Journal of Experimental Psychology:
Learning, Memory, and Cognition
1985, Vol. 11, No. 3, 538-564

0278-7393/85/$00.75

Decision Making Under Uncertainty: A Comparison of Simple Scalability, Fixed-Sample, and Sequential-Sampling Models

Jerome R. Busemeyer
Purdue University

The purpose of this article is to investigate the learning and memory processes involved in decision making under uncertainty. In two different experiments, subjects were given a choice between a certain alternative that produced a single known payoff and an uncertain alternative that produced a normal distribution of payoffs. Initially this distribution was unknown, and in the first experiment it was learned through feedback from past decisions, whereas in the second experiment it was learned by observing sample outcomes. In the first experiment, a response deadline was used to limit the amount of time available for making a decision. In the second experiment, an observation cost was used to limit the number of samples that could be purchased. The mean and variance of the uncertain alternative and the value of the certain alternative were factorially manipulated to study their joint effects on choice probability, choice response time (Experiment 1), and number of observations purchased (Experiment 2). Algebraic-deterministic theories developed for decision making with simple gambles fail to explain the present results. Two new models are developed and tested—fixed- and sequential-sampling models—that attempt to describe the learning and memory processes involved in decision making under uncertainty.

Decision theorists often find it useful to distinguish three classes of situations—decisions made under conditions of certainty, risk, or uncertainty (cf. Luce & Raiffa, 1957, p. 13). Under certainty, each action produces a single (perhaps multidimensional) known outcome. For example, deciding which pair of jeans to buy is a decision under certainty because you can see what you are buying. Under risk, each action produces a set of possible outcomes, and the probability of each outcome is known. State lotteries or gambling games are classic examples of risky decisions, because the probabilities of the outcomes can be calculated. Under conditions of uncertainty, each action again produces a set of possible outcomes, but the probability of each outcome is unknown. For example, deciding to talk with a stranger is an uncertain decision because you have no way to predict how the person will react.

More generally, different decision situations can be ordered in terms of degree of uncertainty (see Luce & Raiffa, 1957, p. 299). One can only speculate about the frequency of real-life situations along this ordering, but one guess is that most decisions are made under partial uncertainty, that is, a decision situation somewhere between risk and complete uncertainty. For example, the decision to build a nuclear energy plant will not be made under complete ignorance, but neither will it be made with perfect knowledge of the probabilities of all outcomes.

One way to make decisions under uncertainty is to base it on past experience with similar situations. For example, suppose that a particular patient has a long history of adverse reactions to different treatments for a serious medical problem. The patient's family physician would have to rely on this

The first experiment was supported by National Institute of Mental Health Grant MH1425706, Larry Jones principle investigator. Douglas Medin generously provided space and equipment to conduct the first study. The second experiment was supported by a grant awarded by Indiana University–Purdue University at Indianapolis.

The author wishes to thank K. Michael Aschenbrenner, Michael Birnbaum, Douglas Medin, Barbara Mellers, Steve Link, Tom Wallsten, and an unknown reviewer for helpful comments on this article.

Requests for reprints should be sent to Jerome R. Busemeyer, Psychological Sciences, Purdue University, West Lafayette, Indiana 47907.

past history to decide how to treat a new emergency call.

Another way to make decisions under uncertainty is to collect new information about the outcomes produced by an action before making a commitment. For example, before a new drug is introduced on the market, a sequence of tests are performed to determine its effectiveness.

Early in the history of behavioral decision theory (see Lee, 1971), researchers actively investigated decision making under both risk and uncertainty. Investigations of risky decision making typically used simple gambles of the form "win 10 dollars with probability .3 or nothing." Investigations of decision making under uncertainty typically used a probability learning task in which the probability of each payoff was initially unknown and had to be learned through trial by trial feedback.

Different theories were developed by researchers studying decisions under risk and uncertainty. Deterministic-algebraic theories such as Edwards's (1962) subjective expected utility theory were developed to describe decision making under risk. Stimulus sampling theories such as Myers and Atkinson's (1964) weak–strong conditioning model were developed to describe decision making under uncertainty.

Recently, Lopes (1983, p. 138) has argued that "after 30 years or more of research on risk, we know a lot about how people make decisions about simple lotteries, but we know remarkably little about decision under uncertainty, possibly because we have not had a good laboratory model of uncertainty." Apparently, this opinion reflects a disenchantment with stimulus-response conditioning theories and a lack of development of cognitive theories of decision making under uncertainty.

The purpose of this article is to investigate the cognitive processes involved in decision making under partial uncertainty. Two experiments are reported. In the first experiment, a probability learning task was used to investigate how memory of past outcomes influences new decisions. In the second experiment, an information purchasing task was used to investigate how new information influences decisions.

Experiment 1

On each trial of Experiment 1, subjects were given a choice between a certain alternative and an uncertain alternative. If the certain alternative was chosen, then a known monetary payoff was delivered. If the uncertain alternative was chosen, then the monetary payoff was randomly selected from a normal distribution with a mean equal to zero. The uncertain alternative was labeled X on some trials and Y on the remaining trials. When the X cue was presented, the payoff was sampled from a distribution with a small standard deviation. When the Y cue was presented, the payoff was sampled from a distribution with a large standard deviation. Initially, subjects did not know the distribution of outcomes produced by the cues X and Y. However, following each choice subjects were given feedback indicating the payoff that would have been received if the uncertain alternative was chosen. Choice responses had to be made within a deadline time limit in order to avoid a severe penalty. A factorial design was constructed by manipulating the variance of the uncertain payoffs (denoted σ^2), the value of the certain payoff (denoted k), and the deadline time limit (denoted L).

In the next section, two decision-making models are developed. The first is a probabilistic extension of the algebraic-deterministic models developed within the risky decision-making paradigm. The second is a probabilistic model based on the idea of memory retrieval.

Simple Scalability Models

Suppose subjects kept track of the long run average utility produced by the cues X and Y. Then after several hundred trials their estimates would converge on the mean or expected utility. If they simply compared the expected utility of the uncertain alternative with the utility of the certain alternative, then the choice process would be deterministic; for a given condition the same alternative would always be chosen. In order to introduce random variation into the choice process, one could hypothesize that the comparison of alternatives is perturbed by factors randomly varying across trials such as the individual's fluctuating state of wealth, randomly occurring patterns of wins and losses, or lapses of attention. Given these random disturbances, this model asserts that the probability of choosing the uncertain alternative is an increasing function of the difference between the expected utility of the uncertain alternative and the utility of the certain alternative. Becker, Degroot, and Marschak (1963a) proposed this model of choice.

The model just mentioned is one example of a general class called "simple scalability models," which were originally developed as probabilistic extensions of deterministic-algebraic models. (See section 7.1 of Luce & Suppes, 1965). According to simple scalability models, each alternative produces an independent utility scale value symbolized as u(UA) for the uncertain alternative, and $u(k)$ for the certain alternative. In general, u(UA) is determined by the probability distribution associated with the cue X or Y, and $u(k)$ is determined solely by the utility of k. Simple scalability models state that the probability of choosing the uncertain alternative is given by $P(\mathrm{UA}) = F[u(\mathrm{UA}), u(k)]$, where F is a strictly increasing function of u(UA) and a strictly decreasing function of $u(k)$.

The Becker et al. (1963a) model is a special case that states that $P(\mathrm{UA}) = F[u(\mathrm{UA}) - u(k)]$, where u(UA) was defined as the mean utility of the uncertain alternative. Other versions are possible by assuming that u(UA) is also determined by the variance, or various percentiles of the distribution for the uncertain alternative.

Predictions. There is one general property that all special cases of simple scalability obey: independence between alternatives (cf. Tversky & Russo, 1969). The *direction* of the effect of manipulating the distribution for the uncertain alternative should be the same for all values of the certain alternative. For example, the pattern of results illustrated in both Figures 1a and 1b conform to the independence between alternatives property. In both figures, the probability of choosing the uncertain alternative is plotted as a function of the certain value, denoted k. (k is plotted from positive to negative so that the curves increase from left to right.) Each line represents a different distribution for the uncertain alternative, one with low variance and another with high variance. Figure 1a represents an *uncertainty seeking* type of individual—increasing the variance increases the tendency to choose the uncertain alternative for all k. Figure 1b represents an *uncertainty averse* type of individual—increasing the variance decreases the tendency to choose the uncertain alternative for all k.

The $\sigma \times k$ crossover interaction shown in Figure 1c violates the independence between alternatives property (see Busemeyer, 1982, Appendix A). Contrary to independence, increasing the variance either increases or decreases P(UA) depending on the certain value, k. However, caution is needed because an unambiguous test of the independence property requires an analysis of individual performance. To see why, suppose we average across the *uncertainty seekers* illustrated in Figure 1a and the *uncertainty avoiders* illustrated in Figure 1b. Although neither type of individual produces a crossover pattern, the average shown in Figure 1c artificially produces the crossover. The present study provides a test of independence at the individual level of analysis.

Fixed-Sample Model

According to the fixed-sample model, the decision process for each choice trial begins by retrieving a fixed number of recent traces from memory. Each memory trace is a record of the numeric value of a past outcome associated with the cue X or Y presented on the current trial. The number of traces retrieved from memory is limited by the deadline time limit. The utilities of the retrieved outcomes are averaged, and this *moving average* is compared with the utility of the certain alternative. Only if the sample mean is greater will the uncertain alternative be chosen.

Choice probability. The probability of choosing the uncertain alternative, denoted P(UA), is an increasing function of the ratio $[\mu(u) - u(k)]/\sigma(\bar{u})$, where $\bar{u}$ denotes the sample mean, $\mu(u)$ is the population mean utility of the uncertain payoffs, $u(k)$ is the utility of the certain payoff, and $\sigma^2(\bar{u})$ is the variance of the sample mean.

Figure 2 illustrates the distribution of the sample mean under the low- and high-variance conditions. The three vertical lines indicate $u(k)$ when the certain value is negative ($k < 0$), zero ($k = 0$), and positive ($k > 0$). P(UA) is represented as the area above each vertical line.

Figure 2 illustrates the effect that variance has on choice probability. When k is positive, the area above the vertical line is greater under the high variance curve. When k is negative, the area above the vertical line is greater under the low variance curve. In sum, the direction of the effect of variance on P(UA) depends on the sign of the certain value, producing a crossover interaction similar to that shown in Figure 1c. Thus the fixed-sample model asserts that the independence between alternatives property should be violated for individual subjects.

This crossover interaction can be derived mathematically as follows. Define n as the number of recalled outcomes (i.e., the sample size of $\bar{u}$), and let $\sigma^2(u)$ be the variance of the utility of the uncertain payoffs. Then the variance of the sample mean equals $\sigma^2(\bar{u}) = \sigma^2(u)/n$. In the present study, the mean of the uncertain payoffs was always zero so that the mean difference $[\mu(u) - u(k)]$ was determined by $-u(k)$. Given these assumptions, P(UA) is an increasing function of the product $\sqrt{n}[-u(k)/\sigma(u)]$. Note that P(UA) increases as k decreases from positive to negative values, and that the slope is determined by the variance. Low variance produces sharp discrimination of the mean difference, yielding a steep slope. High variance produces poor discrimination yielding a flatter slope.

Now consider the effect of increasing the time limit, L. This allows an increase in the number of recalled outcomes, n, which then magnifies all of the slopes shown in Figure 1c. When k is positive, then increasing the time limit increases the probability of correctly[1] choosing the certain alternative. When k is negative, then increasing the time limit increases the probability of correctly choosing the uncertain alternative. More concisely, accuracy always increases as the deadline time limit increases.

