Risk Aversion for Multiple-Prior Expected Utility*

Jan Werner
Department of Economics
University of Minnesota
Minneapolis, MN 55455

May 2005
revised June 2006.

Abstract: The objective of this paper is to identify multiple-prior (maxmin) expected utility functions that exhibit aversion to risk under some probability measure from among the priors. Risk aversion has profound implications on agents' choices and on market prices and allocations. Our approach to risk aversion relies on the theory of mean-independent risk of Werner (2005). We show that a necessary and sufficient condition for risk aversion of concave multiple-rior expected utility under probability measure π is that the set of probability priors be π-stable. The property of π-stability is a new concept. We show that cores of convex distortions of a probability measure have that property. Relative entropy neighborhoods - used in the context of model uncertainty - have it, too, but Euclidean neighborhoods fail to have it. We also show that the existence of a non-trivial unambiguous event precludes risk aversion with respect to any prior.

*I gratefully acknowledge helpful comments of the participants of RUD 2005 in Heidelberg, in particular, Michele Cohen, Simon Grant, Ehud Lehrer, Isaco Meilijson and Klaus Nehring. Supported by the NSF under Grant SES-0099206.

1. Introduction

Multiple-prior (or maxmin) expected utility is the most appealing alternative to expected utility among the, so called, non-expected utilities. Instead of a single probability measure over uncertain states, the decision maker has a set of probability measures. This multiplicity of probability measures reflects her ambiguous information about the states, or uncertainty of her expectations. The decision criterion is the minimum of expected utilities over the set of multiple priors. Taking the minimum reflects the decision maker's concern with the "worst case" scenario. An axiomatic derivation of multiple-prior expected utility has been given by Gilboa and Schmeidler (1989).

One motivation for multiple-prior expected utility comes from the Ellsberg paradox. The pattern of preference over bets on balls drawn from an urn in the Ellsberg experiment is incompatible with expected utility, but can be explained by multiple-prior expected utility. For a single urn with 30 red balls and 60 green or yellow balls with unknown proportions of the two colors, a multiple-prior expected utility with the set of all probability measures that assign probability 1/3 to drawing red ball and arbitrary probabilities (summing up to 2/3) to drawing yellow or green ball leads to the desired pattern of preferences. Another motivation for multiple-prior expected utility comes from applications of robust control theory. Hansen et al (2002) (see also Maccheroni, Marinacci and Rustichini (2005)) show that stochastic robust control with a certain measure of model misspecification can be viewed as decision making with multiple-prior expected utility.

Risk aversion has significant implications on agents' choices and on prices and allocations in security markets. Examples are the positive premium on equilibrium return on the market portfolio of risky securities and the comonotonicity of Pareto-optimal risk sharing rules (see LeRoy and Werner (2001) for a textbook treatment). These results are typically limited to expected utility (see however Chateauneuf, Dana and Tallon (2000)) because the standard theory of risk aversion due to Arrow (1965) and Pratt (1964) is limited to expected utility.

In the paper "Risk and Risk Aversion when States of Nature Matter" (Werner (2005)) we proposed a new theory of risk aversion primarily aimed at applications to non-expected utilities. The theory applies to utility functions that are defined on contingent claims (i.e., random variables) instead of probability distributions.

Many (but not all) multiple-prior expected utility functions are not distribution invariant under any probabilities of states. The basic concept of our theory is the mean-independent risk: for a probability measure π on a finite state space, contingent claim ϵ is a mean-independent risk at another contingent claim z if the conditional expectation $E_\pi(\epsilon|z)$ of ϵ on z equals zero. Utility function U exhibits aversion to mean-independent risk if there exists a probability measure π such that $U(x) \geq U(y)$ whenever y differs from x by a mean-independent risk at π, that is, if $x = z + \epsilon$ and $y = z + \lambda\epsilon$ for some ϵ and z such that ϵ is a mean-independent risk at z and $0 \leq \lambda \leq 1$. Aversion to mean-independent risk under π implies, in particular, that $U(z) \geq U(z+\epsilon)$ whenever ϵ is a mean-independent risk at z. Thus, the agent whose initial position is a contingent claim z rejects a gamble given by the mean-independent risk ϵ. If utility function U is concave, this last condition is equivalent to aversion to mean-independent risk (see Werner (2005)). Arrow (1965) and Pratt (1964) defined risk aversion under the expected utility hypothesis by this condition restricted to risk-free initial claims. Under expected utility, risk aversion in the Arrow-Pratt sense implies rejection of gambles with mean-independent risk. Every utility function that is monotone decreasing with respect to the standard Rothschild-Stiglitz (or stochastic dominance) order of more risky is averse to mean-independent risk. The converse does not hold since aversion to mean-independent risk does not require that utility function be distribution invariant.

