COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 807

KNIGHTIAN DECISION THEORY: PART I

Truman F. Bewley

November 1986

KNIGHTIAN DECISION THEORY: PART I

by

Truman F. Bewley*

Cowles Foundation
Yale University

Abstract

A theory of choice under uncertainty is proposed which removes the completeness assumption from the Anscombe-Aumann formulation of Savage's theory and introduces an inertia assumption. The inertia assumption is that there is such a thing as the <u>status quo</u> and an alternative is accepted only if it is preferred to the <u>status quo</u>. This theory is one way of giving rigorous expression to Frank Knight's distinction between risk and uncertainty.

---

*The author is grateful to Martin Shubik for encouragement and long discussions. He is also grateful to Donald Brown, Edward Leamer, David Schmeidler, Martin Shubik, Sidney Winter, Asad Zaman and Arnold Zellner for calling his attention to relevant literature. This work has been supported by National Foundation Grants SES-8410219 and SES-8605046.

INTRODUCTION

Many years ago, Frank Knight (1921) made a distinction between risk and uncertainty. A random variable is risky if its probability distribution is known, uncertain if its distribution is unknown. He argued that uncertainty in this sense is very common in economic life and he based a theory of profit and entrepreneurship on the idea that the function of the entrepreneur is to undertake investments with uncertain outcome.

From the point of view of Bayesian decision theory, Knight's distinction has no interesting consequences. According to the Bayesian theory decision makers act so as to maximize the expected value of their gain, no matter whether the fluctuations faced are risky or uncertain.

However, Knight's idea does seem to have some intuitive appeal. Ellsberg's (1961, 1963), experiments seem to show that people are repelled by vagueness of probabilities. Bayesian decision theory also has the following disturbing implication. Suppose two decision makers are faced with the same decision problem with the same objectives, constraints and information, but with uncertain outcome. Suppose the objective function is strictly concave and the constraint set convex, so that the prior distribution of a Bayesian maximizer determines a unique decision. Then, if the decision makers do not choose the same decision, one must conclude that they have different prior distributions and so would be willing to make bets with each other about the outcomes. These conclusions strike me as questionable. Betting outside of gambling casinos and race tracks is uncommon, but disagreement over cooperative decisions seems to be part of every-day life. One may try to explain the lack of betting by mutual suspicion that the other decision maker has secretly acquired superior information. I find this argument hard to reconcile with the observation that people usually seem very fond of their own decisions. Knight's ideas suggest another way to explain the absence of bets. In the presence of uncertainty, decisions may not be determinate, and bets may be shunned unless they are very favorable.

In this paper, I propose a rigorous formulation of Knight's somewhat vague ideas. The basic idea is to drop the completeness assumption from Savage's (1954) framework and to add an assumption of inertia. (In fact, I drop completeness from the reformulation of Savage's theory made by Anscombe and Aumann (1963).) The completeness assumption asserts that any two lotteries are comparable in the preference ordering; that is, one is preferred to the other or they are equivalent in the ordering. When this axiom is dropped, one obtains a set of subjective probability distributions rather than a single one. One lottery is preferred to another if its expected value is higher according to all the distributions. The idea of the inertia assumption is that a person never accepts a lottery unless he prefers acceptance to rejection. There is a <u>status quo</u> with which he stays unless an alternative is preferred. Without the inertia assumption, choices between all incomparable alternatives would be arbitrary. I apply the adjective "Knightian" to behavior consistent with the assumptions just described.

Knightian behavior seems to correspond to Knight's intuition about investor behavior. If an individual found a new investment opportunity uncertain and hard to evaluate, he would be unlikely to undertake it, for he would do so only if it had positive expected value for each of a large set of probability distributions. There is a form of aversion to uncertainty which is distinct from the usual risk aversion.

Also, Knightian behavior does not lead to the paradoxical high propensity to bet mentioned earlier. Two Knightian decision makers in disagreement would not be willing to bet with each other on some event unless the minimum probability one decision maker assigned to the event exceeded the maximum probability assigned by the other.

It is important that in Knightian decision theory one cannot predict decisions from knowledge of preferences. The theory can say only which decisions would be undominated by others and would be preferred to the *status quo*. The theory would not be contradicted if a decision maker cannot be persuaded to move from one undominated decision to another, as I shall explain presently.

The indeterminateness of decision may be viewed as a defect of the Knightian theory, since theories should explain as much as possible. But I suspect that indeterminateness, as well as uncertainty aversion and inertia, may turn out to be real and important. If this is so, they will have to be accommodated, perhaps, of course, by a better theory than the one proposed here.

The inertia assumption prevents a Knightian decision maker from making intransitive choices, provided the choices are between lotteries that are to be carried out. Choices between hypothetical choices could, however, be highly intransitive.

What I call the Knightian theory is not original. There exists a sizeable literature on the subject, including papers by Aumann (1962), Smith (1961), Walley (1981, 1982), and Williams (1974). This literature is reviewed in Section 7. What is new in this paper, I believe, is the emphasis on the inertia assumption. This assumption strikes me as crucial, yet I have nowhere found it made explicit or defended.

It is perhaps unfair of me to apply Frank Knight's name to the theory described in this paper. It is not entirely clear what he had in mind, and there is an attractive alternative class of preferences which display uncertainty aversion. These are complete preferences represented by a utility function of the form $u(x) = \min_{\pi \in \Delta} E_\pi x$, where x is a random variable, Δ is a set of probability distributions and E_π is the expectation with respect to π. Such preferences have been characterized axiomatically by Gilboa and Schmeidler (1986). Hereafter, I will refer to them as Gilboa-Schmeidler preferences. Such preferences are probably more convenient than the Knightian theory of this paper for the foundation of statistics, as I explain in Section 7. However, Gilboa-Schmeidler preferences do not lead to the sorts of economic behavior which make Knightian behavior interesting. For instance, unless an individual with Gilboa-Schmeidler preferences has the same utility in all states of the world he may behave much like an individual with preferences obeying the expected utility hypothesis of Savage. The distinction between Gilboa-Schmeidler and Knightian behavior is discussed in Section 5.

I propose the Knightian model of behavior because it helps rationalize many economic phenomena which otherwise seem difficult to explain. Some of these are discussed in the conclusion. The ability of the Knightian model to explain economic phenomena does not make it true in a descriptive sense. Only careful empirical work can establish whether the predictions of the Knightian theory occur with sufficient regularity that the theory may serve as a sound basis for economic analysis. I discuss experimental work and possible experiments in Section 6.

For good reasons, economists tend not to view empirical evidence as sufficient reason for accepting models of individual behavior. Economists want their models also to represent individuals as rational beings. I have viewed my main task in writing this paper to be to convince readers that Knightian behavior is rational, just as rational as behavior generated by the expected utility hypothesis or by Gilboa-Schmeidler preferences.

A person is defined to be rational, I believe, if he does the best he can, using reason and all available information, to further his own interests and values. I argue that Knightian behavior is rational in this sense. However, rationality is often used loosely in another sense, which is that all behavior is rationalizable as serving to maximize some preference. The two senses of rational are in a way converse. The first says that when preference exists, behavior serves it. The second says that all behavior is generated by preferences. The second sense seems to be very unlikely to be true, except by definition. It does not even seem to be useful as a definition. If choice is made the definition of preference, then one is led, like a true sophist, to the conclusion that people always do what they want to do, even when compelled to do things by threats of violence. The first sense of rationality is the one which is important for economic theory, at least as it is presently formulated. One would like to believe that people usually act so as to serve their own economic interests, at least when these interests are clear and do not conflict with other interests. If one identifies the two converse senses of rationality, one needlessly jeopardizes the first sense, since the second sense is probably more likely to be rejected than the first.

Associated with each definition of rationality is a different point of view toward incomplete preference. The view associated with the first definition of rationality is that the preference ordering is a constituent of a model which explains some but not all behavior. Behavior never contradicts the ordering, but not all choices are explained by it, nor are all stated or felt preferences. The model is not contradicted if an individual expresses strong preferences between alternatives which he finds incomparable according to the model. Such unexplained preferences or choices may be erratic and intransitive, but this is no cause for concern. Such behavior does not make the individual irrational, since the intransitive choices are not assumed to be in pursuit of some goal. The individual becomes irrational only if one tries to infer some unchanging goals from his choices or statements. It is because I adopt the point of view just stated that I said earlier that the Knightian theory is not contradicted if an individual shows a preference for one undominated choice over another.

I now turn to the view of incomplete preference associated with the second definition of rationality. This view accepts all stated or revealed preference at face value, but adds a category of incomparability to the categories of indifference and strict preference. That is, an individual may assert that two alternatives are incomparable. Choice behavior cannot distinguish indifference from incomparability. In fact, if one thinks about choice behavior one can quickly convince oneself that incomparability is an empty category. (If an individual chooses x over y, he either will or will not accept a small bribe to reverse his choices.) It is for this reason, I believe, that incompleteness is often referred to in the literature as intransitivity of indifference. A disadvantage of the second view of incompleteness is that it makes all individuals rational by definition.

The obvious way to escape from this tautology is to impose structure on preferences, such as transitivity and monotonicity. But a strong model of this sort is too frequently contradicted by reality, I believe. Is not everyday life full of inconsistent choice and unresolved goal conflicts? One could assert that only economic decisions are assumed to be rational, but this assertion can be justified only by the first definition of rationality.

One could also argue that the concept of preference is operational only if it is identified with choice. However, this is not so. The Knightian theory makes fairly obvious testable predictions. These stem largely from the inertia assumption.

The central problem of this paper is to make the inertia assumption precise and to defend its rationality. The intuitive idea of the assumption is that if a new alternative arises, an individual makes use of it only if doing so would put him in a preferred position. "New" means previously unavailable, and rejection of a new alternative means carrying on with previous plans.

It is not immediately clear how to make sense of this idea. If one adopts the usual point of view of decision theory, one assumes that a decision maker chooses at the beginning of his life an undominated program for his entire lifetime decision tree. What is the <u>status quo</u> or initial position in such a decision tree? If one is defined, why should the decision maker choose only programs preferred to it? It would be equally rational to choose a program incomparable to the initial position.

The answer to the second question is that inertia is not a consequence of rationality. Inertia is an extra assumption which is consistent with rationality.

I present three different versions of inertia. The first, given in Section 2, defines new alternatives to be ones to whose appearance the decision maker had previously assigned probability zero. The initial position is the position planned before the appearance of the new alternatives. The inertia assumption applies to the decision maker's way of reacting to new alternatives when they arrive.

The second approach to inertia defines inertia as a property of the undominated program chosen at the beginning of life. It is assumed that certain choices appearing after the initial period can be identified in a natural way as new, even though they are anticipated. An independence assumption guarantees that choices among new alternatives do not interact with other choices. The inertia assumption is that the chosen program makes use of new alternatives only if any program not doing so would be dominated. It is proved that there exists an undominated program satisfying the inertia assumption. In this sense, inertia is rational. This approach to inertia is presented in Section 3.

Section 4 contains the third approach to inertia. This approach makes a slight concession to bounded rationality in that it recognizes that a decision maker cannot possibly formulate a lifetime plan covering all contingencies. The disadvantage of bounded rationality is that it makes the concept of rational behavior very ambiguous. The best one can do is to imagine what a sensible, self-interested and boundedly rational person might do. I simply tell a plausible story in which inertia may be identified. I assume that the decision maker continually makes approximate plans. The inertia assumption is that these plans are abandoned only if doing so is judged necessary for an improvement. This picture of reality motivates a loose definition of inertia given in Section 4. it is probably the loose definition which should be used when applying the theory.

1. STRUCTURAL THEOREMS

I here present representation theorems for incomplete preferences over gambles. This material is not original. The proofs are contained essentially in Aumann (1962, 1964), Smith (1961), and Walley (1981). In spite of all this previous work, I present my own structural theorems. None of the sources presents the material in a way that is really suitable for my purposes. The presentation here is such that the theory of von Neumann and Morgenstern remains unchanged if the probabilities are known objectively. Incompleteness applies only to gambles over events of unknown probability.

I follows the Anscombe-Aumann (1963) formulation of choice under uncertainty, I have also been much influenced in choosing assumptions by recent papers of Myerson (1979, 1986). I retain essentially all of their assumptions except completeness.

I retain these assumptions not because I believe them but because they do have some normative justification. Experimental evidence seems to show that the assumptions do not describe behavior accurately. However, if lotteries are repeated often and the probabilities are known, their expected values do approximate actual outcomes. In this limited context, at least, it would be foolish to violate the assumptions.

If probabilities are not known, there seems to be no normative justification for completeness. The usual argument against incompleteness is that haphazard choice among incomparable alternatives can be intransitive and so lead to exploitation by a money pump. But any of the inertia assumptions prevents such exploitation, as is explained in Section 5. (A money pump occurs if a person chooses B minus a little money over A, C over B and A over C. If this cycle were repeated, the person would lose a little on each round.)

