the evidence.

To be even more specific, suppose that when $n \geq 10$ tests were observed, subjects in Experiment 1 randomly forgot $m \geq 3$ out of n test results and based their decision on the sign of the partial sum of $n - m$ test results. The probability that the sign of this partial sum matches the sign of the total sum given that $|d(n)| = 1$ is less than .712 (assuming that subjects choose randomly when the partial sum equals zero). However, the proportion of optimal terminal decisions pooled across all 54 sequences in Experiment 1 ending with $|d(N)| = 1$ and $N \geq 10$ equaled $51/54 = .944$, which is significantly higher than expected by the forgetting model.

Joint frequencies of $d(N)$ and N. Table 4 shows the joint relative frequency distribution of the terminal difference, $d(N)$, and the number of tests purchased, N, for each condition and experiment. The first column indicates the cost condition for each experiment. The second column indicates the magnitude of the terminal difference grouped into three intervals: $|d(N)| \leq 1$, $2 \leq |d(N)| \leq 4$, and $|d(N)| \geq 5$. The third column indicates whether the percentages are observed (O) or predicted (P). (Predicted percentages will be discussed later.) The last 11 columns indicate the number of tests purchased grouped into 11 intervals of size three, except for the first interval which only includes $N = 0$. Crossing the three intervals for $|d(N)|$ with the 11 intervals for N produces a 3×11 matrix. There are three matrices for each of the two experiments, and the percentages within each cell are averages across subjects. For example, under the constant cost condition of Experiment 1, subjects stopped with $2 \leq |d(N)| \leq 4$ and $10 \leq N \leq 12$ an average of 10% of all sequences.

Previous research suggests that subjects are willing to stop with less evidence as the number of tests purchased increases, even when the monetary observation costs

TABLE 4

Joint Relative Frequency Distribution of the Number of Observations Purchased and the Terminal Difference

Cond.	$d(N)$		0	1 – 3	4–6	7–9	10–12	13–15	16–18	19–21	22–24	25–27	28–30
			Number of observations purchased										
D	0–1	*O*	0	2	2	1	0	0	0	0	0	0	0
D	0–1	*P*	1	2	1	2	1	1	0	0	0	0	0
D	2–4	*O*	0	9	20	7	7	3	4	2	1	1	1
D	2–4	*P*	0	7	18	12	10	6	4	2	1	1	0
D	⩾5	*O*	0	0	5	10	5	6	3	3	1	2	1
D	⩾5	*P*	0	0	5	8	4	3	1	1	0	0	0
C	0–1	*O*	0	1	0	2	0	1	0	0	0	0	0
C	0–1	*P*	0	1	2	2	1	1	0	0	0	0	0
C	2–4	*O*	0	7	26	10	10	6	5	2	2	1	1
C	2–4	*P*	0	6	23	15	11	6	4	2	1	1	0
C	⩾5	*O*	0	0	5	7	5	3	1	1	0	0	0
C	⩾5	*P*	0	0	5	7	3	3	2	1	1	1	1
I	0–1	*O*	0	1	21	7	0	0	0	0	0	0	0
I	0–1	*P*	0	2	20	6	0	0	0	0	0	0	0
I	2–4	*O*	0	22	45	2	0	0	0	0	0	0	0
I	2–4	*P*	0	18	52	0	0	0	0	0	0	0	0
I	⩾5	*O*	0	0	1	0	0	0	0	0	0	0	0
I	⩾5	*P*	0	0	1	0	0	0	0	0	0	0	0
S	0–1	*O*	0	0	3	6	2	2	0	0	0	0	0
S	0–1	*P*	0	3	6	7	2	1	0	0	0	0	0
S	2–4	*O*	0	17	46	12	8	1	0	0	0	0	0
S	2–4	*P*	0	17	39	15	6	0	0	0	0	0	0
S	⩾5	*O*	0	0	3	1	0	0	0	0	0	0	0
S	⩾5	*P*	0	0	2	2	0	0	0	0	0	0	0
C	0–1	*O*	0	0	6	4	0	0	0	0	0	0	0
C	0–1	*P*	0	3	5	5	1	0	0	0	0	0	0
C	2–4	*O*	0	20	39	11	7	3	2	0	0	0	0
C	2–4	*P*	0	17	36	12	6	2	1	1	0	0	0
C	⩾5	*O*	0	0	3	2	0	0	0	0	0	0	0
C	⩾5	*P*	0	0	2	3	1	1	1	0	0	0	0
F	0–1	*O*	0	0	11	10	0	0	0	0	0	0	0
F	0–1	*P*	0	3	11	12	0	0	0	0	0	0	0
F	2–4	*O*	0	23	46	7	0	0	0	0	0	0	0
F	2–4	*P*	0	19	48	5	0	0	0	0	0	0	0
F	⩾5	*O*	0	0	1	0	0	0	0	0	0	0	0
F	⩾5	*P*	0	0	1	0	0	0	0	0	0	0	0

Note. Relative frequencies were rounded off to two decimal places.

are constant. Whenever subjects terminate with $|d(N)| < 2$ after purchasing more than one observation ($N > 1$), then it is clear that they initially required more evidence than a difference of -1 or $+1$ to make a terminal decision (since they passed it up after each odd serial position within a sequence), but later this same difference was sufficient to make a terminal decision. Subjects in Experiment 1 stopped with $|d(N)| < 2$, $N > 1$, on 6.6, 6.8, and 29.2% of the sequences presented during the decreasing, constant, and increasing cost conditions, respectively. Subjects in Experiment 2 stopped with $|d(N)| < 2$, $N > 1$, on 12.7, 10.7, and 21.7% of the sequences presented during the slow increasing, constant, and fast increasing cost conditions, respectively.

Conditional Statistics

Responses following last observation. According to the optimal model, subjects should never stop and choose R_A immediately following a negative test result, nor should they stop and choose R_B immediately following a positive test result. The tendency to choose R_A immediately following a negative test result or R_B immediately following a positive test result depended on the final difference between the number of positive and negative tests. If the last test was negative but the final difference was positive, or if the last test was positive and the final difference was negative, then *subjects usually decided in agreement with the final difference and contrary to the last test result.*

Table 5 presents the joint frequencies for the eight possible combinations defined by crossing the sign of the final test result, $Z(N)$, the sign of the final difference, $d(N)$, and the two terminal responses. The results are pooled across all sequences for each cost condition and experiment, separately.[2] The first row of the table indicates the sign of the final test result, the second indicates the sign of the final difference, and the third indicates the terminal response. The next three rows labeled D, C, and I, show the frequencies for the decreasing, constant, and increasing cost conditions of Experiment 1. The seventh row shows the totals from Experiment 1. The next three rows show the frequencies for the slow increasing, constant, and fast increasing cost conditions, respectively, and the last row shows the totals for Experiment 2. For example, under the constant cost condition of Experiment 2, a total of 10 sequences ended with a positive test result, a negative terminal difference, and the selection of response R_B.

Overall, subjects terminated with the combinations $[R_A, Z(N) < 0]$ or $[R_B, Z(N) > 0]$ on 166 sequences. On 89.2% of these 166 sequences they ended with either the combination $[d(N) < 0,\ Z(N) > 0,\ R_B]$ or the combination

[2] Due to a memory limitation, the computer program used to store the results of Experiment 1 recorded the patterns only up to and including $n = 8$. The memory limitation was increased to $n = 30$ for Experiment 2. The frequencies for Experiment 1 in Table 5 are, therefore, limited to sequences with $N \leq 8$, which include 97.2% of the sequences for the increasing cost condition, but only about 60% of the sequences for the constant and increasing cost conditions (see Table 2). All sequences were included for the increasing cost condition, and 99.5% were included for the constant cost condition in Experiment 2.

TABLE 5

Frequencies of the Joint Events Formed by Crossing the Sign of the Final Test Result, the Sign of the Final Difference, and the Final Choice Alternative

	Event							
Final test	1	1	1	1	−1	−1	−1	−1
Final difference	1	1	−1	−1	1	1	−1	−1
Final choice	*A*	*B*	*A*	*B*	*A*	*B*	*A*	*B*
D	537	1	2	4	7	1	0	485
C	592	0	1	7	11	0	2	517
I	978	1	7	50	29	5	8	952
Total	2107	2	10	61	47	6	10	1954
S	693	1	2	5	5	0	1	670
C	666	0	2	10	4	0	0	675
F	681	3	0	5	11	0	1	666
Total	2040	4	4	20	20	0	2	2011

Note. The value listed for the final difference indicates only whether the final difference was positive or negative, and it does not reflect the magnitude of the difference. D = decreasing cost, C = constant cost, I = increasing cost, S = slow increasing cost, F = fast increasing cost.

$[d(N) > 0,\ Z(N) < 0,\ R_A]$. This pattern was common across individuals, and was not due to any single subject. In sum, *subjects occasionally* (2%) *stopped and chose an alternative that was contrary to the evidence produced by the last observation. Whenever this occurred, they usually* (89%) *decided in favor of the sum of all the evidence.*

Pattern analysis. After purchasing n tests and observing the particular binary pattern of results $y = [z_1, ..., z_n]$, the subject had to choose between three alternatives—stop and choose R_A, stop and choose R_B, or continue to sample another test. In this section, we analyze the percentage of each choice following various patterns of test results. Two different kinds of analyses are performed. The first is based on the percentages of terminal decisions that followed each pattern produced by the initial four tests. These percentages are averaged across subjects within each condition and are shown in Tables 6A and 6B. The second analysis is based on the percentages of terminal decisions that followed each pattern produced by the four most *recent* test results. These percentages were calculated separately for each subject and condition. They are presented in Appendix A.

