AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

UNIVERSITY
OF
WISCONSIN-
MADISON

86 21
WITHDRAWN

SSRI

Workshop Series

ORDINAL UTILITY
MODELS OF DECISION
MAKING UNDER
UNCERTAINTY

Charles F. Manski

8621

SOCIAL SYSTEMS RESEARCH INSTITUTE

24

Social Systems Research Institute

University of Wisconsin

ORDINAL UTILITY
MODELS OF DECISION
MAKING UNDER
UNCERTAINTY

Charles F. Manski

8621

July, 1986

* This research was supported under National Science Foundation grant SES-8319335. I have benefitted from very lively and useful discussions with Gary Chamberlain, Mark Gertler, Joram Mayshar, Jim Powell, John Rust, and Peter Streufert.

Abstract

This paper studies two models of rational behavior under uncertainty whose predictions are invariant under ordinal transformations of utility. The 'quantile utility' model assumes that the agent maximizes some quantile of the distribution of utility. The 'utility mass' model assumes maximization of the probability of obtaining an outcome whose utility is higher than some fixed critical value. Both models satisfy weak stochastic dominance. Lexicographic refinements satisfy strong dominance.

The study of these ordinal utility models suggests a significant generalization of traditional ideas of riskiness and risk preference. We define one action to be riskier than another if the utility distribution of the latter crosses that of the former from below. The single crossing property is equivalent to a 'minmax spread' of a random variable. With relative risk defined by the single crossing criterion, the risk preference of a quantile utility maximizer increases with the utility distribution quantile that he maximizes. The risk preference of a utility mass maximizer increases with his critical utility value.

1. Introduction

Since the early 1950s, the expected utility model has dominated applied research analyzing rational behavior under uncertainty. At the same time, decision theorists have continued to question whether this one model warrants such exclusive attention. For varying perspectives, see Arrow(1951), Savage(1954), Dreze(1974), Kahneman and Tversky(1979), Fishburn (1981), and Machina(1982).

One position is that the axiom systems implying the expected utility model are so compelling as to make any other mode of behavior under uncertainty nonrational. Deviations from expected utility maximization may occur in practice. An agent made aware of these deviations, however, would correct them.

A second view is that expected utility maximization is a central idealization, but one that is overly restrictive or flawed in particular respects. Feeling thus, numerous researchers have proposed that the expected utility model should be modified or generalized in some manner.

Occasionally, a more catholic opinion is expressed. That is, if actions are characterized by probability measures of outcomes, then we should consider rational any pattern of behavior consistent with the existence of a preference ordering on the space of these probability measures. From this perspective, expected utility maximization is not central to the study of rational behavior under uncertainty. It is simply one tightly specified model. Other models, including some that bear little resemblance to expected utility, may be theoretically enlightening and/or empirically relevant.

ORDINAL UTILITY MODELS: Adopting the last position gives one the freedom to explore ideas that seem useful but are at odds with the expected utility model. This paper is the result of such an exploration. In particular, we ask whether there exist interesting models of rational behavior under uncertainty whose predictions are invariant under ordinal transformations of utility. The predictions of the expected utility model, of course, are invariant with respect to cardinal but not ordinal transformations. Hence, to investigate ordinal utility models, one must be willing to leave expected utility behind.

We find that interesting ordinal utility models do exist (Propositions 1 and 4). One such is the 'quantile utility' model. This assumes that the agent maximizes some quantile of the distribution of utility, for example the median. A second is the 'utility mass' model. Here, the agent maximizes the probability of obtaining an outcome whose utility is higher than some fixed critical value. The quantile utility model appears not to have been proposed previously. The utility mass model dates back at least to Cramer (1930) but seems not to have been studied from an ordinal utility perspective.

LEXICOGRAPHIC REFINEMENTS OF THE MODELS: The quantile utility and utility mass models share one disagreeable feature. Their indifference classes are too large. This emerges from consideration of stochastic dominance.

Both models satisfy weak stochastic dominance, perhaps the most basic criterion of reasonableness (Propositions 2 and 5). That is, if the utility distribution induced by one action weakly dominates that induced by another, then both models predict that the first action is weakly preferred to the other. The problem is that strong dominance does not

always imply strict preference. Rather, there exist cases in which one action strictly dominates another yet the models predict the agent is indifferent between them.

Fortunately, the problem of overly large indifference classes can be resolved. Simple lexicographic refinements of the quantile utility and utility mass models shrink the indifference classes in the right way. In the refined versions of the models, strong dominance does imply strict preference.

RISK PREFERENCE: Demonstrating that ordinal utility is compatible with rational behavior under uncertainty is valuable if for no other reason than that it breaks a conventional wisdom. It turns out that the study of ordinal utility models has a further, serendipitous consequence. That is, it leads to a very significant generalization of traditional ideas of relative riskiness and risk preference.

We define one action to be riskier than another if the induced utility distribution of the latter crosses that of the former from below. And we show that the single-crossing property is equivalent to 'spreading' the less risky random utility in a natural manner (Lemma 2).

The single-crossing criterion for risk comparison is much more general than the accepted characterization of relative risk. There, actions are risk comparable only if the single-crossing property plus three other conditions hold: outcomes must be real-valued, utility must be increasing in outcome, and the actions being compared must induce outcome distributions that have the same mean. Here, the outcome space is unrestricted, utility need only be a measurable function on this space, and the outcome distributions need not be normalized in mean or in any other manner.

With relative risk defined by the single-crossing property, we show that the risk preference of a quantile utility maximizer increases with the utility distribution quantile that he maximizes (Proposition 3). Similarly, the risk preference of a utility mass maximizer increases with his critical utility value (Proposition 6). Heuristically, the higher an agent's characteristic quantile or critical utility value, the more optimistic he is.

We also find that in the case of real valued outcomes, the quantile utility model yields a very simple expression for the risk premium that makes an agent indifferent between an action inducing a distribution of outcomes and one yielding a certain outcome.

DISCLAIMERS: Usually, I would at this point get on with the paper. In reading the literature and in discussing this work with colleagues, however, I have found that the atmosphere surrounding the expected utility model is highly charged. Attitudes seem to depend on an interplay of ideology, intuition, and pragmatism. Given this, it seems prudent to make some explicit disclaimers before beginning.

First, I do not think of the ordinal utility models studied here as dominating the expected utility model. Expected utility, quantile utility, and utility mass should, I feel, be thought of as three tightly specified models. Each illuminates a possible mode of rational behavior under uncertainty.