Incompleteness itself might be thought irrational. But from the point of view of the first definition of rationality mentioned in the Introduction, incompleteness simply limits the criterion for defining rationality. If a rational person is one who acts so as to achieve his objectives, a person without objectives is both rational and irrational, just as any statement is true of an empty set.[1]

I now turn to the representation theorems. In order to make clear the structure of the theory, I first of all present the case in which utility is linear in rewards or payoffs. The utility should be thought of as von Neumann-Morgenstern utility.

The mathematical notation is as follows. S is a finite set of states of nature. If $B \subset S$, R^B is the set of real-valued functions on B. Identify R^B with the obvious subspace of R^S. The function $\Pi^B : R^S \to R^B$ is the natural projection. The symbol e_B denotes the indicator function of B. That is, $e_B(s) = 1$, if $s \in B$, and $e_B(s) = 0$ otherwise. If $x \in R^S$, x_s is the reward or utility in state s. If x and y belong to R^S, $x > y$ means $x_n \geq y_n$, for all n, and $x \neq y$. If π is a probability on S, if $B \subset S$ and $x \in R^S$, then $E_\pi[x|B]$ denotes the expectation of x with respect to π, conditional on B.

---

[1] Aumann (1962, 1964a) has criticized the completeness assumption both as normatively unsound and descriptively inaccurate.

The preference ordering consists of an ordering $\succ_B$ on R^S, for each non-empty subset B of S. The expression $x \succ_B y$ means x is preferred to y if B is known to be true. I write $\succ$ for $\succ_S$. No relation of indifference is assumed. However, x and y are said to be <u>equivalent</u> if for all z and B, $x \succ_B z$ if and only if $y \succ_B z$ and $z \succ_B x$ if and only if $z \succ_B y$.

$\mathcal{X}$ denotes a field of subsets of S of objectively known probability, $q : \mathcal{X} \to [0,1]$ is the objective probability. q is objective in the sense that all observers would be conscious of q and agree to it.

The first assumption expresses the natural relation between the conditional and unconditional preference orderings.

<u>Assumption 1.1</u>. For all non-empty subsets, B, of S and all x and y in R^S, $x \succ_B y$ if and only if $\Pi^B x \succ \Pi^B y$.

The next assumption says simply that more utility is better.

<u>Assumption 1.2</u>. $x > Ey$ implies $x \succ y$.

The following assumption says that $\succ$ is an ordering of strict preference.

<u>Assumption 1.3</u>. $x \succ y \succ z$ implies $x \succ z$ and for no x is $x \succ x$.

The next assumption is of only technical significance.

<u>Assumption 1.4</u>. For all x, $\{y | y \succ x\}$ is open in R^S.

The key structural assumption is the following one. It has an obvious interpretation if one thinks of $\alpha x + (1-\alpha)y$ as the lottery giving x with probability α and y with probability $1-\alpha$.

Assumption 1.5. For all x , y and z , and for all $\alpha \in (0,1)$, $y \succ z$ if and only if $\alpha x + (1-\alpha)y \succ \alpha x + (1-\alpha)z$.

The last assumption asserts that the known probabilities of events in $\mathcal{A}$ are treated as they should be.

Assumption 1.6. For all $A \in \mathcal{A}$, e_A is equivalent to the lottery $q(A)e_S$.

Theorem 1.1. If the $\succ_B$ satisfy assumptions 1.1-1.6, then there is a closed convex set Δ of probabilities on S such that

(i) for all x and y and B , $x \succ_B y$ if and only if $E_\pi[x|B] > E_\pi[y|B]$, for all $\pi \in \Delta$,

(ii) for all $A \in \mathcal{A}$, $\pi(A) = q(A)$, and

(iii) for all $\pi \in \Delta$, $\pi(B) > 0$, for all non-empty subsets, B , of S .

The proof of this theorem appears in the appendix.

In the light of Theorem 1.1, one may define an indifference relation on R^S by $x \sim_B y$ if and only if $E_\pi[x|B] = E_\pi[y|B]$, for all $\pi \in \Delta$. Clearly, if $x \succsim_B y \succsim_B z$, then $x \succsim_B z$. However, the statement " $x \not\succsim_B y$ " does not imply " $x \precsim_B y$." The ordering $\succsim$ is complete on the set of x in R^S such that x is measurable with respect to $\mathcal{A}$.

One may define $\succ$ to be complete if for all $x \in R^S$, the closure of $\{y \in R^S | y \succ x \text{ or } x \succ y\}$ is all of R^S . Clearly, $\succ$ is complete if and only if Δ consists of a singleton and so the expected utility hypothesis applies.

I now turn to the case in which preferences are for lotteries over consequences, so that one must infer the existence of a von Neumann-Morgenstern utility function. Let X be a finite set of consequences. Let Λ be the set of probability distributions in X. Identify $x \in X$ with the probability measure $\delta_x \in \Lambda$ which assigns probability one to x. For a non-empty subset B of S, let $\Lambda^B = \prod_{s \in B} \Lambda$ and let $\Pi^B : \Lambda^S \to \Lambda^B$ be the natural projection. Fix $x^* \in X$ and identify $\lambda \in \Lambda^B$ with the vector $\lambda' \in \Lambda^S$ defined by $\lambda'_s = \lambda_s$, if $s \in B$, and $\lambda'_s = \delta_{x^*}$, if $s \notin B$. With this identification, Λ^B is thought of as a subset of Λ^S. Λ^S is given the usual topology as a subset of a Euclidean space.

Preferences are expressed by orderings $\succ_B$ on Λ^S, where B varies over the non-empty subsets of S. An ordering $\succ_B$ is said to be complete if for each $\lambda \in \Lambda^S$, Λ^S equals the closure of $\{\lambda' | \lambda' \succ_B \lambda \text{ or } \lambda \succ_B \lambda'\}$. If $\succ_B$ is complete, $\lambda \sim_B \lambda'$ means neither $\lambda \succ_B \lambda'$ nor $\lambda' \succ_B \lambda$. Indifference is not the same as the notion of equivalence, defined earlier.

The field $\mathcal{X}$ of subsets of S and the probability $q : \mathcal{X} \to [0,1]$ are as before.

Three additional assumptions are needed. Let

$\Lambda^S_{diag} = \{\lambda \in \Lambda_S | \lambda_s = \lambda_{s'}, \text{ for all } s \text{ and } s'\}$.

<u>Assumption 1.7</u>. The restriction of $\succ_{\{s\}}$ to Λ^S_{diag} is the same for all $s \in S$.

<u>Assumption 1.8</u>. For all $s \in S$, $\succ_{\{s\}}$ is complete.

<u>Assumption 1.9</u>. For any s, $\delta_x \succ_{\{s\}} \delta_{x'}$, for some x and x' in X.

Assumptions 1.2, 1.4 and 1.6 are now replaced by the following.

<u>Assumption 1.2a</u>. $\lambda \succ \lambda'$ if $\lambda \succsim_{\{s\}} \lambda'$, for all s and $\lambda \succ_{\{s\}} \lambda'$, for some s.

<u>Assumption 1.4a</u>. For all λ, $\{\lambda' | \lambda' \succ \lambda\}$ and $\{\lambda' | \lambda' \prec \lambda\}$ are open in Λ^S.

<u>Assumption 1.6a</u>. For any $A \in \mathcal{A}$ and any λ and λ' in Λ, the vector $\lambda_1 \in \Lambda^S$ defined by $\lambda_{1s} = q(A)\lambda + (1 - q(A))\lambda'$, for all s, is equivalent to the vector λ_2 defined by $\lambda_{2s} = \lambda$, if $s \in A$ and $\lambda_{2s} = \lambda'$, if $s \notin A$.

Think of assumptions 1.1, 1.3 and 1.5 as applying to the $\succ_B$ defined on Λ^S rather than R^S.

If $u : X \to R$ and $\lambda \in \Lambda^S$, $E_\lambda u \in R^S$ denotes the vector whose s^{th} component is $\sum_{x \in X} \lambda_s(x)u(x)$.

<u>Theorem 1.2</u>. If the $\succ_B$ satisfy assumptions 1.1, 1.2a, 1.3, 1.4a, 1.5, 1.6a, and 1.7-1.9, then there exists a function $u : X \to R$ and a closed convex set Δ of probabilities in S such that

(i) for all λ, λ' and B, $\lambda \succ_B \lambda'$ if and only if $E_\pi[E_\lambda u|B] > E_\pi[E_{\lambda'} u|B]$, for all $\pi \in \Delta$,

(ii) for all $A \in \mathcal{A}$, $\pi(A) = q(A)$, for all $\pi \in \Delta$, and

(iii) for all $\pi \in \Delta$, $\pi(B) > 0$, for all non-empty subsets B of S.

The proof of this theorem appears in the appendix.

A result of this sort apparently appears in a book being written by Walley. (I have seen a manuscript of only the introduction, Walley (1984).)

Theorems 1.1 and 1.2 are given only in order to show that Knightian behavior is consistent with very strong notions of rationality. From a descriptive point of view, many of the restrictions imposed are not desirable. The assumptions essential for the Knightian theory are an inertia assumption and the structural assumptions given below. Assume that payoffs are in utility and let S, R^S, and so on be as before.

Structural Assumption 1.10. For each non-empty subset B of S, there is a preference ordering $\succ_B$ on R^S such that $x \succ_B y$ if and only if $\Pi^B x \succ_S \Pi^B y$. The ordering $\succ_S$ is transitive and inreflexive and is monotone in the sense that $x \succ_S y$ whenever $x > y$.

Structural Assumption 1.11. There is a set Δ of probabilities on S such that if $x \succ_B y$, then $E_\pi[x-y|B] > 0$, for all $\pi \in \Delta$. Also, $\pi(B) > 0$, for all $\pi \in \Delta$ and for all non-empty subsets B of S.

Notice that it is not assumed that $x \succ_B y$ whenever $E_\pi[x-y|B] > 0$, for all π. This implication seems to be of little interest for economic applications.

The "fatness" of Δ is a measure of the Knightian uncertainty about events in S.

2. BEHAVIORAL ASSUMPTIONS

I now present assumptions relating behavior to preferences. These assumptions are a version of the inertia assumption and the obvious assumption that decision makers make undominated choices. In order to express the inertia assumption rigorously, it is necessary to define a decision problem.

Let the periods of time be $t = 0, 1, \ldots, T$. The description of the environment at time t is $e_t \in E_t$, where E_t is a finite set and E_0 is a singleton. The state or event at time t is $s_t = (e_0, e_1, \ldots, e_t)$. S_t denotes $E_0 \times E_1 \times \ldots \times E_t$ and $S = \bigcup_{t=0}^{T} S_t$ is the event tree. S is ordered naturally by succession. That is, $s_{t+n} = (e_0, \ldots, e_t, e_{t+1}, \ldots, e_{t+n})$ succeeds $s_t = (e_0, \ldots, e_t)$. If $s_t, s_{t+1}, \ldots, s_{t+n}$ are written consecutively, it is implied that s_{t+1} succeeds s_t, and so on. For each $s \in S$, $A(s)$ denotes the set of actions available in state s. Assume that each $A(s)$ is finite. A deterministic program is a function $\underset{\sim}{a}$ giving the action $\underset{\sim}{a}(s) \in A(s)$, for each $s \in S$. If actions are determined by a program $\underset{\sim}{a}$, then the reward in state s_T is $r(\underset{\sim}{a}(s_0), \underset{\sim}{a}(s_1), \ldots, \underset{\sim}{a}(s_T); s_T) = \hat{r}(\underset{\sim}{a}, s_T)$. Rewards are assumed to be measured in utility.

I now describe how programs are compared. It is assumed that the states s_t may be endogenous. That is, their evolution may be influenced by actions. The underlying set of states of nature, Ω, may be described as follows. For each $t = 0, 1, \ldots, T-1$, let $\Sigma_t = \{(s_t, a_t) | s_t \in S_t$ and $a_t \in A(s_t)\}$. Let $\Omega_t = \{\omega_t : \Sigma_t \rightarrow S_{t+1} | \omega_t(s_t, a_t)$ succeeds s_t, for all s_t and $a_t\}$. Then, $\Omega = \Omega_0 \times \Omega_1 \times \ldots \times \Omega_{T-1}$. A deterministic program $\underset{\sim}{a}$ and an $\omega \in \Omega$ together determine a sequence of successive states in S, call it $(s_0, s_1(\underset{\sim}{a},\omega), \ldots, s_T(\underset{\sim}{a},\omega))$.

