TEL-AVIV

9-88

THE FOERDER INSTITUTE FOR ECONOMIC RESEARCH

TEL-AVIV UNIVERSITY

RAMAT AVIV ISRAEL

מכון למחקר כלכלי ע"ש ד"ר ישעיהו פורדר ז"ל

ע"י אוניברסיטת תל־אביב

BEHAVIORALLY CONSISTENT OPTIMAL STOPPING RULES

by

Edi Karni* and Zvi Safra**

Working Paper No.9-88

February, 1 9 8 8

* The Johns Hopkins University; ** Tel-Aviv University

This research was supported by the NSF under grant # SES87-08360 and by funds granted to the Foerder Institute for Economic Research by KEREN RAUCH

FOERDER INSTITUTE FOR ECONOMIC RESEARCH
FACULTY OF SOCIAL SCIENCES,
TEL-AVIV UNIVERSITY, RAMAT AVIV, ISRAEL.

BEHAVIORALLY CONSISTENT OPTIMAL STOPPING RULES

by

Edi Karni and Zvi Safra

1. Introduction

Much of the analysis of optimal stopping rules is based on the assumption that the individuals seek to maximize expected utility[1]. This assumption lends the problem a specific (linear) structure that accounts for the fact that the decisions are dynamically consistent. That is, because of the linearity in the probabilities of the objective functional, the optimal stopping rule as of any given stage in the sampling process agrees with the continuation as of that stage of the optimal stopping rule formulated at the outset.[2]

In recent years, experimental evidence indicating persistent and systematic violations of expected utility theory has inspired the development of new decision theories, in which preferences over risky prospects are not necessarily linear in the probabilities.[3] In the context of sequential decisions, however, such nonlinearity implies that optimal decision strategies are dynamically inconsistent. Thus, a decision maker acting in accordance with one of these theories may find himself in a position in which he chooses one course of action over another knowing that subsequent parts of the chosen course will not be implemented and that the course that will be followed is inferior to the course that was rejected. In the specific context of sequential search, a decision maker may decide to reject a given observation and continue the search on the assumption that he will follow a given stopping rule, knowing that the rule that he will actually follow does not justify the rejection of the given observation. To us, this self-deceiving mode of behavior seems unreasonable. However, ruling it out without assuming expected utility maximizing behavior requires the imposition of some other restriction. We assume, therefore, that, individuals facing sequential decisions restrict their choice to courses of action that they know they will actually follow. This hypothesis is called **behavioral consistency**.[4]

In this paper we apply behavioral consistency to the analysis of optimal stopping rules. As in Karni and Safra (1987b), (1988), we represent a decision maker by a set of agents, one agent at each stage of the sampling process. The agents are assumed to act in their own self-interest and the decision makers behavior is modeled as a subgame perfect equilibrium of the game among the representing agents. Using this procedure we study the implications of the new decision models for the theory of optimal stopping rules. In particular, we examine the differences in the character of the optimal stopping rules resulting from variations in the nature of the objective function for two sequential search models in which a decision maker draws successive observations from a known distribution without recall. In the first model the observations are free but their number is finite. In the second model the number of observations is unbounded but each additional observation involves a fixed cost.

We characterize the behaviorally consistent optimal stopping rules for both models. In the second model, where such rules may not exist, we give sufficient conditions for existence. We show that when the decision maker's preferences are quasi-convex, the optimal stopping rule is qualitatively similar to that obtained under expected utility theory. However, if the preferences are strictly quasi-concave, the behaviorally consistent optimal stopping rules in both models change drastically. A typical stopping rule in such case involves the use of mixed strategies. For instance, if the preferences satisfy a condition introduced by Machina (1982) under the name hypothesis II,[5] then a behaviorally consistent optimal stopping rule is characterized as follows: There exists an interval, say [a,b], such that if the value of the most recent observation is smaller than a, then the observation is rejected and the search continues; if it is higher than b, then the observation is accepted and the search terminates. If the observation is between a and b, the decision maker employs a mixed strategy and rejects the observation with positive probability. This probability decreases monotonically with the value of the observation.

2. Bounded Sampling from a Known Distribution Without Recall

(2.1) Specification of the model and the optimal stopping rules under expected utility theory

Consider the following situation: A motorist driving on a highway is short of gasoline and must refuel in one of the next three gas stations. The motorist passes the station in a sequential order. Once he passes a station he may not return to it. Suppose that the motorist knows the distribution of prices and is interested in minimizing the expected cost of gasoline, what is his optimal stopping strategy?

To analyze this problem we present it formally. A decision maker draws a sequential random sample $X_1, X_2, \ldots$ from a distribution on $\mathbb{R}$. Let F denote the cumulative distribution function and suppose that it is known (i.e., assume that successive observations are independent and identically distributed) and that it has a finite mean. Assume also that the number of observations that may be taken is n, $2 \le n < \infty$, and that, until the permitted number of observations is exhausted, additional observations are free. Finally, suppose that the sampling process is without recall.

At stage j $(1 \le j < n)$ of the sampling process, the decision maker, having observed $X_1 = x_1, X_2 = x_2, \ldots, X_j = x_j$ may accept x_j and terminate the search or continue the search by drawing an additional observation. However, if he had he not stopped earlier, he must now stop and accept the final observation x_n.

Let $u(x_i)$ be the von Neumann-Morgenstern utility assigned to x_i. Under expected utility theory, the decision maker seeks a stopping rule (search strategy) maximizing $E\{u(x_N)\}$, where N denotes the random number of observations taken under a given stopping rule.