Second, ordinal utility models may or may not describe empirical behavior better than the expected utility model does. In particular, the present work was not initiated to resolve the various well-known empirical objections to the expected utility model: Allais' paradox, the co-occurrence of gambling and insurance, anchoring, and so on. To this

author, it seems quixotic to think that any model as simple as expected utility, quantile utility, or utility mass should be able to explain such diverse and complex aspects of behavior under uncertainty. A failure to do so does not render a model valueless.

Third, the quantile utility and utility mass models do not dominate expected utility in ease of use, nor vice versa. The relative tractability of the various models depends on the application. Section 7 discusses this point.

Finally, this paper does not provide a choice theoretic axiomatic foundation for ordinal utility models. Rather, we begin with the ancient notion of preference functionals, propose the quantile and mass functionals as being worthy of attention, and then study the implications for behavior of maximization of these functionals. This approach is similar to that of Machina(1982), who examines maximization of Frechet differentiable preference functionals.

ORGANIZATION OF THE PAPER: Section 2 is technical background. It reviews basic ideas about preference functionals and extends them to lexicographic functional sequences. Section 3 introduces the quantile utility model, establishes its ordinal invariance and stochastic dominance properties, and discusses lexicographic refinements of the model. Section 4 defines relative riskiness by the single crossing property. Then Section 5 studies risk preference in the quantile utility model. The utility mass model is examined in Section 6, which covers the same ground as Sections 3 and 5. Section 7 uses the classical simultaneous search problem as an example of the application of ordinal utility models.

2. Preference Functionals on Utility Distributions

RATIONAL BEHAVIOR UNDER UNCERTAINTY: Let A be a universe of actions and let X be the space of possible outcomes for these actions. Consider an agent who must select an action from a set DCA of feasible actions. Assume that the agent endows X with a σ -algebra Ω and associates with each $a \in A$ a probability measure π_a on the measurable outcome space (X, Ω) . In this setting, we say that behavior is rational if the choice of an action from every DCA is consistent with the existence of a preference ordering on the space Π of probability measures on (X, Ω) .

CONSISTENCY OF PREFERENCES ON Π AND ON X : Strictly speaking, rationality imposes no restrictions on the agent's preference ordering on Π . We shall, however, insist that this ordering be consistent with the agent's ordering of the outcomes X .

To see what is meant by this, let $\rho \succeq \pi$ ($\rho \succ \pi$) denote the event that $\rho \in \Pi$ is weakly (strictly) preferred to $\pi \in \Pi$. Assume that the elements of X are Ω -measurable. Let $y \succeq x$ ($y \succ x$) denote the event that the probability measure with all mass at $\{y\}$ is weakly (strictly) preferred to that with all mass at $\{x\}$.

For each $y \in X$, let $X_y \equiv \{x \in X : y \succeq x\}$ denote the subset of outcomes to which y is weakly preferred. Assume that X_y is Ω -measurable. Then for each $y \in X$ and $\pi \in \Pi$, $\pi(X_y)$ is the probability under π of obtaining an outcome no better than y . For each $\pi \in \Pi$, $[\pi(X_y), y \in X]$ is the set of such probabilities. Let $\Phi^0 \equiv \{[\pi(X_y), y \in X], \pi \in \Pi\}$.

With this as background, we shall say that preferences on Π are consistent with preferences on X if two conditions are satisfied.

First, there should exist a preference ordering on the space Φ^0 . That is, for all $\pi \in \Pi$, $\rho \in \Pi$,

$$(1) \quad [\pi(X_y), y \in X] \succeq [\rho(X_y), y \in X] \iff \pi \succeq \rho.$$

Second, the ordering on Φ^0 should be such that

$$(2) \quad \pi(X_y) \leq \rho(X_y), \forall y \in X \implies [\pi(X_y), y \in X] \succeq [\rho(X_y), y \in X].$$

Condition 1 states that for purposes of decision making, the agent takes the function $[\pi(X_y), y \in X]$ mapping X into $[0,1]$ as a sufficient statistic for the probability measure π mapping Ω into $[0,1]$. It is hard to imagine why the agent should be concerned with π except through $[\pi(X_y), y \in X]$. Condition 2 requires that if π weakly dominates ρ , then π should be weakly preferred to ρ . That dominance should imply preference seems to be universally accepted as an intrinsic characteristic of rational behavior under uncertainty.

UTILITY FUNCTIONS AND PREFERENCE FUNCTIONALS: We now impose two further restrictions on the agent's preference ordering on Π . These are less basic than the consistency conditions (1) and (2) but greatly simplify the analysis of rational behavior under uncertainty.

First, assume that the agent's ordering on X is consistent with the existence of an Ω -measurable utility function u mapping X into R^1 . That is, for all $x \in X$, $y \in X$,

$$(3) \quad u(y) \geq u(x) \iff y \succeq x.$$

For all $t \in \mathbb{R}^1$, the set $\{x \in X: u(x) \leq t\}$ is Ω -measurable.

Second, assume that the ordering on Π is consistent with the existence of a preference functional V mapping the space of utility distributions into $\mathbb{R}^1$. For each $\pi \in \Pi$, let F_π be the probability distribution of utility induced by π . That is, for $t \in \mathbb{R}^1$, $F_\pi(t) \equiv \pi\{x \in X: u(x) \leq t\}$. Let $\Phi \equiv \{F_\pi, \pi \in \Pi\}$ denote the space of utility distributions induced by Π . Then V mapping Φ into $\mathbb{R}^1$ is said to be a preference functional if, for all $\pi \in \Pi$, $\rho \in \Pi$,

$$(4) \quad V(F_\pi) \geq V(F_\rho) \quad \Leftrightarrow \quad \pi \succeq \rho.$$

The existence of $u: X \rightarrow \mathbb{R}^1$ and $V: \Phi \rightarrow \mathbb{R}^1$ reduces the agent's problem of ordering the probability measures Π on (X, Ω) to one of ordering the real numbers $\{V(F), F \in \Phi\}$. Observe that for each $\pi \in \Pi$ and $y \in X$, $F_\pi[u(y)] = \pi(X_y)$. Hence, conditions (3) and (4) imply that the consistency condition (1) holds. The weak dominance Condition 2 holds if the preference functional satisfies the restriction

$$(5) \quad F_\pi[u(y)] \leq F_\rho[u(y)], \quad \forall y \in X \quad \Rightarrow \quad V(F_\pi) \geq V(F_\rho).$$

We note for later use that for all $\pi \in \Pi$, $F_\pi(*)$ can increase only on the set of points $\{u(y), y \in X\}$. Hence,

$$(6) \quad F_\pi[u(y)] \leq F_\rho[u(y)], \quad \forall y \in X \quad \Leftrightarrow \quad F_\pi(t) \leq F_\rho(t), \quad \forall t \in \mathbb{R}^1.$$

So condition (5) is equivalent to

$$(7) \quad F_\pi(t) \leq F_\rho(t), \quad \forall t \in \mathbb{R}^1 \quad \Rightarrow \quad V(F_\pi) \geq V(F_\rho).$$

Within the structure given by conditions (3), (4), and (5), models of rational behavior under uncertainty are distinguished by the conditions they impose on u and V . For example, the expected utility model assumes that the agent's utility function is such that the utility distributions Φ all have finite expectation. Given this, the model assumes that the agent orders Φ by the expectation functional. The quantile utility and utility mass models studied in this paper impose different restrictions on u and V .