Assume that the decision maker's preferences satisfy the assumptions of Theorem 1.1. Since Ω is finite, that theorem applies. Let Δ be the closed convex set of evaluating probabilities on Ω. If $\pi \in \Delta$ and $\underset{\sim}{a}$ is a deterministic program, let $E_\pi \hat{r}(\underset{\sim}{a}) = \sum_{\omega \in \Omega} \pi(\omega) \hat{r}(\underset{\sim}{a}, s_T(\underset{\sim}{a},\omega))$.

The list of objects $(S, A, r, \Omega, \Delta)$ defines a decision problem, which I call P.

A random program is a probability distribution over the set of deterministic programs. Random programs are denoted by γ and deterministic ones by $\underset{\sim}{a}$. Deterministic programs may, of course, be thought of as special cases of random ones. If a decision maker uses a random program γ , he chooses a deterministic program $\underset{\sim}{a}$ according to the probability distribution γ and then uses action $\underset{\sim}{a}(s)$ in each state s at which he arrives.

For a program γ , $\hat{r}(\gamma, s_T)$ denotes $\sum_{\underset{\sim}{a}} \gamma(\underset{\sim}{a})\hat{r}(\underset{\sim}{a}, s_T)$. A program γ dominates program γ' if $E_\pi\hat{r}(\gamma) > E_\pi\hat{r}(\gamma')$, for all $\pi \in \Delta$, where $E_\pi\hat{r}(\gamma) = \sum_{\underset{\sim}{a}} \gamma(\underset{\sim}{a})E_\pi\hat{r}(\underset{\sim}{a})$. A program is undominated if no program dominates it. Because S and the $A(s)$ are finite, an undominated program exists. The sets S and $A(s)$ may be assumed to be infinite, provided enough assumptions are made to guarantee that undominated programs exist.

A new decision problem is said to occur by surprise at time t if a state s'_t occurs which does not belong to S and if associated with s'_t there is a decision problem $P(s'_t) = (S', A', r', \Omega', \Delta')$, where the set S' is a tree with elements $(s'_t, \ldots, s'_T)$. It is assumed that state s'_t and problem $P(s'_t)$ were not anticipated by the decision maker. If he had thought of the possibility that they might occur, he had assigned the possibility probability zero. Assume that to every state $s'_n \in S'$ there naturally corresponds a state $f(s'_n) \in S$. By "naturally," I mean that the description of the environment corresponding to s'_n contains all that is in the description of the environment corresponding to $f(s'_n)$. A program $\underset{\sim}{a}$ for the decision problem P is said to apply to $P(s'_t)$ if $\underset{\sim}{a}(f(s')) \in A'(s')$, for all $s' \in S'$. Suppose the decision maker is following program $\underset{\sim}{a}$ for P when a new problem $P(s'_t)$ occurs by surprise.

If $\underset{\sim}{a}$ applies to $P(s_t')$, then the decision maker is said to adopt a new program γ for $P(s_t')$ if γ does not equal the program for $P(s_t')$ defined by $\underset{\sim}{a}$.

The behavioral assumptions are the following.

Maximality Assumption. In any decision problem, the decision maker's actions are determined by an undominated program.

Inertia Assumption. If any decision problem occurs by surprise, the decision maker changes his program only if the new program dominates the old one in the new problem.

These assumptions imply that if a series of surprise problem changes occur, then each time a change occurs the decision maker chooses an undominated program which differs from his previous program only if the new program dominates the old one. The inertia assumption implies that the initial point with which new alternatives are compared is planned behavior.

It might seem that inertia implies irrationality, for people can be truly surprised only if they assign probability zero to events which in fact do occur. However, there is nothing necessarily irrational in making assumptions about reality which turn out to be false. Such assumptions are irrational only if there is good reason to doubt them.

3. INERTIA AND INDEPENDENCE

I now present a version of the inertia assumption in which the arrival of new alternatives is anticipated, but they are used only if doing so is necessary in order to achieve an undominated program. The new actions are assumed to be distinguished from old ones in a natural way. In order to express this distinction within the model of the previous section, assume that there are finite sets $A_0, A_1, \ldots, A_T$ such that $A(s_t) = A_0 \times A_1 \times \ldots \times A_t$, for all s_t and t. The actions in A_t are new at time t. If $\underset{\sim}{a}$ is a deterministic program, $\underset{\sim}{a}(s_t)$ may be written as $(\underset{\sim}{a}_0(s_t), \ldots, \underset{\sim}{a}_t(s_t))$, where $\underset{\sim}{a}_n(s_t) \in A_n$, for all n. The program $\underset{\sim}{a}_t$ is called the t^{th} component program of $\underset{\sim}{a}$.

Assume also that there are functions $r_0, r_1, \ldots, r_T$ such that $r(\underset{\sim}{a}(s_0), \ldots, \underset{\sim}{a}(s_T); s_T) = r_0(\underset{\sim}{a}_0(s_0), \ldots, \underset{\sim}{a}_0(s_T); s_T)$ $+ r_1(\underset{\sim}{a}_1(s_1), \ldots, \underset{\sim}{a}_1(s_T); s_T) + \ldots + r_T(\underset{\sim}{a}_T(s_T); s_T)$. The function $r_t(\underset{\sim}{a}_t(s_t), \ldots, \underset{\sim}{a}_t(s_T); s_T)$ is written as $\hat{r}_t(\underset{\sim}{a}_t, s_T)$, so that $\hat{r}(\underset{\sim}{a}, s_T) = \sum_{t=0}^{T} \hat{r}_t(\underset{\sim}{a}_t, s_T)$.

It is assumed that for each $t > 0$, there is a point $0_t \in A_t$ such that $r_t(0_t, \ldots, 0_t; s_T) = 0$, for all a_T. The action 0_t corresponds to not using A_t. The t^{th} component program $\underset{\sim}{0}_t$ is defined by $\underset{\sim}{0}_t(s_n) = 0_t$, for all s_n with $n \geq t$.

It will be assumed that the decision maker uses a special kind of program, which I call a behavioral program. Let Z be the set of all $z_t = (a_0, \ldots, a_{t-1}, s_t)$, where $s_t \in S$ and $a_0, \ldots, a_{t-1}$ are actions in preceding states. That is $a_n \in A(s_n)$, for all n, where the s_n precede s_t. For each z_t, the subproblem

$P(z_t) = (S(s_t), A, r', \Omega(z_t), \Delta(z_t))$ is the decision problem obtained by restricting the original problem P to the state s_t and its successors. $P(z_t)$ may depend on all the components of z_t because actions taken before s_t influence rewards. The states of $P(a_0, \ldots, a_{t-1}, s_t)$ are $S(s_t) = \{s \in S | s = s_t \text{ or } s \text{ succeeds } s_t\}$. The rewards for a deterministic program $\underset{\sim}{a}'$ for $P(s_t)$ are $r'(\underset{\sim}{a}'(s_t), \ldots, \underset{\sim}{a}'(s_T); s_T)$ $= r(a_0, \ldots, a_{t-1}, \underset{\sim}{a}'(s_t), \ldots, \underset{\sim}{a}'(s_T); s_T)$. The set of states over which probabilities are defined is $\Omega(a_0, \ldots, a_{t-1}, s_t)$ $= \{\omega \in \Omega | s_t(a_0, \ldots, a_{t-1}, \omega) = s_T\}$ where $s_t(a_0, \ldots, a_{t-1}, \omega)$ is the state at time t determined by $a_0, \ldots, a_{t-1}$ and ω. The set of evaluating probabilities is $\Delta(a_0, \ldots, a_{t-1}, s_t)$ $= \{\pi[\cdot|\Omega(a_0, \ldots, a_{t-1}, s_t)] | \pi \in \Omega\}$. Clearly, $P(z_0) = P(s_0) = P$, where P is the entire decision problem.

A behavioral program β specifies for each $z_t \in Z$ a probability distribution $\beta(z_t)$ over the t^{th} component programs $\underset{\sim}{a}_t$ for $P(z_t)$. A decision maker using a behavioral program β selects a t^{th} component program, $\underset{\sim}{a}_t$, for $P(z_t)$ according to the probability distribution $\beta(z_t)$. The program $\underset{\sim}{a}_t$ determines his choice of actions in A_t until time T. A behavioral program therefore determines actions at every state. If β is a behavioral program and $\pi \in \Delta$, then $E_\pi \hat{r}(\beta)$ is defined in the obvious way. It is not hard to see that for any behavioral program β, there exists a random program γ such that $E_\pi \hat{r}(\beta) = E_\pi \hat{r}(\gamma)$, for all $\pi \in \Delta$.

For any $z_t \in Z$, let $\hat{\beta}(z_t)$ be the behavioral program for $P(z_t)$ defined by β. If $\beta(z_t)$ is not the zero program $\underset{\sim}{0}_t$, let $\hat{\beta}_0(z_t)$ be the behavioral program for $P(z_t)$ which is the same as $\beta(z_t)$ except that the t^{th} component program is $\underset{\sim}{0}_t$.

A behavioral program β is said to have the <u>inertia property</u> if for every z_t either $\beta(z_t)$ is the t^{th} component zero program, $\underset{\sim}{0}_t$, for $P(z_t)$ or $\hat{\beta}(z_t)$ dominates $\hat{\beta}_0(z_t)$ in $P(z_t)$. That is, the decision maker uses actions in A_t only if doing so is advantageous from the point of view of time t.

<u>Inertia Assumption</u>. The decision maker chooses a behavioral program with the inertia property.

Trivial examples show that no undominated program may have the inertia property. However, such programs do exist if an independence assumption is made. Before proceeding, I must define independence. Let $\mathcal{P}_n$, for $n = 1, \ldots, N$, be partitions of Ω. For such n, let $\Delta(\mathcal{P}_n)$ $= \{\pi_n | \pi_n$ is the restriction to the field generated by $\mathcal{P}_n$ of some $\pi \in \Delta\}$. The partitions $\mathcal{P}_n$ are said to be mutually independent with respect to Δ if given any $\pi_n \in \Delta(\mathcal{P}_n)$, for $n = 1, \ldots, N$, there exists $\pi \in \Delta$ such that $\pi(\bigcap_{n=1}^{N} A_n) = \prod_{n=1}^{N} \pi_n(A_n)$, for any sets $A_1, \ldots, A_N$, such that $A_n \in \mathcal{P}_n$, for all n. (It does not follow that the partitions $\mathcal{P}_n$ are mutually independent in the usual sense with respect to every $\pi \in \Delta$.) Functions $g_1, \ldots, g_N$ defined on Ω are independent if the partitions they generate are independent.

I now define the independence assumption. Let $\underset{\sim}{A}$ be the set of all deterministic programs $\underset{\sim}{a}$. In the independence assumption about to be stated, consider $\hat{r}_t$ as the function $h : \Omega \to \{g : \underset{\sim}{A} \to (-\infty,\infty)\}$ defined by $h(\omega)(\underset{\sim}{a}) = \hat{r}_t(\underset{\sim}{a}_t, s_T(\underset{\sim}{a},\omega))$. Similarly, consider s_t as the function $h : \Omega \to \{g : \underset{\sim}{A} \to S_t\}$ defined by $h(\omega)(\underset{\sim}{a}) = s_t(\underset{\sim}{a},\omega)$.

Independence Assumption. The functions $\hat{r}_0, \ldots, \hat{r}_T$ are mutually independent with respect to Δ, and, for all t, the functions $\hat{r}_t$ and s_t are independent with respect to Δ. Finally, for all t, if $\omega \in \Omega$ is such that $\pi(\omega) > 0$ for some $\pi \in \Delta$, then $\hat{r}_t(\underset{\sim}{a}_t, s_T(\underset{\sim}{a},\omega))$ does not depend on $\underset{\sim}{a}_n$, for $n \neq t$, where $\underset{\sim}{a}_n$ to the n^{th} component program of $\underset{\sim}{a}$.

The following assumption is also needed.

Separation Assumption. For each s_T and s'_T in S_T such that $s_T \neq s'_T$, there exists t and a t^{th} component program $\underset{\sim}{a}_t$ such that $\hat{r}_t(\underset{\sim}{a}_t, s_T) \neq \hat{r}_t(\underset{\sim}{a}_t, s'_T)$.

Theorem 3.1. If the above assumptions are satisfied, then there exists a behavioral program which is undominated and has the inertia property.

This theorem is proved in the appendix.

The assumptions of this theorem may no doubt be weakened, but some independence assumption seems necessary.

4. INERTIA AND INCOMPLETE PLANNING

One might like to make use of the inertia assumption in settings where the decision maker would have gained had he anticipated that a particular new alternative might appear and probably would have assigned its appearance positive probability, had he considered the matter previously. For these reasons, one might define inertia loosely as follows.

Loose Inertia Assumption. If the decision maker has not previously planned how to react to particular new alternatives, he does not make use of them unless doing so leads to an improvement from the point of view of the moment when he becomes aware of their existence.