The percentage of terminal responses to the initial patterns was calculated for each subject as follows. Define $f(n, y)$ as the frequency that at least n observations were purchased and pattern y occurred. Define $f(n, y, R_j)$ as the frequency that at least n observations were purchased, pattern y occurred, and the subject's response was R_j. The relative frequency of R_j responses to pattern y was calculated from the

TABLE 6

Relative Frequency of Making a Terminal Decision Conditioned on the Test Pattern

A. Experiment 1

Pattern		d(n)	Decreasing			Constant			Increasing		
			f	P(A)	P(B)	f	P(A)	P(B)	f	P(A)	P(B)
11	3	1	169	2	0	177	1	0	173	1	0
11	1	−1	177	0	1	172	0	1	176	0	1
21	33	2	87	5	0	95	2	0	94	5	0
21	11	−2	93	0	3	96	0	1	96	0	5
22	31	0	79	0	0	81	0	0	77	0	0
22	13	0	83	1	0	76	0	0	79	0	0
31	333	3	51	31	0	57	27	0	52	71	0
31	111	−3	52	0	29	57	0	22	49	0	59
32	331	1	32	1	0	36	0	0	37	1	0
32	113	−1	38	0	1	37	0	0	43	1	0
33	313	1	40	3	0	38	2	0	40	1	0
33	131	−1	41	0	1	36	0	2	40	0	3
34	133	1	42	3	0	39	0	0	38	5	0
34	311	−1	38	0	2	43	0	2	37	0	6
41	3333	4	21	39	0	28	55	0	10	91	0
41	1111	−4	23	0	42	28	0	46	12	0	87
42	3331	2	16	3	0	18	6	0	6	12	0
42	1113	−2	17	0	2	16	0	2	12	0	15
43	3313	2	18	27	0	21	17	0	19	51	0
43	1131	−2	22	0	14	21	0	23	23	0	37
44	3133	2	23	17	0	21	17	0	22	53	0
44	1311	−2	22	0	14	19	0	18	20	0	40
45	1333	2	24	22	0	23	29	0	18	58	0
45	3111	−2	21	0	20	24	0	16	20	0	52
46	3131	0	16	0	0	17	0	0	18	2	2
46	1313	0	18	0	0	17	1	0	20	1	1
47	3113	0	16	1	0	19	0	0	15	0	0
47	1331	0	16	3	0	16	0	2	18	1	2
48	3311	0	15	0	2	14	0	1	17	0	3
48	1133	0	15	0	0	17	0	0	19	2	0

Table continued

TABLE 6—*Continued*

B. Experiment 2

Pattern		$d(n)$	Slow Increasing			Constant			Fast Increasing		
			f	$P(A)$	$P(B)$	f	$P(A)$	$P(B)$	f	$P(A)$	$P(B)$
11	3	1	174	0	0	177	0	0	174	1	0
11	1	−1	176	0	0	173	0	0	177	0	0
21	33	2	94	7	0	97	10	0	94	15	0
21	11	−2	94	0	8	93	0	11	98	0	12
22	31	0	80	0	0	79	0	0	79	0	0
22	13	0	82	0	0	80	0	0	79	0	0
31	333	3	55	49	0	53	54	0	46	69	0
31	111	−3	53	0	43	46	0	59	50	0	66
32	331	1	33	1	0	35	0	0	33	0	0
32	113	−1	34	0	0	36	1	0	37	0	0
33	313	1	38	1	0	38	0	0	41	4	0
33	131	−1	40	0	2	43	0	0	44	0	2
34	133	1	42	1	0	37	4	0	35	2	0
34	311	−1	42	0	1	42	0	1	38	0	1
41	3333	4	17	82	0	20	62	0	19	82	0
41	1111	−4	23	0	78	14	0	60	18	0	93
42	3331	2	16	0	0	11	5	0	7	25	0
42	1113	−2	12	1	10	11	0	8	0	0	0
43	3313	2	18	27	0	18	39	0	19	59	0
43	1131	−2	17	0	24	19	0	37	19	0	56
44	3133	2	20	32	0	22	36	0	20	56	0
44	1311	−2	20	0	31	23	0	42	23	0	59
45	1333	2	24	31	0	19	42	0	20	65	0
45	3111	−2	24	0	38	24	0	45	23	0	59
46	3131	0	18	0	0	16	0	0	20	0	0
46	1313	0	20	0	0	20	0	0	21	0	0
47	3113	0	19	0	0	17	0	0	15	0	0
47	1331	0	18	0	0	17	0	0	14	0	0
48	3311	0	15	0	0	17	0	1	14	0	0
48	1133	0	17	0	0	17	0	0	18	0	0

Note. f = pattern frequency, $P(A)$ = percentage of responses favoring alternative A, $P(B)$ = percentage of responses favoring alternative B. The percentage of reponses favoring continued testing equals $1 - P(A) - P(B)$. All statistics are averages across subjects, including the pattern frequencies. The second column indicates the particular test pattern, where 3 indicates a positive test result, and 1 indicates a negative test result. For example, 331 represents the occurrence of two positive test results followed by a negative test result.

ratio $f(n, y, R_j)/f(n, y)$. The percentage of responses equals 100 times this ratio. Note that for any given n, there are 2^n possible patterns. Consequently, the present analysis is limited to $n \leqslant 4$.

The percentages shown in Tables 6A and 6B are the arithmetic means of the individual percentages. The first column indicates an arbitrary pattern identification number. The second column indicates a particular pattern, where the integer 1 symbolizes the occurrence of a negative test result, and the integer 3 indicates the occurrence of a positive test result. (These two integers are easier to discriminate than + and − packed closely together.) For example, the pattern 3313 represents the sequence $(+1, +1, -1, +1)$ obtained from the first four observations. The third column indicates the difference between the number of positive and negative test results, $d(n)$.

Note that the patterns are organized according to the difference, $d(n)$, and the serial position of the minority outcome. Consider, for example, the patterns labeled 42, 43, 44, and 45. All produced a difference equal to 2 in magnitude. The minority outcome occurs at the fourth position for pattern 42, at the third position for pattern 43, at the second position for pattern 44, and at the first position for pattern 45. A similar serial position ordering was used for the patterns labeled 32, 33, and 34.

Columns 4, 5, and 6 in Table 6A present the mean frequency of each pattern, the mean percentage of R_A responses to each pattern, and the mean percentage of R_B responses to each pattern, respectively, for the decreasing cost condition of Experiment 1. Columns 7, 8, and 9 show these same statistics for the constant cost condition, and columns 10, 11, and 12 present the same statistics for the increasing cost condition of Experiment 1. In Table 6B, columns 4, 5, 6 are for the slow increasing cost condition, columns 7, 8, and 9 are for the constant cost condition, and columns 10, 11, and 12 are for the fast increasing cost condition of Experiment 2.

Three important findings warrant special attention:

(1) Under the constant cost conditions (as well as other conditions), subjects frequently terminated after pattern 43 despite the fact that the last pair of test results in the sequence was non-diagnostic. Assuming constant costs and unlimited number of observations, the optimal model predicts that subjects should continue sampling until a constant critical difference is exceeded. If the critical difference has not been exceeded prior to the non-diagnostic subsequence, then it cannot be exceeded after the non-diagnostic subsequence. Therefore, subjects should never terminate following non-diagnostic subsequences.

(2) Under the increasing cost conditions, subjects frequently terminated with R_A following a negative test or terminated with R_B following a positive test (pattern 42). As noted earlier (Table 5), this is another obvious violation of the optimal model because an observation was purchased at the end that could not change the optimal terminal decision.

(3). There is a strong recency effect. The percentage of terminal decisions increased across the four patterns 42, 43, 44, and 45. Thus, the tendency to stop cannot be described as a function of n and $d(n)$ alone.

The theoretical implications of these findings are discussed in more detail after the presentation of the results of Experiment 3. At present, two explanations can be ruled out. The first explanation is that subjects make occasional random guesses because of lack of attention or memory failure. Consequently, the non-zero precentages of terminal decisions after patterns 42 and 43 are due to random guesses. This conjecture cannot explain the present results for the following four reasons. First, if random guesses were responsible, then these guesses should be equally distributed across all patterns including the non-diagnostic patterns 46, 47, and 48. However, subjects almost never terminated on these particular non-diagnostic patterns. Second, subjects almost always selected the optimal terminal decision. If subjects were guessing at random, then they would occasionally make an error and choose, for example, R_A when the majority of the evidence favored S_B. Third, physical counters were always available during Experiment 2. Subjects only had to look at the video monitor to obtain a perfectly accurate count before each choice. Finally, all subjects were highly practiced and motivated to perform well because their monetary payoff depended on their decisions.

The second explanation is that the deviations result from a couple of ideosyncratic subjects, while the majority of subjects behave according to the optimal model. The individual analyses shown in Appendix A indicate that this possibility is incorrect. Under constant cost conditions, all subjects tended to make a terminal decision immediately after observing the non-diagnostic subsequence occurring at the end of pattern 43. Under fast increasing cost conditions, all but one subject showed a tendency to stop by choosing R_A immediately following a negative test, or stop by choosing R_B immediately following a positive test after observing pattern 42. (The one exception resulted from the fact that subject S9 never permitted pattern 42 to occur in the first place.) Finally, all subjects demonstrated a recency effect. In sum, the previously described conclusions drawn from Table 6 hold for the majority of subjects.

EXPERIMENT 3

Method

Experiment 3 was conducted in 1968 at the Hebrew University of Jerusalem by Rapoport, Kubovy, and Tversky. The subjects in this experiment were tested on several different decision tasks including fixed sample and sequential sampling tasks. The present article is only concerned with the results obtained from the deferred decision task. Readers interested in the results of other decision tasks are referred to Kubovy, Rapoport, and Tversky (1971), and Rapoport and Tversky (1970).

Procedure

Subjects recorded a sequence of four digit numbers representing the heights of individuals (in millimeters) randomly sampled from a population of males (labeled here as S_A) or females (labeled here as S_B). At any stage within a sequence, the subject had a choice between three alternatives—terminate the sequence by deciding that the sequence was being sampled from (a) the male population, (b) the female population, or (c) defer the decision and purchase another observation. The maximum number of observations that could be purchased was limited as described below.

The stimuli were two sets of 500 four-digit numbers constructed so as to provide the best approximation to two normal distributions with means of 1797 for the males and 1630 for the females, and a common standard deviation of 167. Exactly half of the sequences were selected from each population. All the subjects received the same set of sequences in the same order.

The payoff for incorrect decisions was always zero ($v_{AB} = v_{BA} = 0$), and the payoff for a correct decision was the same for both alternatives ($v_{AA} = v_{BB}$). The payoff for a correct decision was equal to 5, 25, or 100 monetary units. The cost for each observation was always equal to 1.0 monetary unit. The maximum number of observations that could be purchased was 10 for both payoff conditions 5 and 25, and 20 for the payoff condition 100.