PARTIAL PREFERENCE FUNCTIONALS AND LEXICOGRAPHIC FUNCTIONAL SEQUENCES:

By (4), an agent endowed with the preference functional V is indifferent between two probability measures of outcomes if and only if, for some $t \in \mathbb{R}^1$, both induce utility distributions in $\Phi_V(t) \equiv \{F \in \Phi : V(F) = t\}$. That is, the sets $\Phi_V(t), t \in \mathbb{R}^1$ are the agent's indifference classes.

In general, preference functionals imply rather large indifference classes. This should not be surprising. After all, a functional attempts to summarize an entire distribution by a single number. Hence, the ability of one functional to discriminate among distributions is inherently limited.

If the sets $\Phi_V(t), t \in \mathbb{R}^1$ seem too large to be indifference classes, one may prefer to interpret them as sets of distributions that are unordered by V . Formally, this amounts to replacing condition (4) with the weaker statement

$$(8) \quad V(F_\pi) > V(F_\rho) \quad \Rightarrow \quad \pi > \rho.$$

When (8) holds but not (4), we say V is a partial preference functional.

Let V_1 be a partial preference functional. Then how does the agent order the subsets of Π that are unordered by V_1 ? One possibility is that the agent is endowed with a sequence $\{V_n, n=1, \dots, \infty\}$ of functionals on Φ . He begins by partially ordering distributions by their V_1 values. For each $t \in \mathbb{R}^1$, he partially orders the distributions in $\Phi_{V_1}(t)$ by their V_2 values. Continuing, he lexicographically shrinks the sets of unordered distributions by considering in turn V_3, V_4 , and so on.

The indifference classes of an agent endowed with a functional sequence are those sets of distributions which remain unordered in the limit of the lexicographic ordering process. These indifference classes may be small or large, depending on the sequence. At the extreme, they are singletons. This occurs if $\{V_n, n=1, \dots, \infty\}$ completely characterizes the distributions in Φ .

Formally, assume that for each $F, G \in \Phi$, $F \neq G$, there exists a finite n such that $V_n(F) \neq V_n(G)$. Then for each $\pi, \rho \in \Pi$, $\pi \neq \rho$,

$$(9) \quad \inf\{n: V_n(F_\pi) > V_n(F_\rho)\} < \inf\{n: V_n(F_\rho) > V_n(F_\pi)\} \Leftrightarrow \pi > \rho.$$

Thus, an agent endowed with a functional sequence that characterizes Φ has a strict preference ordering on Π .

3. The Quantile Utility Functionals

THE QUANTILE UTILITY MODEL: Given any distribution function F on the real line and any $\alpha \in (0,1)$, the α -quantile of F is defined as

$$(10) \quad Q_{\alpha}(F) \equiv [\inf t: F(t) \geq \alpha].$$

We shall say that an agent maximizes quantile utility if, for some fixed $\alpha \in (0,1)$ and for all $F, G \in \mathcal{F}$,

$$(11) \quad Q_{\alpha}(F) \geq Q_{\alpha}(G) \quad \Leftrightarrow \quad F \geq G.$$

The right continuity of distribution functions implies that quantiles always exist and are finite. Hence, quantile utility maximization is always a well-defined decision rule.

In this Section, we establish properties of the quantile utility model that hold for all $\alpha \in (0,1)$. We also consider a lexicographic refinement of the model. In Section 5, we compare the behavior of agents maximizing different quantiles of the utility distribution.

DISTRIBUTIONS AND QUANTILES: For whatever reasons, the basic facts about quantiles are less widely familiar than are, say, those about expectations. We therefore begin by reviewing the relationships between probability distributions and their quantiles.

For each fixed distribution function F , the quantile is a real-valued function on the interval $(0,1)$. Right continuity of distribution functions implies the following relationship between F and its quantiles:

$$(12) \quad F(t) \geq \alpha \iff Q_\alpha(F) \leq t.$$

By the definition of quantiles, $F[Q_\alpha(F)] \geq \alpha$ for all $\alpha \in (0,1)$. If F is continuous at its α -quantile, then $F[Q_\alpha(F)] = \alpha$. But this equality need not hold if $Q_\alpha(F)$ is a point of discontinuity.

For each fixed α , the quantile is a functional on the space of distribution functions on the real line. The following Lemma characterizes how the α -quantile varies on this space.

Lemma 1: Let F and G be any two distribution functions on the real line. Let $\alpha \in (0,1)$. Then

$$(13) \quad \exists t \in \mathbb{R}^1 \text{ s.t. } F(t) < \alpha \leq G(t) \iff Q_\alpha(F) > Q_\alpha(G).$$

PROOF: First prove sufficiency. By (12), $F(t) < \alpha \Rightarrow Q_\alpha(F) > t$ and $G(t) \geq \alpha \Rightarrow t \geq Q_\alpha(G)$. So $Q_\alpha(F) > t \geq Q_\alpha(G)$.

Now prove necessity. Assume there exists no t such that $F(t) < \alpha \leq G(t)$. If there exists t such that $F(t) \geq \alpha > G(t)$, then $Q_\alpha(G) > t \geq Q_\alpha(F)$, by the proof of sufficiency. The remaining possibility is that there exists an $s \in \mathbb{R}^1$ such that $[F(t) < \alpha \cap G(t) < \alpha]$ for $t < s$ and $[F(t) \geq \alpha \cap G(t) \geq \alpha]$ for $t \geq s$. If so, then $Q_\alpha(F) = Q_\alpha(G) = s$.

Q.E.D.

Note that two distribution functions can have the same α -quantiles and yet not have the same values there. That is, $Q_\alpha(F) = Q_\alpha(G)$ does not imply that $F[Q_\alpha(F)] = G[Q_\alpha(G)]$. On the other hand, $Q_\alpha(F) = Q_\alpha(G)$ does imply $F[Q_\alpha(F)] = G[Q_\alpha(G)]$ if both F and G are continuous at the common

α -quantile point.

INVARIANCE OF PREFERENCES TO ORDINAL TRANSFORMATIONS OF UTILITY: Our first substantive task is to show that under the quantile utility model, an agent's preference ordering on Π is invariant with respect to ordinal transformations of his utility function. This simple but central result is given in Proposition 1.