If the stricter version of inertia given in Section 2 is valid from a descriptive point of view, then the above version is probably valid as well. The question to be dealt with is whether the looser version corresponds to rational behavior.

One might be tempted to argue that in the presence of uncertainty, an undominated program may be achieved even if one does not plan for events of low probability. Thus, it would be rational not to plan for new alternatives, if their appearance was thought unlikely. This intuition is valid in some cases. However, that it is not always valid is demonstrated by the following example.

Example. There are three periods, labeled 0, 1 and 2. The problem is to distribute purchasing power between periods 0 and 2. The utility function for expenditures, x, in each period is $\log(1+x)$. Utility is enjoyed only in periods 0 and 2. Income is earned in period 2. There are two states in period 2, states L and H. Income is 3 in state L and 5 in state H. The individual may borrow in period 0 at no interest. The loan must be repaid in period 2. There is uncertainty about the state in period 2. In periods 0 and 1, the individual believes that the probability of state L, π_L, lies in the interval [1/6, 1/3]. Insurance against state L may be offered in period 1. If insurance is offered, two units of account in state L may be had in exchange for one in state H. At time 0, the individual believes that insurance will be offered with probability $\alpha > 0$, where α is small.

Suppose that the individual ignores the possibility that insurance may be offered. If we solve $\max_{x_0}[\log(1 + x_0) + \pi_L \log(1 + x_{2L}) + (1 - \pi_L)\log(1 + x_{2H})]$, subject to $x_{2L} = 3 - x_0$ and $x_{2H} = 5 - x_0$, and with $\pi_L = 1/3$, one obtains $x_0 = 2$, $x_{2L} = 1$, $x_{2H} = 3$. This is therefore an undominated program provided the possibility of insurance is ignored. If this program is used and insurance becomes available in period 1, then buying insurance in period 1 would not lead to a preferred position. For suppose that $a > 0$ units of insurance in state L were purchased. Evaluating the gain with $\pi_L = 1/6$, one finds that it is at most $-\frac{a}{48} < 0$.

Suppose now that at time 0 the individual took into account the possibility that insurance might become available. Suppose he changed his program by borrowing $\varepsilon > 0$ more and buying $\varepsilon > 0$ units of insurance in period 1 if it became available. The derivative of his gain with respect to ε at $\varepsilon = 0$ is

$$\alpha\left(\frac{1}{3} - (1 - \pi_L)\frac{1}{4}\left(\frac{3}{2}\right)\right) + (1-\alpha)\left(\frac{1}{3} - \pi_L\,\frac{1}{2} - (1 - \pi_L)\frac{1}{4}\right) ,$$

which is positive for any $\alpha > 0$ and any $\pi_L \in [1/6, 1/3]$. Thus, if $\varepsilon > 0$ is small, this change leads to an unambiguous gain. It is not rational to ignore the possibility that insurance may become available, no matter how small α may be.

If one had chosen the initial program by maximizing with respect to some $\pi_L < 1/3$, then it would have been rational to plan not to buy insurance, if α were sufficiently small. This observation might tempt one to restrict attention to undominated programs which were optimal with respect to some intermediate prior distribution (in the relative interior of Δ). However, such a restriction would conflict with the inertia assumption. In a general decision problem with an initial point, the only undominated program dominating the initial point may be one which is optimal with respect to a prior distribution near to the boundary of the set Δ of prior distributions.

One is thus pushed toward bounded rationality in looking for arguments to defend the rationality of the loose inertia assumption. Bounded rationality certainly makes sense in the context of lifetime decision planning. It is obviously impossible to specify in advance a complete lifetime decision problem. Powerful computers would not help overcome this limitation, since the limiting factor is imagination, not computational capacity.

If complete forward planning is impossible, it makes sense to change one's mind from time to time and not to act according to plan. Inertia is the refusal to change plans unless doing so leads to an improvement.

The disadvantage of bounded rationality is that the concept of rationality itself becomes ambiguous. In trying to describe rational behavior, the best one can do is to imagine what a wise person might do in trying to advance his own interests. Such a person might well show inertia. Inertia may sound conservative and boring, but it can simplify life by reducing the frequency of changes in plans and by eliminating from consideration new alternatives which may arise. There seems to be little more that one can say.

Some insight may be gained into the meaning of the loose inertia assumption by trying to express it slightly more formally in the context of bounded rationality. Suppose the decision maker is faced with a decision problem $(S, A, r, \Omega, \Delta)$ of the sort described in Section 2. If the problem were much too large and complex to be solved completely, a sensible decision maker might organize his thinking by solving a simple approximation to the problem at each stage s_t. Let $M(s_t) = (S', A', r', \Omega', \Delta')$ be the approximating model used in state s_t. We can imagine that the decision maker could achieve coherence between current and future behavior by specifying a function f_{s_t} at state s_t. The function f_{s_t} would assign to any model $M(s_{t+n})$, for $n \geq 0$, a program for $M(s_{t+n})$. The program $f_{s_t}(M(s_t))$ should be maximal in $M(s_t)$. This program would determine action at state s_t. The f_{s_t} could correspond to rules of thumb or standardized procedures for reacting to situations. Of course, f_{s_t} could also simply specify the program maximal with respect to some fixed prior distribution, if one could be specified in advance.

<u>Inertia Assumption</u>. If $f_{s_t}(M(s_{t+1}))$ is maximal in $M(s_{t+1})$, then $f_{s_{t+1}} = f_{s_t}$. Otherwise, $f_{s_{t+1}}(M(s_{t+1}))$ dominates $f_{s_t}(M(s_{t+1}))$ in $M(s_{t+1})$.

5. DIAGRAMMATIC ILLUSTRATION

The relation between the Savage and Knightian theories may be seen easily in a diagrammatic representation of the case with two states. Let $S = \{1,2\}$ and label the abscissa and ordinate with the payoffs, r_1 and r_2, in states 1 and 2, respectively. Payoffs are in units of utility. In the Savage case, the preference ordering is represented by indifference curves which are parallel straight lines with slope $-\pi_1(1 - \pi_1)^{-1}$, where π_1 is the subjective probability of state 1. π_1 is defined by the relation $(\pi_1, \pi_1) \sim (1,0)$. (See Figure 1.)

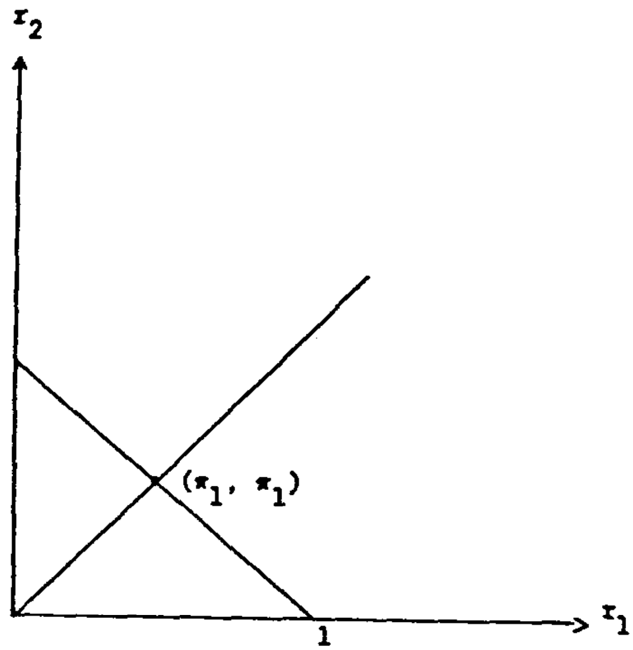

FIGURE 1

According to the Knightian theory of Theorem 1.1, preferences are defined by two families of parallel straight lines with slopes $-\bar{\pi}_1(1 - \bar{\pi}_1)^{-1}$ and $-\underline{\pi}_1(1 - \underline{\pi}_1)^{-1}$, respectively, where $0 < \underline{\pi}_1 \leq \bar{\pi}_1 < 1$.

A point y is preferred to x if and only if y is above the two lines through x . The preference ordering is complete if and only if $\underline{\pi}_1 = \bar{\pi}_1$. (See Figure 2.) If the Knightian theory is that of assumptions 1.10 and 1.11, then one can assert only that if y is preferred to x , then y lies above the two lines through x .

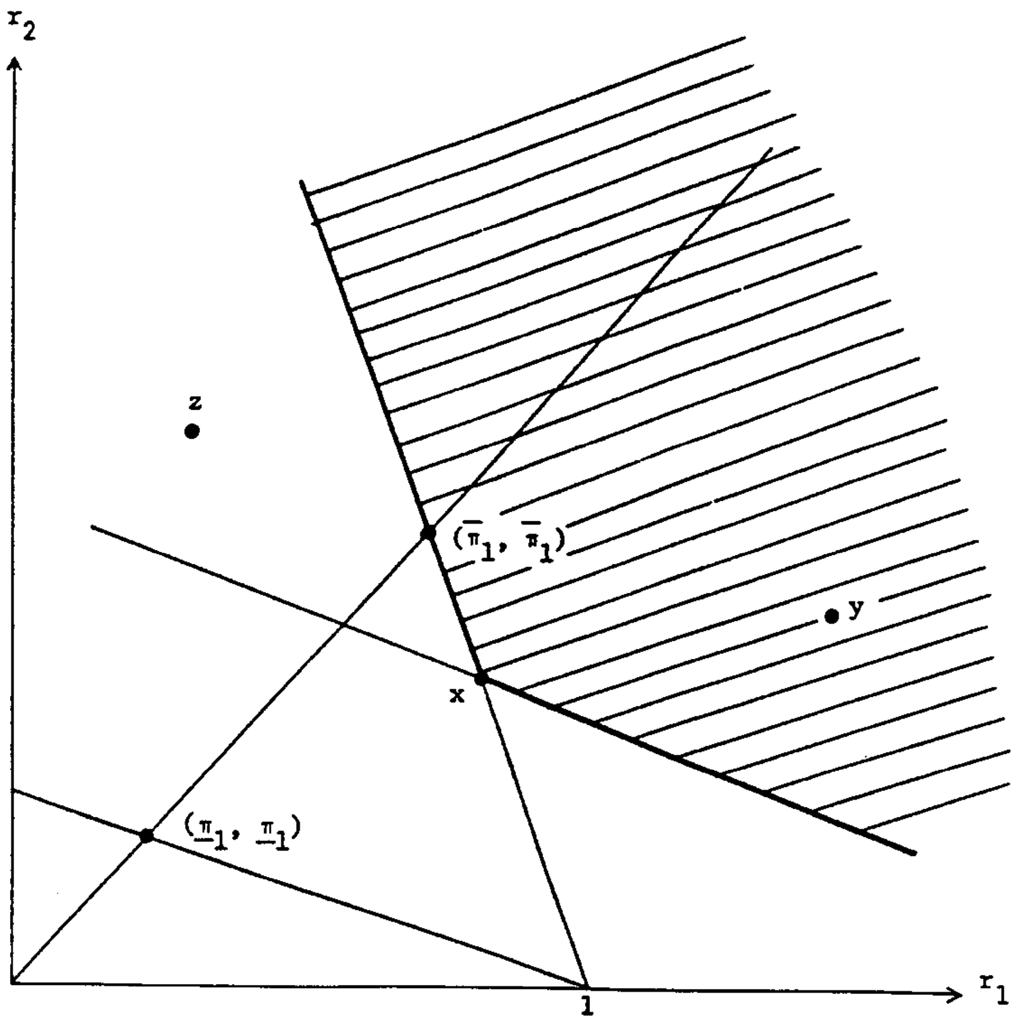

FIGURE 2

The numbers $\bar{\pi}_1$ and $\underline{\pi}_1$ may be thought of as upper and lower probabilities for state 1, respectively. $\bar{\pi}_1$ is defined by $(a,a) \succsim (1,0)$ if and only if $a > \bar{\pi}_1$. $\underline{\pi}_1$ is defined by $(a,a) \prec (1,0)$ if and only if $a < \underline{\pi}_1$. That is, the decision maker is willing to receive \$a in exchange for \$1 in state 1 if and only if $a > \bar{\pi}_1$. Similarly, he is willing to give \$a in exchange for \$1 in state 1 if and only if $a < \underline{\pi}_1$. For any $\varepsilon > 0$, $\underline{\pi}_1 - \varepsilon$ and $\bar{\pi}_1 + \varepsilon$ are possible bidding and asking prices, respectively, for \$1 in state 1. If $\underline{\pi}_1 \leq a \leq \bar{\pi}_1$, then $(a,0)$ and $(1,0)$ are not comparable. Similar definitions may be given for upper and lower probabilities for state 2, $\bar{\pi}_2$ and $\underline{\pi}_2$, respectively. They satisfy $\bar{\pi}_2 = 1 - \underline{\pi}_1$ and $\underline{\pi}_2 = 1 - \bar{\pi}_1$.