The optimal stopping rule for this problem is well known and we state it here without proof.[6] Let $v_1 \le v_2 \le v_3 \le \ldots \le v_n$ be a sequence of numbers such that v_j represents the expected utility from taking at least one additional observation and then proceeding optimally when there are j observations that may be taken before the sampling must end. The optimal stopping rule in this case is to continue the sampling if $x_j < v_{n-j}$, otherwise to accept x_j and to terminate the sampling. The expected utility from following this procedure is v_n. Furthermore, under the expected utility hypothesis the optimal stopping rule is dynamically consistent in the sense that the values $\{v_i\}_{i=1}^n$ remain unchanged throughout the process. Thus, the optimal continuation of the search strategy as of any given stage coincides with the continuation of the optimal stopping rule as specified at the outset. When the decision maker's preferences do not satisfy the independence axiom of expected utility theory, the optimal stopping rule is not dynamically consistent in the aforementioned sense. To illustrate this point and to set the stage for the analysis of behaviorally consistent optimal stopping rules we discuss a simple example.

(2.2) **Dynamic inconsistency: An example**

Consider the search model of the preceding subsection. Let $n = 2$ and suppose that F is the c.d.f. of the uniform distribution over the set $\{x_1, x_2, x_3\}$ of money prizes, where $x_1 < x_2 < x_3$. Suppose that the decision maker's preference relation on the space L of cumulative distribution functions on $\mathbb{R}$ is representable by a preference functional $V: L \to \mathbb{R}$ that satisfies first-order stochastic dominance. We describe the optimal stopping problem in terms of a decision tree (see figure 1 below) denoting decision nodes by □, chance nodes by ○ and the terminal nodes by •. Next to each terminal node we indicate the prize awarded if the process terminates at that node. Following chance nodes, the number along the branches indicate the probability of the event represented by the branch.

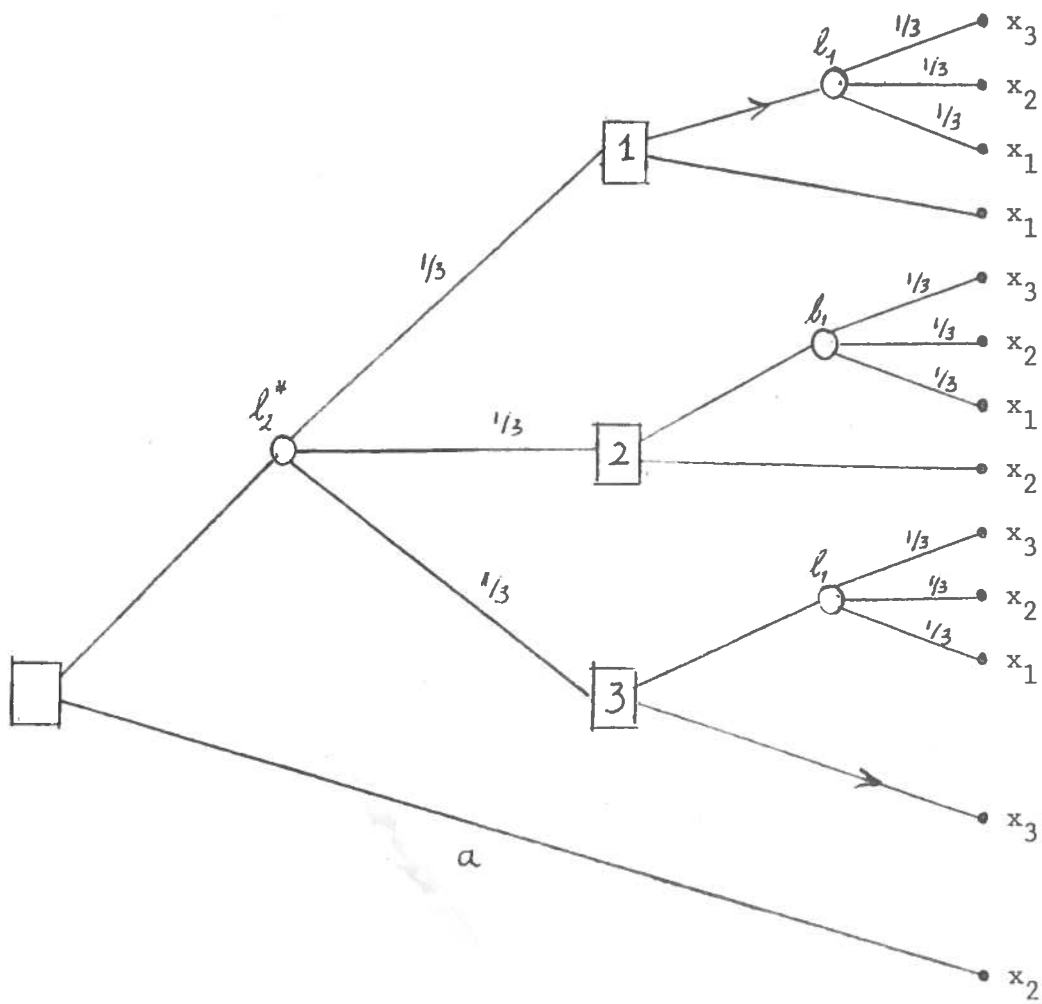

Figure 1

Figure 1 describes the situation in which the decision maker's first draw results in the observation $X_1 = x_2$. This places him at the initial decision node, facing the choice between accepting x_2 and continuing the search. Accepting x_2 corresponds to choosing the lower branch, marked by a, which leads to a terminal node and the reward x_2. Rejection of x_2 in favor of continuation of the sampling process corresponds to moving to the first chance node, ℓ_2^*.