Proposition 1: Let $u: X \rightarrow R^1$ be a utility function. Let $m: R^1 \rightarrow R^1$ be any strictly increasing function; thus $m[u(*)]$ is an ordinal transformation of u . For each $\pi \in \Pi$, let F_π denote the distribution of u induced by π and let H_π be the induced distribution of $m[u(*)]$. Then for all $\alpha \in (0,1)$ and all $\pi, \rho \in \Pi$,

$$(14) \quad Q_\alpha(F_\pi) \geq Q_\alpha(F_\rho) \iff Q_\alpha(H_\pi) \geq Q_\alpha(H_\rho). \blacksquare$$

PROOF: Let F be any distribution on R^1 , let u be a random variable distributed F , and let H be the distribution of $m(u)$. For all $t \in R^1$,

$$(15) \quad H[m(t)] = F(t).$$

It follows from this and from the definition of quantiles that

$$(16) \quad Q_\alpha(H) = m[Q_\alpha(F)].$$

for all $\alpha \in (0,1)$. Relationship (14) follows immediately.

Q.E.D.

DOMINANCE: Next we verify that the quantile utility model satisfies the weak dominance condition (7). We also characterize the model's strong dominance properties. Proposition 2 gives these results.

Proposition 2: Let $u: X \rightarrow \mathbb{R}^1$ be a utility function, let $\alpha \in (0,1)$, and let $\pi, \rho \in \Pi$. Assume that $F_\pi(t) \leq F_\rho(t)$, $\forall t \in \mathbb{R}^1$. Then

- (a) $Q_\alpha(F_\pi) \geq Q_\alpha(F_\rho)$.
- (b) $Q_\alpha(F_\pi) > Q_\alpha(F_\rho)$ if $F_\pi[Q_\alpha(F_\rho)] < \alpha$.
- (c) $Q_\alpha(F_\pi) = Q_\alpha(F_\rho)$ if $F_\pi[Q_\alpha(F_\rho)] \geq \alpha$. ■

PROOF: (a) By assumption, there exists no t such that $F_\rho(t) < \alpha \leq F_\pi(t)$.

By Lemma 1 then, $Q_\alpha(F_\rho)$ cannot exceed $Q_\alpha(F_\pi)$.

(b) By assumption, $F_\pi[Q_\alpha(F_\rho)] < \alpha$. But $\alpha \leq F_\pi[Q_\alpha(F_\pi)]$. Hence,
 $Q_\alpha(F_\pi) > Q_\alpha(F_\rho)$.

(c) $F_\pi[Q_\alpha(F_\rho)] \geq \alpha \Rightarrow Q_\alpha(F_\pi) \leq Q_\alpha(F_\rho)$. But Part (a) showed that
 $Q_\alpha(F_\pi) \geq Q_\alpha(F_\rho)$. Hence, $Q_\alpha(F_\pi) = Q_\alpha(F_\rho)$.

Q.E.D.

Part (a) of Proposition 2 shows that the quantile utility model satisfies a basic criterion of reasonableness. An agent maximizing quantile utility weakly prefers one probability measure of outcomes to another if the distribution of utility induced by the first weakly dominates that induced by the second. On the other hand, parts (b) and (c) of the Proposition point out a weakness of the model. That is, strong dominance does not always imply strict preference.

A distribution F on the real line is said to strongly dominate another distribution G if F weakly dominates G and there exists some $t \in \mathbb{R}^1$ such that $F(t) < G(t)$. Assume that the utility distribution F_π

strongly dominates the distribution F_ρ . It is reasonable to predict that π should then be strictly preferred to ρ , not just weakly preferred.

The quantile utility model does not guarantee this. Parts (b) and (c) of Proposition 2 show that given weak dominance, strict preference holds only in the presence of the additional condition $F_\pi[Q_\alpha(F_\rho)] < \alpha$. This condition implies strong dominance, since $F_\rho[Q_\alpha(F_\rho)] \geq \alpha$, but not vice versa. That is, it is possible that F_π strongly dominates F_ρ yet $F_\pi[Q_\alpha(F_\rho)] \geq \alpha$. If so, then an α -quantile maximizer is indifferent between π and ρ .

Figures 1 and 2 provide a pair of illustrations. In Figure 1, $F(t)$ lies strictly below $G(t)$ at every $t \in \mathbb{R}^1$ except at point s , where they are tangent. Both F and G are continuous distributions. In this setting, quantile utility maximizer strictly prefers F to G unless the quantile maximized is $F(s)=G(s)$. Then he is indifferent between F and G .

In Figure 2, $F(t)$ lies strictly below $G(t)$ everywhere. Both F and G are discontinuous at the point s and continuous elsewhere. In this setting, a quantile utility maximizer strictly prefers F to G if the quantile maximized is above $F(s)$ or below $g(s)$, the limit of $G(t)$ as t approaches s from the left. He is indifferent if the quantile maximized is in the interval $[g(s), F(s)]$.

THE SEQUENTIAL QUANTILE UTILITY MODEL: The fact that strong dominance does not guarantee strict preference indicates that the indifference classes of the quantile utility model are too large to be reasonable. The problem is that an agent maximizing the α -quantile of utility is unconcerned with the shapes of utility distributions above and below their α -quantile points. Among the set of distributions whose

α -quantiles equal a common value t , the agent does not compare the shapes of these distributions conditional on the event $u \leq t$ or on the event $u > t$. Instead, he is indifferent among them all.

Given this, one may prefer to view quantile utility not as a preference functional but as a partial preference functional. If so, one reinterprets the indifference classes of the quantile utility model as unordered sets of utility distributions. One may then consider the possibility that these sets are ordered by some lexicographic process.

An obvious candidate process is a sequential quantile utility model. Here, we assume that the agent is endowed with a sequence of quantiles. He begins by partially ordering distributions by their α_1 -quantiles. For each $t \in \mathbb{R}^1$, he partially orders the distributions whose α_1 -quantile equals t by α_2 -quantile, and so on. If the sequence $\alpha_1, \alpha_2, \dots, \infty$ characterizes the distributions in Φ , the agent strictly orders Π . In particular, for $\pi, \rho \in \Pi$, $\pi \neq \rho$,

$$(17) \quad \inf[n: W_{\alpha_n}^{\pi}(F_{\pi}) > W_{\alpha_n}^{\rho}(F_{\rho})] < \inf[n: W_{\alpha_n}^{\rho}(F_{\rho}) > W_{\alpha_n}^{\pi}(F_{\pi})] \Leftrightarrow \pi > \rho.$$

Under the sequential quantile utility model, preferences continue to be invariant to ordinal transformations of utility and to satisfy the weak dominance condition. Moreover, if the quantile sequence characterizes Φ , then strong dominance implies strict preference.