Choice response time. According to the fixed-sample model, increasing the number of outcomes recalled, n, increases the number of operations required to compute the sample mean. As a consequence, the time required to estimate the sample mean and ultimately make a choice response is an increasing function of the sample size. However, increasing the sample size also increases the probability of choosing the correct alternative. In sum, a

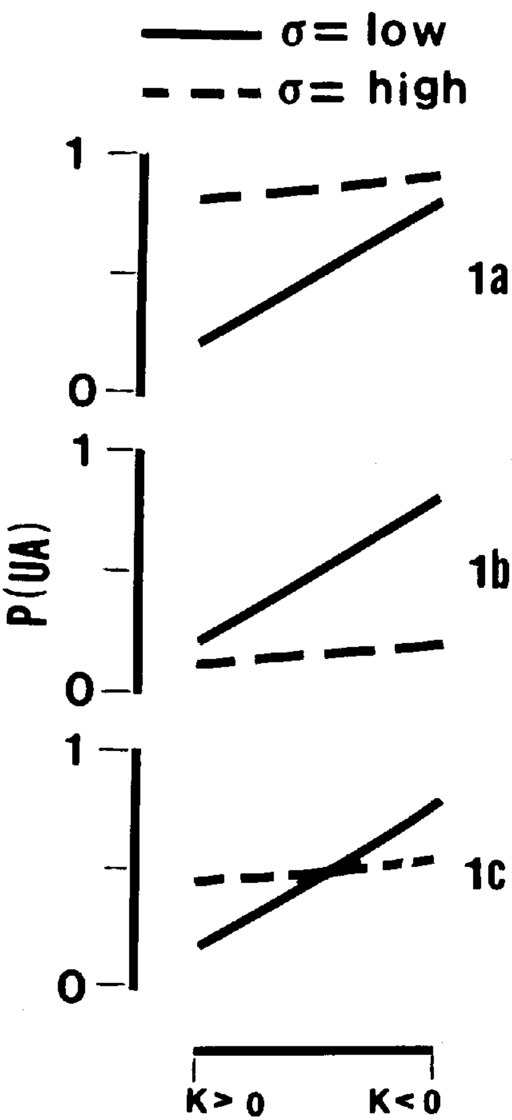

Figure 1. Fictitious results indicating different possible patterns of the uncertain alternative variance (σ^2) by certain value (k) interaction. [P(UA) equals the probability of choosing the uncertain alternative.]

[1] For this binary choice task, the term *correct alternative* is a label for the alternative producing the larger expected utility. The term *accuracy* is equivalent to the probability of choosing the correct alternative.

basic property of the fixed-sample model is the *speed-accuracy* trade-off, that is, faster responses to the short deadline time limit yield lower accuracy (cf. Swensson & Thomas, 1974).

Method

Subjects. Six psychology students (4 seniors and 2 graduate students, 2 males and 4 females) from the University of Illinois–Champaign volunteered to participate for 15 1-hr daily sessions. Subjects were paid according to their performance, and they earned an average of $3.25 per session.

Apparatus. Each subject was tested individually in a quiet room. The experiment was computer controlled, stimuli were presented on a video terminal driven by an Apple microcomputer, and response times were recorded in milliseconds using a Mountain Hardware clock.

Design and procedure. Monetary payoffs for the uncertain alternative were generated according to a normal distribution with a mean of zero, and a standard deviation of 5 or 50 units (in units of $.01) depending on the stimulus condition. The payoff for the certain alternative equaled either −3, 0, or +3 units, and the deadline time limit was either 1, 2, or 3 s. The variance of the uncertain alternative, the payoff values for the certain alternative, and the time limit conditions were combined according to a 2 × 3 × 3 factorial design. Failure to beat the deadline on any trial produced a loss equal to 25 units. Pilot research indicated that a simple reaction time required .25 s.

The protocol for a typical trial is illustrated as follows:

Deadline = 1 s

(.5-s delay)

Ready

(.5-s delay, then screen clears)

Choose X or −3

(response timing begins)
(subject types a response and selects the certain alternative)
(response timing ends)

X = −45 on this trial
Your pay equals −3 on this trial

(subject initiates new trial by typing the return key)
(screen clears)

As the display illustrates, each trial started with a message indicating the deadline requirement, after a delay the screen cleared, and then the choice stimulus was presented. In the display, X is the cue for the uncertain alternative and the certain value equals −3. The low- or high-variance condition was indicated by the cues X or Y, where X signaled the high variance and Y the low variance for half the subjects, and the reverse pairing was used for the other half. Feedback followed each response, which indicated both the payoff that would have been received if uncertain alternative was chosen (e.g., X = −45), and the actual payoff received based on the selected alternative (e.g., −3, because the certain alternative was chosen in the example).

The actual distributions associated with the cues X and Y were strictly stationary and independently distributed across trials. However, subjects were given no prior information about the payoffs produced by each cue. After the experiment, many subjects remarked that the means seemed to vary across sessions in an unsystematic fashion. In fact, the sample means did vary across sessions, even though the population means were always zero.

Each subject received six blocks per session. The deadline requirement was constant within a block of 60 trials, and varied across blocks in a counterbalanced order. The values of σ and k were selected at random with equal probabilities for any given trial. The spatial position of the alternatives remained constant within a given session, but alternated from left to right across sessions.

Initially, subjects were told that their final pay would be determined by their accumulated wins and losses across all sessions. Because it was possible under the high-variance condition to win or lose a dollar on each

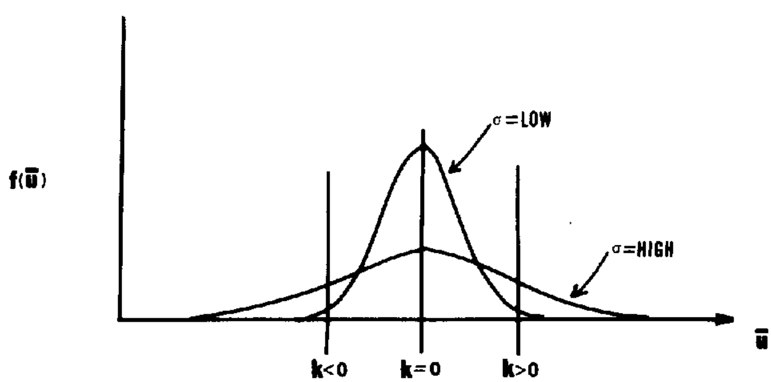

Figure 2. Illustration of the effect of uncertain alternative variance (σ^2) on the probability of choosing the uncertain alternative, according to the fixed-sample model.

trial, there was considerable incentive to perform well. They were not told their accumulated wins until the end of the experiment. However, after the experiment was concluded, they were paid their cumulative wins and losses, and they were also given a bonus for completing the experiment.

After completing the experiment, subjects were asked to provide preference judgments for each of the six choice pairs produced by combining the two uncertainty cues, X and Y, with the three certain values, k. The preference scale was a 10 cm line anchored at the left, center, and right with the labels "uncertain alternative is 10 units better," "equal," or "certain alternative is 10 units better." Subjects placed a tick mark anywhere along the scale to indicate the direction and strength of preference. Subjects were also asked to construct separate relative frequency histograms for each of the uncertainty cues X and Y using five categories of their own choosing.

Results

Choice probability.[2] The proportion of trials that the uncertain alternative was chosen, denoted $\bar{P}$(UA), was estimated for each subject by pooling across Sessions 3 to 15. The first two sessions were treated as practice, and there were no significant training effects following the first two sessions.

Figure 3 is a plot of $\bar{P}$(UA) averaged across subjects, as a function of the certain value, k, with a separate curve for each level of variance, σ^2, and a separate panel for each time limit, L. The expected total number of choice responses per data point in Figure 3 equals 1,560.

First consider the $\sigma \times k$ interaction which is illustrated within each panel of Figure 3. This is precisely the interaction predicted by the fixed-sample model, and it directly contradicts the independence between alternatives property implied by the simple scalability models. All 6 subjects produced this crossover pattern, and so the crossover shown in Figure 3 is not the result of averaging across different types of subjects. A repeated measures analysis of variance (ANOVA) indicated that the $\sigma \times k$ interaction was significant, $F(2, 10) = 14.98$, $p < .001$.

Next consider the effect of the time limit, L, on the $\sigma \times k$ interaction. This three-way interaction can be seen by comparing the slopes across panels in Figure 3. Recall that the fixed-sample model asserts that increasing the time limit should increase all slopes. Contrary to this prediction, increasing the time limit actually decreased the slope for the high-variance condition, although the slope for the low-variance condition increased as predicted. The three-way interaction pattern shown in Figure 3 was statistically significant, $F(4, 20) = 8.3$, $p < .001$, and this pattern was consistent across all 6 subjects.

Table 1 illustrates the effects of the deadline time limit and variance on accuracy. The observed proportion correct within each cell of Table 1 was calculated by averaging $\bar{P}$(UA) when k was negative, with $1 - \bar{P}$(UA) when k was positive. The table shows that under the low-variance condition, accuracy consistently increased as the time limit increased. But under the high-variance condition, accuracy consistently decreased as the time limit increased.

Judged preference. Figure 4 is a plot of the mean preference as a function of the certain value, k, with a different curve for each level of variance, σ^2. Negative scale values indicate increasingly stronger preferences favoring the certain alternative, whereas positive scale values indicate increasingly stronger preferences favoring the uncertain alternative. Figure 4 is a replication of the $\sigma \times k$ interaction shown in Figure 3, but based on a single judged preference from each subject rather than a proportion pooled across several hundred choice trials. All 6 subjects produced the crossover interaction indicated by Figure 4, and the $\sigma \times k$ interaction effect was statistically significant, $F(2, 10) = 10.72$, $MS_e = 22.35$, $p < .01$.

Mean choice response time.[3] Mean choice response times ($\overline{\text{CRT}}$) for each subject were obtained by averaging the latencies across Sessions 3 to 15. Figure 5 is a plot of these latencies in seconds averaged across subjects, as a function of the sure thing value, k, with a separate curve for each level of variance, σ^2, and a separate panel for each deadline time limit, L. Note that $\overline{\text{CRT}}$ represents the latencies pooled across both responses.

Recall that the fixed-sample model asserts that mean choice response time is an increas-

[2] The analysis of variance was computed on the arcsine transformed proportions in order to achieve homogeneous variance.

[3] The analysis of variance was computed on choice speed (the reciprocal of the latencies) in order to achieve homogeneous variance.

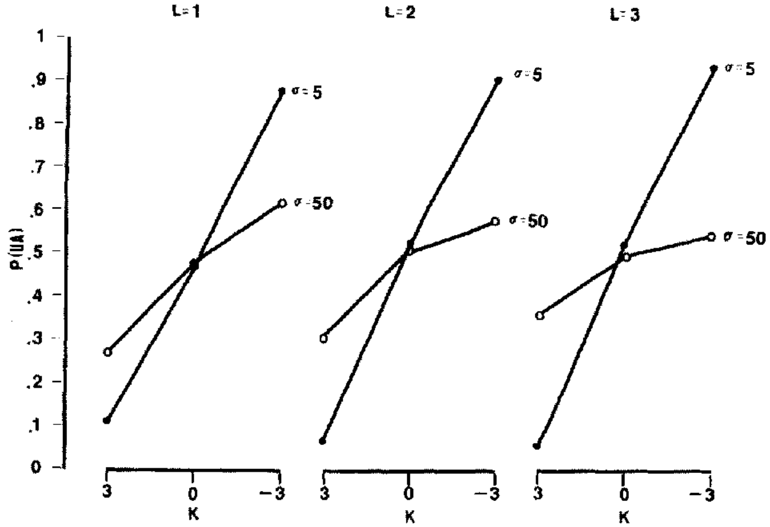

Figure 3. Probability of choosing the uncertain alternative, [$\bar{P}$(UA)] plotted as a function of the certain value (k), with a different curve for each level of variance (σ^2) and a different panel for each level of deadline time limit (L).

ing function of sample size, n, and sample size increases with the time limit, L. Figure 5 clearly supports this expectation, and the effect is statistically significant, $F(2, 10) = 44.9$, $p < .001$.

The fixed-sample model also asserts that a speed-accuracy trade-off relation should always occur when the deadline time limit is increased, i.e., faster responses to the short deadlines should produce lower accuracy. Note that this speed-accuracy trade off relation failed to occur under the high-variance condition—faster responses produced greater accuracy when performance to the short and long time limits are compared. The predicted speed-accuracy trade-off relation was confirmed for the low-variance condition.