Each subject participated in a total of 32 one-hour sessions, which were distributed across a variety of different types of decision tasks. The deferred decision task occurred on session 11 for payoff condition 25, session 13 for payoff condition 5, and session 16 and 25 for payoff condition 100. The experiment took place across a 2-month period with five sessions per week. Forty sequences were presented during each session.

Subjects were run simultaneously in one large room. Communication among the subjects was not permitted. The stimuli were presented by a slide projector and the responses were recorded manually. The number of points gained was computed at the end of each sequence.

Subjects

The subjects were seven volunteers—four male and three females—all first year psychology majors at the Hebrew University of Jerusalem. Subjects were paid in proportion to the number of points earned. The mean payoff was about $4.00 per hour.

Results of Experiment 3

Percentage of Correct Terminal Decisions

Table 7 shows the percentage of correct decisions separately for each subject, condition, and state. The first two columns show the percentage of correct decisions separately for each state. The last two columns show the mean and standard deviation of the number of observations purchased.

The percentage of correct decisions tends to increase as the mean number of observations purchased increases. On the basis of a simple linear regression of the percentage of correct decisions (averaged across states) on the mean number of observations purchased, each additional observation produces a 5% increase in percent correct ($R^2 = .82$).

The percentage of correct decisions for states A and B were quite different under payoff condition 5, but the direction of the bias varied across subjects. Subjects purchased very few observations under payoff condition 5, and the fluctuations in percent correct were due to chance. The percentage of correct decisions were nearly equal under payoff conditions 25 and 100. The mean number of observations purchased increased as the magnitude of the terminal payoff increased. The standard deviation was an increasing function of the mean.

TABLE 7

Percentage Correct for Each State, Mean and Standard Deviation of the Number of Observations Purchased

Payoff Condition	Subj.	$P[R_A \mid S_A]$	$P[R_B \mid S_B]$	Mean N	Std. N
5	1	0.75	0.45	0.05	0.22
25	1	0.80	0.80	5.68	2.43
100	1	0.95	0.95	8.55	2.82
5	2	0.75	0.45	0.28	0.71
25	2	0.80	0.90	4.48	2.19
100	2	0.95	0.90	8.12	4.16
5	3	0.45	0.55	0.05	0.31
25	3	0.80	0.75	3.20	1.36
100	3	0.90	1.00	6.30	2.83
5	4	0.55	0.70	1.08	0.91
25	4	0.80	0.90	4.93	1.33
100	4	0.80	1.00	8.28	2.86
5	5	1.00	0.00	0.00	0.00
25	5	0.80	0.80	4.70	1.86
100	5	0.90	1.00	7.30	3.55
5	6	0.70	0.65	1.08	0.26
25	6	0.75	0.85	2.53	0.77
100	6	0.95	0.95	4.85	2.25
5	7	0.75	0.75	1.88	0.84
25	7	0.80	0.80	3.48	1.47
100	7	0.90	1.00	5.55	2.10

Responses following the Last Observation

One of the most intruiging findings from the first two experiments was the fact that subjects occasionally made a terminal decision favoring an alternative that the majority of evidence supported, immediately after observing evidence against that same alternative (see Table 5). The following analysis was performed to determine whether a similar result occurred with normally distributed stimuli.

The evidence engendred by the observation $Z(n) = z$ was defined as the log odds, $s(n) = \ln[f_A(z)/f_B(z)]$, where f_i is the density of $Z(n)$ given state S_i and ln is the natural logarithm. Assuming an independent and identically distributed sequence of normally distributed random variables, $s(n)$ is proportional to the deviation score $X(n) = [Z(n) - 1713.5]$, where 1713.5 is the average of the two state means. Recall that N is the total number of observations purchased on a given sequence. The evidence produced by the last observation is said to favor state S_A if $X(N) > 0$ and to favor state S_B if $X(N) < 0$. The log posterior odds is said to favor S_A if $d(N) = X(1) + \cdots + X(N) > 0$, and the log posterior odds favors S_B if $d(N) < 0$.

Table 8 shows the frequencies of the eight possible combinations obtained by crossing the sign of the evidence produced by the last observation, the sign of the log posterior odds, and the terminal response pooled across subjects for each payoff condition. The first row indicates whether the evidence from the last observation favored S_A $(+1)$ or S_B (-1), the second row indicates whether the log posterior odds favored S_A $(+1)$, or S_B (-1), and the third row indicates the terminal response. The next three rows show the frequencies for payoff conditions 5, 25, and 100, respectively. The last row shows the total frequencies for each combination. For example, under payoff condition 25, subjects stopped and chose R_A when the log posterior odds favored state A but the most recent observation favored state B on 11 sequences.

TABLE 8

Frequencies of the Joint Events Formed by Crossing the Sign of the Final Test Result, the Sign of the Final Log Posterior Odds, and the Final Choice Alternative

	Event							
Final test	1	1	1	1	−1	−1	−1	−1
Final log odds	1	1	−1	−1	1	1	−1	−1
Final choice	*A*	*B*	*A*	*B*	*A*	*B*	*A*	*B*
5	52	2	1	1	0	2	7	54
25	129	4	4	12	11	6	1	113
100	130	2	6	5	12	1	1	123
Total	311	8	11	18	23	9	9	290

Note. The values listed for the final test and final log odds indicate only the sign and not the magnitude of final test and log odds. The numbers 5, 25, and 100 on the far left indicate the payoff conditions.

The combinations $[X(N) < 0, R_A]$ or $[X(N) > 0, R_B]$ occurred on 58 sequences (8.5%). (This percentage is a bit higher than the first experiment, but the difference may be partly due to the binary categorization of the continuous distribution of evidence.) The combinations $[X(N) < 0, d(N) > 0, R_A]$ or $[X(N) > 0, d(N) < 0, R_B]$ occurred on 71% of these 58 sequences. In conclusion, the results of Experiment 3 replicate and strengthen one of the most interesting findings of Experiments 1 and 2–subjects occasionally stop and choose an alternative that is contrary to the evidence produced by the most recently experienced observation. Whenever this occurs, they usually decide in favor of the log posterior odds. This finding was obtained with both binomial and normally distributed observations.

Discussion

Results from all three experiments refute the optimal model of deferred decision making behavior. Considering both (a) the difficulty of computing the optimal stopping rule and (b) the limited information processing capacity of human decision makers, this conclusion is not very surprising. We now turn to the evaluation of seven psychological models that are more limited in their information processing requirements. These seven models are discussed one at time, beginning with a brief description of each stopping rule followed by an empirical evaluation of each model.

For the first six models, parameter free qualitative tests are possble under fairly general assumptions. Specifically, we assume that the stopping criteria vary systematically as a function of the prior probabilitites, the costs of the observations, and the terminal payoffs. The criteria are also permitted to differ across subjects. Furthermore, the criteria may fluctuate stochastically across sequences for a given subject and payoff condition. However, the criteria are assumed to remain constant within a given sequence of observations.

The last model to be considered allows the criteria to vary within a sequence of observations. In the latter case, quantitative tests based on specific parameters are required to assess lack of fit.

1. *Fixed Sample Model*

A fixed sample size, M', is selected prior to each sequence. A total of M' observations are purchased, independent of the observed test results. Alternative R_A is chosen if the log posterior odds is greater than some criterion, β, and R_B is chosen if the log posterior odds is less than β. For example, in the first two experiments, subjects may sample $M' = 4$ tests and then make a terminal decision depending on whether $d(4)$ is greater or less than $\beta = 0$. Pitz *et al.* (1969), Swenson and Thomas (1974), and Stone (1960) have discussed this model.

The fixed sample model predicts that the decision to stop purchasing observations is independent of the observed pattern of test results. Consequently, this model can account for the fact that subjects occasionally stop on differences of zero or one in magnitude (Table 4), and the fact that subjects stop and choose one alternative immediately after experiencing evidence that favors the contrast alternative (Table 5).

A problem arises from the prediction that the probability of termination is equal for all sequences of identical length. Referring to Table 6, the probability of termination following the highly diagnostic pattern 41 should be equal to that for non-diagnostic patterns such as 46, 47, and 48. Similarly, the probability of stopping after pattern 31 should be equal to that for patterns 32, 33, and 34. It is clear that this prediction is incorrect.

2. *Constant Bound Random Walk Model*

Observations are purchased until either (a) the log posterior odds is greater than or equal to a fixed upper bound, α, in which case response R_A is chosen, or (b) the log posterior odds is less than or equal to a lower bound, $-\beta$, in which case response R_B is chosen. For example, in the first two experiments, subjects may continue purchasing observations until $d(n) \geqslant \alpha = 3$ or until $d(n) \leqslant -\beta = -2$. This model was proposed by Audley and Pike (1965), Laming (1968), Pitz *et al.* (1969), and Stone (1960). Link and Heath (1975) proposed a similar model except that the cumulative evidence does not necessarily equal the log posterior odds.

The constant bound random walk model predicts that subjects will never terminate on sequences that end with a positive log posterior odds that is less than or equal to a positive log posterior odds occurring earlier in the sequence. Nor should they terminate on sequences that end with a negative log posterior odds that is greater than or equal to a negative log posterior odds occurring earlier in the sequence. Referring to Table 6, a terminal decision should never occur following patterns 42, 43, 46, 47, and 48 simply because the last two tests do not increase the magnitude of the log posterior odds.

This model can explain why subjects are more likely to stop after diagnostic patterns such as 31 and 41 as compared to non-diagnostic patterns such as 46, 47, and 48 in Table 6. It also explains why subjects are more likely to terminate after patterns 44 and 45 than pattern 42, despite the fact that terminal differences are equal. The log posterior odds increases in magnitude at the end for patterns 44 and 45, but it decreases at the end for pattern 42.

The problem with this model is that subjects should never stop immediately after observing a non-diagnostic subsequence. Referring to Table 6, subjects should never stop on the patterns numbered 42 and 43, but all subjects frequently do stop on these patterns. Furthermore, Pitz *et al.* (1969) have noted that according to the random walk model, subjects should never stop with a terminal difference equal to zero or one after purchasing more than one test. As can be seen in Table 4, this prediction is clearly incorrect.