4. Riskiness

THE SINGLE-CROSSING CRITERION: To define risk preference, we need to choose a criterion by which one utility distribution will be termed riskier than another. Henceforth, we say that $G \in \Phi$ is riskier than $F \in \Phi$ if F crosses G from below. That is, G is riskier than F if the following holds:

$$(18) \quad \exists s \in R^1 \text{ s.t. } F(t) \leq G(t), \forall t < s \quad \cap \quad F(t) \geq G(t), \forall t > s.$$

The single-crossing property is preserved under ordinal transformations of the agent's utility function. To see this, let $\pi, \rho \in \Pi$ and assume that F_π crosses F_ρ from below. Let $m: R^1 \rightarrow R^1$ be any strictly increasing function. Let H_π and H_ρ be the induced distributions of the ordinally transformed utility function $m[u(*)]$. Then (15) and (18) imply that

$$(19) \quad H_\pi(t) \leq H_\rho(t), \forall t < m(s) \quad \cap \quad H_\pi(t) \geq H_\rho(t), \forall t > m(s).$$

Figure 3 depicts a pair of continuous utility distributions satisfying the single-crossing property. In Figure 4, G is continuous and F is degenerate. Observe that if F is degenerate, then F crosses every distribution G from below. Let F have all its mass at $s \in R^1$. Then whatever distribution G is,

$$(20) \quad 0 = F(t) \leq G(t), \forall t < s \quad \cap \quad 1 = F(t) \geq G(t), \forall t \geq s.$$

Note the following curiosity. If F is degenerate at s_1 and G is degenerate at s_2 , then F crosses G from below at s_1 and G crosses F from below at s_2 .

RISKINESS IN THE EXPECTED UTILITY MODEL: The single-crossing criterion is considerably more general than the relative riskiness criterion used in applications of the expected utility model. There, outcome measure $\rho \in \Pi$ is said to be riskier than $\pi \in \Pi$ if the following conditions hold:

$$(21a) \quad X \subset R^1$$

$$(21b) \quad x \geq y \iff u(x) \geq u(y)$$

$$(21c) \quad \exists s \in R^1 \text{ s.t. } \pi(x < t) \leq \rho(x < t), \forall t < s \quad \cap \quad \pi(x < t) \geq \rho(x < t), \forall t > s.$$

$$(21d) \quad \int x d\rho = \int x d\pi.$$

See Rothschild and Stiglitz (1970).

Conditions (21a) and (21b) imply that the utility function $u: X \rightarrow R^1$ transforms X ordinally. Thus, $u(x) \equiv x$ is itself a representation of the ordinal utility function. With u specified in this way, each $\pi \in \Pi$ and its induced utility distribution F_π are the same. So the expected-utility analysis of riskiness restricts attention to decision problems in which outcomes and ordinal utilities are synonymous. Given this, Condition (21c) is the same as the single-crossing property (18).

The remaining condition, (21d), states that if two probability measures have different expectations, then they are noncomparable in riskiness. This restriction is necessary in the context of the expected utility model but has no relevance in ordinal utility models of decision making.

MINMAX SPREADS: Some of the appeal of condition (21) as a criterion for risk comparison derives from its equivalence to the operation of adding unbiased noise to a random variable. That is, given (21a) and (21b), (21c) and (21d) hold if and only if there exists a bivariate random variable (δ, ϵ) such that δ is distributed π , $E(\epsilon|\delta)=0$, and $\delta+\epsilon$ is distributed ρ . See Rothschild and Stiglitz (1970).

It is of interest to report that the single-crossing condition (18) is itself equivalent to a noise-increasing operation, provided only that F equals G at the crossing point s . Given that $F(s)=G(s)$, Lemma 2 shows that F crosses G from below if and only if there exists a bivariate random variable (δ, ϵ) such that δ is distributed F and the 'minmax spread' $\min(\delta, \epsilon) * 1[\delta \leq s] + \max(\delta, \epsilon) * 1[\delta > s]$ is distributed G . It seems appropriate to say that the minmax spread operation increases the noisiness of δ . This operation makes small values of δ smaller and large ones larger, small and large being defined relative to s .

Lemma 2: (a) Let (δ, ϵ) be a bivariate random variable, with δ distributed F . Fix $s \in \mathbb{R}^1$. Let G denote the distribution of $\min(\delta, \epsilon) * 1[\delta \leq s] + \max(\delta, \epsilon) * 1[\delta > s]$. Then F crosses G from below at s and $G(s)=F(s)$.

(b) Let F cross G from below at some $s \in \mathbb{R}^1$. Assume that $F(s)=G(s)$. Then there exists a bivariate random variable (δ, ϵ) such that δ is distributed F and $\min(\delta, \epsilon) * 1[\delta \leq s] + \max(\delta, \epsilon) * 1[\delta > s]$ is distributed G . ■

PROOF: (a) Let $t \leq s$. Then

$$G(t) = \Pr[\delta \leq s, (\delta \vee \epsilon) \leq t] = \Pr(\delta \leq t) + \Pr(\delta \leq s, \epsilon \leq t) - \Pr(\delta \leq t, \epsilon \leq t) \geq F(t).$$

Moreover, $G(s)=F(s)$. Let $t>s$. Then

$$1-G(t) = \Pr[\delta>s, (\delta\vee\epsilon)>t] = \Pr(\delta>t) + \Pr(\delta>s, \epsilon>t) - \Pr(\delta>t, \epsilon>t) > 1-F(t).$$

(b) The proof is constructive. Let (δ, ϵ) be a bivariate random variable whose joint distribution is

$$\Pr(\delta\leq r, \epsilon\leq t) = \min[F(r), G(t)], \quad \forall r, t \in \mathbb{R}^1.$$

This is a valid bivariate distribution. Ord(1972) reports that it first appears in Frechet(1951). The marginal distribution of δ is F and that of ϵ is G . The probability that (δ, ϵ) falls in a rectangle $(r_0, r_1] \times (t_0, t_1]$ is

$$\begin{aligned} \Pr(r_0 < \delta \leq r_1, t_0 < \epsilon \leq t_1) &= \Pr(\delta \leq r_1, \epsilon \leq t_1) + \Pr(\delta \leq r_0, \epsilon \leq t_0) \\ &\quad - \Pr(\delta \leq r_1, \epsilon \leq t_0) - \Pr(\delta \leq r_0, \epsilon \leq t_1) \\ &= \min[F(r_1), G(t_1)] + \min[F(r_0), G(t_0)] \\ &\quad - \min[F(r_1), G(t_0)] - \min[F(r_0), G(t_1)]. \end{aligned}$$