Denote by δ_{x_i} the cumulative distribution function that assigns the entire probability mass to x_i (i.e., $\delta_{x_i}(x) = 0$ for $x < x_i$ and $\delta_{x_i}(x) = 1$ for $x \geq x_i$). Then, at the initial decision node the decision maker must choose between δ_{x_2} and ℓ_2^*, where ℓ_2^* is the lottery induced by the optimal continuation of the sampling process. By stochastic dominance the decision maker will choose ℓ_1 if in the final stage he is at decision node 1, and x_3 if he is at decision node 3. Suppose that $V(\frac{1}{3}\ell_1 + \frac{1}{3}\delta_{x_2} + \frac{1}{3}\delta_{x_3}) < V(\frac{2}{3}\ell_1 + \frac{1}{3}\delta_{x_3})$. Then $\ell_2^* = \frac{2}{3}\ \ell_1 + \frac{1}{3}\delta_{x_3}$. It seems, therefore, that the decision between stopping and accepting x_2 at the initial decision node and continuing the search depends on whether $V(\delta_{x_2})$ is larger or smaller than $V(\ell_2^*)$. Suppose that $V(\ell_2^*) > V(\delta_{x_2})$, then the decision maker will reject x_2 and continue the search. This indeed, is the only consideration when the decision maker is an expected utility maximizer. For in this case, $V(\frac{1}{3}\ell_1 + \frac{1}{3}\delta_{x_2} + \frac{1}{3}\delta_{x_3}) < V(\frac{2}{3}\ell_1 + \frac{1}{3}\delta_{x_3})$ if and only if $V(\ell_1) > V(\delta_{x_2})$. Consequently, at decision node 2 in the final stage the decision maker rejects x_2 in favor of one additional observation. In general, however, the first inequality may hold even though $V(\ell_1) < V(\delta_{x_2})$. In other words, without the additional structure imposed by the independence axiom of expected utility theory, the optimal decision at node 2 of the final stage may be to accept x_2 as the outcome. This, however, is at variance with the continuation of the optimal search strategy formulated at the initial decision node, thus suggesting that this decision rule is dynamically inconsistent.

Suppose that $V(\frac{1}{3}\ell_1 + \frac{1}{3}\ell_{x_2} + \frac{1}{3}\ell_{x_3}) < V(\delta_{x_2})$. If the decision maker recognizes that, should he find himself at node 2 of the final decision stage he will stop the search and accept x_2, he should then accept x_2 at the initial stage. On the other hand, if he believes that he will follow the strategy that induced ℓ_2^*, then he will reject x_2 at the initial stage. Both modes of behavior are possible. The second, however, is self-deceiving, because at the final stage, if he is at decision node 2, acting rationally, the decision maker will accept x_2. A decision maker who rejects x_2 at the initial stage on the presumption that he will follow ℓ_2^* and then, at decision node 2 in the final stage accepts x_2, is said to be **behaviorally inconsistent**. A decision-maker who chooses his initial move optimally, subject to the restriction that at the final stage he will follow the optimal course of action as of that stage, is said to follow a **behaviorally consistent** optimal stopping rule.

Clearly, the optimal stopping rules induced by the two behavioral modes are different. In particular, behaviorally inconsistent optimal stopping rules require the continuation of the sampling process whenever the behaviorally consistent rules do, but not vice versa.

(2.3) **Behaviorally consistent optimal stopping rules: A definition**

To formalize the notion of behaviorally consistent optimal stopping rules we represent the decision maker by a set of agents, one agent at each **stage** of the sampling process. We assume that each agent acts in his own self-interest. Finally, we formulate the game among the agents and define a behaviorally consistent optimal stopping rule to be a subgame perfect equilibrium of this game.

Let $\{1,...,n\}$ be the set of agents representing the decision maker at the n stages of the search process in the natural order. Each agent, in his turn, must choose between accepting the most recent observation or drawing another observation, thereby passing the decision on to the next agent. A **strategy** for agent j is a function s_j: $X \rightarrow [0,1]$, where $X = [\underline{x}, \bar{x}]$ is the set of prizes, and $s_j(x)$ is the probability that agent j accepts the outcome $x \in X$. For each agent, the set of strategies S is identical to the set of all the functions from X to $[0,1]$. Let S^j be the Cartesian product representing the set of strategies of agents $j, j+1,...,n$. Elements of S^j are denoted by s^j. Given $s \in S^1$, for each j let ℓ_j^s be the lottery on X induced by $s^j \in S^j$. Then the payoff of agent j, $j = 1,...,n$ is given by $V(\ell_j^s)$. Notice that the agents share the payoff function V and that this is common knowledge.