3. *Fixed Forgetting Model*

Morgan and Robertson (1980) described a version of the random walk model with memory limited to the m most recent observations. Subjects continue purchasing tests until the log odds summed across the m most recent tests exceeds an upper criterion or crosses below a lower criterion. For example, in the first two experiments, sampling may continue until the sum of the last $m = 4$ tests exceeds 2 in magnitude.

This model also predicts that subjects should never stop on patterns 42 and 43 in Table 6. This is true for any value of m. As noted in the preceding paragraph, this prediction is clearly incorrect.

4. *Hybrid Random Walk and Fixed Sample Model*

Pitz *et al.* (1969) proposed the following hybrid rule: the terminal decision is determined by a constant bound random walk strategy as long as the number of observations sampled is less than some self imposed limit, M'. However, if a decision was not reached after M' samples, then the random walk process is aborted, and a decision is based on a fixed sample decision rule. For example, in the first two experiments, sampling may continue until either the magnitude of $d(n)$ equals a critical difference of 3 or until the number purchased equals $M' = 4$, whichever comes first.

A simple test of this model can be performed by considering sequences ending with a non-diagnostic subsequence (patterns 42, 43, 46, 47, 48 in Table 6). For these sequences, the hybrid model predicts that subjects terminate only if the number of tests purchased exceeds the fixed sample size M'. In this case, the hybrid model makes exactly the same predictions as the pure fixed sample model—the probability of terminating after pattern 43 should equal that for patterns 46, 47, and 48. As can be seen in the tables, this is clearly incorrect.

5. *Horse Race or Accumulator Model*

A separate evidence accumulator is used for each alternative. A positive accumulator sums the log odds corresponding to the subsequence of observations that produced positive evidence favoring state S_A. A negative accumulator sums the log odds corresponding to the subsequence of observations that produced negative evidence favoring state S_B. Observations are purchased until either (a) the positive accumulator is greater than or equal to positive criterion α, in which case R_A is chosen, or (b) the negative accumulator is less than or equal to a negative criterion $-\beta$, in which case R_B is chosen. For example, in the first two experiments, testing may continue until either $\alpha = 4$ positive test results occur or until $\beta = 3$ negative test results occur in any order. Audley and Pike (1965), Pike (1968), Pitz *et al.* (1969), and Vickers (1979) have described versions of this model under various names such as accumulator model, counter model, recruitment model, and "world series" model.

An unbiased accumulator model cannot explain the fact that subjects occasionally stop on a log posterior odds equal to zero, since this event is impossible if $\alpha = \beta$. To account for the fact that subjects occasionally stopped with a difference equal to zero in the study by Pitz *et al.* (1969) or in our first two experiments, it is necessary to assume that a bias ($\alpha \neq \beta$) was present on some sequences.

The accumulator model can explain why subjects occasionally stop on patterns ending with a non-diagnostic subsequence such as pattern 43 in Table 6. For example, subjects would stop on pattern 3313 if the criteria were set to $\alpha = \beta = 3$, so that the positive counter would accumulate three positive test results for the first time on the fourth observation.

Note that the accumulator model implies that the terminal response always agrees with the sign of the log odds produced by the last observation. Thus, the accumulator model predicts that response R_A should never ocur if the terminal log posterior odds is positive but the evidence from the last observation is negative. Similarly, the response R_B should never occur if the terminal log posterior odds is negative but the evidence from the last observation is positive. The results in Tables 5 and 8 indicate that this prediction is systematically violated for both binomial and normal distributions.

Referring to Table 6, the accumulator model predicts that subjects should never stop and choose R_A following the pattern 3331, and the response R_B should never occur following the pattern 1113. This prediction is frequently violated under the increasing cost conditions for pattern number 42.

The accumulator model also predicts an ordering of the cumulative frequency distributions shown in Table 2. If observation costs only influence the criterion bounds, then the cumulative frequencies for one of the cost conditions should always exceed the cumulative frequencies for another cost condition (see Appendix B). For example, if the cumulative frequencies for the decreasing cost condition exceed the cumulative frequencies for the constant cost condition at $n < 5$, then this dominance relation should also be maintained for all $n > 5$. As can be seen in Table 2, the cumulative distributions crossover at $n = 5$ which contradicts the predicted distribution ordering property.

6. *Simple and Cumulative Runs Models*

There are two different versions of the runs model. The simple version (cf. Audley & Pike, 1965; Estes, 1960) assumes that subjects continue to purchase observations until either (a) the log odds of each of the α most recent observations are all positively signed, in which case response R_A is chosen, or (b) the log odds of each of the β most recent observations are all negatively signed, in which case response R_B is chosen. For example, in the first two experiments, testing may continue until a consistent run of $\alpha = 3$ positive tests occur, or until a consistent run of $\beta = 2$ consecutive negative tests occur.

This model predicts that a terminal decision should never occur following a run length which is less than or equal to earlier run lengths of the same sign. Referring to Table 6, a terminal decision should never occur following patterns 43, 46, or 47. The results for the pattern 43 indicate that this prediction is wrong.

The cumulative runs models states that subjects continue purchasing observations until either (a) the posterior log odds is consistently positive on the α most recent observations, in which case R_A is chosen, or (b) the posterior log odds is consistently negative on the β most recent observations, in which case response R_B is chosen. For example, in the first two experiments, testing may continue until $d(n)$ is consistently positive on the last $\alpha = 4$ observations, or until $d(n)$ is consistently negative on the last $\beta = 3$ observations.

The cumulative runs model can easily explain why subjects stop on patterns such as 42 or 43 in Table 6. For example, if $\alpha = \beta = 4$ and pattern 3331 is observed, then $d(1) > 0$, $d(2) > 0$, $d(3) > 0$, and $d(4) > 0$ which meets the criterion for stopping and choosing R_A. Furthermore, they would not stop on the patterns 46, 47, and 48 since $d(n)$ is not consistently positive or negative.

The problem with this model is that it predicts that subjects should be equally likely to stop on patterns 41, 42 and 43, and more likely to stop on pattern 42, than 44 and 45. However, Table 6 indicates that the tendency to stop on pattern 42 was far below that for patterns 41, 43, 44, and 45. Also, the cumulative runs model predicts that subjects should never terminate with a difference equal to zero after purchasing one test, but the results reported by Pitz *et al.* (1969), and to a lesser extent the present study, indicate that occasionally they are willing to do so.

7. *Myopic Decision Rules*

Recall that the optimal model prescribes purchasing additional observations until the expected loss of making a terminal decision is less than the expected loss after purchasing one or more additional observations. This decision rule requires planning many steps ahead—until the termination of the task—because the possibility of purchasing any number of additional observations must be considered (see Rapoport & Burkheimer, 1971).

There is strong evidence suggesting that the decision maker's planning horizon is severely limited (Rapoport, 1966). Suppose that the decision maker considers only a small number of steps (say m steps) and uses the following stopping rule: If the expected loss of making a terminal decision on the basis of the current information is less than the expected loss of making a terminal decision after purchasing at most m more observations, then stop; otherwise, purchase another observation. For example, if $m = 1$, then this rule prescribes purchasing another observation if and only if the expected loss of making a terminal decision on the basis of the current information is more than the expected loss of making a terminal decision after purchasing just one more observation. Of course, m may be larger than one step, but m must be fairly small ($m \leqslant 3$) to make predictions that differ from those generated by the optimal rule for each cost condition.

The problem with this decision rule is that it always prescribes stopping with a terminal difference smaller than $m + 1$ in magnitude. For example, if $m = 1$ then this rule prescribes stopping after purchasing the very first observation (see Appendix C). However, the results shown in Table 4 indicate that terminal differences larger than 4 in magnitude frequently occurred, which is contrary to the supposition that m is fairly small.

The following myopic decision rule minimizes future planning: If the expected loss of making a terminal decision after purchasing n observations is less than or equal to the sum of the costs of $n+1$ observations, then a terminal decision is made; otherwise another observation is purchased. To be more specific, let $P[S_i \mid d(n)=d]$ be the posterior probability of state S_i conditioned on the observed difference $d(n)=d$. If R_A is chosen, then the expected loss equals

$$J_A(n, d) = P[S_B \mid d(n)=d] \cdot [v_{AA} - v_{BA}] + \sum c(k), \qquad k=1, n$$

and if R_B is chosen then the expected loss equals

$$J_B(n, d) = P[S_A \mid d(n)=d] \cdot [v_{BB} - v_{AB}] + \sum c(k), \qquad k=1, n.$$

According to the myopic rule, another observation is purchased if and only if $J(n, d) = \min[J_A(n, d), J_B(n, d)] > c(1) + \cdots + c(n+1)$. Following the decision to stop purchasing observations, R_A is chosen over R_B if $J_A(n, d) < J_B(n, d)$, R_B is chosen over R_A if the inequality is reversed, and the choice is made randomly if the two expected losses are equal.

Under the conditions of Experiments 1 and 2, the myopic stopping rule can be reformulated as follows (see Appendix D). For $c(n+1) > 0$, define $r(n) = (v_{ii} - v_{ij})/c(n+1)$ as the ratio of the cost of an incorrect terminal decision to the cost of the next test. Then according to the myopic stopping rule, another observation is purchased if and only if $|d(n)| < \delta(n)$, where

$$\delta(n) = \begin{cases} \ln[r(n) - 1]/\ln(p/q) & \text{for} \quad r(n) > 1, \\ 0 & \text{for} \quad r(n) \leqslant 1, \end{cases} \tag{1}$$

$p = P[Z(n) = +1 \mid S_A]$, and $q = 1 - p$. (Recall that $Z(n)$ is the nth test result in a sequence of observations.) Following the decision to stop purchasing observations, R_A is chosen if $d(n) > 0$, R_B is chosen if $d(n) < 0$, and the choice is made randomly if $d(n) = 0$.

The ratio $r(n)$ is not defined for $c(n+1) = 0$. This problem is circumvented by postulating a subjective cost of waiting to make a terminal decision. The longer one has to delay the terminal response, the more impatient one becomes, especially when the stopping time is uncertain (cf. Osuna, 1985). Define $w(n)$ as the increment to the cumulative waiting cost produced by purchasing the nth test. Although this function is unknown (except that $w(0) = 0$), we approximated it by the linear function, $w(n) = b \cdot n$, (where $b > 0$ is a scaling constant that expresses the cost of waiting in monetary units). The subjective cost of purchasing the nth observation is defined as $c'(n) = c(n) + w(n)$. Hereafter, the stopping criterion, $\delta(n)$, is defined by Eq. 1 with $c'(n)$ substituted for $c(n)$.