Table 1
Proportion Correct[a] *for Each Combination of Deadline Time Limit and Variance Condition*

	Deadline time limit		
Variance condition	1	2	3
$\sigma = 5$	.885	.920	.937
$\sigma = 50$	.676	.633	.593

[a] Proportion correct for each cell was calculated by averaging $\bar{P}$(UA) when $k = -3$, with $1 - \bar{P}$(UA) when $k = +3$. Each cell is based on approximately 3,120 observations.

Another interesting finding in Figure 5 is that mean choice time increased as the certain value, k, decreased from positive to negative. Both the main effect of the certain value and the Certain Value × Time Limit interaction effect were statistically significant, $F(2, 10) = 9.31$, $p < .01$; $F(4, 20) = 5.19$, $p < .05$. None of the effects due to the variance manipulation produced statistically significant effects on mean choice response time.

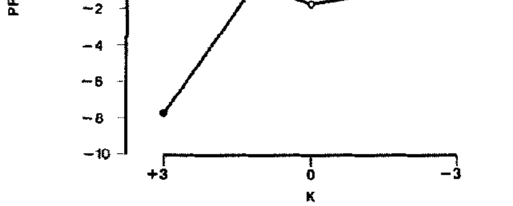

Figure 4. Mean preference rating plotted as a function of certain value with a different curve for each level of variance. (Negative ratings indicate preference for the certain alternative, and positive ratings indicate preference for the uncertain alternative.)

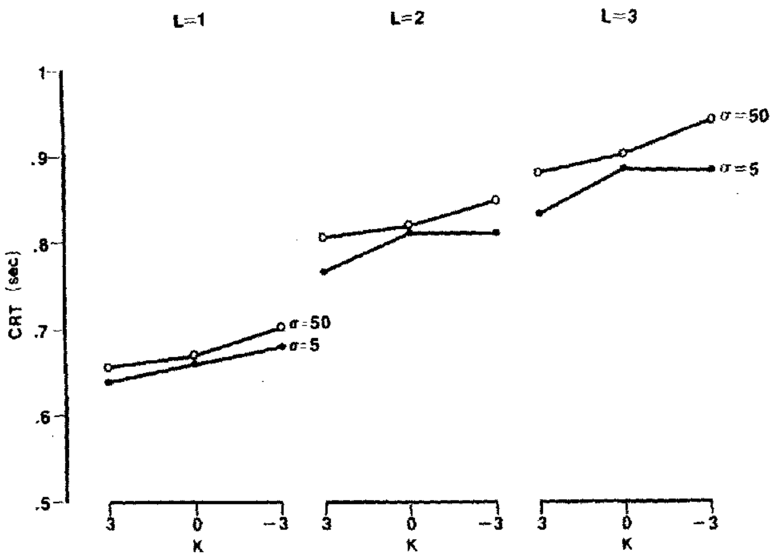

Figure 5. Mean choice response time ($\overline{CRT}$) in seconds plotted as a function of the certain value (k), with a different curve for each level of variance (σ^2) and a different panel for each level of deadline time limit (L).

Judged cumulative probabilities. The judged cumulative probabilities averaged across subjects for each variance condition are shown in Figure 6, and they are compared with the cumulative normal distribution. If the judged distributions were perfectly accu-

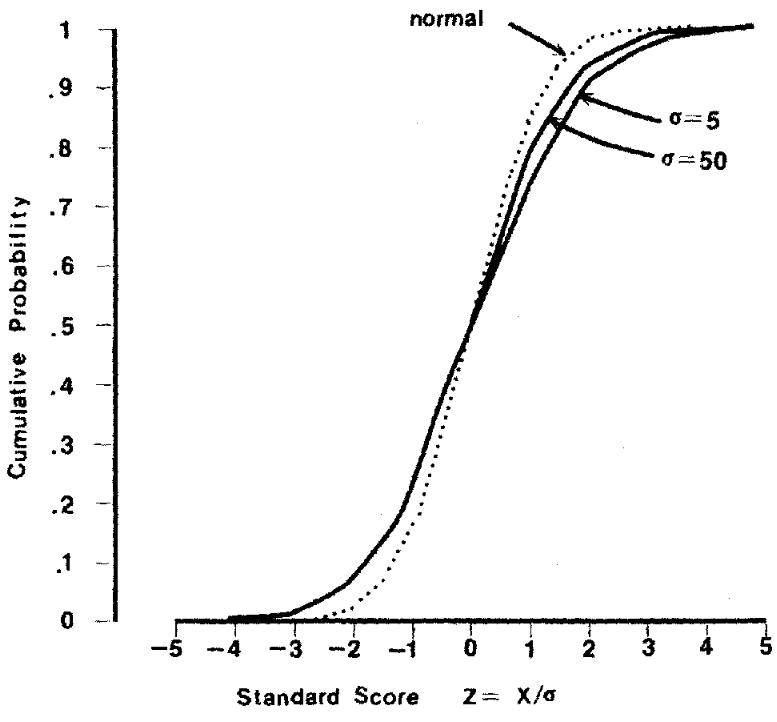

Figure 6. Mean judged cumulative distributions for the low- and high-variance conditions, as compared with the normal distribution.

rate, then all three curves would coincide. The means of the judged distributions equal .185 and −3.36 for the low- and high-variance conditions, which in standard units equals .037 and −.067, respectively. The standard deviations of the judged distributions equal 7.36 and 67.58 for the low- and high-variance conditions, respectively.

Discussion

Summary of major results. Probability of choosing the uncertain alternative increased as the certain value decreased, and the magnitude of this effect was inversely related to the variance, producing a $\sigma \times k$ crossover interaction (see Figure 3). This effect was replicated with judged preference (see Figure 4). Increasing the time limit increased accuracy under the low-variance condition, but it decreased accuracy under the high-variance condition (see Table 1).

Mean choice response time also increased as the certain value decreased (see Figure 5). Increasing the deadline time limit produced a large increase in mean response time. Variance did not have a reliable effect on mean response time.

The finding that accuracy decreased as the time limit increased under the high-variance condition may have important practical implications. Sensitivity to the *safe* alternative seems higher under short deadlines. Apparently, when discrimination is very poor, additional time to think about the choice decreases sensitivity to the sure thing, and accuracy suffers.

Relation to prior research. An earlier study by Myers, Suydam, and Gambino (1965) reported a $\sigma \times k$ interaction similar to that shown in Figure 1c. However, they used a between subjects design, and only reported averages across subjects. Thus their results could have been artificially produced by averaging across uncertainty seekers and uncertainty avoiders, as shown in Figures 1a and 1b. The present study resolved this difficulty by demonstrating that the choice proportions and single preference judgments of individuals produced the crossover interaction pattern.

The decrease in accuracy with increased time limit for the high-variance condition was a surprising result. However, it is interesting to note that a similar finding occurred in an earlier study by Irwin, Smith, and Mayfield (1956), although they did not emphasize this result in their report. Subjects in that study were asked to decide whether the mean of a deck of cards was less than or greater than zero, based on sample sizes of 10 or 20 cards. Their results indicated that when discriminability was low (because of a small mean difference and a large variance as in the present study), accuracy in detecting positive mean differences decreased with increased sample size. Together the results of the present study and Irwin et al. suggest that this finding is reliable and robust across tasks.

A possible difference between decision making under risk versus partial uncertainty is illustrated by comparing the present results with an earlier study by Ben Zur and Breznitz (1981). Subjects in the latter study were required to choose between two gambles within a deadline time limit. Each gamble produced one of two possible outcomes, and the probability of each outcome was displayed upon a request from the subject. The two gambles within a pair had equal expected values, but the variance differed. Ben Zur and Breznitz reported a decrease in the probability of choosing the high variance gamble within a pair as the time limit decreased from 32 to 8 s.

The present study did not replicate this effect. In order to compare the present results with those reported by Ben Zur and Breznitz (1981), the expected value of the alternatives must be equated, which is true only for the condition with $k = 0$. As can be seen in Figure 3 for $k = 0$, $\bar{P}(\mathrm{UA})$ remained at .5 across all deadline conditions.

Ben Zur and Breznitz (1981) concluded that their results were caused by increased attention to the displayed loss probability under the short time limit. In the present study, outcome probabilities were not displayed, and subjects had to rely on recall of past outcomes. Apparently, the tendency to recall negative or positive outcomes was not differentially influenced by the deadline time limit.

Simple scalability models. Deterministic-algebraic models of risky decision making (e.g., subjective expected utility theory) cannot be used to describe the quantitative aspects of choice probability without making additional assumptions. The most natural way to extend these models is the class of probabilistic models called *simple scalability* models. Simple scalability models continue to be popular among applied researchers (business, marketing, and consumer behavior) because each choice alternative can be assigned a single scale value independent of other alternatives, which can easily be estimated by commonly available scaling programs.

Despite the popularity of these models, there are strong reasons for rejecting this entire class as a representation of decision making under certainty, risk, or uncertainty. Tversky (1972) has shown that this class is inadequate for describing decision under certainty because the similarity between the outcomes produced by each alternative strongly influences choice. Becker, DeGroot, and Marschak (1963b) have shown that this class is inadequate for describing choice under risk because the known correlation among outcomes produced by each alternative influences choice. Finally, the present study has shown that this class is inadequate for describing choice under uncertainty because the variance of the difference among outcomes produced by each alternative influences choice. In all three cases, independence between alternatives was violated, which is a fundamental property of simple scalability models.

Fixed-sample model. Violations of the independence property are evidenced by the crossover interactions within each panel of Figure 3. This crossover interaction can easily be explained by a memory retrieval model. Presumably, subjects base their decision on a sample of recalled outcomes, and compare this sample mean (rather than a population mean) with the certain value. The variability of this difference influences discriminability similar to a signal detection task. Formally, the variance of the sample mean difference divides the mean utility of both alternatives, which causes them to be interlocked rather than independent.

Of course there are other possible explanations for the Uncertain Variance × Certain Value crossover interaction. The expected loss ratio model (Fishburn, 1976; Myers et al., 1965) and the weak–strong conditioning model (Myers & Atkinson, 1964) are possibilities. However, both of these models are inadequate for different reasons. The expected loss ratio model provides no explanation for the dynamics of the underlying decision process, and cannot explain the systematic effects of deadline and sure thing value on choice response time. Although the weak–strong conditioning model does provide a mechanism for describing mean choice response time, it does not provide any mechanism for explaining the effects of information processing variables such as the deadline time limit. The fixed-sample model generates predictions for choice probability and response time as a function of deadline time limit, but it occasionally makes incorrect predictions.

A basic property of the fixed-sample model is the predicted speed-accuracy trade-off relation—increasing the time limit increases the number of recalled outcomes (i.e., the sample size), which then causes an increase in mean response time and accuracy. Although the predicted speed-accuracy trade-off was obtained under the low-variance condition, it was violated under the high-variance condition.

One could argue that sample size does not necessarily increase with the time limit, instead the time limit simply increases the mean time to execute each of the operations involved in computing the sample mean and the comparison with the sure thing value. However, this explanation does not save the fixed-sample model. In order to explain the speed-accuracy trade-off effect obtained under the low variance, it would be necessary to assume that slower mean execution times produce fewer computational errors. But the reduction of computational errors should also occur for the high-variance condition.

Another problem for the fixed-sample model is the fact that mean choice time was slower for negative than for positive certain values. This could be explained by assuming that subjects waited for the choice stimulus to appear, and then selected a larger sample size (i.e., recalled more past outcomes) when the certain value was negative. This implies that the factors influencing the choice of sample size are more complicated than originally assumed.

Correlated bias. One way to revise the fixed-sample model is to assume that subjects choose the uncertain alternative only if the sample mean difference, $\bar{u} - u(k)$, is greater than some bias factor, b. This bias factor may be positively correlated with the certain value so that $b > 0$ when $k > 0$, and $b < 0$ when $k < 0$. Without any further assumptions the predictions remain unchanged, but if the bias is assumed to decrease in magnitude towards zero as the deadline time limit increases, then the fixed-sample model can explain the pattern of results in Table 1.