Definition 1: A behaviorally consistent optimal stopping rule, s, is an element of S^1 such that for all $j = 1,\ldots,n$, and for all $x \in X$

$$V(s_j(x)\delta_x + (1-s_j(x))\ell^s_{j+1}) \geq V(\alpha\delta_x + (1-\alpha)\ell^s_{j+1}) \ \forall\alpha \in [0,1].$$

(2.4) **Behaviorally consistent optimal stopping rules: Characterization**

The nature of the behaviorally consistent optimal stopping rules depends on the underlying preferences. To study this relationship we introduce the following notations and definitions. Let $G_1 \equiv F$ and, for all j, $N-1 > j \geq 0$,

$$G_{n-j}(\cdot;s^{j+1}) = \int_{-\infty}^{\infty} [s_{j+1}(x)\delta_x + (1-s_{j+1}(x))G_{n-j-1}(\cdot;s^{j+2})]dF(x).$$

Theorem 1: *If V is quasi-convex, then any behaviorally consistent optimal stopping rule $s \in S^1$ is characterized as follows: there exists an n tuple $(v_n,\ldots,v_1) \in X^n$ defined by*

$$V(\delta_{v_{n-j}}) = V(G_{n-j}(s^{j+1})), \quad j = 0,1,\ldots,n-1$$

such that for all $j = 1,2,\ldots,n-1$

$$s_j(x) = \begin{cases} 1 & \text{if } x > v_{n-j} \\ 0 & \text{if } n < v_{n-j} \end{cases}$$

and $s_j(x) \in [0,1]$ when $x = v_{n-j}$.

Proof: (by induction). Consider agent n-1 and let x be the most recent observation. Let $s_{n-1}(x)$ be the value of α that maximizes $V(\alpha\delta_x + (1-\alpha)G_1)$ (because at the last stage any observation must be accepted, continuation yields the distribution G_1). By quasi-convexity

$$s_{n-1}(x) = \begin{cases} 1 & \text{if } V(\delta_x) > V(G_1) = V(\delta_{v_1}) \\ \\ 0 & \text{if } V(\delta_x) < V(G_1) = V(\delta_{v_1}) \end{cases}$$

and $s_{n-1}(x) \in [0,1]$ when $V(\delta_x) = V(\delta_{v_1})$.

Suppose the conclusion is valid for $j+1$ and consider agent j. Given s^{j+1} and $x \in X$, let $s_j(x)$ be such that for all $\alpha \in [0,1]$,

$$V(s_j(x)\delta_x + (1 - s_j(x))G_{n-j}(s^{j+1})) \geq V(\alpha\delta_x + (1-\alpha)G_{n-j}(s^{j+1})).$$

By quasi-convexity,

$$s_j(x) = \begin{cases} 1 & \text{if } V(\delta_x) > V(G_{n-j}(s^{j+1})) = V(\delta_{v_{n-j}}) \\ \\ 0 & \text{if } V(\delta_x) < V(G_{n-j}(s^{j+1}) = V(\delta_{v_{n-j}}) \end{cases}$$

and $s_j(x) \in [0,1]$ if $V(\delta_x) = V(\delta_{v_{n-1}})$. The conclusion of the theorem follows from first-order stochastic dominance. □

To simplify the exposition if the preferences are strictly quasi-convex, we adopt the convention that $s_j(x) = 1$ whenever $x = v_j$.

Remark: A subset of quasi-convex preferences that have received special attention in the literature is the set of preferences satisfying a property called betweenness, i.e., for any F and H in L such that $V(F) > V(H)$, $V(F) > V(\alpha F + (1-\alpha)H) > V(H)$ for all $\alpha \in (0,1)$.[7] The characterization of behaviorally consistent optimal stopping values in this case is the same as in the case of strictly quasi-convex preferences with the exception that if $x = v_{n-j}$ for some j, $s_j(x)$ may take any value in the interval [0,1]. In this case we shall use the convention that $s_j(x)$ is chosen so as to maximize $V(G_{n-j+1}(s^j))$.[8]

Qualitatively speaking, behaviorally consistent optimal stopping rules in the case where V is quasi-convex are the same as the optimal stopping rules under expected utility theory. This is not the case, however, when V is strictly quasi-concave. In this case a behaviorally consistent optimal stopping rule $s \in S^1$ takes the following form: At each stage j of the process there are two values $\underline{v}_{n-j}$ and $\bar{v}_{n-j}$ such that

$$s_j(x) = \begin{cases} 1 & \text{if } x \geq \bar{v}_{n-j} \\ 0 & \text{if } x \leq \underline{v}_{n-j}. \end{cases}$$

and $s_j(x) \in [0,1]$ if $x \in [\underline{v}_{n-j}, \bar{v}_{n-j}]$. s_j is continuous and if V satisfies the following additional condition, then s_j is also monotonic on X.

Hypothesis II:[9] Let V be Frechet differentiable and denote by u_F its local utility function at F. Then for all F and G in L,

$$F >_1 G \implies -\frac{u_F''(x)}{u_F'(x)} \geq -\frac{u_G''(x)}{u_G'(x)} \quad \forall x \in X,$$

where $>_1$ denotes the relation of the first-order stochastic dominance.

Hypothesis II states that the Arrow-Pratt measure of absolute risk aversion of the local utility function increases with shifts towards stochastically dominating distributions.