The stopping rule described above is deterministic; for a fixed pattern of test results, the probability of making a terminal decision will be zero or one. The results in Table 6 indicate that variability must be introduced into the decision process. For $\delta(n) > 0$, it may be difficult to evaluate the difference $z = |d(n)| - \delta(n)$. In this case, the probability of stopping was represented by logistic function

$$F(z) = [1 + \exp(-\Theta \cdot z)]^{-1}, \qquad \text{with} \quad z = |d(n)| - \delta(n).$$

The coefficient, $\Theta > 0$, determines the degree of discriminability. Nearly perfect discriminability can be achieved as a special case by setting Θ to a very large number.

If the cost of the next test is more than the loss expected by random guessing ($\delta(n) = 0$), then there is little doubt that it is time to stop testing. (In fact, subjects almost never purchased another test under these conditions.) Therefore, the probability of stopping was set to 1.0 for $\delta(n) = 0$.

There is one last consideration. The impatience of the subject may wax and wane during a session. For example, subjects may become more impatient towards the end of a session or the end of a week. This heterogeneity was taken into account by defining b (the scaling constant for waiting cost) as a random variable. The distribution of this random variable was approximated by a binomial distribution over a set of 26 values ($\frac{0}{10}, \frac{1}{10}, ..., \frac{25}{10}$). A binomial distribution was chosen because it produces a skewed bounded distribution that depends on only one parameter (the mean, denoted μ). The upper bound of the set ($b = 2.5$) was chosen to be large relative to the observation costs (e.g., if $b = 2.5$ then $w(10) = 25$ cents, which is a large cost for a single test).

In sum, the myopic decision rule, requires estimation of two parameters—the discriminability coefficient, Θ, and the mean of the distribution of the waiting cost scaling constant, μ. The discriminability coefficient was assumed to remain constant across all conditions. However, the mean of the scaling constant was allowed to vary across cost conditions because of variation in the expected waiting time for different conditions (cf. Osuna, 1985). Subjects expected that a decision would be reached within a small number of tests with the fast increasing cost conditions, but they were unable to predict how long they would have to wait for the decreasing or constant cost condtions.

Test of the myopic model. The model described above was evaluated in two phases. First, the parameters were estimated by fitting the model to joint relative frequency distributions presented in Table 4, separately for each individual. Then, predictions for the relative frequencies shown in Table 6 were calculated using the same parameters estimated from Table 4. The frequencies in the two tables are not independent, but the results in Table 4 are pooled over patterns, while the results in Table 6 are conditioned on specific patterns. The details of the model fitting procedure are given in Appendix D.

The discriminability parameter for all subjects and conditions was fixed at $\Theta = 2.0$. The means of waiting cost scaling parameter, μ, for each subject and condition are shown in Table 9. Note that while it was necessary to allow μ to vary across conditions in Experiment 1, it was possible to fix μ across conditions in Experiment 2 without substantially reducing the fit of the model. Further discussion of these parameters is delayed until the predictions are discribed.

The predictions for the joint frequency distribution of $d(N)$ and N can be evaluated by comparing the two adjacent rows for observed (denoted O) and predicted (denoted P) frequencies in Table 4. Although the fit is not perfect, no large systematic deviations are apparent. Note that the model generates non-zero percentages of stopping on a difference less than or equal to one in magnitude after $N > 1$ similar to the observed results. At the same time, the model correctly predicts that large terminal differences frequently occur under certain cost conditions. Also note that the systematic changes in the joint frequencies across observation cost conditions in Experiment 2 are accurately reproduced, even though both parameters were fixed across conditions in this experiment. (In other words, only two parameters were required to fit all of the joint frequencies in Experiment 2.)

Table 10 shows the observed (denoted O) and predicted (denoted P) percentages of sequences that subjects stopped after observing a particular pattern on the first four tests. The first column indicates the type of pattern (the pattern number refers to the same patterns listed in Table 6), and the last six columns indicate the observation cost conditions for Experiments 1 and 2. Because of the symmetry of the results, the data were averaged across adjacent patterns associated with the same

TABLE 9

Parameter Estimates of the Mean of the Distribution of the Scaling Coefficients for the Subjective Cost of Waiting

		Experimental Condition		
Expt.	Subj.	1	2	3
1	1	0.25	0.25	0.22
1	2	0.95	1.02	0.77
1	3	0.18	0.20	0.82
1	4	0.18	0.30	0.72
1	5	0.18	0.12	0.72
1	6	0.33	0.37	0.62
2	5	0.85	0.85	0.85
2	7	0.18	0.18	0.18
2	8	0.57	0.57	0.57
2	9	0.72	0.72	0.72

Note. For Experiment 1, conditions 1, 2, and 3 refer to the decreasing, constant, and increasing cost conditions; for Experiment 2, conditions 1, 2, and 3 refer to the slow, constant, and fast increasing cost conditions, respectively.

TABLE 10

Predicted and Observed Relative Frequencies of Stopping and Making a Terminal Decision Conditioned on the Preceding Pattern of Test Results

		Expt. 1			Expt. 2		
Pattern		*D*	*C*	*I*	*S*	*C*	*F*
41	*O*	41	51	89	80	61	88
41	*P*	42	43	94	74	73	85
42	*O*	3	4	14	6	7	13
42	*P*	6	7	36	20	19	29
43	*O*	21	20	44	26	38	58
43	*P*	8	9	43	27	26	37
44	*O*	16	18	47	32	39	58
44	*P*	9	9	44	28	28	38
45	*O*	21	23	55	35	44	62
45	*P*	9	9	44	28	28	38
46	*O*	0	1	3	0	0	0
46	*P*	0	0	2	1	1	2
47	*O*	2	1	2	0	0	0
47	*P*	0	0	2	1	1	2
48	*O*	1	1	3	0	1	0
48	*P*	0	0	2	1	1	2

Note. The pattern numbers correspond to the pattern numbers shown in Table 6. O = observed percentage, P = predicted percentage.

labels in Table 6. For example, the third pair of rows in Table 10 show the observed and predicted percentage of sequences on which subjects stopped after observing pattern number 43, i.e., after observing either pattern 3313 or pattern 1131.

Table 10 shows that predictions for highly diagnostic (41) or non-diagnostic (46, 47, 48) patterns are fairly accurate. Also, the model correctly predicts that subjects will occasionally stop on pattern 43 under constant and decreasing cost conditions, and more frequently stop on both patterns 42 and 43 under increasing cost conditions. However, it is also apparent that there are systematic discrepancies. For example, although the model correctly predicts that the percentage for pattern 42 is less than that for pattern 43, the predicted difference is much smaller than the observed difference. Although the model correctly predicts that the percentages for patterns 43, 44, and 45 are approximately equal, the predicted values are systematically lower than the observed values.

Are the discrepancies shown in Table 10 due to the myopic decision rule per se or the axiliary assumptions? To explore this possibility, we varied the assumption concerning the subjective cost of waiting (by comparing models with $w(n) = b$, $w(n) = b \cdot n$, or $w(n) = b \cdot n^2$), but the linear function ($w(n) = b \cdot n$) produced the best overall results. In addition, the model was fit to the data shown in Table 10 using new parameters. The results of this analysis indicated that the most serious discrepancies were eliminated by increasing the discriminability parameter (e.g., setting $\alpha = 4$). However, increasing the discriminability substantially reduced the fit to the results in Table 4. Thus, the problem does not seem to be fitting the results in Table 10, but rather the model is unable to account for the results in both tables simultaneously.

Another problem for the model is the variation of the mean of the waiting cost scaling constant, μ, across conditions in Experiment 1 for subjects 3–6. Although the means for the decreasing and constant cost conditions were approximately equal, both of these means tended to be smaller than the mean for the increasing cost condition. Possibly, the increased mean for the increasing cost condition indicates that the stopping criteria decrease in magnitude more rapidly than predicted on the basis of observation costs alone.

In sum, the myopic decision rule captures several of the qualitative features of the data, yet there are systematic discrepancies from the model. Possibly with some subjects, planning beyond the next observation occurred on some of the sequences. However, the myopic model may be considered better than the other psychological models for two reasons. First, all of the other models can be rejected on the basis of *parameter free* predictions that are disconfirmed by the known facts, but the parameter free predictions of the myopic decision model remain consistent with the known facts. Second, the myopic decision model provides a simple way to describe how changes in observation costs influence the stopping rule, whereas the other psychological models require ad hoc assumptions that relate changes in observation costs to the parameters of the stopping rule.

CONCLUSION

This research investigated psychological models of deferred decision making. A new method of evaluating these models, called pattern analysis, proved to be very diagnostic. The fixed sample, constant bound random walk, hybrid fixed sample-constant bound random walk, fixed forgetting, horse race or accumulator, simple and cumulative runs models all generate parameter free predictions that were falsified by the results of three experiments. A myopic stopping rule provided a better account of these results. The essential idea of the latter model is that subjects continue purchasing observations until the expected loss of a terminal decision after purchasing n observations is less than the sum of costs of $n+1$ observations. The shortsighted nature of this stopping rule allows one to wastefully purchase observations that have no impact on the final decision.

Are the conclusions of the present study relevant to more complex extralaboratory settings? Although a complete answer to this question must await further research, two preliminary answers are worth mentioning. First, any psychological theory of deferred decision making general enough to be applicable to complex settings should be able to account for behavior in simple laboratory settings, provided that subjects are sufficiently well motivated (as they were in the present studies). Second, Young, Fried, Hershey, Eisenberg, and Williams (1985) reported results similar to the present findings that were based on decisions made by physicians who were asked to answer a questionnaire containing complex realistic medical scenarios. Young *et al.* found that physicians would occasionally fail to recommend a treatment for a disease before observing a piece of evidence, but later decide to recommend treatment after observing a piece of evidence that actually lowered the probability that the disease was present. This result is similar to the present finding that subjects occasionally make a terminal decision in favor of one disease immediately following evidence that favors the alternative disease.