Aversion to mean-independent risk, if it holds for multiple-prior expected utility, identifies a probability measure in the set of priors such that the agent's preferences exhibit patterns discussed above under this probability measure. Thus the role of probabilities bears some similarity to Machina and Schmeidler's (1992) probabilistic sophistication. Probabilistic sophistication requires that preferences (or utility function) be distribution invariant under some probability measure on states. Aversion to mean-independent risk requires that the utility preserve the order induced by mean-independent risk under some probability measure.

Some multiple-prior expected utilities have a specific reference probability measure. Such are multiple-prior expected utilities obtained as representations of rank-dependent expected utilities of Quiggin (1982) and Yaari (1987) with convex distortions of an (objective) probability measure. Sets of priors in multiple-prior expected utilities of Hansen et al (2002) are neighbourhoods of a probability measure. These reference measures are of special relevance in the question of aversion to mean-independent risk.

The objective of this paper is to provide characterization of multiple-prior expected utilities that are averse to mean-independent risk. More precisely, we identify conditions on the set of probability priors (on a finite state space) and the von Neumann-Morgenstern utility function that guarantee aversion to mean-independent risk for some probability measure from the set of priors. The condition on the set of priors is called π-*stability* and it requires that for every probability measure P in the set of priors and every partition of states F, a probability measure that coincides with P on elements of partition F and has conditional probabilities of π within each element of F lies in the set of priors. Our main result, Theorem 1, states that concave multiple-prior expected utility exhibits aversion to mean-independent risk under π if and only if the set of priors is π-stable.

Which sets of priors are π-stable for some probability measure π in the set? Sets obtained in the representation of rank-dependent expected utilities with convex distortion of probability measure π are π-stable. [These are cores of convex distortions of π]. Euclidean neighborhoods of a probability measure π in the probability simplex are not π-stable unless π is the uniform probability. Most remarkably, neighborhoods of π in the relative entropy distance are π-stable, for arbitrary π. The latter are the sets of priors proposed by Hansen et al (2002), see also Maccheroni, Marinacci and Rustichini (2005).

The set of priors in the Ellsberg experiment is not π-stable for any probability π in the set. The reason is the existence of a unambiguous non-trivial event – red ball drawn – to which all priors assign the same probability. We shaw that, in general, the existence of non-trivial unambiguous event precludes mean-independent risk aversion.

The paper is organized as follows: In Section 2 we provide definition and discussion of multiple-prior expected utilities with a finite state space. In Section 3 we introduce the notion of π-stable set of probability measures and prove our main result on risk aversion of multiple-prior expected utility. We also provide a characterization of π-stable sets of priors. Examples of π-stable sets including relative entropy neighborhoods are discussed in Section 4. Section 5 is about risk aversion and existence of unambiguous events.

2. Multiple-Prior Expected Utility and Risk Aversion

There is a finite set $S = \{1, \ldots, S\}$ of states of nature (with $S > 1$.) The set of all (additive) probability measures on the set of all subsets of S is denoted by $\mathcal{M}$, and can be identified with the unit simplex Δ in $\mathcal{R}^S$. The subset of strictly positive probability measures is denoted by $\mathring{\mathcal{M}}$. Any S-dimensional vector $x = (x_1, \ldots, x_S) \in \mathcal{R}^S$ is called *contingent claim*. The expected value $\sum_{s=1}^S P(s)x_s$ of x under a probability measure $P \in \mathcal{M}$ is denoted by $E_P(x)$; the expected utility $\sum_{s=1}^S P(s)v(x_s)$ of x under P is denoted by $E_P[v(x)]$.