Theorem 2: *Let V be a strictly quasi-concave utility functional.*

(a) *If $s \in S^1$ is a behaviorally consistent optimal stopping rule of V then for each $j = 1,\ldots,n-1$, s_j is a continuous function of X to [0,1] and there exists an interval $[\underline{v}_{n-j}, \bar{v}_{n-j}]$ such that $s_j(x) = 1$ for $x \geq \bar{v}_{n-j}$* and *$s_j(x) = 0$ for $x \leq \underline{v}_{n-j}$.*

(b) *If, in addition, V satisfies Hypothesis II, then s_j is a monotonic nondecreasing function.*

Proof: For each j the choice of $s_j(x)$ satisfies

$$V(s_j(x)\delta_x + (1-s_j(x))G_{n-j}(s^{j+1})) \geq V(\alpha\delta_x + (1-\alpha)G_{n-j}(s^{j+1})) \quad \forall\alpha \in [0,1].$$

Continuity of s_j follows from the differentiability and strict quasi-concavity of V. Clearly, Supp $G_{n-j}(s^{j+1}) \subset$ Supp F and, by stochastic dominance, $s_j(x) = 1$ if $x \geq z$ $\forall z \in$ Supp F and $s_j(x) = 0$ if $x \leq z$ $\forall z \in$ Supp F. Define

$$\bar{v}_{n-j} = \text{Min } \{x \mid s_j(x) = 1 \text{ for all } y \geq x\}$$

and

$$\underline{v}_{n-j} = \text{Max } \{x \mid s_j(x) = 0 \text{ for all } y \leq x\}.$$

This completes the proof of part (a).

Suppose that V satisfies Hypothesis II and let $\bar{\alpha} = s_j(x)$. If $\bar{\alpha} \in (0,1)$, then

$$(1) \quad 0 = \frac{d}{d\alpha} V(\alpha\delta_x + (1-\alpha)G_{n-j}(s^{j+1}))\Big|_{\bar{\alpha}} = \int U_A d(\delta_x - G_{n-j}(s^{i+1})),$$

where $A = \bar{\alpha}\delta_x + (1-\bar{\alpha})G_{n-j}(s^{j+1})$. Take $y > x$. Clearly $B \equiv \bar{\alpha}\delta_y + (1-\bar{\alpha})G_{n-j}(s^{j+1}) >_1 A$. By Hypothesis II this implies that the certainty equivalent of $G_{n-j}(s^{j+1})$ given U_B is smaller than x which by (1) is the certainty equivalent of $G_{n-j}(s^{j+1})$ given U_A. Thus,

$$\int U_B d(\delta_x - G_{n-j}(s^{j+1})) \geq 0.$$

But this implies

$$\frac{d}{d\alpha}V(\alpha\delta_y + (1-\alpha)G_{n-j}(s^{j+1}))\Big|_{\bar{\alpha}} = \int U_B d(\delta_y - G_{n-j}(s^{j+1}))$$

(2)

$$= \int U_B d(\delta_y - \delta_x) + \int U_B d(\delta_x - G_{n-j}(s^{j+1})) \geq 0,$$

where we use the fact that $\int U_B d(\delta_y - \delta_x) \geq 0$ by first-order stochastic dominance. Hence, the optimal value of α for y, i.e., $s_j(y)$, is larger than $\bar{\alpha}$.

If $\bar{\alpha} = 1$ then $\int U_A d(\delta_x - G_{n-j}(s^{j+1})) \geq 0$. Again, by Hypothesis II, $\int U_B d(\delta_x - G_{n-j}(s^{j+1})) \geq 0$ and, as before, $s^j(y) = 1$ for $y > x$.

If $\bar{\alpha} = 0$ then, by definition, $s_j(y) \geq 0$ for $y > x$.

□

3. Unbounded Sampling From a Known Distribution

(3.1) The specification of the model and the optimal stopping rule under expected value theory

Consider the problem of sampling from a known distribution when the number of observations that may be taken is unbounded, but there is a fixed positive cost, c, per observation. This problem has the property that, if the search strategy remains the same, the lottery induced by a continuation of the sampling as of any decision node is a leftward translation by the amount c of the lottery induced by the same decision at the preceding node. It is well known that if the objective of the search is the maximization of the expected monetary gain then, when the optimal stopping rule exists, it has the following simple characterization: At each stage of the sampling process the decision maker compares the value of the most recent observation to a number, v^*, representing the maximal expected gain among all stopping rules. If the value of the most recent observation exceeds v^*, then the decision maker terminates the search and accepts this observation, otherwise he continues the search.[10] A by-product of our analysis of general preference functionals is the establishment of the validity of this characterization of the optimal stopping rules for expected utility functionals. The economic interpretation of this model, however, is not obvious, as unbounded search with positive cost requires unbounded wealth. One may interpret the optimal stopping rule derived from this model as a convenient approximation of behavior in search situations when the number of observations is finite but large.

(3.2) A preliminary observation concerning the nature of the behaviorally consistent optimal stopping rules

Because the number of observations that may be taken is unbounded, the analysis of behaviorally consistent optimal stopping rules cannot be based on backward induction. The key fact about the problem at hand, however, is the recurrence of the same choice problem at each stage of the sampling process. This implies that **if the behavioral mode is the same throughout the process then, independent of the exact nature of the behavioral mode, the optimal decision rule must be the same at each decision node.** In other words, whether we assume that the decision maker is behaviorally consistent or not, a decision to continue the sampling at each stage places him in essentially the same position that he was in at the outset of the process. The fact that by that time he incurred a cost is irrelevant, as this cost is sunk.

To reveal the consequences of this observation we introduce the following notation. Denote by S^{ω} the infinite Cartesian product of S. Let $(s_j, s^{j+1}) \in S \times S^{\omega}$ be a continuation strategy as of stage j. We denote by $\hat{s} = (\hat{s}_1, \hat{s}_2, \ldots)$ the element of S^{ω} that consists of the first entry of (s_1, s^2), $(s_2, s^3), \ldots, (s_j, s^{j+1}), \ldots$. Defined in this manner, $\hat{s}$ is the **realized stopping rule**. An element $s \in S^{\omega}$ such that $s_j = s_1$, $j = 1,2,\ldots$ is a **constant stopping rule**. Finally, we use the term **behavioral mode** to describe the way a decision maker perceives the stopping rule as of any given stage in the sampling process. Thus, behavioral consistency is a behavioral mode according to which the perceived optimal stopping rule as of any given decision is the realized stopping rule as of that node. Formally, a behavioral mode at a decision node in stage j is an element of S^{ω} representing the **assumption** of the agent at stage j about the strategies that will be followed by all the agents starting from agent j+1. A **constant behavioral mode** is a behavioral mode such that $s^j = s^{j+1}$ for all j.