Notice that if $x_1 > 0$ and $x_2 < 0$, then $x > 0$ if and only if $\underline{\pi}_1 x_1 + \bar{\pi}_2 x_2 > 0$. Gains are weighted by the upper probability and losses by the lower probability. This conservative weighting of gains and losses results in uncertainty aversion.

When there are only two states, preferences are defined entirely by upper and lower probabilities. When there are more than two states, it may not be possible to derive preferences from upper and lower probabilities, for the set Δ of subjective probabilities may be "round."

According to the inertia assumption of Section 2, if x in Figure 2 is the initial point, then a point such as y would be chosen instead of x , if y were offered by surprise as an alternative to x . A point such as z would not be chosen over x . After y is chosen, it becomes the new initial point. Thus, the inertia assumption would prevent intransitive choices among successive surprise alternatives. However, if various alternatives were offered in some sequence which was foreseen or at least thought possible, some successive choices might be intransitive. However, the maximality or undominatedness of the program guiding these choices would prevent the occurrence of a money pump. If either of the other two versions of inertia are assumed, a money pump is impossible for the same reasons.

According to the inertia assumption of Section 3, the initial point necessarily plays a role in choice only if the events 1 and 2 are independent of all else of significance in the decision maker's life. If they are independent of the rest of his life, then we can say only that any choices made must be preferred to zero. It will not necessarily be the case that each choice becomes the new initial point.

If a set, such as C in Figure 3, is made available and if the initial point is zero, then any point on the boundary of C between A and B could be chosen. Two decision makers with the same preferences might choose

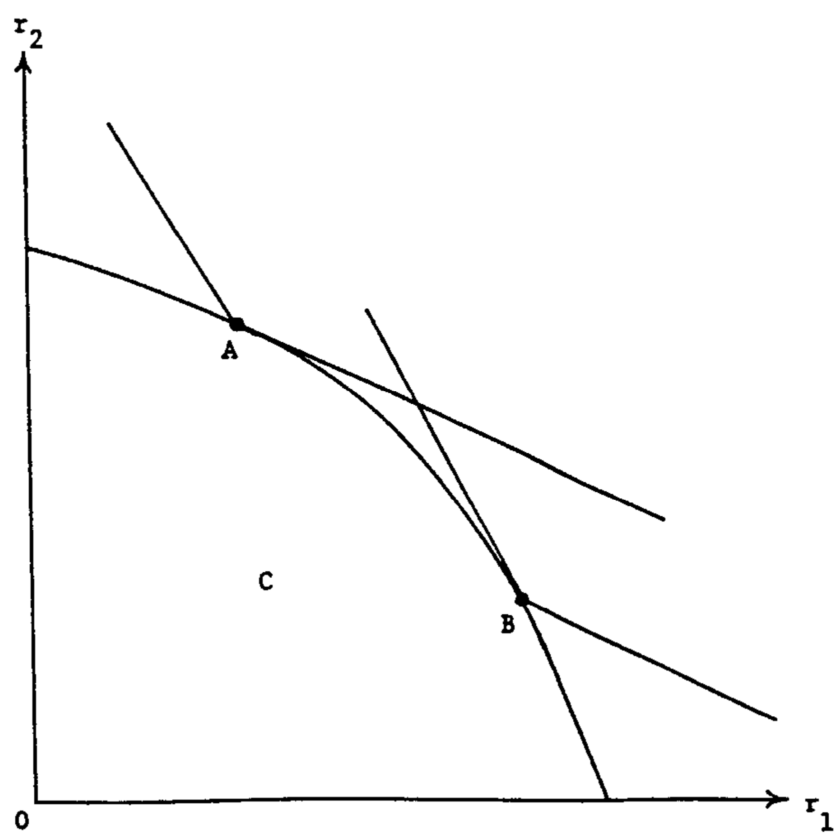

FIGURE 3

different points along this frontier. The inertia assumptions of Section 1 or 3 imply that once a choice was made, the decision maker would not want to move to another point along the frontier AB . Two decision makers making different choices would not want to make side bets with each other unless an upper probability of one decision maker were less than the corresponding lower probability of the other decision maker.

Uncertainty aversion could discourage mutual insurance. Let the rewards r_s now be measured in units of one commodity and suppose that utility is concave in r_s . Then, the Edgeworth box diagram for the case of two states and two traders could be as in Figure 4. Assume that the initial endowment point, ω , is also the initial point. The sets of points preferred to the initial point are denoted P_1 and P_2 , respectively. Even if

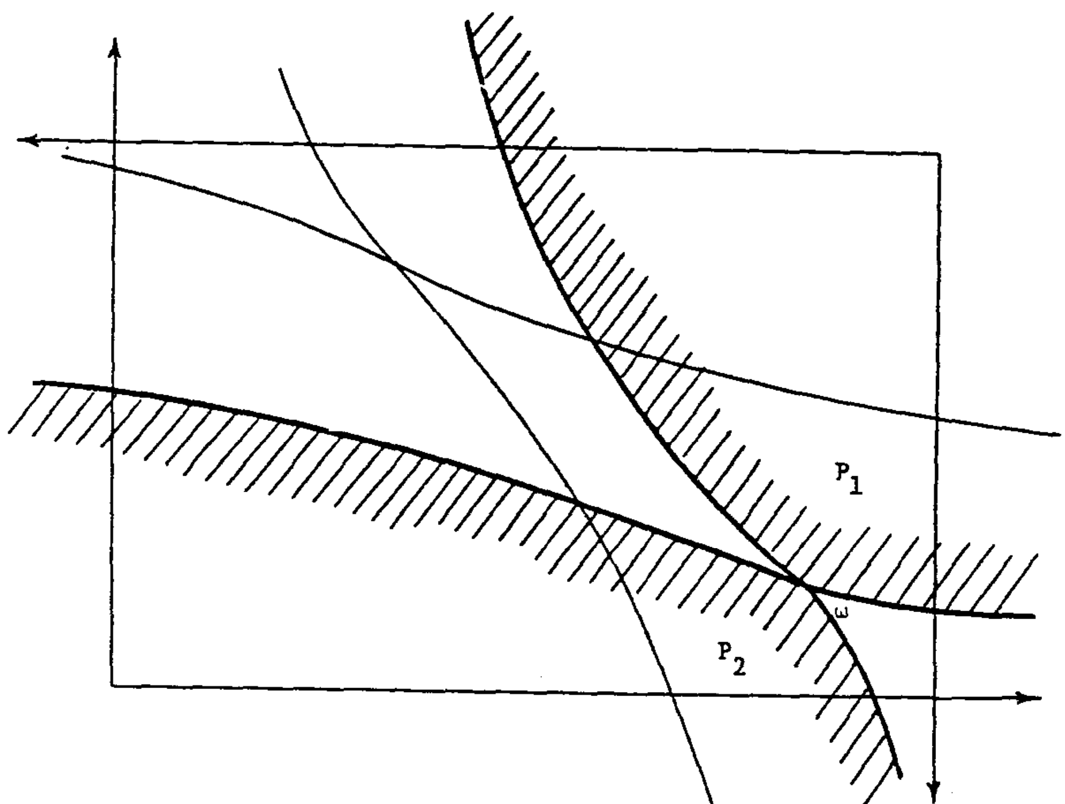

FIGURE 4

endowments were very asymmetric and the preferences of the traders were the same, there might be no trade in insurance. Nevertheless, the equilibrium would be Pareto optimal. I now compare the interpretation of Knight presented in this paper with the competing one of Gilboa and Schmeidler (1986) mentioned in the Introduction. Recall that Gilboa-Schmeidler preferences are complete and represented by the utility function of the form $u(x) = \min_{\pi \in \Delta} E_{\pi} x$. Here, $x : S \to (-\infty,\infty)$ is a gamble over a set of states S with rewards in utility, and Δ is a set of probability distributions over S . These preferences display uncertainty aversion since preferred sets are convex. However, these preferences otherwise have implications quite different from those described in Section 1. People with Gilboa-Schmeidler preferences would be very apt to buy insurance. The Edgeworth box diagram

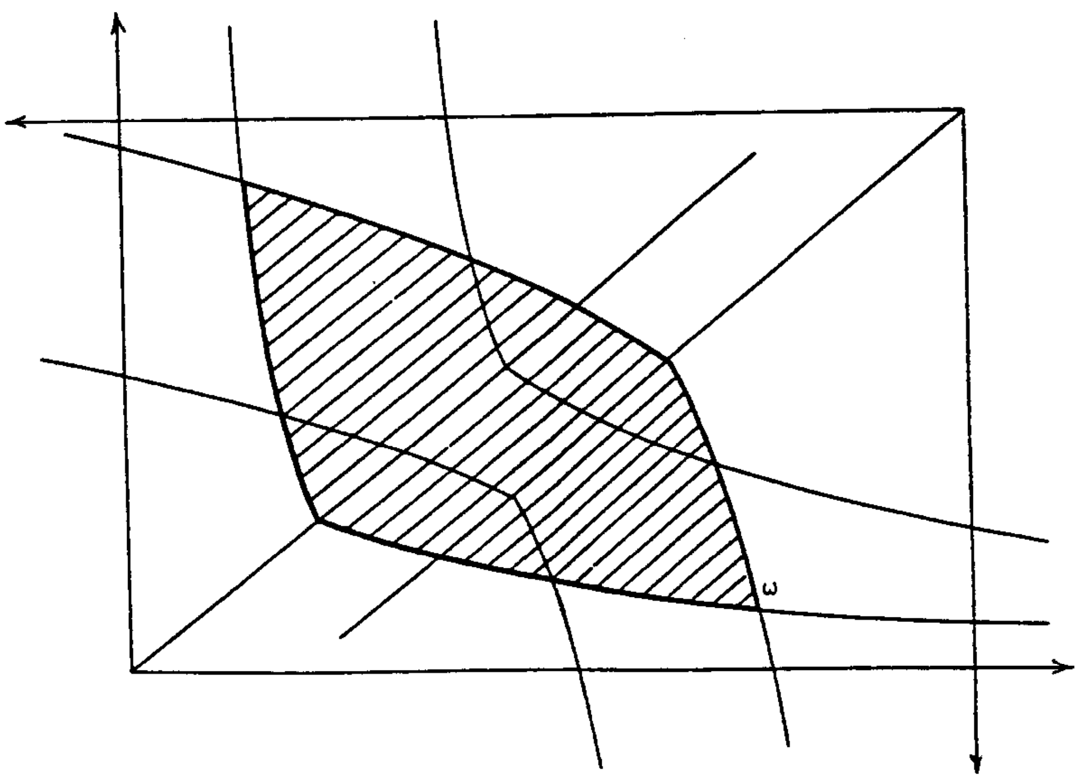

FIGURE 5

corresponding to Figure 4 would be as in Figure 5. Trades in the shadowed area Pareto dominate the initial endowment point.

If two decision makers with the same Gilboa-Schmeidler preferences were offered a set C as in Figure 3, then they would choose the same point. If two Gilboa-Schmeidler decision makers choose different points from C, they would be willing to make side bets with each other after the choices were made, as is illustrated in Figure 6. The decision maker with the solid indifference curves would choose point A. The decision maker with dotted indifference curves would choose point B.