APPENDIX A

Individual analyses of responses to patterns. The results in Table A were calculated for each subject in Experiments 1 and 2 as follows. Recall that $y = [z_1, \ldots, z_n]$ represents a particular pattern of test results obtained from the first n observations. The pattern $y^* = [z_{n-3}, z_{n-2}, z_{n-1}, z_n]$ represents the results of the four most *recent* tests in the pattern y. Define $f(n, y^*)$ as the frequency that at least n observations were purchased and the pattern y^* occurred on the four most recent tests. Define $f(n, y^*, R_j)$ as the frequency that at least n observations were purchased, y^* occurred, and the response R_j was chosen. The relative frequency of a terminal decision to the pattern y^* was obtained from the ratio $\sum f(n, y^*, R_j)/\sum f(n, y^*)$, where the summation ranges from $n = 4$ to $n = 8$. The percentage of responses equals 100 times the relative frequency.

Pooling across n was needed to obtain frequencies large enough to provide reasonable estimates for individual analyses. The analyses were limited to $n = 8$ because few sequences extended beyond $n = 8$ in Experiment 2.

The first column in Table A indicates the subject number, and the fourth column is the difference between the number of positive and negative tests based on the last four observations. The remaining columns are defined exactly the same as in Table 6.

APPENDIX B

The purpose of this appendix is to derive a cumulative distribution ordering property from the horse race model. When the test results are binomially distributed, as they were in Experiments 1 and 2, then the horse race stopping rule can be defined as follows. Let $A(n)$ and $B(n)$ be the number of positive and negative test

TABLE A

Relative Frequency of Stopping Conditioned on the Test Pattern Appearing on the Four Most Recent Tests

Subj.	Pattern		d	Decreasing			Constant			Increasing		
				f	P(A)	P(B)	f	P(A)	P(B)	f	P(A)	P(B)
1	1	3333	4	83	41	0	90	42	0	24	83	0
1	1	1111	−4	95	0	38	82	0	44	42	0	76
1	2	3331	2	65	3	0	47	4	0	21	19	5
1	2	1113	−2	60	0	0	63	0	2	30	3	27
1	3	3313	2	69	16	0	68	10	0	46	39	0
1	3	1131	−2	82	0	4	76	0	9	62	2	31
1	4	3133	2	86	13	0	73	14	0	54	52	0
1	4	1311	−2	100	0	12	85	0	15	66	0	53
1	5	1333	2	93	15	0	88	30	0	47	57	0
1	5	3111	−2	101	0	12	98	0	21	53	0	53
1	6	3131	0	71	0	0	64	0	3	67	3	16
1	6	1313	0	78	1	0	73	1	0	66	26	2
1	7	3113	0	76	0	0	62	0	0	52	6	2
1	7	1331	0	67	0	0	67	1	0	46	4	9
1	8	3311	0	69	0	0	63	0	0	43	0	19
1	8	1133	0	59	2	0	71	3	0	47	28	0
2	1	3333	4	17	59	0	19	79	0	13	100	0
2	1	1111	−4	21	0	76	24	0	79	12	0	100
2	2	3331	2	18	0	0	18	6	0	7	0	0
2	2	1113	−2	15	0	7	20	0	5	9	22	0
2	3	3313	2	56	41	0	49	31	0	45	58	0
2	3	1131	−2	56	0	27	55	0	40	45	0	56
2	4	3133	2	59	59	0	63	44	0	51	75	0
2	4	1311	−2	48	0	50	66	0	50	53	0	77
2	5	1333	2	54	52	0	69	57	0	41	71	0
2	5	3111	−2	60	0	52	54	0	50	38	0	61
2	6	3131	0	59	0	8	49	0	4	38	0	24
2	6	1313	0	55	9	0	40	7	0	38	13	3
2	7	3113	0	34	3	0	45	0	0	30	10	0
2	7	1331	0	47	0	9	45	0	0	42	0	26
2	8	3311	0	41	0	12	36	0	17	31	0	32
2	8	1133	0	43	7	0	43	14	0	42	19	0

Table continued

TABLE A—*Continued*

				Decreasing			Constant			Increasing		
Subj.	Pattern		d	f	$P(A)$	$P(B)$	f	$P(A)$	$P(B)$	f	$P(A)$	$P(B)$
3	1	3333	4	91	42	0	109	51	0	10	100	0
3	1	1111	−4	100	0	49	105	0	53	10	0	100
3	2	3331	2	58	0	0	61	0	0	6	0	0
3	2	1113	−2	82	0	0	56	0	2	9	0	22
3	3	3313	2	88	5	0	91	2	0	43	56	0
3	3	1131	−2	88	0	2	82	0	2	35	6	49
3	4	3133	2	99	10	0	96	13	0	43	72	0
3	4	1311	−2	106	0	11	88	0	9	33	0	76
3	5	1333	2	94	12	0	81	15	0	46	78	0
3	5	3111	−2	99	0	10	97	0	22	40	0	65
3	6	3131	0	75	0	0	90	0	1	38	3	21
3	6	1313	0	77	0	0	84	1	0	44	36	2
3	7	3113	0	80	0	0	73	0	0	36	25	0
3	7	1331	0	82	0	0	82	0	0	33	3	12
3	8	3311	0	78	0	3	64	0	0	40	0	27
3	8	1133	0	71	1	0	72	3	0	36	17	0
4	1	3333	4	89	37	0	67	52	0	19	100	0
4	1	1111	−4	104	0	34	92	0	49	15	0	100
4	2	3331	2	59	0	0	60	0	0	10	20	0
4	2	1113	−2	54	0	0	63	0	0	18	0	33
4	3	3313	2	80	7	0	74	7	0	36	58	0
4	3	1131	−2	71	0	4	80	0	7	44	0	50
4	4	3133	2	87	10	0	78	15	0	48	65	0
4	4	1311	−2	91	0	9	81	0	10	39	0	59
4	5	1333	2	77	14	0	80	20	0	41	76	0
4	5	3111	−2	85	0	12	81	0	17	41	0	71
4	6	3131	0	68	0	1	72	0	1	40	5	38
4	6	1313	0	58	2	0	71	4	0	38	18	8
4	7	3113	0	59	0	0	70	1	0	28	21	14
4	7	1331	0	69	0	0	61	0	0	29	7	14
4	8	3311	0	60	0	5	70	0	1	34	0	41
4	8	1133	0	58	0	0	74	3	0	32	28	6

Table continued

TABLE A—*Continued*

Subj.	Pattern		d	Decreasing f	Decreasing $P(A)$	Decreasing $P(B)$	Constant f	Constant $P(A)$	Constant $P(B)$	Increasing f	Increasing $P(A)$	Increasing $P(B)$
5	1	3333	4	97	24	0	137	21	0	18	72	0
5	1	1111	−4	109	0	23	141	0	13	14	0	79
5	2	3331	2	67	0	0	85	0	0	10	40	0
5	2	1113	−2	53	0	2	79	0	3	19	0	32
5	3	3313	2	59	14	0	93	3	0	32	66	0
5	3	1131	−2	78	0	1	78	0	3	43	0	70
5	4	3133	2	74	14	0	98	13	0	45	62	0
5	4	1311	−2	77	0	13	80	0	9	46	0	59
5	5	1333	2	84	17	0	112	13	0	43	74	0
5	5	3111	−2	73	0	14	87	0	5	40	0	72
5	6	3131	0	55	0	2	85	2	1	33	15	24
5	6	1313	0	70	0	0	87	1	0	38	13	3
5	7	3113	0	78	3	0	77	0	0	30	17	10
5	7	1331	0	63	3	0	62	0	2	34	9	24
5	8	3311	0	73	0	0	73	0	0	25	4	36
5	8	1133	0	57	2	0	86	2	0	29	24	7
6	1	3333	4	60	52	0	70	64	0	26	96	0
6	1	1111	−4	110	0	43	85	0	52	37	0	86
6	2	3331	2	45	2	0	56	7	0	9	11	11
6	2	1113	−2	61	2	2	57	0	2	22	0	5
6	3	3313	2	67	22	0	72	31	0	26	62	4
6	3	1131	−2	70	0	16	71	0	23	45	0	56
6	4	3133	2	69	12	0	80	21	0	50	50	0
6	4	1311	−2	85	0	13	71	0	25	42	0	60
6	5	1333	2	79	29	0	93	37	0	45	69	0
6	5	3111	−2	81	0	19	69	0	12	48	0	56
6	6	3131	0	58	0	2	50	0	4	39	3	13
6	6	1313	0	59	3	0	45	2	0	54	13	6
6	7	3113	0	68	3	0	67	1	1	31	23	13
6	7	1331	0	75	1	8	67	0	1	47	4	30
6	8	3311	0	56	0	4	66	0	9	44	2	14
6	8	1133	0	86	5	0	58	2	0	35	14	3

Table continued

TABLE A—*Ccontinued*

Subj.	Pattern		d	Decreasing f	Decreasing P(A)	Decreasing P(B)	Constant f	Constant P(A)	Constant P(B)	Increasing f	Increasing P(A)	Increasing P(B)
5	1	3333	4	52	100	0	41	76	0	0	0	0
5	1	1111	−4	40	0	97	29	0	76	6	0	100
5	2	3331	2	23	0	4	23	4	0	4	75	0
5	2	1113	−2	28	4	0	27	0	0	0	0	0
5	3	3313	2	50	40	0	57	56	2	30	97	0
5	3	1131	−2	58	0	38	53	0	51	30	0	100
5	4	3133	2	60	65	0	48	54	0	32	94	0
5	4	1311	−2	64	0	48	72	0	53	36	0	92
5	5	1333	2	66	52	0	47	47	0	27	100	0
5	5	3111	−2	56	0	48	65	0	49	39	0	95
5	6	3131	0	38	0	21	56	0	18	38	0	29
5	6	1313	0	45	11	0	49	14	0	34	35	0
5	7	3113	0	53	15	2	47	19	0	18	28	6
5	7	1331	0	41	0	22	46	0	15	33	0	45
5	8	3311	0	36	0	25	40	0	13	25	0	20
5	8	1133	0	56	18	0	52	21	0	29	17	0
7	1	3333	4	86	43	0	104	49	0	71	79	0
7	1	1111	−4	103	0	58	85	0	47	63	0	79
7	2	3331	2	71	0	0	58	0	0	33	3	6
7	2	1113	−2	59	0	0	63	0	0	48	2	0
7	3	3313	2	94	6	0	78	9	0	59	27	0
7	3	1131	−2	78	0	4	87	0	6	74	0	32
7	4	3133	2	89	28	0	93	13	0	83	35	0
7	4	1311	−2	105	0	11	91	0	14	59	0	31
7	5	1333	2	88	14	0	88	25	0	84	45	0
7	5	3111	−2	96	0	17	102	0	21	63	0	40
7	6	3131	0	111	0	2	73	0	1	70	0	13
7	6	1313	0	89	2	0	82	1	0	81	9	0
7	7	3113	0	86	1	0	69	0	0	58	2	0
7	7	1331	0	72	0	0	70	0	0	51	0	12
7	8	3311	0	67	0	1	72	0	3	46	0	20
7	8	1133	0	73	4	0	62	5	0	61	13	0