Proposition: *If a decision maker has a constant behavioral mode, then the realized optimal stopping rule of an unbounded sampling process from a known distribution with constant positive cost per observation is a constant stopping rule.*

Proof: Consider the situation facing a decision maker with constant behavioral mode at stage j of the sampling process. Let s_j be his optimal strategy at stage j given the behavioral mode s^{j+1}. Then, at stage j+1, the decision maker is in precisely the same position that he was in at the preceding stage because the behavioral mode is constant, i.e., $s^{j+1} = s^{j+2}$. This implies that the optimal strategy s_{j+1} can be chosen so that $s_{j+1} = s_j$. Consequently, the realized stopping rule $(s_1, s_2,...)$ is constant.

□

A behaviorally consistent decision maker recognizes that the realized stopping rule is constant. Consequently, by definition, we have the following:

Corollary: *Under the hypothesis of the proposition a behaviorally consistent optimal stopping rule is the optimal stopping rule among all the constant stopping rules.*

(3.3) **Characterization of behaviorally consistent optimal stopping rules**

Consider a decision maker who, at stage j of the sampling process and having observed x, must choose between accepting x and continuing the sampling. Let G denote the lottery induced by the continuation of a behaviorally consistent optimal stopping rule. Because this rule is constant, G is independent of j.

If the decision maker's utility functional is quasi-convex then, given the conditions of Theorem 1, his optimal stopping rule requires that at stage j he rejects x if $x < v_{n-j}$ (where v_j is the certainty equivalent of the lottery induced by optimal continuation strategy), otherwise he accepts x. In the present case the lottery induced by the optimal continuation strategy, G, is constant. Therefore, behaviorally consistent optimal stopping rules, s, have a simple form, that is, there exists v^* defined by the equation $V(\delta_{v^*}) = V(G)$ such that for all $j = 1,2,\ldots$

$$s_j(x) = \begin{cases} 1 & \text{if } x \geq v^* \\ 0 & \text{if } x < v^*. \end{cases}$$ [11]

By a similar argument and using Theorem 2 it is easy to see that if the decision maker's preference relation on L is strictly quasi-concave, then the characterization of behaviorally consistent optimal stopping rules, s, is as follows: For each $j = 1,2,\ldots$ there exist numbers $\underline{v}^*, \bar{v}^* \in X$, $\underline{v}^* \leq \bar{v}^*$ such that $s_j(x) = 1$ for $x \geq \bar{v}^*$, $s_j(x) = 0$ for $x \leq \underline{v}^*$ and $s_j(x) \in [0,1]$ for $x \in (\underline{v}^*, \bar{v}^*)$. Furthermore, if V satisfies Hypothesis II, then s_j is monotonic non-decreasing.

Finally, notice that the characterization of the optimal stopping rule for quasi-convex preferences include preference relations that are representable by expected utility theory as a special case.

(3.4) **Existence of behaviorally consistent optimal stopping rules**

We turn now to the issue of the existence of behaviorally consistent optimal stopping rules. In particular, we show that if the observations are taken from a continuous distribution on a finite interval, then behaviorally consistent optimal stopping rules exist for any utility functional.[12]

Theorem 3: *Let* V *be a continuous utility functional and assume that* F *is continuous with Supp* $F = [\underline{x}, \bar{x}]$. *If* c *is small enough, then there exists a behaviorally consistent optimal stopping rule* $s \in S^{\omega}$ *such that*

$$s_j(x) = s_{j+1}(x) \quad \text{for all} \quad j = 1,2,\ldots .$$

Proof: Let $(G+c)$ be a distribution defined by $(G+c)(x) \in G(x-c)$, i.e., $G+c$ is a translation of G to the right by the amount c. Consider a decision maker at stage j of the sampling process. Given the continuation strategy s^{j+1} a decision to continue the search induces the lottery $G(s^{j+1}) \in L$. By the Proposition, $s^{j+1} = s^j$ for all j, therefore we can write that $G(s^j) = G$ for $j = 1,2,\ldots$. Hence, playing the strategy s_j at j yields the distribution

$$\int_{\underline{x}}^{\bar{x}} [(1-s_j(x))G + s_j(x)\delta_x] dF(x).$$

However, this is the lottery induced by a decision to continue at stage j-1 not taking into account the cost involved in obtaining an additional observation. Thus, G is defined by the equation

$$(3) \qquad (G+c) = \int_{\underline{x}}^{\bar{x}} [(1-s_j(x))G + s_j(x)\delta_x]dF(x).$$

(a) Suppose that V is quasi-convex. We have already established (see the preceding subsection) that s_j is characterized as follows: There exists $v \in [\underline{x}, \bar{x}]$ and $s_j(x) = 1$ for $x \geq v$, $s_j(x) = 0$ otherwise. Hence, given v there exists a unique distribution $G = G(v)$ that solves equation (3), or equivalently, given the nature of s in this case,

$$(4) \qquad G(x-c) = F(v)G(x) + \max\{F(x) - F(v), 0\}.$$

However, given G, a decision maker acting optimally at stage j will continue the sampling if and only if the most recent observation x is such that $V(\delta_x) < V(G)$. Let v′ be defined by $V(G) = V(\delta_{v'})$. Then, by stochastic domnance $v' \leq \bar{x}$.