Briefly, the accuracy obtained under the high-variance condition may be entirely due to the correlated bias because discriminability is so very poor. As the deadline increases and the bias decreases in magnitude, then accuracy would decrease for the high-variance condition. This does not happen under the low-variance condition because discriminability is very high, and the increase in discriminability produced by the increased sample size dominates the effect of the decrease in bias.

One problem with this explanation is the ad hoc assumption that the bias decreases in magnitude with increasing time limit. During the introduction to Experiment 2, a sequential-sampling model is proposed which provides an explanation for the decreasing magnitude of the bias.

Conclusion. The results of Experiment 1 indicate that the deterministic-algebraic models used to describe decision making under risk with known outcome probabilities cannot be directly applied to decision making under uncertainty with outcomes learned from past experience. The reason is that when individuals rely on past experience, they do not assign a utility scale value independently to each alternative and compare these values. Instead, it seems that they compare the estimated value of each alternative, and these estimates are based on a small number of outcomes retrieved from memory. Because of the probabilistic nature of the memory retrieval process, the estimated value of each alternative fluctuates from one decision to the next. This variability influences discriminability which in turn influences choice behavior.

Experiment 2

The second experiment used an information purchasing task to investigate decision making under uncertainty. Subjects were asked to choose between an uncertain alternative and a certain alternative. The payoff for the uncertain alternative was generated by randomly selecting a lottery ticket from a fictitious urn, and the values of the tickets in the urn were normally distributed. Prior to each choice trial, a new urn was randomly selected, so the population mean and variance of the urn was initially unknown. Information about this distribution was learned by a sequence of requests for sample observations. On each request, the computer sampled a ticket from the urn, displayed the value of the ticket, and then returned the ticket to the urn. Although an unlimited number of observations could be requested, each one cost a fixed amount. In sum, a single choice trial consisted of a sequence of sample observations, followed by a terminal choice of either the certain or the uncertain alternative.

One can think of the present task as a discrimination problem—the subject has to determine on the basis of sample observations whether the mean of the uncertain alternative is greater or less than the certain value. From this perspective, the present task is similar to the "expanded judgment" task investigated by Irwin and Smith (1956; 1957).

The second experiment was designed to determine whether or not the decreased accuracy with increased sample size found in Experiment 1 (see Table 1) could be replicated by manipulating observation cost rather than a deadline time limit. A four-way factorial design was constructed by manipulating the mean of the uncertain alternative, denoted $\mu(x)$, the variance of the uncertain alternative, denoted $\sigma^2(x)$, the certain value, denoted k, and the cost of purchasing observations, denoted C.

Two decision-making models are described next—a fixed-sample model similar to that described in Experiment 1, and a sequential-sampling model. The sequential-sampling model is worth considering for three reasons. One is that under certain ideal conditions, it is the optimal strategy for minimizing expected losses (see Edwards, 1965). The second is that sequential-sampling models are useful for describing psychophysical discrimination (see Laming, 1968; Link & Heath, 1975; Stone, 1960; Vickers, 1979). The third is that sequential-sampling models have been applied to decision making under certainty by Aschenbrenner, Albert, and Schmalhofer (1984), to decision making under risk by Petrusic and Jamieson (1978), and to decision making under uncertainty by Busemeyer (1982). The model developed by Busemeyer (1982) was restricted to binary outcomes, whereas the present development can be applied to either discrete or continuous outcome distributions.

Fixed-Sample Model

One possible decision strategy is to first decide how many observations to buy, and then choose the uncertain alternative only if the average of the sample is greater than the certain value. This is an inefficient strategy for the information purchasing task. Unlike, Experiment 1, there is no limit on the number of observations that can be purchased, and the number purchased can vary depending on the informativeness of the sample. For example, suppose the subject decides to purchase only a small number of observations, and the difference between the sample mean and the certain value is very small. In this case, the sample is uninformative, and it may be worthwhile to continue sampling. Now suppose the subject decides to purchase a large number of observations, but discovers after the first few observations that the sample mean is much larger than the certain value. In this case, the small sample is very informative, and it may not be worthwhile to continue sampling costly observations.

There is a simple way to test whether or not subjects are using a fixed-sample strategy. If the sample size is selected before any observations are purchased, then the actual number purchased should be independent of the mean and the variance of the uncertain alternative distribution. This is because the distribution properties of the uncertain alternative will not be known when the sample size is selected.

Sequential-Sampling Model

According to the sequential-sampling model, each observation produces an increment in preference. This increment is defined as the difference between the utility of the observed value sampled from the uncertain alternative and the utility of the certain value. Positive increments increase preference for the uncertain alternative, and negative increments increase preference for the certain alternative. If the first increment is insufficient to evoke a choice response, then another observation is purchased, and a new increment is produced. The new increment is added to the previous increment to produce a new cumulative preference. This accumulation of increments continues until a positive criterion is exceeded, evoking the choice of the uncertain alternative, or until a negative criterion is exceeded, evoking the choice of the certain alternative.

The criterion are selected by the subject prior to sampling. Three factors are assumed to influence the selection of criterion bounds. The first factor is an individuals attitude toward uncertainty. If the subject tends to avoid uncertainty, then the magnitude of the certain alternative bound will be smaller than the uncertain alternative bound. The opposite relation holds if the subject tends to approach uncertainty. The second factor is information available before sampling begins. For example, knowledge of the certain value may influence the magnitudes of each bound. The third factor is the observation cost. Under high cost, the bounds must be close to zero so that very few observations are purchased; but under the low cost, the boundaries can be farther apart, allowing a stronger preference to accumulate before making a final decision.

A mathematical model of this process is illustrated in Figure 7. Sequential sampling produces a sequence of utilities where $u(1)$ is the utility of the first sample, $u(j)$ is the utility of the j^{th} sample, and $u(N)$ is the utility of the last sample before making a final decision. A comparison of the j^{th} sample with the utility of the certain alternative produces an increment, $d(j) = [u(j) - u(k)]$. The cumulative preference after the first sample is $d(1)$, and the cumulative preference after $j + 1$ samples equals $D(j + 1) = D(j) + d(j + 1)$. For example, in Figure 7, $d(3)$ is the increment and $D(3) = D(2) + d(3)$ is the cumulative preference following the third sample. The solid line in Figure 7 is a sample path of the cumulative preference for a single choice trial. The dotted line in Figure 7 represents the mean of the cumulative pref-

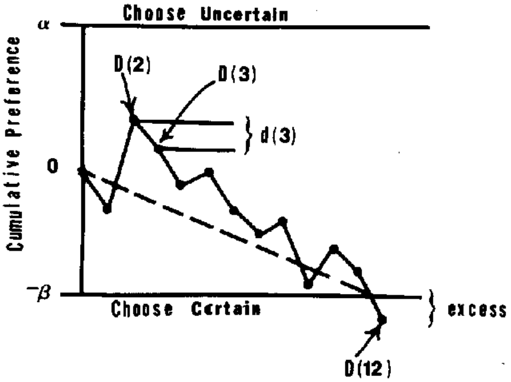

Figure 7. The sample path of the cumulative preference for the sequential-sampling model.

erence averaged across trials under identical stimulus conditions.

The mean of the increments, symbolized as $\mu(d)$, equals $[\mu(u) - u(k)]$, where $\mu(u)$ is the mean utility of the uncertain alternative. In Figure 7, the mean increment is represented by the negative slope of the dotted line, indicating that for this example the mean of the uncertain alternative is less than the certain value. As can be seen in the figure, the mean increment determines the direction and rate of the sample path. The variance of the increments is simply equal to the variance of the utility of the uncertain alternative, $\sigma^2(u)$, because $u(k)$ is a constant.

The criterion for choosing the uncertain alternative is symbolized as α and the criterion for choosing the certain alternative is symbolized as $-\beta$. The bounds are drawn as the upper- and lower-horizontal lines in Figure 7. For example, Figure 7 indicates that the cumulative preference eventually exceeded the criterion for the certain alternative after observing $N = 12$ samples.

Choice probability. Probability of choosing the uncertain alternative, $P(\text{UA})$, is determined by the probability that the cumulative preference exceeds the uncertain criterion before the certain criterion. Although the hypothesized decision process is very simple, the behavioral properties that it produces are quite complex. The derivation of $P(\text{UA})$ is given in Appendix A, and only the qualitative properties are described here.

Choice probability is a function of three parameters—a measure of discriminability (θ), a measure of response bias (δ), and a measure of the distance between the bounds (A). More specifically, the distance between bounds is the sum of the criterion magnitudes, $A = (\alpha + \beta)$. Increasing the distance between bounds increases the average strength of preference required to reach a decision.

The discriminability parameter θ is determined by the distribution of the increments $d(j)$. If the increments are normally distributed, then $\theta = 2\mu(d)/\sigma^2(u)$, that is, θ is proportional to the mean increment divided by the variance of the increments.

The parameter δ is called the *relative bias*, and it is defined by the ratio $\delta = (\beta - \alpha)/A$. In other words, it reflects the difference between the magnitude of the certain bound and the magnitude of the uncertain bound, relative to the total distance between the bounds. When $\delta = +1$, the uncertain alternative is always chosen; when $\delta = 0$, then the initial preference is midway between the bounds; and when $\delta = -1$, the certain alternative is always chosen. It is important to note that a constant absolute bias, $(\beta - \alpha)$, will have a smaller effect on the relative bias if the total distance between the bounds, A, is increased.

First consider the predictions when the bias is assumed to equal zero. Under these conditions, $P(\text{UA})$ is an increasing function of the product $A[\mu(d)/\sigma^2(u)]$. Thus, $P(\text{UA})$ increases as the mean increment increases from negative to positive. The slope of this function decreases as the variance increases. Finally, all slopes increase as the observation cost decreases due to the increase in in the total distance, A. In sum, accuracy (see Footnote 1) always increases as the observation cost decreases. Note that these predictions are similar to those generated by the fixed-sample model (with no bias) described in Experiment 1.

Now relax the assumption that the bias is fixed at zero. Suppose that subjects adjust the certain criterion magnitude, β, depending on the certain value, k, producing a correlated bias. More specifically, assume that the certain bound equals $(\beta - b_k)$ where b_k is positive when $k > 0$, b_k is zero when $k = 0$, and b_k is negative when $k < 0$. In other words, subjects have an initial tendency to favor the certain alternative when the certain value is positive, and they have an initial tendency to disfavor the certain alternative when the certain value is negative.

In this case, accuracy may decrease as the observation cost decreases for the following reason. When the variance is very high, then discriminability is near zero, and the cumulative preference wanders without any systematic direction. Despite the fact that discriminability is low, accuracy may be high due to the correlated bias. Under the high-cost condition, the correct criterion bound is close to the starting point, and the cumulative preference is very likely to hit the correct bound before wandering off in a random direction. However, under the low-cost condition, the criterion bounds are far apart, and the correlated bias has very little effect. The cumulative preference has ample opportunity to wander off randomly before hitting either the correct or the incorrect criterion bound.

Note that these predictions are similar to the fixed-sample model with correlated bias described in the discussion of Experiment 1. However, the sequential-sampling model provides a simple explanation for the decreasing effectiveness of the correlated bias as the sample size increases. The decreased effectiveness is due to the property that choice probability is a function of the relative bias, $\delta_k = [(\beta - b_k) - \alpha]/[\alpha + (\beta - b_k)]$. It is assumed that α and β increase as the observation cost decreases, but b_k remains constant. As α and β increase, the average sample size increases, and δ_k converges towards a single value independent of b_k.

If the correlated bias hypothesis is the correct explanation for the decreased accuracy with increased sample size observed in Experiment 1 (see Table 1), then this result should only occur when the certain value is manipulated. It should not occur if the certain value is fixed at zero. This suggests that two different patterns of results should occur in Experiment 2, depending on whether the mean of the uncertain alternative or the certain value is manipulated. If the uncertain alternative mean is fixed at zero as in Experiment 1, then the same pattern of results shown in Table 1 should occur in Experiment 2, with the time limit variable replaced by observation cost. However, if the certain value is fixed at zero, then a different pattern should occur. Decreasing the observation cost should always increase accuracy under both variance conditions.