Multiple-prior or maxmin expected utility takes the form

$$\min_{P \in \mathcal{P}} E_P[v(x)], \tag{1}$$

for some utility function $v : \mathcal{R} \to \mathcal{R}$ and some convex and closed set $\mathcal{P} \subset \mathcal{M}$ of probability measures. If the set $\mathcal{P}$ consists of a single probability measure, then multiple-prior expected utility reduces to the standard expected utility. In the other polar case, if $\mathcal{P}$ is the set of all probability measures $\mathcal{M}$, then the multiple-prior expected utility reduces to the Wald's criterion $\min_{s \in S} v(x_s)$. Gilboa and Schmeidler's (1989) axiomatization of multiple-prior expected utility (1) does not specify the set of multiple priors beyond closedness and convexity.

The multiple-prior expected utility with linear utility,

$$\min_{P \in \mathcal{P}} E_P(x), \tag{2}$$

has been extensively studied in the context of coherent measures of risk (see Föllmer and Schied (2002)).

Multiple-prior expected utility (1) with concave v is, of course, concave. It is not differentiable unless the set of priors is a singleton (and v is differentiable), that is, unless it is an expected utility. We shall consider its superdifferential. We recall that the superdifferential of a concave function $U : \mathcal{R}^S \to \mathcal{R}$ at $x \in \mathcal{R}^S$ is the set $\partial U(x)$ consisting of all vectors $\phi \in \mathcal{R}^S$ that satisfy $U(y) \leq U(x) + \phi(y - x)$ for every $y \in \mathcal{R}^S$.

If function v is differentiable, then the superdifferential of utility function (1) at x is (see Aubin (1998))

$$\{\phi \in \mathcal{R}^S : \phi_s = v'(x_s)P(s) \text{ for some } P \in \mathcal{P}_v(x)\}, \tag{3}$$

where $\mathcal{P}_v(x)$ denotes the subset of priors on which the minimum expected utility is attained. That is

$$\mathcal{P}_v(x) = \{\bar{P} : E_{\bar{P}}[v(x)] = \min_{P \in \mathcal{P}} E_P[v(x)]\}. \tag{4}$$

For the linear multiple-prior expected utility (2) the subset of minimizing priors is

$$\mathcal{P}(x) = \{\bar{P} \in \mathcal{P} : E_{\bar{P}}(x) = \min_{P \in \mathcal{P}} E_P(x)\}, \tag{5}$$

and the superdifferential of (2) coincides with $\mathcal{P}(x)$.

Multiple-prior expected utilities are often not distribution invariant under any probability measure on states. That is, there may not exist a probability measure such that the utility of a contingent claim depends only on its probability distribution.[1] The Rothschild-Stiglitz (or stochastic dominance) theory of risk and risk aversion applies only to distribution invariant utilities, and hence cannot be used. We shall use the concepts of mean-independent risk and aversion to mean-independent risk of Werner (2005).

Consider a probability measure $\pi \in \mathring{\mathcal{M}}$. Contingent claim $\epsilon \in \mathcal{R}^S$ is a *mean-independent risk at* $z \in \mathcal{R}^S$ if $E_\pi(\epsilon|z) = 0$. For two contingent claims $x, y \in \mathcal{R}^S$ with the same expectation, $E_\pi(x) = E_\pi(y)$, x *differs from* y *by mean-independent risk* if there exist $z, \epsilon \in \mathcal{R}^S$ and $0 \leq \lambda \leq 1$ such that ϵ is a mean-independent risk at z, and $x = z + \epsilon$ and $y = z + \lambda\epsilon$.

Utility function U on $\mathcal{R}^S$ is *averse to mean-independent risk* if there exists a probability measure π such that $U(y) \geq U(x)$ whenever x differs from y by mean-independent risk. Every utility function that is decreasing with respect to the relation of Rothschild-Stiglitz more risky is averse to mean-independent risk (Werner (2005), Theorem 2.1). The converse is not true. Every concave expected utility is averse to mean-independent risk.

Aversion to mean-independent risk is closely related to preference for conditional expectations. Utility function U on $\mathcal{R}^S$ exhibits *preference for conditional expectations* under π if $U(E_\pi(x|F)) \geq U(x)$, for every $x \in \mathcal{R}^S$ and every partition

[1] An exception is when the set of priors is a convex distortion of a probability measure. Then, the multiple-prior expected utility is a rank-dependent expected utility, and it is distribution invariant under the underlying probability measure. Marinacci (2002) studies distribution invariance of multiple-prior expected utilities on a continuum state space.