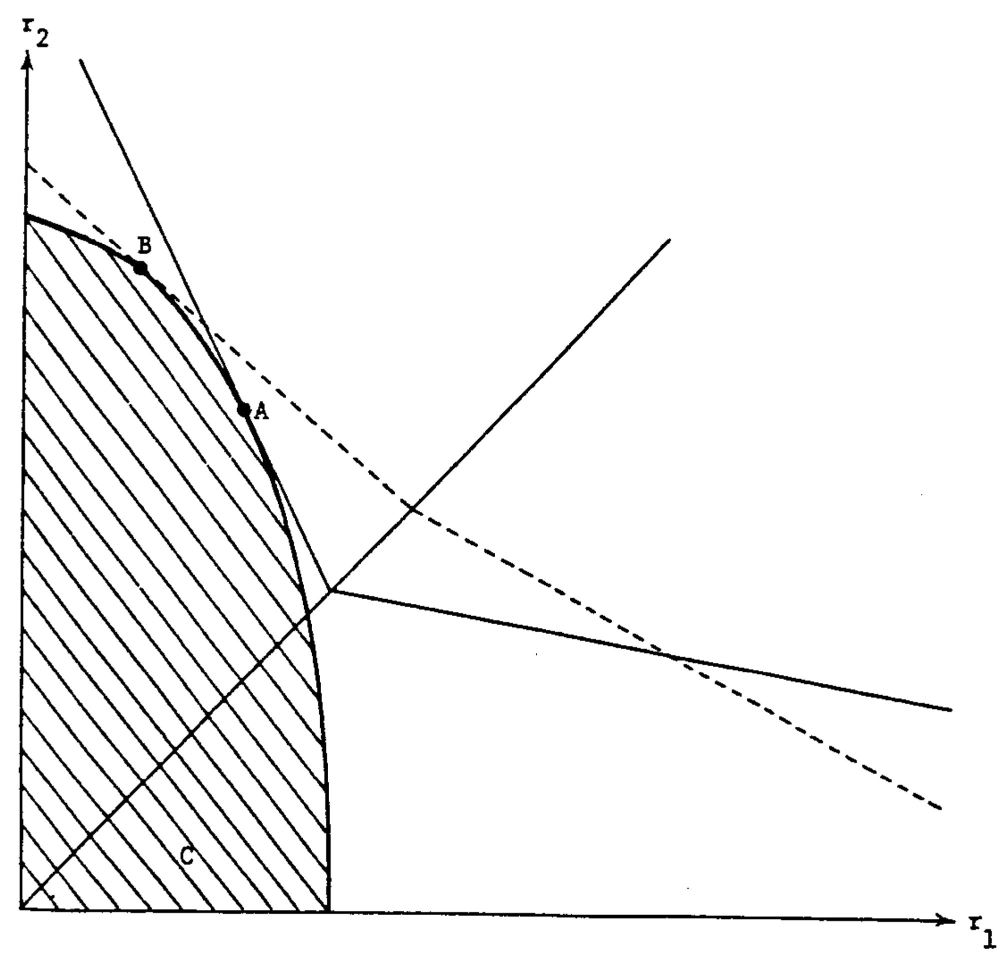

6. RELATION TO EXPERIMENTAL EVIDENCE

The expected utility hypothesis seems to be rejected systematically by experimental evidence. (Surveys of the experimental literature on the subject may be found in Machina (1982, 1983) and Schoemaker (1982).) Since the von Neumann-Morgenstern theory is a special case of the Knightian theory of Theorems 1.1 and 1.2, that theory is rejected too. However, the essence of the Knightian theory has little to do with the expected utility hypothesis. The Knightian theory is captured by the inertia assumption and structural assumptions 1.10 and 1.11. These do not imply the expected utility hypothesis. But they do imply the essential phenomena of uncertainty aversion and inertia. These phenomena may have some chance of being verified experimentally.

It is perhaps encouraging that the essential phenomena have little to do explicitly with probabilities or the calculus of probabilities. Probability is foreign to most people's everyday experience. It requires training, after all, simply to get used to the elementary concepts of expected value, independence and conditioning. Since lotteries with known probabilities are encountered rarely by most people, the law of large numbers does not justify the axioms of von Neumann-Morgenstern. Because decision making under Knightian uncertainty is a large part of life, there may be some grounds for hope that people react more systematically to it than they do to lotteries with explicit probabilities.

One of the implications of inertia and uncertainty aversion is that bid prices for insurance of an uncertain event may be systematically less than asking prices, even if the insured event results in a loss to the bidder and not to the asker or seller. Many explanations may be given for this bid-ask spread. Game and information theoretic explanations may be found in Leamer (1985a). Still another explanation may be found in Einhorn and Hogarth (1985).[2] The same paper reports experimental work which tends to confirm the existence of bid-ask spreads on insurance of an event of vague probability, even when the event causes a loss to the bidder.

One might imagine that Ellsberg's (1961) experiments lend support to the Knightian theory. However, the choices among the alternatives he offered would be indeterminate according to the theory presented here, so that his experiments neither confirm nor contradict the theory. Ellsberg's (1961, 1963) experiments are, however, consistent with preferences of the Gilboa-Schmeidler type discussed in the previous section. This fact is an advantage of such preferences. These preferences, however, would not explain a gap between bid and ask prices for insurance against a loss suffered by the bidder.

One can imagine simple experiments designed to test the Knightian theory. For instance, subjects could be shown a photograph of someone whose age is verifiable but unknown to the subject. The subject could be offered a sequence of lotteries whose outcome would depend on the true age. Once a lottery was accepted, each new lottery should be offered as an alternative to the one previously accepted. The sequence of choices offered should not depend on choices made, and this fact should be made clear to the subject. The last lottery accepted should be paid off at the end of the session when the true age was revealed. Some of the payoffs must be negative. In order to induce participation, it might be necessary to pay a fixed sum in

---

[2] I owe this reference to Sidney Winter.

addition. Payoffs should be large enough to interest the subjects but so small that utility could be assumed to be linear in payoffs.

Structural assumptions 1.10 and 1.11 and the loose inertia assumption of Section 4 imply that if any lottery is chosen, it should have a positive worth according to a set of prior probability distributions. Also, if new lotteries were accepted in succession, each should be at least as valuable as the previously accepted one according to the same probability distributions. The experiment could determine whether behavior was consistent with these assertions.

If the strict inertia assumption of Section 2 is assumed, one could assert only that each lottery chosen should have positive value according to all the prior probability distributions. One could not assert that each new alternative chosen should dominate the previous choice, for one could not assure that each new alternative was a surprise. The subject would surely expect a sequence of alternatives to be offered. Hence, behavior in the experiment would be largely as predicted by the Gilboa-Schmeidler theory. Only if choices seemed to be irregular or indeterminate could one assert that the experimental results favored the Knightian theory over that of Gilboa-Schmeidler.

It could be difficult to design an experiment which could distinguish clearly Gilboa-Schmeidler preferences from incomplete preferences obeying the strict inertia assumption of Section 2. It is hard to imagine how one could generate true surprise so as to manipulate the initial position. It seems that one would have to take the initial positions as given and seek subjects with different initial positions with respect to some events. For instance, one could offer to buy or sell small but real insurance contracts on some event of vague probability from which certain subjects would suffer financially. An example would be job loss. The object would be to determine if potential sufferers were willing to pay as much as non-sufferers were willing to accept for such contracts. Clearly, any such experiment would be fraught with ambiguities.

7. RELATED LITERATURE

There are two bodies of literature very closely related to what has been presented here, one in economics and done chiefly by Robert Aumann, and another in statistics.

Robert Aumann (1962, 1964a, b) studied the representation of incomplete preferences on what he called mixture spaces. Among other things, he gave conditions on an ordering such that a linear utility function u represents it in the sense that $x \succ y$ implies $u(x) > u(y)$. His work is described in Fishburn (1970), Chapter 9, and has been extended to infinite dimensional spaces by Kannai (1963). None of these authors related the work to Knightian uncertainty. The typical interpretation made of this work by economists seems to be that incompleteness is consistent with the expected utility hypothesis. (See, for example, Yaari (1985).) The representation theorems of Section 1 in this paper are essentially interpretations of Aumann's theorems.

The body of literature in statistics consists of papers by Smith (1961), Williams (1976) and Walley (1981, 1982, 1984). Smith, among other things, presents in an informal way Theorem 1.1 of this paper. His work is formalized and elaborated in the papers by Williams and Walley. All these authors are interested mainly in upper and lower probabilities as tools of statistical analysis. These may be derived as follows from preferences obeying Theorem 1.1. If $A \subset S$, the upper probability of S is $\bar{p}(A) = \{\pi(A) | \pi \in \Delta\}$. The lower probability of A is $\underline{p}(A) = \bar{p}(S \backslash A) = \min\{\pi(A) | \pi \in \Delta\}$. Because of the authors' interest in probabilities, their presentation is not in a form convenient for economic interpretation, so that I could not replace the structural theorems of Section 1 by citations of their work. These authors tend to focus on the set $K = \{x \in R^S | x \succ 0\}$ rather than on the entire preference ordering. The set K is referred to as the set of desirable or acceptable gambles. Assumptions are made such that K is a convex cone not intersecting $R_-^S = \{x \in R^S | x \leq 0\}$ or such that the convex hull of K does not intersect R_-^S. It is taken for granted that no gamble would be accepted unless it were preferred to zero. Thus, the inertia assumption is implicit. In fact, the authors in statistics seem not to make any use of the incompleteness of preferences. Williams (1974) does not even mention incompleteness. It seems to me that Gilboa-Schmeidler preferences are the most appropriate foundation for the use in statistics of upper and lower probability.[3] A normatively sound theory of choice should be enough for the foundations of statistics. Why should statisticians care about the descriptive accuracy of the theory of choice they use? Why saddle statistics with the ambiguities associated with incomplete preferences?

There is a large statistical literature on upper and lower probabilities, which I do not cite. I mention, however, that there is a philosophical literature which uses upper and lower probabilities to characterize

---

[3]Leonid Hurwicz (1951) has made such a suggestion.

beliefs and discusses how a rational person ought to relate his beliefs to information or evidence. General sources in this area are Levi (1980, 1984) and Shafer (1978). Leamer (1985b) has argued that upper and lower probabilities should be used to present econometric conclusions.

8. CONCLUSION

Even though one cannot be sure that uncertainty aversion and inertia are facts of life, one can speculate about their role in economic life. I intend to indulge in such speculations in future papers for I believe that Knightian decision theory may explain many puzzling economic phenomena. I here sketch some possible insights.

One of the consequences of the incompleteness of preferences and uncertainty aversion is that uncertainty can make very simple programs be undominated. Apparently excessively simple economic behavior becomes rational when seen from a Knightian point of view. Examples of such behavior are the use of mark-up pricing rules in retail firms (see Cyert and March (1963), Chapter 7 and Baumol and Stewart (1971)), and the frequent lack of diversification of individual investment portfolios (see Blume and Friend (1978)). The same arguments can rationalize the behavioral routines discussed in Nelson and Winter's (1982) theory of the firm.

As has already been mentioned in Section 5, uncertainty aversion and inertia can explain reluctance to buy or sell insurance when the probability of loss is ambiguous. Thus, Knightian behavior may explain the absence of many markets for insurance and forward contracts.

Uncertainty aversion and inertia can be used to give a rigorous presentation of Knight's theory of the entrepreneur in terms of a general equilibrium model. The entrepreneurs are those with fatter cones of preferences or smaller sets of subjective probabilities.

Knightian decision theory may also offer a possible solution to the vexing question of how to explain wage rigidity, layoffs, and rigid long-term contracts in general. In fact, I was led to Knightian decision theory by exasperation and defeat in trying to deal with these questions using the concepts of asymmetric information and risk aversion.

I try to give an intuitive explanation of the connection between wage contracts and Knightian decision theory. I believe that this explanation has something in common with the ideas of Oliver Williamson (1975, 1985).

Imagine the employees of a firm and the firm's owners as being locked in a long-term relationship, with the employees being capable of collective action. Leave aside the question of where that capacity comes from. (One can resort to the theory of repeated games in order to make strikes subgame perfect.) Suppose the employees and owners have already agreed on a criterion for fair division of the benefits of their relationship. Imagine that the business prospects of the firm and the value marginal products of labor and capital are hard to assess. Suppose the two sides have the same objective information about all relevant matters. If they did not have the same information, one would imagine that it would be to their mutual advantage to share it, since in models of bargaining asymmetric information leads to a Pareto loss. In the context of a long-term relationship, there ought to be little reluctance to share information, since it should be possible to punish either side for taking unfair advantage of shared information. Because of the Knightian uncertainty, it would not necessarily be clear what the agreed-on criterion of fairness implied, even though there was no asymmetric information about observables. Even if the two sides were honest, they could disagree and no outsider would be able to say who was right. Because of the ambiguity, there would be room for posturing and falsification of one's opinion.

If the only issue at stake were fairness, an arbiter might be used, if one could be found who understand the complex business situation. But more than fairness might be at issue, for if the wage were too high in the opinion of the owners, then they might be discouraged from investing in the firm, which would in turn be against the interests of the employees. But, again, there may be vagueness about the relation between investment and the wage.

This vagueness cannot be resolved by arbitration, for it is important that the wage be acceptable to the owners. The vagueness must be resolved by bargaining. But bargaining is meaningless without some loss that can be imposed to prevent posturing by either side. The role of strikes, lockouts, and other bargaining costs may be to prevent posturing and to achieve incentive compatibility in the bargaining process. There is asymmetric information about each side's judgment as to an appropriate outcome.

The greater the vagueness, the greater the potential punishments needed and the more likely it is that they will be imposed. Thus, bargaining costs increase with vagueness. For this reason, it is valuable to have contracts be simple; simpler contracts being easier to evaluate. Simplicity may imply wage rigidity and lack of indexation.

The fact that contracts are long-term may be explained by bargaining costs resulting from ambiguity. Suppose one had weekly bargaining for weekly contracts. The incentive for posturing might be as large in bargaining for short-term contracts as for long-term ones. Business conditions and opinions about them probably change slowly. Therefore, if one side could succeed in conveying a false impression of its own opinion, this success would be of value in many future periods. For instance, if the owners once persuaded their employees that low wages were necessary for the health of the firm, the workers would be likely to stay convinced for some time. Thus, the costs of bargaining for each short-term contract could be as high as the costs of bargaining for a long-term contract. It could, therefore, be advantageous to have long-term contracts. The same argument can apply, of course, to any model with bargaining costs which result from asymmetric information.