Now let $\eta \equiv \min(\delta, \epsilon) * 1[\delta \leq s] + \max(\delta, \epsilon) * 1[\delta > s]$. Then for $t \leq s$,

$$\begin{aligned} \Pr(\eta \leq t) &= \Pr(\delta \leq t) + \Pr(\delta \leq s, \epsilon \leq t) - \Pr(\delta \leq t, \epsilon \leq t) \\ &= F(t) + \min[F(s), G(t)] - \min[F(t), G(t)]. \end{aligned}$$

For $t > s$,

$$\Pr(\eta > t) = \Pr(\delta > t) + \Pr(\delta > s, \epsilon > t) - \Pr(\delta > t, \epsilon > t)$$

$$\begin{aligned}
&= \{1-F(t)\} + \{1+\min[F(s), G(t)]-G(t)-F(s)\} \\
&\quad - \{1+\min[F(t), G(t)]-G(t)-F(t)\}.
\end{aligned}$$

The assumptions that F crosses G at s and that $F(s)=G(s)$ imply that $F(t) \leq G(t) \leq F(s)$ for $t \leq s$ and that $F(s) \leq G(t) \leq F(t)$ for $t > s$. Hence, η is distributed G .

Q.E.D.

5. Risk Preference in the Quantile Utility Model

Let $\Psi \equiv \{(F, G) \in \Phi \times \Phi : F \text{ crosses } G \text{ from below}\}$. We shall say that one agent is more risk preferring than another if, for all $(F, G) \in \Psi$, the first agent weakly prefers G to F whenever the second does.

Within the class of quantile utility models, an agent's preferences are determined by his ordinal utility function and by his characteristic quantile. Consider two agents whose quantiles differ. Proposition 3 shows that the agent maximizing the higher quantile is the more risk-prefering.

Proposition 3: Let $(F, G) \in \Psi$ cross at $s \in \mathbb{R}^1$. Then

$$(22a) \quad Q_{\alpha}(F) \geq Q_{\alpha}(G) \text{ if } \alpha \leq G(s).$$

$$(22b) \quad Q_{\alpha}(F) > Q_{\alpha}(G) \text{ if } \alpha \leq G(s) \text{ and } F[Q_{\alpha}(G)] < \alpha.$$

$$(22c) \quad Q_{\alpha}(F) = Q_{\alpha}(G) \text{ if } \alpha \leq G(s) \text{ and } F[Q_{\alpha}(G)] \geq \alpha.$$

$$(23a) \quad Q_{\alpha}(G) \geq Q_{\alpha}(F) \text{ if } \alpha > G(s).$$

(23b) $Q_\alpha(G) > Q_\alpha(F)$ if $\alpha > G(s)$ and $G[Q_\alpha(F)] < \alpha$.

(23c) $Q_\alpha(G) = Q_\alpha(F)$ if $\alpha > G(s)$ and $G[Q_\alpha(F)] \geq \alpha$. ■

PROOF: By the single-crossing property and the right continuity of probability distributions,

(24) $t < s \Rightarrow F(t) \leq G(t). \quad t \geq s \Rightarrow F(t) \geq G(t).$

Let $\alpha \leq G(s)$. Then $Q_\alpha(G) \leq s$. By (24), $F(t) \leq G(t) < \alpha$ for all $t < Q_\alpha(G)$. It follows that $Q_\alpha(F) \geq Q_\alpha(G)$. So (22a) holds. Given this, (22b) and (22c) may be proved in the same manner as parts (b) and (c) of Proposition 2.

Let $\alpha > G(s)$. Then $Q_\alpha(G) > s$. By (24), $F(t) \geq G(t) \geq \alpha$ for all $t \geq Q_\alpha(G)$. It follows that $Q_\alpha(F) \leq Q_\alpha(G)$. So (23a) holds. Given this, (23b) and (23c) may be proved in the same manner as parts (b) and (c) of Proposition 2.

Q.E.D.

The Proposition shows that given any pair of utility distributions (F, G) such that F crosses G at s , a quantile utility maximizer weakly prefers the less risky distribution F if $\alpha \leq G(s)$ and weakly prefers the more risky one G otherwise. Thus, risk preference is weakly increasing in the quantile the agent maximizes.

Risk preference does not necessarily strictly increase with quantile. In Figure 3, $G(s) < \alpha_1 < \alpha_2$. An agent maximizing the α_1 -quantile strictly prefers G to F but one maximizing the α_2 -quantile is indifferent.

RISK PREFERENCE IN THE EXPECTED UTILITY MODEL: The present notion of

risk preference differs in several respects from the one familiar in the expected utility literature. We have already pointed out that the single-crossing condition does not restrict attention to mean-preserving spreads of real-valued outcomes. Here, we note three further distinctions.

First, we do not characterize agents as risk-averse, risk-neutral, or risk-seeking. In some applications, it may be useful to select a particular quantile, say the median, as defining risk neutrality. There is, however, no compelling theoretical reason to do so.

Second, the risk preference of a quantile utility maximizer is defined without reference to his preference ordering on X . Consider two agents endowed with utility-quantile pairs (u_1, α_1) and (u_2, α_2) respectively. We say that agent two is more risk-preferring than agent one if $\alpha_2 > \alpha_1$, whatever u_1 and u_2 may be. In contrast, the expected utility analysis of riskiness presumes that the two agents have the same ordinal utility function, given by condition (21b). There, risk preference is characterized by the agent's cardinal transformation of x into $u(x)$.

It is important to understand why, in the present analysis, agents with different preference orderings on X can be compared in terms of risk preference. We define riskiness as a property of utility distributions rather than outcome measures. An agent's ordering of X determines the utility distribution F_π induced by a given $\pi \in \Pi$. Once F_π is specified, however, its derivation is irrelevant.

Third, the present notion of risk preference is global rather than local. If $\alpha_2 > \alpha_1$, then the second agent is always more risk-preferring than the first. In contrast, in the expected utility model, the relative risk preference of two agents may vary. For example, it may be that u_2 is convex where u_1 is concave and vice versa. If so, then the

second agent is more risk-preferring than the first in some settings and less risk-preferring in others.

RISK PREMIUMS: Assume that Conditions (21a) and (21b) hold and that F is a degenerate utility distribution. In this special but important case, we can obtain a result much sharper than Proposition 3. Specifically, we can characterize an agent's risk preference by the size of the risk premium that would make him indifferent between F and any other distribution G . It turns out that the risk premium of a quantile utility maximizer depends in a very simple manner on the mass point of F , on the form of G , and on the quantile maximized.