Number of samples purchased. The predictions for the mean number of samples purchased, denoted $E(N)$, are derived in Appendix A, and only the qualitative properties are discussed here. Intuitively, $E(N)$ is determined by the average value of the final cumulative preference divided by the mean increment. (This is analogous to a measure of the distance traveled divided by rate of travel.)

If the relative bias is zero, and discriminability is held constant, then $E(N)$ is an increasing function of the total distance, A. When the relative bias is zero, and all factors other than the mean increment are held constant, then $E(N)$ is a decreasing function of the magnitude of the mean increment. When the mean increment is zero, and all factors other than the relative bias are held constant, then $E(N)$ is a decreasing function of the magnitude of the relative bias. Finally, a large increase in the variance reduces $E(N)$ if the total distance, relative bias, and mean increment are held constant. Note that these predictions contrast sharply with those of the fixed-sample model.

Method

Subjects. Six psychology students (5 graduate and 1 senior, 1 male and 5 females) from Indiana University–Purdue University at Indianapolis volunteered to participate for 1 practice and 10 experimental 1-hr daily sessions. Subjects were paid according to their performance, and they earned an average of \$3.75 per session. Five students verbally expressed interest throughout the experiment. One subject (S4) became a little discouraged midway due to initially low performance, but she completed the entire experiment.

Apparatus. Same as Experiment 1.

Design and procedure. The uncertain alternative generated monetary payoffs (in units of \$.02) according to a normal distribution with a mean of +10, 0, or −10 and a standard deviation of 10 or 20 depending on the condition. The certain alternative was set equal to −5, 0, or +5, and the cost of sampling each observation from the computer equaled 0 or 2 units depending on the condition. Factorial combination of the uncertain alternative mean, uncertain alternative variance, certain value, and observation cost produced $3 \times 2 \times 3 \times 2 = 36$ conditions, and all 36 conditions were presented in a new random order within each block of 36 trials. Each session consisted of three such blocks.

The protocol for a typical trial is illustrated below:

1. Stimulus:

   Lottery Batch # 2176
   Observation Cost = 2 Certain Value = 5
   Make Your Choice

2. Subject requests a sample observation by typing a "1."
3. Screen clears.
4. Stimulus:

   Lottery Ticket # 51328 Was Worth
   Value = 23

5. Delay .5 s
6. Screen clears
7. Program returns to line 1
   (This cycle continues until a final choice is made.)
8. Subject requests a terminal choice by typing a "2" or "3."
9. Screen clears
10. Subject initiates a new trial by pushing the return key.

As the display illustrates, each trial began with the random selection of an urn or batch of lottery tickets, indicated by a new random number (e.g., #2176). The observation cost (e.g., $C = 2$) and the certain value (e.g., $k = 5$) were also indicated at the start of each trial. If the subject requested an observation from the computer (by typing the number 1), then the computer randomly selected a new lottery ticket (e.g., # 51328) and printed the value of the ticket (e.g., Value = 23). After a .5-s delay, the screen cleared, and the subject was again faced with the original choice problem, but now having more information about the uncertain alternative. This information request cycle continued until the subject eventually selected a terminal choice, the certain alternative (typing the number 2), or the uncertain alternative (typing the number 3). Following a terminal choice, a new trial was initiated.[4]

Subjects were told that their pay would equal the accumulated wins and losses produced by their terminal choices minus the cost of sampling observations. Subjects were never shown the payoffs produced by the terminal choice on each trial. Instead they were shown the total accumulated pay after each block of trials. This procedure was used to prevent trial by trial sequential effects due to runs of large wins or losses.

At the beginning of the experiment, subjects were informed (using a graphical display) that the distribution of lottery tickets was normal, but the mean and standard deviation of the distribution would be randomly selected for each choice trial from independent uniform distributions. They were told that the mean could range from −100 to 100, and the standard deviation could range from 1 to 50. This instruction was used to introduce a great deal of uncertainty. During a postexperimental interview, subjects indicated that the mean did appear to follow a uniform distribution. Also during this interview, subjects were asked to construct a relative frequency histogram based on their overall experience across all 10 sessions, using the same method as described in Experiment 1. In general the subjective histograms closely approximated the shape of the normal distribution.

Results

After the first practice session, there were no interactions with training, and the responses were averaged across the 10 experimental sessions. Although there were individual differences (noted later), averages across individuals provided an accurate representation of the qualitative pattern produced by the majority of subjects. Figures 8a and 9a (described in detail later) provide a graphical illustration of the pattern of results for the two performance measures—proportion of trials that the uncertain alternative was chosen (denoted $\bar{P}$(UA) in Figure 8a), average number of observations purchased on each trial (denoted $\bar{N}$ in Figure 9a). Each mean is based on 180 observations. Figures 8b and 9b illustrate the predictions generated by the sequential-sampling model. A detailed theoretical analysis of the results is presented in the discussion section.

Choice probability.[5] Figure 8a (top) presents the observed $\bar{P}$(UA) as a function of the programmed mean difference $[\mu(x) - k]$, produced by each of nine combinations of uncertain alternative mean, $\mu(x)$, and certain

[4] Subjects were also asked to provide confidence ratings following each terminal decision, but these results will not be reported in detail. In general, confidence ratings decreased as the observation cost increased. Holding cost constant, confidence ratings decreased as the number purchased increased. However, the rate of decrease in confidence with increased number purchased was smaller for the high-cost condition.

[5] The analysis of variance was computed on the arcsine transformed proportions. It should be noted that the ideal method for evaluating the sequential-sampling model would be to consider the response probability conditioned on the cumulative preference. However, there are two problems with this approach. First, this method requires estimating the sample path of the cumulative preference, $D(j)$, (see the solid line in Figure 7). Although the exact sequence of numeric values which determine the sample path can be recorded, the sample path underlying the decision process is unobservable. The second problem is that it would be impractical to estimate the probabilities conditioned on each possible sequence of numeric values. This method would be feasable if the outcomes were generated by a discrete (e.g., binary) rather than continuous (e.g., normal) distribution because this would reduce the number of outcome sequences to a small number. Because the purpose of Experiment 2 was to replicate and extend Experiment 1, the normal distribution was used.

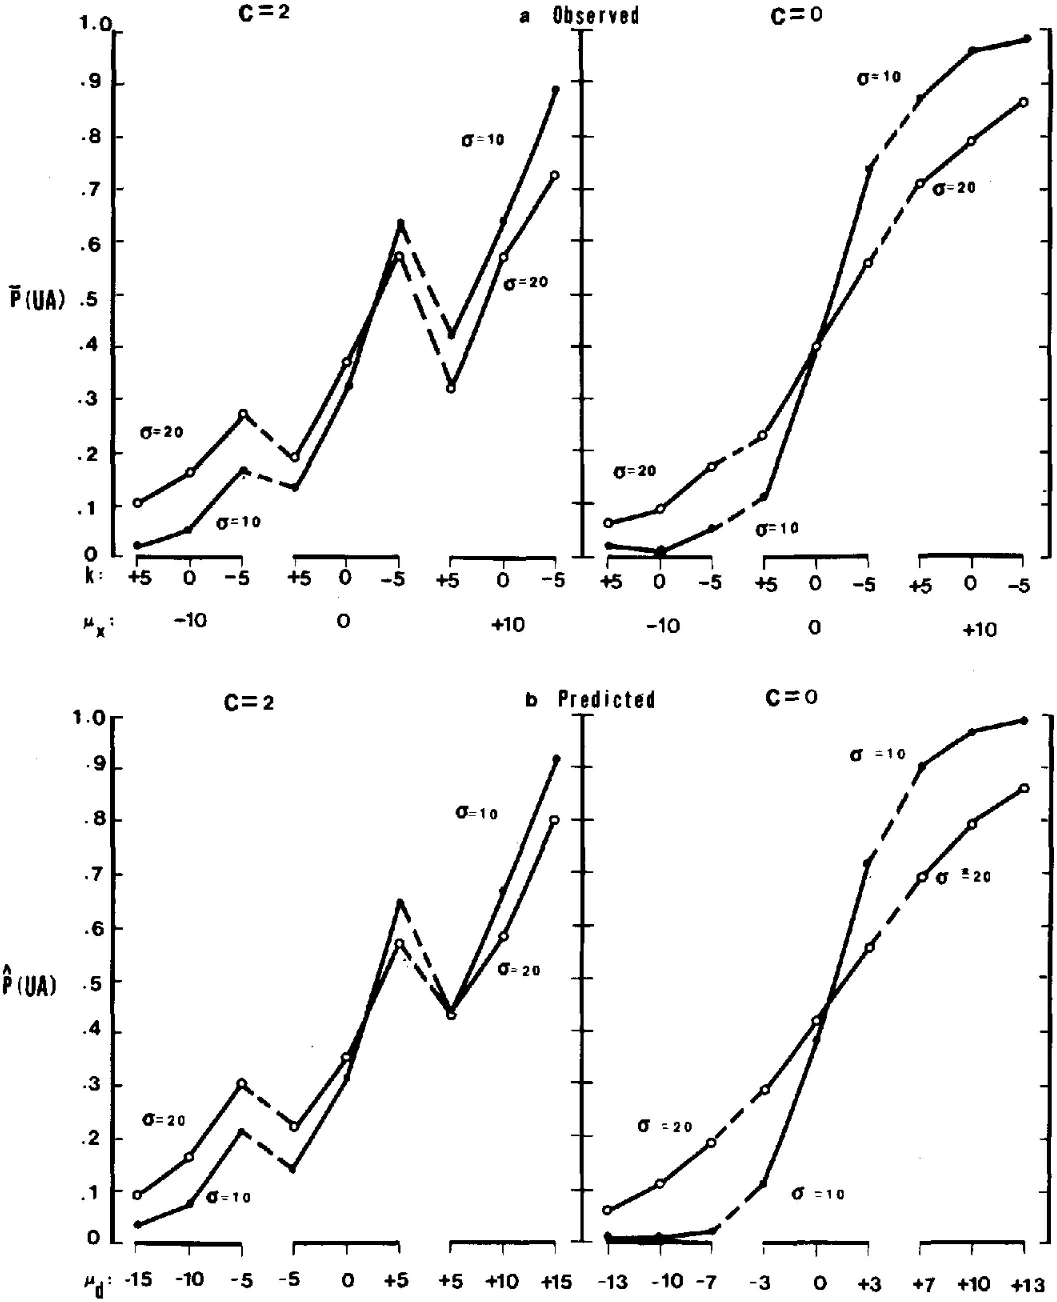

Figure 8. Probability of choosing the uncertain alternative [$\bar{P}$(UA)] plotted as a function of the mean difference, with a different curve for each level of variance (σ^2) and a different panel for each level of observation cost (C). [The top figure (8a) shows observed results, and the bottom figure (8b) shows the predictions of the sequential-sampling model.]