Since F is continuous, then by equation (4), G is pointwise continuous in v. By continuity of V, v′ is continuous in G. Let $h = [\underline{x}, \bar{x}] \rightarrow (-\infty, \bar{x}]$ be the composition of these continuous functions.

If $v = \bar{x}$ then, by equation (4), $G(x-c) = G(x)$. Since $G(\bar{x}) = 1$ it follows that G has a unit mass at $-\infty$. Thus, h(v) is infinitely small and

$h(\bar{x}) < \bar{x}$. If $v = \underline{x}$ then, by equation (4), $G(x-c) = F(x)$ or $G = F - c$. If c is sufficiently small so that a $V(F-c) = V(\delta_{\underline{x}})$, then $h(\underline{x}) > \underline{x}$. By continuity of h there exists a $v^* \in [\underline{x}, \bar{x}]$ such that $h(v^*) = v^*$. The strategy

$$s_j(x) = \begin{cases} 1 & \text{if } x \geq v^* \\ 0 & \text{otherise} \end{cases}$$

is behaviorally consistent optimal stopping rule for this case.

(b) Suppose that V is strictly quasi-concave. Let $Q = \{s:[\underline{x}, \bar{x}] \rightarrow [0,1]\}$ be endowed with the product topology. Then, by compactness of $[0,1]$ and the Tychonoff Theorem, Q is compact.[13] We show next that Q is a metric space.

Q is Hausdorff. To see this, let s and s' be two distinct points in Q. Thus, $\exists\, x_0 \in [\underline{x}, \bar{x}]$ such that $s(x_0) \neq s'(x_0)$. But $[0,1]$ is a metric space, therefore it is Hausdorff. Hence, there exist disjoint open sets U and U' in $[0,1]$ such that $s(x_0) \in U$, $s'(x_0) \in U'$. Let

$$P(U,x_0) = \{h \mid h \in Q,\ h(x_0) \in U\}$$
$$P(U',x_0) = \{h \mid h \in Q,\ h(x_0) \in U'\}.$$

Then, $s \in P(U,x_0)$ $s' \in P(U',x_0)$ and $P(U,x_0) \cap P(U', x_0) = \phi$, which implies that Q is Hausdorff. But every compact Hausdorff space is normal. Hence Q is normal. By the Urysohn Metrization Theorem, Q is metrizable. Clearly, Q is convex.

Next, we define a continuous function $h: Q \to Q$. From equation (3) we obtain

$$(5) \qquad G(y) = G(y+c)\left[1 - \int_{\underline{x}}^{\bar{x}} s_j(x)dF(x)\right] + \int_{y+c}^{\bar{x}} s_j(x)dF(x).$$

Since $G(\bar{x}) = 1$ we have $G(y) = 1$ for $y \in [\bar{x}-c, \bar{x}]$, and clearly G is uniquely defined. Furthermore, if $s_j^n \xrightarrow[n\to\infty]{} s_j$ pointwise, then $G(s_j^n) \xrightarrow[n\to\infty]{} G(s_j)$ at every $y \in [\underline{x}, \bar{x}]$, and the function $G: Q \to L$ is continuous.

Next, for every $H \in L$ let $\alpha(H) \in Q$ be such that for any given $x \in [\underline{x}, \bar{x}]$

$$V(\alpha(H)(x)\delta_x + (1 - \alpha(H)(x))H) \geq V(\alpha(x)\delta_x + (1-\alpha(x))H)$$

for all $\alpha(x) \in [0,1]$. By continuity of V, α is continuous on L. Now we define $h = \alpha \circ G$. Hence, $h: Q \to Q$ is continuous. By a theorem of Schauder (see Smart [1974], p.15), it has a fixed point. This fixed point is the fixed strategy that defines the behaviorally consistant optimal stopping rule.

□

4. Concluding Remarks

Our analysis of behaviorally consistant optimal stopping rules is intended to deal with conceptual issues arising when nonlinear preferences are admitted. It was not our intention to offer an exhaustive study of sequential sampling models. We took two models as a vehicle to present the main principles involved. For the same reason we did not present a comparative statics analysis of these models, even though such an analysis may be of interest in its own right.

It is interesting to note a certain analogy between the characterization of the optimal stopping rules obtained here and the equilibrium strategies of behaviorally consistant bidders in ascending bid auctions. In particular, our analysis of ascending bid auctions with behaviorally consistant bidders (Karni and Safra [1988]) concluded that: (a) A Bayesian-Nash equilibrium exists if preferences are quasi-concave but not necessarily strictly quasi-concave; (b) when a bidder's preference relation on L satisfies the betweenness assumption (i.e., they are both quasi-concave and quasi-convex), then there exists a behaviorally consistent dominant strategy. Furthermore, this strategy requires that the bidder stays in the auction as long as the price falls short of his value of the auctioned object and that he withdraws from the auction as soon as the price exceeds this value; (c) when the bidder's preferences on L are strictly quasi-concave, then his equilibrium strategy is characterized as follows: There exist values, say $\underline{v}$ and $\bar{v}$, $\underline{v} < \bar{v}$; the bidder stays at the auction as long as the price is lower than $\underline{v}$ and he is never in the auction when the price exceeds $\bar{v}$. In between he plays a mixed strategy, i.e., stays at the auction with positive probability.