In the presence of vagueness, bargaining costs could also be reduced by resolving disagreements according to some commonly recognized formula. For instance, the current wage could be the previous one plus some adjustment for changes in productivity and the cost of living. Such an argument can explain wage rigidity from contract to contract.

APPENDIX

Proof of Theorem 1.1

Let $K = \{x \in R^S | x \succ 0\}$. I show that for x and y in R^S , $x \succ y$ if and only if $x - y \in K$. By assumption 1.5, for all x , y and z in R^S and $\alpha \in (0,1)$, $y \succ z$ if and only if $(1-\alpha)x + \alpha y \succ (1-\alpha)x + \alpha z$. Letting $x = 0$, one obtains $y \succ z$ if and only $\alpha y \succ \alpha z$. Hence, $y - z \succ 0 = z - z$ if and only if $\frac{y-z}{2} \succ \frac{z-z}{2}$, which by assumption 1.5 is true if and only if $y \succ z$.

I next show that K is a convex, open cone containing $R^S_+ \backslash \{0\} = \{x \in R^S | x > 0\}$. K is convex, for suppose that x and y belong to K and $\alpha \in (0,1)$. $\alpha x + (1-\alpha)y \succ \alpha x$, by assumption 1.5, since $y \succ 0$. Since $x \succ 0$, $\alpha x \succ 0$. Hence, by transitivity $\alpha x + (1-\alpha)y \succ 0$ or $\alpha x + (1-\alpha)y \in K$. K is a cone, for let $x \in K$ and $t > 0$. If $t < 1$, $tx \in K$ by what has already been proved. If $t > 1$, $x = t^{-1}(tx) \in K$ only if $tx \in K$, by what has been proved. By assumption 1.2, K contains $R^S_+ \backslash \{0\}$. By assumption 1.4, K is open. By assumption 1.3, $0 \notin K$.

By the Minkowski separation theorem, the set $\Delta = \{\pi : S \to [0,1] | \sum_s \pi_s = 1 \text{ and } \pi \cdot x > 0, \text{ for all } x \in K\}$ is non-empty and $K = \{x \in R^S | \pi \cdot x > 0, \text{ for all } \pi \in \Delta\}$. Consider the π in Δ to be probability measures on S .

If B is a non-empty subset of S , $e_B \in K$ and so $\pi(B) = \pi \cdot e_B > 0$, for all $\pi \in \Delta$.

This proves part (iii) of the theorem.

By assumption 1.1, for x and y in R^S and B a non-empty subset of S, $x \succ_B y$ if and only if $\Pi^B x \succ \Pi^B y$, which is true if and only if $E_\pi[x|B] > E_\pi[y|B]$, for all $\pi \in \Delta$. This proves part (i) of the theorem.

In order to prove part (ii), let $A \in \mathcal{A}$ and let $\varepsilon > 0$ be arbitrarily small. By assumptions 1.2 and 1.6, $(q(A) + \varepsilon)e_S \succ e_A$, so that for all $\pi \in \Delta$, $0 < \pi \cdot ((q(A) + \varepsilon)e_S - e_A) = q(A) + \varepsilon - \pi(A)$. Similarly, $(q(A) - \varepsilon)e_S \prec e_A$, so that $\pi(A) > q(A) - \varepsilon$. Therefore, $\pi(A) = q(A)$.

Q.E.D.

Proof of Theorem 1.2

It follows from assumption 1.5 that

A.1) if $\lambda \succ \lambda'$ and if $0 \le \alpha < \beta \le 1$, then

$$\beta\lambda + (1-\beta)\lambda' \succ \alpha\lambda + (1-\alpha)\lambda' .$$

By assumption 1.7, the orderings $\succ_{\{s\}}$, for $s \in S$, induce a unique order $\succ_0$ on Λ. By assumption 1.1, $\succ_0$ satisfies assumptions 1.3, 1.4a, 1.5 and 1.9. Hence, statement A.1 applies to $\succ_0$. By assumption 1.8, $\succ_0$ is complete. Let $\sim_0$ be the indifference relation associated with $\succ_0$. Since statement A.1 and assumptions 1.3 and 1.4a apply to $\succ_0$, the relation $\succsim_0$ is transitive.

Let $\underline{x}$ and $\bar{x}$ in X is such that $\underline{x} \precsim_0 x \precsim_0 \bar{x}$, for all $x \in X$. By assumption 1.5 $\underline{x} \precsim_0 \lambda \precsim_0 \bar{x}$, for all $\lambda \in \Lambda$. By assumption 1.9, $\underline{x} \prec_0 \bar{x}$. By statement A.1 and assumption 1.4a, for each $\lambda \in \Lambda$, there is a unique $u(\lambda) \in [0,1]$ such that $u(\lambda)\bar{x} + (1 - u(\lambda))\underline{x} \sim_0 \lambda$. From assumption 1.5, it follows that if $0 < \alpha < 1$, then $u(\alpha\lambda + (1-\alpha)\lambda') = \alpha u(\lambda) + (1-\lambda)u(\lambda')$. Hence, $u(\lambda) = \sum_{x \in X} \lambda(x)u(x)$.

If $\lambda \in \Lambda^S$, let $U(\lambda) \in \Lambda^S$ be the vector defined by

$U(\lambda)_s = u(\lambda_s)\delta_{\bar{x}} + (1 - u(\lambda_s))\delta_{\underline{x}}$, for all $s \in S$.

Lemma. If λ and λ' belong to Λ^S , then $\lambda \succ \lambda'$ if and only if $U(\lambda) > U(\lambda')$.

Proof. Suppose that $\lambda \succ \lambda'$. By assumptions 1.2a and 1.4a and statement A.1, it is possible to choose $\lambda'' \in \Lambda^S$ such that $\lambda'' \succ \lambda'$ and $\lambda_s \succsim_{\{s\}} \lambda''_s$, for all s , and $\lambda_s \succ_{\{s\}} \lambda''_s$, for some s . By assumption 1.2a, $U(\lambda) \succ \lambda''$. Therefore, $U(\lambda) > \lambda'$. A similar argument proves that $U(\lambda) \succ U(\lambda')$.

The same sort of argument proves that $\lambda \succ \lambda'$ if $U(\lambda) > U(\lambda')$.

Q.E.D.

The proof of Theorem 1.1 may now be applied to the orderings $\succ_B$ restricted to $\{U(\lambda) | \lambda \in \Lambda^S\}$, which is isomorphic to a subset of R^S . The proof must be modified slightly because now utility levels vary over $[0,1]$, whereas in Theorem 1.1 they varied over $(-\infty,\infty)$.

Q.E.D.

Proof of Theorem 3.1

For any t , let $\mathcal{S}_t = \{S(s_t) \cap S_T | s_t \in S_t\}$ be the partition of S_T generated by information available at time t . Let $\mathcal{P}_t$ be the partition of S_T generated by r_t . That is, $\mathcal{P}_t$ is generated by the function h : $h : S_T \to \{g : \prod_{n=t}^{T} A_t \to (-\infty,\infty)\}$ defined by $h(s_T)(a_t, \ldots, a_T) = r_t(a_t, \ldots, a_T; s_T)$. For $n = t, \ldots, T$, let $\mathcal{P}_{tn}$ be the partition $\{E | \text{for some } E' \in \mathcal{S}_n, E = \cup\{E'' \in \mathcal{P}_t | E'' \cap E' \neq \emptyset\}\}$. $\mathcal{P}_{tn}$ represents information about the function r_t available at time n . Clearly, $\mathcal{P}_{tT} = \mathcal{P}_t$. By the independence of s_t and $\hat{r}_t$ assumed in the independence assumption,

$\mathcal{P}_{tt} = \{S_T\}$. Since $\mathcal{P}_{t,n+1}$ refines $\mathcal{P}_{tn}$, for all n , the partitions $\mathcal{P}_{tt}, \mathcal{P}_{t,t+1}, \ldots, \mathcal{P}_{tT}$ from a tree. Call this tree $\mathcal{T}_t$.

Let P_t be the (T-t+1)-period decision problem $(\mathcal{T}_t, A_t, r'_t, \Omega'_t, \Delta_t)$ defined as follows. $A_t(E) = A_t$, for every $E \in \mathcal{T}_t$. If $E \in \mathcal{P}_{tT}$ and $a_n \in A_n$, for $n = t, \ldots, T$, then $r'_t(a_t, \ldots, a_T: E) = r_t(a_t, \ldots, a_T; s_T)$, for some $s_T \in E$. $\Omega'_T = \prod_{n=t}^{T-1} \Omega'_{tn}$, where $\Omega'_{tn} = \{\omega : \Sigma'_{tn} \to S_{t+1}\}$ and $\Sigma'_{tn} = \{(E_n, a_n) | E_n \in \mathcal{P}_{tn} \text{ and } a_n \in A_t\}$. For $\omega' \in \Omega'_t$ and for a program $\underset{\sim}{a}$ for P_t , let $E_T(\underset{\sim}{a}, \omega')$ be a unique member of $\mathcal{P}_t$ reached if $\underset{\sim}{a}$ is used and the state of nature is ω' . The mapping from ω' to the function $E_T(\cdot, \omega')$ is one to one.

Observe that any deterministic program $\underset{\sim}{a}_t$ for P_t may also be thought of as defined on $\bigcup_{n \geq t} S_t$, and so may be thought of as a t^{th} component program for the decision problem P . That is, if $n \geq t$, think of $\underset{\sim}{a}_t$ as assigning action $\underset{\sim}{a}_t(E(s_n))$ to $s_n \in S_n$, where $E(s_n)$ is the member of $\mathcal{P}_{tn}$ containing s_n .

I now define the sets Δ_t . Let $\Omega' = \{\omega \in \Omega |$ for all t, $\hat{r}_t(\underset{\sim}{a}_t, s_T(\underset{\sim}{a},\omega))$ does not depend on $\underset{\sim}{a}_n$, for $n \neq t\}$. By the independence assumption, $\pi(\Omega') = 1$, for all $\pi \in \Delta$. For each t , let $\mathcal{P}'_t$ be the partition of Ω' generated by the function $\hat{r}_t$. By the independence assumption, the partitions $\mathcal{P}'_t$ are independent with respect to Δ . Let $\Delta(\mathcal{P}'_t) = \{\pi_t | \pi_t$ is the restriction of π to the field generated by $\mathcal{P}'_t$, where $\pi \in \Delta\}$. The partition $\mathcal{P}'_t$ may be identified with Ω'_t . That is, $E \in \mathcal{P}'_t$ corresponds to the $\omega' \in \Omega'_t$ satisfying $s_T(\underset{\sim}{a},\omega) \in E_T(\bar{\underset{\sim}{a}}_t, \omega')$, for all deterministic programs $\bar{\underset{\sim}{a}}_t$ for P_t , where ω is any element of E , and $\underset{\sim}{a} = (\underset{\sim}{a}_0, \ldots, \underset{\sim}{a}_T)$ is any program for the decision problem P satisfying

$\underset{\sim}{a}_t = \bar{\underset{\sim}{a}}_t$. Let $\Delta_t = \Delta(\mathcal{P}'_t)$, where $\mathcal{P}'_t$ is identified with Ω'_t .

The following notation is applied to the problem P_t . If $\underset{\sim}{a}$ is a deterministic program for P_t and $E_T \in \mathcal{P}_t$, then $\hat{r}'_t(\underset{\sim}{a}, E_T) = r'_t(\underset{\sim}{a}(E_t), \ldots, \underset{\sim}{a}(E_T); E_T)$. If $\pi \in \Delta_t$, then $E_\pi \hat{r}'_t(\underset{\sim}{a}) = \sum_{\omega' \in \Omega'_t} \pi(\omega') \hat{r}'_t(\underset{\sim}{a}, E_T(\underset{\sim}{a}, \omega'))$. If γ is a random program for P_t , then $\hat{r}'_t(\gamma, E_T)$ and $E_\pi \hat{r}'_t(\gamma)$ are defined in the obvious ways. The functions $\hat{r}'_t$ should not be confused with the functions $\hat{r}_t$ defined in Section 3.

I now choose a behavioral program for the decision problem P . First of all, I select an undominated program $\bar{\gamma}_t$ for each P_t . If the zero program $\underset{\sim}{0}_t$ is undominated in P_t , let $\bar{\gamma}_t = \underset{\sim}{0}_t$. Otherwise, let $\bar{\gamma}_t$ be any undominated program for P_t which dominates $\underset{\sim}{0}_t$. If $\underset{\sim}{a}_t$ is in the support of $\bar{\gamma}_t$, then $\underset{\sim}{a}_t$ may be thought of as a t^{th} component program for the problem P as well. It is therefore also a t^{th} component program for any subproblem $P(z_t)$, where $P(z_t)$ is defined as in Section 3. Thus, $\bar{\gamma}_t$ defines a random t^{th} component program for each $P(z_t)$. Let β be the behavioral program defined by $\beta(z_t) = \bar{\gamma}_t$, for all z_t . It must be shown that β satisfies the inertia assumption and is maximal.