Let F have all its mass at the point $s \in \mathbb{R}^1$. Under the quantile utility model, the risk premium is the value $p_\alpha(s, G) \in \mathbb{R}^1$ that solves the equation

$$(25) \quad u[s - p_\alpha(s, G)] = Q_\alpha(G).$$

Recall that under (21a) and (21b) we can, without loss of generality, let $u(x) = x$ be the ordinal utility function. This done, the risk premium has the explicit form

$$(26) \quad p_\alpha(s, G) = s - Q_\alpha(G).$$

Thus, for every s , G , and α , there exists a unique, finite risk premium. With utility measured in units of x , this premium equals the mass point of F minus the α -quantile of G .

QUANTILES AS BEST PREDICTORS OF UTILITY: To conclude this discussion, we call attention to a fact about quantiles that may enhance the intuitive

appeal of Proposition 3. It is well-known that if u is a real-valued random variable distributed F , then the α -quantile of F is a best predictor of u , in the sense of solving the problem

$$(27) \quad \min_{t \in \mathbb{R}^1} \int [(1-\alpha)|u-t|*1[u < t] + \alpha|u-t|*1[u > t]] dF.$$

That is, the α -quantile minimizes expected loss under the asymmetric absolute loss function that weighs overpredictions to underpredictions in the ratio $(1-\alpha)/\alpha$.

Thus, we may interpret an α -quantile utility maximizer as summarizing each $F \in \mathcal{F}$ by the best predictor $Q_\alpha(F)$ and ranking utility distributions in order of this best predictor. The higher α is, the less heavily the agent weighs overpredictions relative to underpredictions. In other words, the higher α is, the more the agent is concerned with the upper tails of utility distributions relative to their lower tails.

Assume that distribution F crosses distribution G from below. Then F has a more favorable lower tail than G and a less favorable upper tail. So the best predictor property of quantiles suggests that the agent should prefer G if α is high enough and F otherwise. That is precisely what Proposition 3 shows.

Note that expected utility also has a best predictor interpretation: the expectation of u minimizes expected square loss. Problem (27) suggests a generalization of the expected utility model. That is, one might view expected utility as a member of the class of preference functionals that minimize expected loss under asymmetric square loss functions. These solve the problems

$$(28) \min_{t \in \mathbb{R}^1} \int [(1-\alpha)(u-t)^2 * 1[u < t] + \alpha(u-t)^2 * 1[u > t]] dF$$

for $\alpha \in (0,1)$. The solutions to (28) have been studied recently by Newey and Powell (1986), who term them 'expectiles'.

5. The Utility Mass Functionals

THE UTILITY MASS MODEL: Given any distribution F on the real line and any $\tau \in \mathbb{R}^1$, we shall define the τ -mass of F by

$$(29) \quad R_{\tau}(F) \equiv 1 - F(\tau).$$

That is, the τ -mass of F is the probability under F of obtaining a realization larger than τ . We shall say that an agent maximizes utility mass if, for some fixed $\tau \in \mathbb{R}^1$ and for all $F, G \in \Phi$,

$$(30) \quad R_{\tau}(F) \geq R_{\tau}(G) \iff F \geq G.$$

Maximization of utility mass is always a well-defined decision rule.

Marschak (1950) reported that the utility mass model had been proposed by Cramer (1930). Arrow (1951) noted that the mass functional can be written in the form of an expectation, that is

$$(31) \quad R_{\tau}(F) = \int 1[u > \tau] dF.$$

From this, he concluded that the utility mass model is a special case of the expected utility model, the cardinal utility function being the

indicator function $1[u(*) > \tau]$. As a cardinal utility function, the indicator function appears rather uninteresting. This may account for the fact that the utility mass model seems not to have been scrutinized closely by researchers working within the expected utility paradigm.

In fact, the utility mass model differs considerably from the generic expected utility model. The preference ordering of an agent maximizing utility mass is invariant to ordinal transformations of u and τ . It is also invariant to ordinal transformations of the indicator function $1[u(*) > \tau]$. In its ordinality and in other respects, the properties of the utility mass model are very similar to those of the quantile utility model. This should not be surprising; the quantile and mass functionals are closely related.

In what follows, we establish the main properties of the utility mass model. We cover much the same ground as Sections 3 and 4 did for the quantile utility model. The present discussion, however, is much briefer. Results that parallel Propositions 1, 2, and 3 are immediate consequences of the definition of the utility mass functional.

INVARIANCE OF PREFERENCES TO ORDINAL TRANSFORMATIONS OF UTILITY: This basic property of the utility mass model follows immediately from (15). Proposition 4 states the result.

Proposition 4: Let $u: X \rightarrow \mathbb{R}^1$ be a utility function. Let $m: \mathbb{R}^1 \rightarrow \mathbb{R}^1$ be any strictly increasing function. For $\pi \in \Pi$, let F_π denote the distribution of u induced by π and let H_π be the induced distribution of $m[u(*)]$. Then for all $\tau \in \mathbb{R}^1$ and all $\pi, \rho \in \Pi$,

$$(32) \quad R_\tau(F_\pi) \geq R_\tau(F_\rho) \iff R_{m(\tau)}(H_\pi) \geq R_{m(\tau)}(H_\rho). \blacksquare$$

DOMINANCE: The utility mass model satisfies the weak dominance condition, trivially. As with the quantile utility model, strong dominance does not always imply strict preference. Proposition 5 gives these results.

Proposition 5: Let $u: X \rightarrow \mathbb{R}^1$ be a utility function, let $\tau \in \mathbb{R}^1$ and let $\pi, \rho \in \Pi$. Assume that $F_\pi(t) \leq F_\rho(t)$, $\forall t \in \mathbb{R}^1$. Then

- (a) $R_\tau(F_\pi) \geq R_\tau(F_\rho)$.
- (b) $R_\tau(F_\pi) > R_\tau(F_\rho)$ if $F_\pi(\tau) < F_\rho(\tau)$.
- (c) $R_\tau(F_\pi) = R_\tau(F_\rho)$ if $F_\pi(\tau) = F_\rho(\tau)$. ■

THE SEQUENTIAL UTILITY MASS MODEL: Utility mass indifference classes are clearly too large to be reasonable. Among the set of utility distributions whose τ -masses equal a common value, the agent is unconcerned with the shapes of these distributions conditional on $u \leq \tau$ or on $u > \tau$. In particular, he is indifferent among distributions whose τ -mass equals zero and also among those with τ -mass one. These are, respectively, the classes of distributions bounded above and below by τ .

It is probably best to consider utility mass a partial preference functional. Then distributions with equal τ -utility mass are unordered rather than equivalent. Such distributions may be ordered by a lexicographic process.