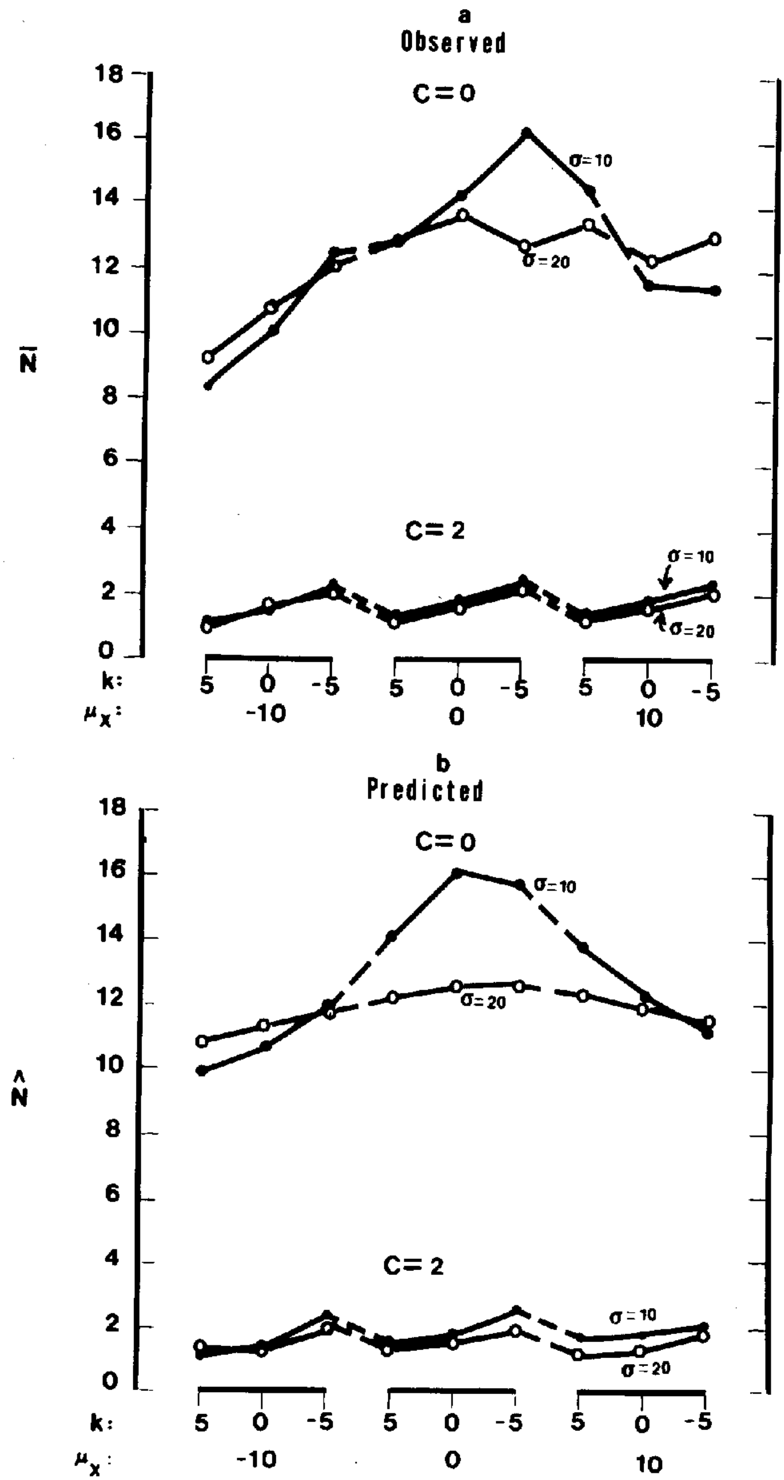

Figure 9. Mean number of observations purchased ($\bar{N}$) plotted as a function of the mean difference, with a different curve for each variance and observation cost condition. [The top figure (Figure 9a) shows observed results, and the bottom figure (Figure 9b) shows predictions for the sequential sampling model.]

Table 2
Probability of Choosing the Uncertain Alternative for Conditions With a Programmed Mean Difference of +5 or −5

			$C = 2$		$C = 0$	
$\mu(x)$	k	$[\mu(x) - k]$	Observed	Predicted	Observed	Predicted
−10	−5	−5	.22	.26	.11	.11
0	+5	−5	.16	.18	.17	.20
0	−5	+5	.60	.61	.65	.64
+10	+5	+5	.37	.44	.79	.79

Note. Each choice proportion was obtained by averaging $\bar{P}$(UA) across the low- and high-variance conditions. Predictions were derived from the sequential-sampling model.

value, k. The panels on the left and right display the responses under the high- ($C = 2$) and low- ($C = 0$) observation cost conditions, respectively. The solid and open dots indicate the choice proportions for the low- and high-variance conditions, respectively. Three major patterns are apparent in Figure 8a:

1. The most obvious feature of Figure 8a is the jagged sawtooth pattern obtained under the high-cost condition ($C = 2$), and the smooth ogival pattern obtained under the low-cost condition ($C = 0$). Under the high-cost condition, a decrease in P(UA) occurred whenever the certain value changed from $k = -5$ to $k = +5$ (see the dashed lines for $C = 2$). Under the low-cost condition, an increase in P(UA) occurred whenever the certain value changed from $k = -5$ to $k = +5$ (see the dashed lines for $C = 0$).

Four of the 6 subjects produced the sawtooth pattern under the high cost and the ogival pattern under the low cost. Of the 2 that failed to show this pattern, 1 subject (S2) produced sawtooth patterns, and the other subject (S4) produced ogival patterns under both cost conditions. A repeated measures ANOVA indicated a statistically significant Uncertain Alternative Mean × Certain Value × Observation Cost interaction, $F(4, 20) = 4.05$, $p < .05$.

Table 2 analyzes this interaction effect in more detail. It presents P(UA) for four different combinations of $\mu(x)$ and k, averaged across variance conditions. The first two combinations result in a constant mean difference equal to $[\mu(x) - k] = -5$, and the last two combinations result in a constant mean difference equal to $[\mu(x) - k] = +5$. The first pair of columns indicate the experimental conditions, the next pair of columns indicate the probabilities obtained under the high-cost condition, and the last pair of columns indicate the probabilities obtained under the low-cost condition. Note that changing from $k = -5$ to $k = +5$ decreased $\bar{P}$(UA) under the high cost, but this same manipulation increased $\bar{P}$(UA) under the low cost. The results in Table 2 suggest that the certain value is more effective under the high cost, and the uncertain alternative mean is more effective under the low-cost condition.

2. A Variance × Mean Difference interaction occurred under both cost conditions. This can be seen by comparing the curve under the low-variance condition (solid dots) with the curve under the high-variance condition (open dots) within each panel of Figure 8a. Note that for both cost conditions, the high-variance curve lies above when the mean difference is negative, but the low-variance curve lies above when the mean difference is positive. This interaction is similar to the $\sigma \times k$ crossover interaction observed in Experiment 1. Theoretically, this interaction was due to a reduction in discriminability under the high-variance condition. All 6 subjects produced this crossover pattern, and the effect was statistically significant.

3. Table 3 provides a detailed analysis of the effects that observation cost had on accuracy. First consider the top part of Table 3, which only involves conditions with the uncertain alternative mean equal to zero, as in Experiment 1. Each of the four Variance × Observation Cost cells contain the proportion

of correct responses calculated by averaging $\bar{P}$(UA) when $k = -5$, with $1 - \bar{P}$(UA) when $k = +5$.

The results presented in the top half of Table 3 replicate the pattern of results from Experiment 1. Accuracy increased as the observation cost decreased under the low-variance condition (.75 vs. .82). Accuracy decreased as the observation cost decreased under the high-variance condition (.69 vs. .66).

Next consider the bottom half of Table 3, which only involves conditions with the certain value equal to zero. The proportion correct within each cell was obtained by averaging the $\bar{P}$(UA) when $\mu(x) = +10$, with $1 - \bar{P}$(UA) when $\mu(x) = -10$. As can be seen in the bottom half of Table 3, accuracy increased as the observation cost decreased for both the low- and the high-variance conditions.

The results shown in Table 3 agree with the predictions of the correlated bias hypothesis proposed in the discussion of Experiment 1 and the introduction to Experiment 2. A more detailed analysis is given in the discussion.

Table 3
Proportion of Correct Responses for Each Combination of Variance and Observation Cost Condition

Variance condition	$C = 2$	$C = 0$
	$\mu(x) = 0$	
$\sigma = 10$		
observed	.75	.82
predicted	.75	.80
$\sigma = 20$		
observed	.69	.66
predicted	.67	.63
	$k = 0$	
$\sigma = 10$		
observed	.80	.98
predicted	.79	.98
$\sigma = 20$		
observed	.71	.85
predicted	.71	.84

Note. The proportion within each cell of the top table was calculated by averaging $\bar{P}$(UA) given $\mu(x) = 0$ and $k = -3$ with $[1 - \bar{P}$(UA)] given $\mu(x) = 0$ and $k = +3$. The proportion with each cell of the bottom table was calculated by averaging $\bar{P}$(UA) given $k = 0$ and $\mu(x) = +10$ with $[1 - \bar{P}$(UA)] given $k = 0$ and $\mu(x) = -10$. Predictions were derived from the sequential-sampling model.

Four subjects produced the same pattern of results as shown in top half of Table 3, and all 6 produced the pattern shown in the bottom half of Table 3. Two subjects (S2 and S4 again) failed to produce the pattern in top half of Table 3. For both of these subjects, accuracy consistently increased as observation cost decreased under both variance conditions.

Number of samples purchased. Figure 9a (top) presents the observed $\bar{N}$ as a function of the programmed mean difference, $[\mu(x) - k]$. The pair of curves located on the top half of the figure represent the results for the low-cost condition, and the pair of curves located on the bottom half of the figure represent the results for the high-cost condition. The lines connected by the solid and open dots represent the low- and high-variance conditions, respectively. Two major patterns are apparent in Figure 9a:

1. The most salient feature of Figure 9a is the difference in the functions relating $\bar{N}$ to the mean difference for the low- and high-cost conditions. When the cost was high, $\bar{N}$ was primarily determined by the certain value, and $\bar{N}$ increased as certain value decreased from +5 to −5. Five of the 6 subjects produced this pattern of results for the high-cost condition (S4 produced an inverted U). A repeated measures ANOVA performed on the high cost data produced a significant main effect due to the certain value, $F(2, 10) = 4.6$, $MS_e = 1.951$, $p < .01$.

The effect of the certain value on $\bar{N}$ under the high-cost condition is similar to the mean choice time results from Experiment 1 under the short ($L = 1$) deadline. This result suggests that the criterion bounds were more extended when the certain value was negative.

2. Under the low-cost condition, $\bar{N}$ generally followed an inverted U-shape function of the mean difference as predicted by the sequential-sampling model. The inverted U shape was much flatter under the high variance as compared to low-variance condition. All 6 subjects produced the inverted U-shape pattern shown in Figure 9a. A repeated measures ANOVA performed on the low cost data produced a significant Uncertain Alternative Mean × Uncertain Alternative Variance × Certain Value interaction effect, $F(4, 20) = 5.26$, $MS_e = 1.4125$, $p < 01$.

Recall that the fixed-sample model predicted no effect of the uncertain alternative mean and variance on the number sampled. Contrary to this prediction, the mean and variance had an effect on the number purchased for all 6 subjects.

Discussion

Summary of major results. Probability of choosing the uncertain alternative was an increasing function of the difference between the uncertain alternative mean and the certain value (see Figure 8). Under the high-observation cost condition the shape of this function was sawtooth, suggesting that the certain value had a much stronger effect on choice than the uncertain alternative mean when sampling was severely restricted. Under the low-cost condition, the shape of this function was smooth and ogival, suggesting that the uncertain alternative mean had a much stronger effect on choice than the certain value when sampling was unrestricted.

The slope of the function relating the mean difference to choice probability was much steeper under the low- as compared with the high-variance condition. This produced a Variance × Mean Difference interaction, extending the $\sigma \times k$ interaction reported in Experiment 1. Theoretically, discriminability was reduced by increasing the variance.

On the one hand, when the uncertain alternative mean was fixed at zero, then decreasing the observation cost increased accuracy under the low variance, but decreased accuracy under the high variance, replicating the pattern obtained in Experiment 1 (compare Table 1 and the top half of Table 3). On the other hand, when the certain value was fixed at zero, then decreasing the observation cost always increased accuracy under both low- and high-variance conditions (see the bottom half of Table 3). Both patterns (Table 3) were exactly in accord with the predictions of the correlated bias hypothesis.

Under the high-cost condition, the mean number of observations purchased increased as the certain value decreased, similar to the choice response time results of Experiment 1. Under the low-cost condition, the number purchased was an inverted U-shape function of the mean difference. The inverted U-shape function was flatter under the high- as compared with the low-variance condition (see Figure 9). The strong influence of the uncertain alternative mean and variance on the number purchased rules out the fixed-sample model.