One cannot fail to see the analogy between these behaviorally consistent equilibrium bidding strategies and the behaviorally consistent optimal stopping rules.

FOOTNOTES

1 See DeGroot (1970) ch.13 for a mathematical analysis and Lippman and McCall (1976) for a survey of the literature emphasizing the economic issues.

2. Karni and Safra (1987a) show, in the context of ascending bid auctions, that the linearity of the expected utility functional is a necessary and sufficient condition for dynamic consistency. This conclusion applies here as well.

3. For a survey of the evidence see MacCrimmon and Larsson (1979) and Machina (1982). The new theories of decision making under risk include: Kahnman and Tversky (1979), Machina (1982), Quiggin (1982), Fishburn (1983), Chew and MacCrimmon (1979), Chew (1983), Yaari (1987), and Dekel (1986).

4. The idea of behavioral consistency was suggested by Strotz (1956) as a way of solving the problem of dynamic inconsistency that may be encountered in optimal consumption plans. The issue of the existence of behaviorally consistent optimal consumption plans was discussed in Peleg and Yaari (1973) and more recently in Goldman (1980). In Karni and Safra (1987b), (1988) we analyzed the equilibrium of ascending bid auction games with behaviorally consistent bidders.

5. A formal definition is given in the sequel.

6. For a detailed analysis see DeGroot (1970), pp.331-333.

7. See Fishburn (1983), Chew and MacCrimmon (1979), Chew (1983), and Dekel (1986).

8. This convention captures the notion that, when indifferent between alternative courses of action, an agent acts so as to improve the position of the other agents in the set.

9. This condition is due to Machina (1982), whose terminology is used here. Machina showed that this condition is consistent with Allais-type experimental evidence, as well as other violations of expected utility theory. **Frechet differentiability** of $V: L \rightarrow \mathbb{R}$ means that at each $F \in L$ we have

$$V(G) - V(F) = \int u_F d(G-F) + o(\|G-F\|).$$

The function $u_F:\mathbb{R}\rightarrow\mathbb{R}$ is called the **local utility** function of V at F.

10. For details see DeGroot [1970], pp.333-34. Notice that the value of continuation in this model remains unchanged so that if an observation is rejected then the option of recalling it, even if it exists, is never utilized.

11. The lottery G is defined by

$$G(x-c;v^*) = F(v^*)G(x;v^*) + \max\{F(x)-F(v^*);0\}.$$

For further details see the proof of Theorem 3 in the text.

12. Our analysis generalizes certain aspects of the results reported in DeGroot (1970). In particular, our analysis imposes less severe restrictions on the structure of the decision maker's preferences. In DeGroot, existence was merely shown for utility functionals that are linear in the probabilities and the prizes.

13. Convergence according to this topology is pointwise.

REFERENCES

Chew, S.H. (1983). "A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox," **Econometrica** 51, 1061-92.

Chew, S.H. and K.MacCrimmon (1979). "Alpha-Nu Choice Theory: A Generalization of Expected Utility Theory," unpublished manuscript, University of British Columbia.

DeGroot, H. (1970). **Optimal Statistical Decision.** New York: McGraw Hill.

Dekel, E. (1986). "An Axiomatic Characterization of Preference under Uncertainty: Weakening the Independence Axiom," **Journal of Economic Theory** 40, 304-18.

Fishburn, P.S. (1983). "Transitive Measurable Utility," **Journal of Economic Theory** 31, 298-317.

Goldman, S.M. (1980). "Consistent Plans", **Review of Economic Studies**, 47, 533-37.

Kahneman, D., and A.Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk," **Econometrica** 47, 263-91.

Karni E., and Z.Safra (1987a). "Dynamic Consistency and Revelations in Auctions and the Structure of Preferences". Working Paper No.87-W14 The Johns Hopkins University.

——————————— (1987b). "Ascending Bid Auction Games: A Nonexpected Utility Analysis," Working Paper No.20-87, The Foerder Institute for Economic Research, Tel-Aviv University.

Karni, E., and Z.Safra (1988). "Ascending Bid Auctions with Behaviorally Consistent Bidders". Working Paper No.1-88, The Foerder Institute for Economic Research, Tel-Aviv University.

Lippman, S.A. and J.J.McCall (1976). "The Economics of Job Search: A Survey: **Economic Inquiry**, 14, 155-89, 347-68.

McCrimmon, K.R. and S.Larsson (1979). "Utility Theory: Axioms versus Paradoxes" in **Expected Utility Hypotheses and the Allais Paradox**, M.Allen and O.Hagen (eds.) Dordbrecht, Holland, D.Reidel Publishing Co.

Machina, M.J. (1982). "'Expected Utility' Analysis without the Independence Axiom," **Econometrica** 50, 277-323.

Peleg, B., and M.Yaari (1973). "On the Existence of a Consistent Course of Action when Tastes are Changing," **Review of Economic Studies** 40, 391-401.

Quiggin, J. (1982). "A Theory of Anticipated Utility," **Journal of Economic Behavior and Organization** 3, 324-43.

Smart, D.R. (1974). **Fixed Point Theorems**, Cambridge University Press.

Strotz, R.H. (1956). "Myopia and Inconsistency in Dynamic Utility Maximization," **Review of Economic Studies** 23, 165-80.

Yaari, M.E. (1987). "The Dual Theory of Choice under Risk," **Econometrica** 55, 95-115.