F of the set of states S. Equivalently, U exhibits preference for conditional expectations if the agent always rejects a mean-independent risk, that is, $U(z) \geq U(z+\epsilon)$ for every $\epsilon, z \in \mathcal{R}^S$ such that ϵ is mean-independent risk at z. It follows from Theorem 5.1 in Werner (2005) that, if U is quasi-concave, then it is averse to mean-independent risk under π if and only if it exhibits preference for conditional expectations under π. Theorem 6.1 in Werner (2005) provides a characterization of utility functions that are mean-independent risk averse and concave in terms of superdifferentials: Concave utility function U is averse to mean-independent risk if and only if for every x there exists $\phi \in \partial U(x)$ such that if $x_s = x_{s'}$, then $\frac{\phi_s}{\pi(s)} = \frac{\phi_{s'}}{\pi(s')}$.

3. π-Stable Sets of Priors and Risk Aversion.

We first introduce the concept of a π-stable set of priors. Let $\pi \in \mathring{\mathcal{M}}$ be a strictly positive probability measure. For every partition F of states and every probability measure $P \in \mathcal{M}$, we define another probability measure $P_F^\pi \in \mathcal{M}$ by

$$P_F^\pi(A) = \sum_{i=1}^{k} \pi(A|A_i) P(A_i) \tag{6}$$

for every $A \subset S$, where sets A_i's are elements of partition F, $i = 1, \ldots, k$. Probability measure P_F^π coincides with P on elements of partition F, but has conditional probabilities of π within each element of partition F.

Two elementary properties of probability measure P_F^π will be repeatedly used:

Lemma 1: *For every $\pi \in \mathring{\mathcal{M}}$, $P \in \mathcal{M}$, $x \in \mathcal{R}^S$, and every partition F*

(i) $E_{P_\pi^F}(x|F) = E_\pi(x|F)$,

(ii) if x is F-measurable, then $E_{P_\pi^F}(x) = E_P(x)$.

Set of probability measures $\mathcal{P}$ is said to be π-*stable* if $P_F^\pi \in \mathcal{P}$ for every $P \in \mathcal{P}$ and every partition F. Note that, if $\mathcal{P}$ is π-stable, then $\pi \in \mathcal{P}$.

The simplest examples of π-stable sets are the singleton set $\{\pi\}$ and the set of all possible priors $\mathcal{M}$. The latter is π-stable for every $\pi \in \mathring{\mathcal{M}}$. Bounds on probabilities of states give rise to π-stable sets of priors. The set

$$\mathcal{P}_l = \{P \in \mathcal{M} : P(s) \geq \gamma_s, \ \forall s\}, \tag{7}$$

where $\gamma_s \in [0,1]$ are lower bounds on probabilities (such that $\sum_s \gamma_s \le 1$), is π-stable for probability measure π defined by $\pi(s) = \gamma_s/\gamma$ where $\gamma = \sum_s \gamma_s$. Note that $\mathcal{P}_l = \{P \in \mathcal{M} : P \ge \gamma\pi\}$, and also $\mathcal{P}_l = \gamma\pi + (1-\gamma)\Delta$. The set

$$\mathcal{P}_u = \{P \in \mathcal{M} : P(s) \le \lambda_s, \ \forall s\}, \tag{8}$$

where $\lambda_s \in [0,1]$ are upper bounds on probabilities (such that $\sum_s \lambda_s \ge 1$), is π-stable for π defined by $\pi(s) = \lambda_s/\lambda$ where $\lambda = \sum_s \lambda_s$, It holds $\mathcal{P}_u = \{P \in \mathcal{M} : P \le \lambda\pi\}$ and $\mathcal{P}_u = \lambda\pi + (1-\lambda)\Delta$.

Since the intersection (and the union) of any two π-stable sets is π-stable, it follows that the (order) interval of probabilities

$$[\gamma\pi, \lambda\pi] = \{P \in \mathcal{M} : \ \gamma\pi \le P \le \lambda\pi\}, \tag{9}$$

where $\gamma \le 1 \le \lambda$, is π-stable.

Two important classes of π-stable sets - cores of convex distortions and neighborhoods in statistical measures of distance - will be discussed in Sections 4 and 5. We state now our main theorem.