Relation to prior research. Irwin and Smith (1956, 1957) investigated the effects of mean and variance on number of observations required to determine whether the mean of a deck of cards was positive or negative. Similar to the present results, they found that the number of observations requested was an inverted U-shape function of the card deck mean. Unlike the present results, they found that increasing the variance always increased the number of observations required to reach a decision independent of the mean of the deck.

There are several differences between the present study and those of Irwin and Smith (1956, 1957). This includes the form of the stimulus distribution, selection of mean and variance parameters, amount of training, payoff procedure, task instructions, and subject populations. Further research is needed to isolate the particular cause for the different results.

Theoretical analysis. One may question whether or not the sequential-sampling model can explain the complicated pattern of results observed in Experiment 2. To answer this question, it was necessary to obtain quantitative fits. The major assumptions used to fit the model are described in the next few paragraphs, and the details are given in Appendix A. In brief, eight parameters were required to fit the 36 observed choice proportions in Figure 8. These same eight parameters were then used to generate predictions for the number of observations purchased shown in Figure 9. The predictions for number purchased were adjusted by fitting a slope and intercept parameter separately for each observation cost condition. The fitted results are shown in Tables 4 and 5 and Figures 8b and 9b. The first four columns of Tables 4 and 5 indicate the experimental condition, the next two columns indicate the derived parameters,[6] and the last two pairs

[6] Only eight parameters were fit to the data. The derived parameters were calculated from these eight. The

Table 4
Observed $\bar{P}(UA)$ and Predicted $\hat{P}(UA)$ Probabilities of Choosing the Uncertain Alternative, Observed $\bar{N}$ and Predicted $\hat{N}$ Number of Observations Purchased

Experimental condition				Derived parameters		Observed and predicted values			
C	$\sigma(x)$	$\mu(x)$	k	$A\theta$	δ	$\bar{P}$(UA)	$\hat{P}$(UA)	$\bar{N}$	$\hat{N}$
0	10	−10	+5	−9.6	−.3	.02	.001	8.3	9.9
0	10	−10	0	−7.2	−.3	.01	.004	10.0	10.6
0	10	−10	−5	−4.8	−.3	.05	.02	12.5	11.9
0	10	0	+5	−2.4	−.3	.11	.11	12.9	14.1
0	10	0	0	0	−.3	.39	.38	14.3	16.1
0	10	0	−5	2.4	−.3	.74	.72	16.3	15.7
0	10	+10	+5	4.8	−.3	.87	.90	14.6	13.8
0	10	+10	0	7.2	−.3	.96	.97	11.6	12.3
0	10	+10	−5	9.6	−.3	.98	.99	11.5	11.2
0	20	−10	+5	−3.02	−.3	.06	.06	9.2	10.8
0	20	−10	0	−2.26	−.3	.09	.11	10.8	11.3
0	20	−10	−5	−1.52	−.3	.17	.19	12.1	11.8
0	20	0	+5	−.76	−.3	.23	.29	12.9	12.2
0	20	0	0	0	−.3	.40	.42	13.8	12.5
0	20	0	−5	.76	−.3	.56	.56	12.8	12.5
0	20	+10	+5	1.52	−.3	.71	.69	13.5	12.3
0	20	+10	0	2.26	−.3	.79	.79	12.4	11.9
0	20	+10	−5	3.02	−.3	.86	.86	13.1	11.5

Note. Predictions were derived from the sequential-sampling model. Global fit index R^2 = .996 for $\hat{P}$(UA) and .708 for $\hat{N}$; Global fit index RMS = .022 for $\hat{P}$(UA) and 1.02 for $\hat{N}$.

of columns present the observed and predicted values. Included in the note for Table 4 are the two global fit indices for each performance measure—R^2 is the percentage of variance predicted by the model, and RMS is the root mean square error.

The first question is how does the model explain the interaction shown in Table 2? According to the correlated bias hypothesis, subjects select a larger or smaller certain criterion magnitude depending on the certain value. This bias has a very strong effect on choice probability under the high-cost conditions, because the criterion bounds are close to the starting point of the cumulative preference. Thus the results shown under the C = 2 column of Table 2 were due to the strong influence of the correlated bias.

Under the low-cost condition, subjects maintain a correlated bias, but they also increase the total distance between the bounds to allow stronger preferences to accumulate before reaching a decision. As noted earlier, a large increase in the total distance practically eliminates the effect of the correlated bias. This analysis implies that the cell differences under column C = 0 in Table 2 should have been in the same direction as those found under column C = 2, but reduced nearly to zero in magnitude. On the contrary, the observed differences in each cell under the column C = 0 are in the opposite direction.

The results in Table 2 suggest that the certain value was less effective than the uncertain alternative mean for determining choice under the low observation cost condition. One reason may be that subjects occasionally forget to subtract the utility of certain value during preference accumulation,

derived parameter $A\theta$ represents the discriminability times the total distance between the bounds, and δ is the relative bias. The measures A and θ are not presented separately because they are not uniquely identified, but the product is uniquely identified. See Appendix A for more details.

particularly after a large number of observations. Define p as the probability that the subject forgets to subtract the utility of the certain value when computing the increment $d(j)$. Then the mean increment equals $\mu(d) = [\mu(u) - (1 - p)u(k)]$. For simplicity, assume that the utility function is linear, so that the mean increment can be expressed as $\mu(d) = [\mu(x) - (1 - p)k]$, where $\mu(x)$ is the programmed mean for the uncertain alternative, and k is the certain value. Setting $p = 0$ for the high observation cost and setting $p = (1/3)$ for the low observation cost produces the mean increments shown on the abcissa of Figure 8b. Note that the mean increments, $\mu(d)$, are in perfect rank order agreement with the observed choice probabilities for the low observation cost condition in Figure 8.

The second question is how does the model explain the crossover interaction between the mean difference and the variance shown in Figure 8? This is easily explained as the result of a reduction in the discriminability parameter, θ, which is the ratio of the mean increment divided by the variance of the increment. In other words, preferences bounce up and down haphazardly under large variances, but they cumulate systematically in the correct direction under small variances.

The third question is how does the model explain the interaction shown in Table 3. First consider the top half of Table 3 with the uncertain alternative mean equal to zero. In this case, the correct decision depends on the sign of the certain value. By adjusting the certain criterion magnitude according to the certain value, accuracy can be improved. However, this improvement is largely eliminated when the total distance between bounds is large. Thus, decreasing the observation cost has two opposing effects. On the one hand, it reduces the effect of the correlated bias which reduces accuracy and, on the other hand, it increases the average sample size which increases accuracy. When the variance is large and discrimination is very poor, the

Table 5
Observed $\bar{P}(UA)$ and Predicted $\hat{P}(UA)$ Probabilities of Choosing the Uncertain Alternative, Observed $\bar{N}$ and Predicted $\hat{N}$ Number of Observations Purchased

Experimental condition				Derived parameters		Observed and predicted values			
C	$\sigma(x)$	$\mu(x)$	k	$A\theta$	δ	$\bar{P}$(UA)	$\hat{P}$(UA)	$\bar{N}$	$\hat{N}$
2	10	−10	+5	−2.66	−.94	.02	.03	1.1	1.1
2	10	−10	0	−2.0	−.72	.05	.07	1.5	1.4
2	10	−10	−5	−1.32	−.3	.16	.21	2.2	2.4
2	10	0	+5	−.88	−.94	.13	.14	1.3	1.5
2	10	0	0	0	−.72	.32	.31	1.8	1.8
2	10	0	−5	1.32	−.3	.63	.64	2.4	2.6
2	10	+10	+5	.88	−.94	.42	.44	1.4	1.7
2	10	+10	0	2.0	−.72	.64	.66	1.9	1.9
2	10	+10	−5	4.0	−.3	.89	.91	2.3	2.0
2	20	−10	+5	−1.42	−.94	.10	.09	1.0	1.2
2	20	−10	0	−1.06	−.72	.16	.16	1.5	1.3
2	20	−10	−5	−.70	−.3	.27	.30	2.0	1.9
2	20	0	+5	−.46	−.94	.19	.22	1.2	1.3
2	20	0	0	0	−.72	.37	.35	1.6	1.5
2	20	0	−5	.70	−.3	.57	.57	2.2	1.9
2	20	+10	+5	.46	−.94	.32	.43	1.2	1.2
2	20	+10	0	1.06	−.72	.57	.58	1.6	1.3
2	20	+10	−5	2.14	−.3	.72	.80	2.1	1.9

Note. Predictions were derived from the sequential-sampling model. Global fit index $R^2 = .977$ for $\hat{P}$(UA) and .808 for $\hat{N}$; Global fit index RMS = .0375 for $\hat{P}$(UA) and .188 for $\hat{N}$.

correlated bias has a larger effect on accuracy than does the increased sample size. The opposite is true when the variance is small. This explains the crossover interaction in the top half of Table 3.

Next consider the bottom half of Table 3, with the certain value fixed at zero. In this case the bias does not vary, and only the distance between bounds changes as the observation cost decreases. Increasing the total distance increases the average number of samples purchased, and accuracy improves for both variance conditions.

The last question is how does the model explain the different pattern of results obtained for the mean number purchased under the low- and high-observation cost conditions (see Figure 9)? Under the high cost, the criterion magnitudes are very small and the number purchased is largely determined by the adjustment of the certain criterion magnitude correlated with the certain value. Under the low-cost conditions, the criterion magnitudes are very large, and the mean number purchased is largely determined by the mean increment or average rate of movement towards each criterion bound.

In sum, by comparing the predicted with the observed patterns in Figures 8 and 9, and Tables 2 and 3, it is clear that the model provides a fairly complete description of the qualitative pattern of results. Although there are some minor quantitative deviations, these do not appear severe enough to warrant rejecting the model.

Individual differences. Four subjects produced the pattern of results illustrated in Figure 8a. However, 2 subjects produced inconsistent patterns—subject S2 produced saw tooth patterns, whereas subject S4 produced ogival patterns in both observation cost conditions. The behavior of subject S4 can be explained by assuming that this subject was unbiased across all conditions. The behavior of subject S2 is more difficult to explain. Previously, it was assumed that the adjustment of the certain criterion bound was constant, whereas the relative bias, δ, varied across observation cost conditions. One explanation for the pattern produced by subject S2 is that this subject maintained a constant relative bias, δ, by allowing the adjustment of the certain criterion magnitude to vary across observation cost conditions. If the relative bias remains constant across observation cost conditions (as assumed for both S2 and S4), then the model asserts that accuracy always increases as the observation cost increases because of the increased total distance. This was exactly the pattern of results obtained from subjects S2 and S4.

Conclusion. When individuals are given the opportunity to learn the outcomes produced by an uncertain action, and the number of observations are limited only by the cost of sampling, then they tend to use a sequential- rather than a fixed-sample decision process. According to the sequential-sampling model, individuals continue sampling observations until their cumulative preference exceeds a criterion bound for one of the alternatives. The criterion boundaries are determined by individual attitudes towards uncertainty, by prior information such as the value of the certain alternative, and by the observation cost. This model provided a fairly complete explanation for the complex pattern of results observed in the second experiment.

Several assumptions were required to explain the results. Most of the assumptions are not new, but are similar to those commonly made by all sequential-sampling theories of decision making (e.g., the effects of mean and variance on discriminability, and the effects of observation costs on the criterion magnitudes). Two assumptions were new—the correlated bias hypothesis was needed to explain the fact that the certain value had a stronger effect than the uncertain alternative mean when sampling was severely limited, and the forgetting hypothesis was needed to explain the fact that the uncertain alternative mean had a stronger effect than the certain value when sampling was unlimited. A simpler set of assumptions may be able to explain part of the results, but it appears that both the correlated bias and the forgetting hypothesis are needed to explain the entire pattern for both choice probability and number of observations purchased.