One such is the sequential utility mass model. This assumes that the agent is endowed with a sequence of mass points. He begins by partially ordering distributions by their τ_1 -masses. For each $p \in [0, 1]$, he partially orders the distributions whose τ_1 -mass equals p by τ_2 -mass,

and so on. If the sequence $\tau_1, \tau_2, \dots, \infty$ characterizes Φ , the agent strictly orders Π . That is, for $\pi, \rho \in \Pi$, $\pi \neq \rho$,

$$(33) \quad \inf[n: R_{\tau_n}(F_\pi) > R_{\tau_n}(F_\rho)] < \inf[n: R_{\tau_n}(F_\rho) > R_{\tau_n}(F_\pi)] \Leftrightarrow \pi > \rho.$$

The properties of the sequential utility mass model are like those of the sequential quantile utility model. Preferences are invariant to ordinal transformations of utility and satisfy the weak dominance condition. If the mass point sequence characterizes Φ , then strong dominance implies strict preference.

RISK PREFERENCE: As before, let the relative riskiness of distributions be defined by the single-crossing condition (18). If two agents are both utility mass maximizers, risk preference increases weakly with the point τ whose utility mass is maximized. This follows immediately from the definition of the mass functional. Proposition 6 states the result.

Proposition 6: Let $(F, G) \in \Psi$ cross at $s \in \mathbb{R}^1$. Then

$$(34a) \quad R_\tau(F) \geq R_\tau(G) \text{ if } \tau < s.$$

$$(34b) \quad R_\tau(F) > R_\tau(G) \text{ if } \tau < s \text{ and } F(\tau) < G(\tau).$$

$$(34c) \quad R_\tau(F) = R_\tau(G) \text{ if } \tau < s \text{ and } F(\tau) = G(\tau).$$

$$(35a) \quad R_\tau(G) \geq R_\tau(F) \text{ if } \tau \geq s.$$

$$(35b) \quad R_\tau(G) > R_\tau(F) \text{ if } \tau \geq s \text{ and } F(\tau) > G(\tau).$$

$$(35c) \quad R_\tau(G) = R_\tau(F) \text{ if } \tau \geq s \text{ and } F(\tau) = G(\tau). \blacksquare$$

6. An Example

Accepting that ordinal utility models have some appeal in principle, we should ask how easily they may be applied. In particular, are the quantile utility and utility mass models more or less tractable than expected utility maximization? There is no uniform answer. The relative tractability of different models varies across applications.

Compound lotteries, that is mixtures, are analyzed more readily with the expected utility model. The expectation is a linear functional: hence, the law of iterated expectations holds. Quantiles and mass functionals do not iterate: so compound lotteries do not reduce in any nice way to simple lotteries.

Problems involving the order statistics of a random sample of observations are treated more easily by ordinal utility models. The expectation of an order statistic depends in a complex manner on the sample size and on the sampling distribution. The quantiles and masses of an order statistic are simple functions of these factors. We give an example.

THE SIMULTANEOUS SEARCH PROBLEM: Ordinal utility models are particularly well-suited to the analysis of the familiar simultaneous search problem. Let H be a probability distribution on the positive real half-line. Assume that H is strictly increasing on its support and that H is continuous. Let w be a random variable distributed H .

We consider an agent who can make independent draws of w , at a cost per draw of c , where c is in units of w . Let w_i denote the i^{th}

realization of w . Define $w_0 = 0$. Conditional on making n draws, the agent's realized utility is, up to an ordinal transformation,

$$(36) \quad x(n) \equiv \max(w_i, 0=1, \dots, n) - cn.$$

His decision problem is to choose a number of draws.

To characterize decision rules, we need the probability distribution of $x(n)$, denoted F_n . For $n=0$, F_n is degenerate with all its mass at zero. For n positive and $t \in \mathbb{R}^1$,

$$(37) \quad F_n(t) = \prod_{i=1}^n \Pr(w_i \leq t+cn) = H(t+cn)^n.$$

By assumption, H is continuous, strictly increasing on its support. It follows from (37) that F_n is continuous, strictly increasing on its own support; that is, the support of H translated cn units to the left.

For each $\tau \in \mathbb{R}^1$, the τ -utility mass of F_n is

$$(38) \quad R_\tau(F_n) = 1 - H(\tau+cn)^n.$$

For each $\alpha \in (0,1)$, the α -quantile of F_n is the unique value of t solving the equation $H(t+cn)^n = \alpha$. That is,

$$(39) \quad Q_\alpha(F_n) = H^{-1}(\alpha^{1/n}) - cn.$$

An agent maximizing τ -utility mass or α -quantile utility selects a positive integer n that maximizes (38) or (39), respectively. He then compares this trial maximum with the null option of making no draws.

Whichever ordinal utility model holds, analysis of the maximization

problem is straightforward. In contrast, analysis under the expected utility model is difficult. There, one must contend with the fact that the expectation of the maximum of n independent draws from H depends in a complex manner on n and on H .

Figure 1

Figure 2

Figure 3

Figure 4

References

- Arrow, K.(1951), "Alternative Approaches to the Theory of Choice in Risk Taking Situations", Econometrica 19, 404-437.
- Cramer, H.(1930), "On the Mathematical Theory of Risk", Forsakring-saktiebolaget Skandias Festskrift, Stockholm: Centraltryckeriet, 7-84.
- Dreze, J.(1974), "Axiomatic Theories of Choice, Cardinal Utility and Subjective Probability: A Review", in J. Dreze(ed), Allocation Under Uncertainty: Equilibrium and Optimality, London: MacMillan.
- Fishburn, P.(1981), "Subjective Expected Utility: A Review of Normative Theories", Theory and Decision 13, 139-199.
- Fréchet, M.(1951), "Sur Les Tableaux de corrélation dont les marges sont donnees", Annals de Universite de Lyon A, Series 3, 14, 53-77.
- Kahneman, D. and Tversky, A.(1979), "Prospect Theory: An Analysis of Decision Under Risk", Econometrica 47, 263-291.
- Machina, M.(1982), "'Expected Utility' Analysis Without the Independence Axiom", Econometrica 50, 277-323.
- Marschak, J.(1950), "Rational Behavior, Uncertain Prospects, and Measurable Utility", Econometrica 18, 111-141.
- Newey, W. and Powell, J.(1986), "Asymmetric Least Squares Estimation and Testing", Econometrica, forthcoming.
- Ord, J.(1972), Families of Frequency Distributions, Griffin's Statistical Monographs & Courses No.30, New York: Hafner.
- Rothschild, M. and Stiglitz, J.(1970), "Increasing Risk I: A Definition", Journal of Economic Theory 2, 225-243.
- Savage, L.(1954), The Foundations of Statistics, New York: Wiley.