Table continued

TABLE A—*Continued*

				Decreasing			Constant			Increasing		
Subj.	Pattern		*d*	*f*	*P*(*A*)	*P*(*B*)	*f*	*P*(*A*)	*P*(*B*)	*f*	*P*(*A*)	*P*(*B*)
8	1	3333	4	27	93	0	31	87	0	18	89	0
8	1	1111	−4	27	0	96	22	0	91	18	0	94
8	2	3331	2	16	6	0	25	4	0	19	16	5
8	2	1113	−2	21	0	5	25	0	8	15	0	13
8	3	3313	2	42	38	0	45	29	0	36	39	0
8	3	1131	−2	51	0	24	54	0	26	33	0	30
8	4	3133	2	53	38	0	65	42	0	54	44	0
8	4	1311	−2	70	0	49	59	0	51	53	0	62
8	5	1333	2	60	52	0	58	41	0	61	61	0
8	5	3111	−2	48	0	48	56	0	45	44	0	50
8	6	3131	0	62	2	3	35	0	3	36	3	11
8	6	1313	0	61	7	0	44	11	0	48	13	4
8	7	3113	0	51	4	0	47	9	2	39	5	0
8	7	1331	0	47	6	0	45	0	2	44	7	16
8	8	3311	0	39	0	15	46	2	7	39	0	23
8	8	1133	0	42	17	0	37	14	0	48	21	0
9	1	3333	4	12	100	0	7	86	0	0	0	0
9	1	1111	−4	3	0	100	4	0	100	0	0	0
9	2	3331	2	7	0	0	6	17	0	0	0	0
9	2	1113	−2	11	0	18	3	0	0	0	0	0
9	3	3313	2	46	41	0	33	61	0	35	83	3
9	3	1131	−2	23	0	57	47	0	45	34	0	88
9	4	3133	2	54	52	0	57	47	0	48	58	0
9	4	1311	−2	41	0	56	48	0	48	37	0	73
9	5	1333	2	53	70	0	51	76	0	37	89	0
9	5	3111	−2	63	0	84	47	0	91	41	0	98
9	6	3131	0	41	0	5	42	0	0	42	2	14
9	6	1313	0	43	5	0	57	0	2	45	0	2
9	7	3113	0	30	0	0	56	0	2	31	0	16
9	7	1331	0	47	0	0	51	0	6	35	9	0
9	8	3311	0	43	0	14	52	0	6	30	0	13
9	8	1133	0	40	5	0	41	10	0	26	15	4

results, respectively, that have been observed after purchasing n observations, and note that $n = A(n) + B(n)$ so that $B(n) = n - A(n)$. A terminal decision favoring R_A occurs after purchasing N observations if and only if $A(N) = \alpha$ and $B(n) < \beta$, and a terminal decision favoring R_B occurs if and only if $B(N) = \beta$, and $A(N) < \alpha$.

The probability that the total number purchased, N, is less than or equal to n, denoted $P[N \leqslant n]$, will be a function of the stopping criteria α and β. The proposition to be proved asserts:

If a horse race stopping rule is used, then $P[N \leqslant n \mid \alpha, \beta]$ is a non-increasing function of α and β.

The proof is straightforward.

$$\begin{aligned} P[N > n \mid \alpha, \beta] &= P[A(n) < \alpha] \cdot P[B(n) < \beta \mid A(n) < \alpha] \\ &= P[A(n) < \alpha] \cdot P[A(n) > n - \beta \mid A(n) < \alpha] \\ &= P[n - \beta < A(n) < \alpha] = \sum P[A(n) = r], \end{aligned}$$

where the summation ranges from $n - \beta + 1$ to $(\alpha - 1)$. In the expression above $P[A(n) = r] = [n!/((n-r)!\, r!)] \cdot p^r \cdot q^{(n-r)}$, for $0 \leqslant r \leqslant n$, and zero otherwise, where p is the probability of a positive test result and $q = 1 - p$. Note that the parameters α and β enter into only the range of sumation. Increasing either α or β extends this range, which cannot decrease $P[N > n \mid \alpha, \beta]$. (See Proposition 9.9 on page 285 of Townsend and Ashby (1983) for a related derivation.)

Appendix C

The purpose of this appendix is to show that if the decision maker's planning horizon is limited to m steps, then the observed terminal differences must be less than or equal to m in magnitude. In the first two experiments, the test results $Z(1), \ldots, Z(n)$ were generated by a Bernoulli sequence with $P[Z(n) = +1 \mid S_A] = p$, $P[Z(n) = +1 \mid S_B] = 1 - p = q$, and $P[S_A] = P[S_B] = .5$. Define LR as the ratio $LR = p/q$. The posterior probability of state S_A conditioned on $d(n) = d$ equals

$$P[S_A \mid d] = 1/[1 + LR^{-d}],$$

and the posterior probability of state S_B given $d(n) = d$ equals

$$P[S_B \mid d] = 1/[1 + LR^{d}].$$

The probability of observing $Z(n+1) = z$ conditioned on the previous difference, $d(n) = d$, equals

$$P[z \mid d] = [LR^{d+z} + 1]/[(LR^d + 1) \cdot (LR^z + 1)].$$

For Experiments 1 and 2, $v_{ii} = 0$ and $v_{ij} = -v$ for $i \neq j$, so the loss after purchasing n

tests equals $L = -v_{ij} + \sum c(k)$, $k = 1, n$. Whenever a terminal decision is made, the decision maker chooses the response that minimizes the expected loss. The expected value of the loss given that n tests were purchased, $d(n) = d$ was observed, and the response R_A was chosen equals

$$E[L \mid n, d, R_A] = v \cdot P[S_B \mid d] + \sum c(k), \qquad k = 1, n.$$

Similarly,

$$E[L \mid n, d, R_B] = v \cdot P[S_A \mid d] + \sum c(k), \qquad k = 1, n.$$

The decision makes will choose the smallest of the two expected losses shown immediately above, and this minimum, denoted $J[n, d(n)]$, can be expressed as follows:

$$J[n, d(n)] = v/[1 + LR^{|d(n)|}] + \sum c(k), \qquad k = 1, n.$$

First consider a one-step stopping rule ($m = 1$). Suppose that n observations have already been purchased and the difference $d(n)$ was observed. If the decision maker plans to purchase another test, then the expected loss of making a terminal decision after purchasing one more test (but before observing the result) equals

$$E\{J[n+1, d(n+1)] \mid d(n)\} = \sum P[z \mid d(n)] \cdot J[n+1, d(n+1)],$$

where the sum ranges across the two possible values of z. Whenever $d(n)$ is non-zero, then $E\{J[n+1, d(n+1)] \mid d(n)\} = J[n+1, d(n)]$. When $d(n) = 0$, then $E\{J[n+1, d(n+1)] \mid d(n)\} = J[n+1, 1]$. If the decision maker decides to stop as soon as $J[n, d(n)] \leqslant E\{J[n+1, d(n+1)] \mid d(n)\}$, then a terminal decision will be reached as soon as $|d(n)| = 1$. This will occur after the very first observation.

Next consider a two-step stopping rule ($m = 2$). In this case the decision maker compares the expected loss of making a terminal decision based on the current information with the expected loss of making a terminal decision after purchasing either one or two more observations. Suppose that $n \geqslant 2$ observations have been purchased and $d(n) \geqslant 2$. Consider what could happen if a decision to purchase observation $Z(n+1)$ is made. The difference would change to either $d(n+1) = d(n) + 1$ or $d(n+1) = d(n) - 1$, and in either case, the second observation $Z(n+2)$ will not be chosen because $E\{J[n+2, d(n+2)] \mid d(n+1)\} = J[n+2, d(n+1)] \geqslant J[n+1, d(n+1)]$ given that $d(n+1) > 0$. Therefore, the decision maker would plan not to purchase observation $Z(n+2)$. But now the problem is reduced to the one-step stopping rule described earlier, and recall that in this case the next observation is purchased if and only if $d(n) = 0$. Assuming that $d(n) \geqslant 2$, observation $Z(n+1)$ would not be purchased either. A similar argument can be made if one assumed that $d(n) \leqslant -2$. Therefore, if a two-step rule is used, then a terminal decision would be made at or before the occurrence of $|d(n)| = 2$.

The argument for the $m = 3$ step rule follows the same line of reasoning as was used with the $m = 2$ step rule. Suppose that $n \geqslant 3$ observations have been purchased and $d(n) \geqslant 3$. Consider what could happen if a decision to purchase both observations $Z(n+1)$ and $Z(n+2)$ is made. Note that $d(n+2)$ must exceed zero, and in this case, $E\{J[n+3, d(n+3)] \mid d(n+2)\} = J[n+3, d(n+2)] \geqslant J[n+2, d(n+2)]$, and therefore observation $Z(n+3)$ will not be chosen. But now the problem is reduced to the two-step rule described earlier. Since $d(n) \geqslant 2$, neither $Z(n+1)$ or $Z(n+2)$ will be chosen. A similar argument can be made if one assumes that $d(n) \leqslant -3$. Therefore, if a three-step rule is used, then a terminal decision would be made at or before the occurrence of $|d(n)| = 3$.

This argument can be applied recursively to show that for any small integer m, the m-step rule implies that the terminal difference will never exceed m in magnitude.

APPENDIX D

The purpose of this appendix is to describe how the myopic decision rule was tested with the results shown in Tables 4 and 6. The rule asserts that the decision maker stops testing as soon as the expected loss from error falls below the cost of the next test. If a terminal decision is made, then the response producing the smallest expected loss is selected.