I first show that β has the inertia property. Suppose that $\bar{\gamma}_t \neq \underset{\sim}{0}_t$, so that $\bar{\gamma}_t$ dominates $\underset{\sim}{0}_t$ in P_t . Let $z_t \in Z$ and let $\hat{\beta}(z_t)$ and $\hat{\beta}_0(z_t)$ be as in the definition of the inertia property. Let $\pi \in \Delta$ and let π_t be the restriction of π to the field generated by $\mathcal{P}'_t$. Then, $E_\pi[\hat{r}(\hat{\beta}(z_t)) - \hat{r}(\hat{\beta}_0(z_t)) | z_t \text{ occurs}] = E_{\pi_t}[\hat{r}'(\bar{\gamma}_r) - \hat{r}'(\underset{\sim}{0}_t)] > 0$, provided z_t occurs with positive probability.

I now show that for each t , there is $\bar{\pi}_t \in \Delta_t$ such that for any random program γ for P_t , $E_{\bar{\pi}_t} \hat{r}'_t(\gamma) \leq E_{\bar{\pi}_t} \hat{r}'_t(\bar{\gamma}_t)$. Let C_t be the convex hull of $\{h(\underset{\sim}{a}) : \Omega'_t \to (-\infty, \infty) | \underset{\sim}{a} \text{ is a deterministic program for } P_t\}$, where

$h(\underset{\sim}{a})(\omega') = \hat{r}'_t(\underset{\sim}{a}, E_T(\underset{\sim}{a}, \omega'))$. Let $K_t = \{x : \Omega'_t \to (-\infty,\infty) | E_\pi x > E_\pi \hat{r}'_t(\bar{\gamma}_t)$, for all $\pi \in \Delta_t\}$. Since $\bar{\gamma}_t$ is maximal for P_t , $K_t \cap C_t = \emptyset$. Therefore, by the Minkowski separation theorem, there is $\bar{\pi}_t \in \Delta_t$ such that $E_{\bar{\pi}_t} x > E_{\bar{\pi}_t} c$, for all $x \in K_t$ and $c \in C_t$. Hence, $E_{\bar{\pi}_t} c \leq E_{\bar{\pi}_t} \hat{r}'_t(\bar{\gamma}_t)$, for all $c \in C_t$.

By the independence assumption, there is $\bar{\pi} \in \Delta$ such that the restriction of $\bar{\pi}$ to the field generated by $\mathcal{P}'_t$ is $\bar{\pi}_t$, for each t and the partitions $\mathcal{P}'_0, \ldots, \mathcal{P}'_t$ of Ω' are mutually independent under $\bar{\pi}$.

I now show that β is undominated.[4] I must show that if γ is any random program for P , then $E_{\bar{\pi}}\hat{r}(\gamma) \leq E_{\bar{\pi}}\hat{r}(\beta)$. Clearly, it is sufficient to show that $E_{\bar{\pi}}\hat{r}(\underset{\sim}{a}) \leq E_{\bar{\pi}}\hat{r}(\beta)$, for any deterministic program $\underset{\sim}{a}$. Let $\underset{\sim}{a} = (\underset{\sim}{a}_0, \ldots, \underset{\sim}{a}_T)$ be fixed. Clearly, $E_{\bar{\pi}}\hat{r}(\underset{\sim}{a}) = \sum_{t=0}^{T} E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a})$, where $E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a}) = \sum_{\omega\in\Omega'} \bar{\pi}(\omega)\bar{r}_t(\underset{\sim}{a}_t, s_T(\underset{\sim}{a},\omega))$. Also, $E_{\bar{\pi}}\hat{r}(\beta) = \sum_{t=0}^{T} E_{\bar{\pi}}\hat{r}_t(\bar{\gamma}_t) = \sum_{t=0}^{T} E_{\bar{\pi}_t} \hat{r}'_t(\bar{\gamma}_t)$. Therefore, it is sufficient to prove that $E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a}) \leq E_{\bar{\pi}_t} \hat{r}'_t(\bar{\gamma}_t)$, for any t .

Let t be fixed and let $\mathcal{P}'_t$ be the partition of Ω' defined earlier. Let $\mathcal{P}^c_t$ be the join of the partitions $\mathcal{P}'_n$, for $n \neq t$. Then,
$E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a}) = \sum_{B\in\mathcal{P}'_t} \sum_{A\in\mathcal{P}^c_t} \bar{\pi}(A \cap B)E_{\bar{\pi}}[\hat{r}_t(\underset{\sim}{a}) | A \cap B]$. Let $A \in \mathcal{P}^c_t$ and $B \in \mathcal{P}'_t$. By the separation assumption, $s_T(\underset{\sim}{a},\omega)$ is the same for all $\omega \in A \cap B$. Let $s_{A\cap B} = s_T(\underset{\sim}{a},\omega)$, for any $\omega \in A \cap B$. By the independence property of $\bar{\pi}$, $\bar{\pi}(A|B) = \bar{\pi}(A \cap B)(\bar{\pi}_t(B))^{-1} = \bar{\pi}(A)$. Therefore,
$E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a}) = \sum_{B\in\mathcal{P}'_t} \bar{\pi}_t(B) \sum_{A\in\mathcal{P}^c_t} \bar{\pi}(A)\hat{r}_t(\underset{\sim}{a}_t, s_{A\cap B})$.

[4] What follows is really an application of Blackwell's (1964) principle of irrelevant information. (See Whittle (1983), p. 6.)

For $A \in \mathcal{P}_t^c$, define the deterministic program $\underset{\sim}{a}_A$ for P_t as follows. Let $E_T \in \mathcal{P}_t$. By the separation assumption, there is a unique $s_T \in E_T \cap \{s_T(\underset{\sim}{a},\omega) | \omega \in A\}$. If E_n is the member of $\mathcal{P}_{tn}$ which precedes E_T, let $\underset{\sim}{a}_A(E_n) = \underset{\sim}{a}_t(s_n)$, where s_n precedes s_T, for $n = t, \ldots, T$. Recall that Ω_t' may be identified with $\mathcal{P}_t'$. If $\omega' = B \in \mathcal{P}_t' = \Omega_t'$, then $\hat{r}_t'(\underset{\sim}{a}_A, E_T(\underset{\sim}{a}_A, \omega')) = \hat{r}_t(\underset{\sim}{a}_t, s_{A\cap B})$. It now follows that $\sum_{A\in\mathcal{P}_t^c} \bar{\pi}(A)\hat{r}_t(\underset{\sim}{a}_t, s_{A\cap B})$, considered as a function from $\Omega_t' = \mathcal{P}_t'$ to $(-\infty,\infty)$, belongs to the set C_t defined earlier when defining $\bar{\pi}_t'$. Therefore, by the defining property of $\bar{\pi}_t$, $E_{\bar{\pi}_t} \hat{r}(\bar{\gamma}_t) \geq \sum_{B\in\mathcal{P}_t'} \bar{\pi}_t(B) \sum_{A\in\mathcal{P}_t^c} \bar{\pi}(A)\hat{r}_t(\underset{\sim}{a}_t, s_{A\cap B})$ $= E_{\bar{\pi}}\hat{r}_t(\underset{\sim}{a})$, as was to be proved.

REFERENCES

Anscombe, F. J. and R. J. Aumann (1963). "A Definition of Subjective Probability," Annals of Mathematical Statistics, 34, pp. 199-205.

Aumann, Robert J. (1962). "Utility Theory without the Completeness Axiom," Econometrica, 30, pp. 445-462.

__________ (1964). "Subjective Programming," in Human Judgments and Optimality, edited by Maynard W. Shelly, II and Glenn L. Bryan. New York: John Wiley & Sons, pp. 217-242.

__________ (1964b). "Utility Theory without the Completeness Axiom: A Correction," Econometrica, 32:1-2, pp. 210-212.

Baumol, William J. and Maco Stewart (1971). "On the Behavioral Theory of the Firm," in The Corporate Economy, edited by Robin Marris and Adrian Wood. Cambridge, MA: Harvard University Press.

Blackwell, David (1964). "Memoryless Strategies in Finite-Stage Dynamic Programming," Annals of Mathematical Statistics, 35, pp. 863-865.

Blume, Marshall E. and Irwin Friend (1978). The Changing Role of the Individual Investor. New York: John Wiley & Sons.

Cyert, Richard M. and James March (1963). The Behavioral Theory of the Firm. Englewood Cliffs, NJ: Prentice-Hall.

Einhorn, Hillel J. and Robin M. Hogarth (1985). "Ambiguity and Uncertainty in Probabilistic Inference," Psychological Review, 92, pp. 433-461.

Ellsberg, Daniel (1961). "Risk, Ambiguity, and the Savage Axioms," Quarterly Journal of Economics, 75, pp. 643-669.

__________ (1963). "Reply," Quarterly Journal of Economics, 77, pp. 336-347.

Fishburn, Peter, (1970). Utility Theory for Decision Making. New York: John Wiley & Sons.

Gilboa, Itzhak and David Schmeidler (1986). "Maxmin Expected Utility with a Non-Unique Prior," Working Paper No. 16-86, Foerder Institute for Economic Research, Faculty of Social Sciences, Tel-Aviv University, Ramat-Aviv, Israel.

Hurwicz, Leonid (1951). "Some Specification Problems and Applications to Econometric Models," Econometrica, 19, pp. 343-344.

Kannai, Yakar (1963). "Existence of a Utility in Infinite Dimensional Partially Ordered Spaces," Israel Journal of Mathematics, 1, pp. 229-234.

Knight, Frank H. (1921). Risk, Uncertainty and Profit. Boston and New York: Houghton Mifflin.

Leamer, Edward E. (1985a). "Bid-Ask Spreads for Subjective Probabilities," in Bayesian Inference and Decision Techniques, edited by P. Goel and A. Zellner. New York: Elsevier.

__________ (1985b). "Econometric Metaphors," Paper presented at the Fifth World Congress of the Econometric Society, August, 1985.

Levi, Isaac (1980). The Enterprise of Knowledge. Cambridge, MA: MIT Press.

__________ (1984). Decisions and Revisions. Cambridge, England: Cambridge University Press.

Machina, Mark (1982). "Expected Utility Analysis without the Independence Axiom," Econometrica, 50, pp. 277-323.

__________ (1983). "The Economic Theory of Individual Behavior toward Risk: Theory, Evidence and New Directions," Technical Report No. 433, Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford, California.

Myerson, Roger B. (1979). "An Axiomatic Derivation of Subjective Probability and Evaluation Functions," Theory and Decision, 11, p. 339-352.

__________ (1986). "Axiomatic Foundations of Bayesian Decision Theory," Discussion Paper No. 671, Center for Mathematical Studies in Economics and Management Science, Northwestern University, Evanston, IL.

Nelson, Richard R. and Sidney Winter (1982). An Evolutionary Theory of Economic Change. Cambridge, MA: The Belknap Press.

Savage, Leonard (1954). The Foundations of Statistics. New York: John Wiley & Sons.

Schoemaker, Paul J. H. (1982). "The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations," Journal of Economic Literature, 20, pp. 529-563.

Shafer, Glenn (1976). A Mathematical Theory of Evidence. Princeton, NJ: Princeton University Press.

Smith, Cedric A. B. (1961). "Consistency in Statistical Inference and Decision," Journal of the Royal Statistical Society, Series B, 22, pp. 1-25.

Walley, Peter (1981). "Coherent Lower (and Upper) Probabilities," Department of Statistics, University of Warwick, England.

__________ (1982). "The Elicitation and Aggregation of Beliefs," Department of Statistics, University of Warwick, England.

__________ (1984). Rationality and Vagueness, Chapter I, unpublished book manuscript.

Whittle, Peter (1983). Optimization over Time, Vol. II. New York: John Wiley & Sons.

Williams, P. M. (1976). "Indeterminate Probabilities," in Formal Methods in the Methodology of Empirical Sciences, edited by Marian Przelecki, Klemens Szaniawski and Ryszard Wojcicki. Dordrecht, Holland: D. Reidel Publishing Co.

Williamson, Oliver (1975). Markets and Hierarchies. New York: The Free Press.

__________ (1985). The Economic Institutions of Capitalism. New York: The Free Press.

Yaari, Menahem E. (1985). "On the Role of 'Dutch Books' in the Theory of Choice under Risk," 1985 Nancy L. Schwartz Memorial Lecture, J. L. Kellogg Graduate School of Management, Northwestern University.