General Discussion

This article began by making a distinction between decision making under risk (perfect knowledge of outcome probabilities) and decision making under uncertainty (outcomes probabilities initially unknown and learned through experience). Although the latter problem seems more realistic, the majority of past research has been concerned with decision making under risk. Two experiments were reported which investigated the cognitive processes involved in decision making under uncertainty—the first used a probability learning task, and the second used an information purchasing task. What has been learned from these experiments? Three conclusions are discussed as follows:

First, two new empirical findings were discovered. One finding was the crossover interaction shown in Tables 1 and 3. Intuitively, one would expect that if a decision maker has more time to decide, or observes more sample outcomes, then the likelihood of choosing the correct alternative (see Footnote 1) should increase. However, if the choice is between a nonzero sure thing alternative, and an alternative that produces a wide dispersion of outcomes centered at zero, then the opposite tends to occur—the likelihood of choosing the correct alternative decreases.

Another finding was the crossover interaction shown in Table 2. If a decision must be based on a small number of observations sampled from an uncertain alternative, then the value of a certain alternative is more effective than the mean of the uncertain outcomes for determining choice. The opposite tends to be true when the decision is based on a large number of observations.

The second conclusion is that many of the algebraic-deterministic models developed within the risky decision-making paradigm cannot be applied to decision making under uncertainty in any simple manner.[7] These theories typically assume that each alternative is assigned a single scale value, such as a subjective expected utility. Simple scalability models are probabilistic versions of these algebraic-deterministic models. In both of the reported experiments, a fundamental property of simple scalability, the independence between alternatives property, was violated by all subjects.

The third conclusion is that decision making under uncertainty involves learning and memory processes when the outcome probabilities are learned through experience. There are two ways that an individual can decide under uncertainty: one is to recall outcomes of previous decisions under similar circumstances, and another is to observe new outcomes sampled from the unknown distribution before making a final decision. In either case, a fixed- or sequential-sampling strategy could be used. If a fixed-sample strategy is used, the number of past outcomes recalled or the number of new outcomes observed is determined first. Then a decision is made by selecting the alternative producing the larger average utility. If a sequential-sampling strategy is used, the decision maker continues recalling past outcomes or continues observing new sample outcomes until a cumulative preference exceeds a critical bound for one of the alternatives. The cumulative preference is formed by comparing the utilities of the outcomes produced by each alternative, and summing these individual comparisons.

What determines whether a fixed- or sequential-sampling strategy will be selected? In the first experiment, the use of a deadline time limit was likely to encourage a fixed-sample strategy for the following reason. Presumably, subjects believe that the probability of making the correct decision is increased by recalling more past outcomes. Although the number retrieved is limited by the deadline, no loss is produced by recalling as many past outcomes as possible within this limit. Thus the maximum number possible is recalled. In the second experiment, the number of observations that could be purchased was unlimited, but each one cost a fixed amount. This procedure would tend to encourage a sequential-sampling strategy because the subject must simultaneously maximize accuracy and minimize the number purchased. Thus subjects will continue sampling if the cumulative preference is weak in order to increase accuracy, but they will stop sampling if the cumulative preference is strong in order to minimize the observation cost.

The results of Experiment 2 clearly indicate that all subjects adopted a sequential-sampling strategy. The results of Experiment 1 do not provide a clear cut way to determine whether

[7] One exception is the additive difference model (Tversky, 1969) which is similar to the sequential-sampling model.

a fixed- or sequential-strategy was used. Both models are capable of accounting for the observed pattern of results if a correlated bias is assumed.

In summary, the present research suggests that more attention should be given to theories of decision making that emphasize learning and memory retrieval (e.g., the fixed- and sequential-sampling models) rather than concentrating exclusively on deterministic-algebraic theories (e.g., subjective expected utility models). There are a number of questions about the nature of the decision processes that remain to be answered. First, how are preferences accumulated when a sequential-sampling strategy is assumed? The present analysis proposed one idea, but others are possible such as the "accumulator" model described by Vickers (1979). Second, what factors influence the selection of criterion bounds for the sequential-sampling strategy? For example, suppose the observation costs rapidly increased across samples rather than remaining constant as in Experiment 2. According to the optimal model (Rapoport & Burkheimer, 1971), the criterion bounds decrease in magnitude after each observation is purchased as a consequence of rising observation costs. Finally, how could the model be extended to more than two alternatives? Ratcliff (1978) has proposed one approach, but there are several other possibilities. The answer to these questions must await further research on decision making under uncertainty.

References

Aschenbrenner, K. M., Albert, D., & Schmalhofer, F. (1984). Stochastic choice heuristics. *Acta Psychologica, 56,* 153–166.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963a). Stochastic models of choice behavior. *Behavioral Science, 8,* 41–55.

Becker, G. M., DeGroot, M. H., & Marschak, J. (1963b) Probabilities of choice among very similar objects. *Behavioral Science, 8,* 306–311.

Ben Zur, H., & Breznitz, S. J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica, 47,* 89–104.

Busemeyer, J. R. (1982). Choice behavior in a sequential decision making task. *Organizational Behavioral and Human Performance, 29,* 175–207.

Cox, D. R., & Miller, H. D. (1965). *The theory of stochastic processes.* London: Methuen.

Edwards, W. (1962). Subjective probabilities inferred from decisions. *Psychological Review, 69,* 109–135.

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.

Fishburn, P. C. (1976). Binary choice probabilities between gambles: Interlocking expected utility. *Journal of Mathematical Psychology, 14,* 99–122.

Irwin, F. W., & Smith, W. A. S. (1956). Further tests of theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology, 52,* 345–348.

Irwin, F. W., & Smith, W. A. S. (1957). Value, cost, and information as determiners of decision. *Journal of Experimental Psychology, 54,* 229–232.

Irwin, F. W., Smith, W. A. S., & Mayfield, J. (1956). Tests of two theories of decision in an "expanded judgment" situation. *Journal of Experimental Psychology, 51,* 261–268.

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

Lee, W. (1971). *Decision theory and human behavior.* New York: Wiley.

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

Lopes, L. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 133–144.

Luce, R. D., & Raiffa, H. (1957). *Games and decisions: Introduction and critical survey.* New York: Wiley.

Luce, R. D., & Suppes, P. (1965). Preference, utility, and subjective probability. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.), *Handbook of Mathematical Psychology, Vol. 3.* New York: Wiley.

Myers, J. R., & Atkinson, R. C. (1964). Choice behavior and reward structure. *Journal of Mathematical Psychology, 1,* 170–204.

Myers, J. R., Suydam, M. M., & Gambino, B. (1965). Contingent gains and losses in a risk taking situation. *Journal of Mathematical Psychology, 2,* 363–370.

Petrusic, W. M., & Jamieson, D. G. (1978). Relation between probability of preferential choice and time to choose changes with practice. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 474–482.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85,* 59–108.

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

Stein, W. E., & Rapoport, A. (1978). A discrete time model for detection of randomly presented stimuli. *Journal of Mathematical Psychology, 17,* 110–137.

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

Swensson, R. G., & Thomas, R. E. (1974). Fixed and optimal stopping models for two choice discrimination times. *Journal of Mathematical Psychology, 11,* 213–236.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review, 76,* 31–48.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review, 79,* 281–299.

Tversky, A., & Russo, J. E. (1969). Substitutability and similarity in binary choices. *Journal of Mathematical Psychology, 6,* 1–19.

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

Appendix A

The predictions generated by the sequential-sampling model are based on the theory of random walk stochastic processes (Cox & Miller, 1965, chap. 2 & 3). There are two methods for deriving these predictions, and both methods provide only approximations. The first method requires one to assume that the excess over the boundaries (see the difference between $D(12)$ and the lower bound in Figure 7) is small relative to the criterion magnitude, so that it can be ignored. This is not a reasonable assumption for the present study (because of the small mean number of samples purchased in Experiment 2 under the high-cost condition), so an alternative approximation method was used. Stein & Rapoport (1978) recommended using a finite state Markov chain to approximate the continuous state Markov process of the random walk $D(j)$. This approach was realized by discretizing the continuous interval $[-\beta, \alpha]$ into m equally spaced intervals, and constructing a transient state matrix that approximated the transition densities $f[D(j+1)|D(j)] = f[d(j)]$ for the normal distribution. The theory for finite state Markov chains with absorbing states $(-\infty, -\beta)$ and $(\alpha, +\infty)$ was then used to calculate $P(\text{UA})$ and $E(N)$.

The width of each discrete state was defined as $w = (\alpha + \beta)/m$. A set of m discrete preference states was constructed from the original continuous preference states by defining the discrete states as $D_i = \alpha - (i - .5)w$. The probability of a transition from state D_i to state D_j was set equal to $T_{ij} = F[(i-j)w + .5w] - F[(i-j)w - .5w]$, where F is the cumulative normal distribution function with mean $\mu(d)$ and variance $\sigma^2(u)$. Let $\mathbf{T} = [T_{ij}]$ represent the $m \times m$ transient state probability matrix.

The probability of exceeding α from state D_i equals $r_{i1} = 1 - F[(i - .5)w]$, and the probability of exceeding β from state D_i equals $r_{i2} = F[(i - m - .5)w]$. Let $\mathbf{R} = [r_{i1}\ r_{i2}]$ represent the $m \times 2$ response probability matrix.

Define $\boldsymbol{\pi}$ as the initial probability row vector. It was assumed that the cumulative preference always starts at zero so that all elements of $\boldsymbol{\pi}$ were set to zero except for the element corresponding to the zero preference state which was set to unity. $P(\text{UA})$ can be obtained from the matrix equation

$$[P(\text{UA})\ 1 - P(\text{UA})] = \boldsymbol{\pi}(\mathbf{I} - \mathbf{T})^{-1}\mathbf{R},$$

where $\mathbf{I}$ is the identity matrix. The mean number of samples purchased can be calculated from the matrix equation

$$E(N) = \boldsymbol{\pi}(\mathbf{I} - \mathbf{T})^{-2}\mathbf{RJ},$$

where $\mathbf{J}$ is a two-element column vector with each element set to unity.

For both cost conditions, the mean increment was set to $\mu(d) = \mu(x) - (1 - p)k$, where $\mu(x)$ is the programmed mean of the uncertain alternative. For the high cost p was set to zero, and for the low cost p was set to 1/3. The variance of the utility of the uncertain alternative was set equal to the programmed variance.

The criterion bounds were estimated from the choice data using the following assumptions. For both cost conditions, two uncertain criterion bounds were estimated—one for each variance condition. Note that changes in the criterion magnitude with changes in variance may simply reflect the possibility that the variance of the utilities was not equal to the programmed variance. This is due to the fact that the variance of the increments and the criterion magnitudes trade off, so changes in one parameter can be offset by changes in another.

The treatment of the certain criterion bound depended on the observation cost condition. For the high-cost condition, the certain criterion magnitude was set equal to the uncertain criterion bound plus an adjustment factor, b_k, which depended on the certain value k. Under the low-cost condition, it was assumed that for all practical purposes the effect of the correlated bias was reduced to zero. Thus the relative bias was assumed to be constant, and this was achieved by forcing the certain criterion magnitude to be proportional to the uncertain criterion magnitude.

In sum, a total of three parameters were estimated under the low-cost condition: one relative bias, and two uncertain criterion bounds. A total of five parameters was estimated from the choice proportions obtained under the high cost: two uncertain criterion bounds, and three adjustments for the certain criterion bound.

The parameters were estimated by minimizing the sum of squared error between the model predictions and the sample means shown in Figure 8a. A hill climbing grid search algorithm was used with a fairly course grid. The fits could probably be improved with a more efficient search algorithm.

The parameter estimates obtained from the choice proportions were then used to generate predictions for the mean number of samples purchased. However, the predicted values were related to the observed values by an arbitrary multiplicative constant because the mean increment is only defined up to an arbitrary constant on the basis of the choice data. Thus a linear regression equation was used to generate the predictions reported in Table 4. A separate linear regression equation was used for the low- and high-observation cost conditions, because the mean increment changed across these two conditions.

The discriminability parameter θ can be multiplied by an arbitrary constant, and the criterion bounds can be divided by this same constant to produce identical predictions for choice probability. Thus only the product $A\theta$ and the relative bias δ can be unambiguously interpreted. These parameter values are reported in Table 4.

Received June 24, 1984
Revision received November 30, 1984 ■