In Appendix C, it was shown that

$$E[L \mid n, d(n), R_A] = v/[1 + LR^{d(n)}] + \sum c(k), \qquad k = 1, n$$

$$E[L \mid n, d(n), R_B] = v/[1 + LR^{-d(n)}] + \sum c(k), \qquad k = 1, n.$$

If a terminal decision is made, then the response producing the smaller of the two expected values shown above is chosen. Note that $E[L \mid n, d(n), R_A] < E[L \mid n, d(n), R_B]$ iff $d(n) > 0$, and the inequality is reversed iff $d(n) < 0$. So the response R_A is chosen if $d(n) > 0$, and R_B is chosen if $d(n) < 0$. When $d(n) = 0$, the two expected values are equal, and the choice is assumed to be random. The expected loss of making a terminal decision will be equal to the minimum of the two expected losses shown above, and this minimum is denoted $J[n, d(n)]$. It was noted in Appendix C that $J[n, d(n)] = v/[1 + LR^{|d(n)|}] + \sum c(k)$, $k = 1, n$.

The myopic stopping rule states that a terminal decision is made as soon as

$$J[n, d(n)] \leqslant c(1) + \cdots + c(n+1),$$

$$\Rightarrow v/[1 + LR^{|d(n)|}] \leqslant c(n+1),$$

$$\Rightarrow LR^{|d(n)|} \geqslant r(n) - 1,$$

where $r(n) = v/c(n+1)$ for $c(n+1) > 0$. Define $\delta(n) = \ln[r(n) - 1]/\ln(p/q)$ for $r(n) > 1$ and $\delta(n) = 0$, otherwise. Then a terminal decision will be made as soon as $|d(n)| \geq \delta(n)$.

The stopping rule described above does not include the subjective cost of waiting since $c(n)$ simply represents the monetary costs. The subjective cost of purchasing the nth observation is defined as $c'(n) = c(n) + b \cdot n$. Hereafter, the criterion bound, $\delta(n)$, is defined as before with $c'(n)$ substituted for $c(n)$.

The predictions for Table 4 were obtained as follows. The decision process may be modeled by a time-varying Markov chain. The state space can be partitioned into two sets, a set of transient states and a single absorbing state. The absorbing state represents the decision to stop and make a terminal decision.

The set of transient states is defined on the set of integers $\{-m, -m+1, ..., -1, 0, 1, ..., m\}$. Each transient state represents the conjunction of (a) the value of the difference, $d(n)$, and (b) the decision to continue testing after observing this difference. The upper limit, m, is chosen to be larger than the largest possible value of $\delta(n)$. For the present studies, $m = 10$ was sufficiently large.

The $(2m+1)$ row vector, $\mathbf{P}(n)$, represents the probability distribution across the transient states after observing n test results. $\mathbf{P}_i(n) = P[d(n) = i$ and $N > n]$, where N is a random variable representing the total number of observations that are purchased on each sequence.

The $(2m+1)\,X(2m+1)$ transition matrix for the transient states, denoted $\mathbf{T}(n)$, can be factored into two matrices—$\mathbf{T}(n) = \mathbf{D} \cdot \mathbf{G}(n)$. The $(2m+1)\,X(2m+1)$ matrix $\mathbf{D}$ is time invariant, and it contains the transition probabilities $D_{ij} = P[d(n+1) = j \mid d(n) = i] = P[Z(n) = j - i \mid S_A]$.

The $(2m+1)\,X(2m+1)$ matrix $\mathbf{G}(n)$ is a time variable diagonal matrix. Each off-diagonal element is zero, and each diagonal element equals $G_i(n) = P[N > n \mid N \geq n, d(n) = i]$. If $\delta(n) > 0$, then $G_i(n) = F[\delta(n) - |i|]$, and $G_i(n) = 0$ otherwise, where $F[z]$ is the logistic function defined in the discussion section.

The initial starting vector, $\mathbf{P}(0)$, was set to zero except for state $d(0) = 0$, which was normally set to 1.0. However, if $\delta(0) = 0$, then the entire vector was set to zero (i.e., no observations were purchased). The probability distribution for $n > 0$ was calculated iteratively by the matrix product $\mathbf{P}(n) = \mathbf{P}(n-1) \cdot \mathbf{D} \cdot \mathbf{G}(n)$.

The $(2m+1)$ row vector $\mathbf{Q}(n)$ contains the desired joint probabilities $Q_i(n) = P[N = n, d(n) = i]$. These absorbtion probabilities can be obtained from the matrix product $\mathbf{Q}(n) = \mathbf{P}(n-1) \cdot \mathbf{D} \cdot [\mathbf{I} - \mathbf{G}(n)]$, where $\mathbf{I}$ is an identity matrix.

A total of 26 sets of probability distributions were calculated, one set for each of the 26 possible values $\{\frac{0}{10}, \frac{1}{10}, ..., \frac{25}{10}\}$ of the subjective waiting cost scaling constant, b. A mixture distribution, $\mathbf{Q}(n) = \sum W_b \cdot \mathbf{Q}(n \mid b)$, was computed by setting $W_b = (25!)/[(25-k)! \cdot k!] \cdot \Pi^k (1-\Pi)^{25-k}$, where $k = 10 \cdot b$ and $\Pi = \frac{\mu}{2.5}$. The relative frequencies shown in Table 4 were computed from this mixture distribution.

The two parameters, Θ and μ, were estimated separately for each subject by minimizing the χ^2 lack of fit criterion computed from the predicted and observed relative frequencies, separately for each subject. (The $-2 \cdot \log$ likelihood ratio χ^2 formula described on p. 737 of Hays (1973) was used. However, due to the strong dependence of the observations across sequences, it is not possible to assume that this lack of fit index has a χ^2 distribution.)

The predictions for Table 6 were calculated using the same parameters estimated from Table 4 as follows. The probability of stopping immediately after observing the pattern $y = (z_1, z_2, z_3, z_4)$ can be obtained from

$$P[N = n \mid N > n-1, y] = 1 - (P[N > n \mid y]/P[N > n-1 \mid y]).$$

For $\delta(0) > 0$,

$$P[N > n \mid y] = G_{d_1}(1) \cdot G_{d_2}(2) \cdots G_{d_n}(n),$$

where $d_n = z_1 + \cdots + z_n$, and $G_i(n)$ is the ith diagonal element of the matrix $\mathbf{G}$ that was used to fit the results in Table 4.

Acknowledgments

The authors thank F. Greg Ashby and Richard Schweickert for their comments on an earlier draft of the paper. Requests for reprints should be sent to Jerome R. Busemeyer, Psychological Sciences, Purdue University, West Lafayette, IN 47907. This research was supported by NSF Grant BNS 8710103.

References

Audley, R. J., & Pike, A. R. (1965). Some alternative stochastic models of choice. *British Journal of Mathematical and Statistical Psychology*, **21**, 183–192.

De Groot, M. H. (1970). *Optimal statistical decisions.* New York: McGraw–Hill.

Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction time, and human information processing. *Journal of Mathematical Psychology*, **2**, 312–329.

Estes, W. K. (1960). A random walk model for choice behavior. In K. J. Arrow, S. Karlin & P. Suppes (Eds.), *Mathematical methods in the social sciences.* Stanford, CA: Stanford Univ. Press.

Hays, W. L. (1973). *Statistics for the social sciences.* N.Y.: Holt, Rinchart, Winston.

Kubovy, M., Rapoport, A., & Tversky, A. (1971). Deterministic vs. probabilistic strategies in detection. *Perception and Psychophysics*, **9** 427–429.

Laming, D. R. (1968). *Information theory and choice reaction time.* New York: Wiley.

Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, **40**, 77–105.

Morgan, B. J. T., & Robertson, C. (1980). Short term memory models for choice behavior. *Journal of Mathematical Psychology*, **21**, 30–52.

Osuna, E. E. (1985). The psychological cost of waiting. *Journal of Mathematical Psychology*, **29**, 82–105.

Pike, A. R. (1968). Latency and relative frequency of response in psychological discrimination. *British Journal of Mathematical and Statistical Psychology*, **21**, 161–182.

Pitz, G. F., & Barret, H. R. (1969). Information purchase in a decision task following the presentation of free information. *Journal of Experimental Psychology*, **82**, 410–414.

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

Rapoport, A. (1966). A study of human control in a stochastic multistage decision task. *Behavioral Science*, **11**, 18–30.

RAPOPORT, A., & BURKHEIMER, G. J. (1971). Models of deferred decision making. *Journal of Mathematical Psychology*, **8**, 508–538.

RAPOPORT, A., & TVERSKY, A. (1970). Choice behavior in an optimal stopping task. *Organizational Behavior and Human Performance*, **5**, 105–120.

SANDERS, A. F., & TER LINDEN, W. (1967). Decision making during paced arrival of probabilistic information. *Acta Psychologica*, **27**, 170–177.

SHANTEAU, J. C. (1970). An additive model for sequential decision making. *Journal of Experimental Psychology*, **85**, 181–191.

STONE, M. (1960). Models of choice reaction time. *Psychometrika*, **25**, 251–260.

SWENSSON, R. G., & THOMAS, R. E. (1974). Fixed and optional stopping models for two-choice discrimination times. *Journal of Mathematical Psychology*, **11**, 213–236.

TOWNSEND, J. T., & ASHBY, F. G. (1983). *Stochastic modeling of elementary psychological processes.* London: Cambridge Univ. Press.

VICKERS, D. (1979). *Decision processes in perception.* New York: Academic Press.

VIVIANI, P. (1979). Choice reaction times for temporal numerosity. *Journal of Experimental Psychology: Human Perception and Performance*, **5**, 157–167.

WALLSTEN, T. S. (1968). Failure of predictions from subjectively expected utility theory in a Bayesian task. *Organizational Behavior and Human Performance*, **3**, 239–252.

WALLSTEN, T. S., & SAPP, M. M. (1977). Strong ordinal properties of an additive model for sequential processing of probabilistic information. *Acta Psychologica*, **41**, 225–253.

YOUNG, M. J., FRIED, L. S., HERSHEY, S. C., EISENBERG, J. M., & WILLIAMS, S. V. (1985). Precise probability estimates and decision theory. Paper presented at the 7th Annual Society for Medical Decision Making Meeting in Pacific Grove, CA.

RECEIVED: August 14, 1985