Theorem 1: *Suppose that v is concave, $\mathcal{P}$ is closed and convex, and $\pi \in \mathring{\mathcal{M}}$. The following conditions are equivalent:*

(i) *multiple-prior expected utility (1) is averse to mean-independent risk under π,*

(ii) *$\mathcal{P}$ is π-stable,*

(iii) *for every x there exists $P \in \mathcal{P}(x)$ such that*

$$\text{if } x_s = x_{s'}, \quad \text{then} \quad \frac{P(s)}{\pi(s)} = \frac{P(s')}{\pi(s')}. \tag{10}$$

Proof: We first prove that (ii) implies (i). Since $\mathcal{P}$ is convex and v is concave, utility function (1) is concave. By Theorem 5.1 (Werner (2005)) it suffices to show that it exhibits preference for conditional expectation under π.

Consider an arbitrary partition F. Using Lemma 1, we have

$$E_P[v(E_\pi(x|F))] = E_{P_F^\pi}[v(E_{P_F^\pi}(x|F))]. \tag{11}$$

Applying conditional Jensen's inequality, we obtain

$$E_{P_F^\pi}[v(E_{P_F^\pi}(x|F))] \geq E_{P_F^\pi}[v(x)]. \tag{12}$$

Combining (11) and (12) and applying the minimum over $\mathcal{P}$ to both sides, we obtain

$$\min_{P\in\mathcal{P}} E_P[v(E_\pi(x|F))] \geq \min_{P\in\mathcal{P}} E_{P_F^\pi}[v(x)]. \tag{13}$$

Since $\mathcal{P}$ is π-stable, it follows that the right-hand side of inequality (13) is greater than $\min_{P\in\mathcal{P}} E_P[v(x)]$. This shows that multiple-prior expected utility (1) exhibits preference for conditional expectation under π.

Next we prove that (i) implies (iii). Since function v is concave, it is differentiable except for at most countable many points on the real line. This implies that for every $x \in \mathcal{R}^S$, there exists $y \in \mathcal{R}^S$ and $\gamma > 0$ such that $\gamma x_s = v(y_s)$ and v is differentiable at y_s for every s. Using Theorem 6.1 in Werner (2005), condition (i) implies that there exists $P \in \mathcal{P}_v(y)$ such that

$$\text{if } y_s = y_{s'}, \quad \text{then } \frac{P(s)}{\pi(s)} = \frac{P(s')}{\pi(s')}. \tag{14}$$

We note that $\mathcal{P}_v(y) = \mathcal{P}(\gamma x)$ and $\mathcal{P}(\gamma x) = \mathcal{P}(x)$. Further, $y_s = y_{s'}$ if and only if $x_s = x_{s'}$. Taking all these into account, (14) implies (iii).

Last, we prove that (iii) implies (ii). Since the superdifferential of linear multiple-prior expected utility (2) at x is $\mathcal{P}(x)$, it follows from Theorem 6.1 in Werner (2005) and condition (iii) that utility function (2) is averse to mean-independent risk under under π. Suppose by contradiction that $\mathcal{P}$ is not π-stable, so that $\bar{P}_F^\pi \notin \mathcal{P}$ for some $\bar{P} \in \mathcal{P}$ and some partition F. By the separation theorem, there exists $\hat{x} \in \mathcal{R}^S$ such that

$$E_{\bar{P}_F^\pi}(\hat{x}) < \min_{P\in\mathcal{P}} E_P(\hat{x}) \tag{15}$$

By Lemma 1, $E_{\bar{P}_F^\pi}(\hat{x}) = E_{\bar{P}}[E_\pi(\hat{x}|F)]$. Further, since $\bar{P} \in \mathcal{P}$, we obtain from (15) that

$$\min_{P\in\mathcal{P}} E_P[E_\pi(\hat{x}|F)] < \min_{P\in\mathcal{P}} E_P(\hat{x}) \tag{16}$$

This contradict preference for conditional expectations under π of linear multiple-prior expected utility (2). The latter is equivalent to mean-independent risk aversion. This contradiction concludes the proof. □

The most important part of Theorem 1 is the equivalence of conditions (i) and (ii). Condition (iii) should be viewed as a criterion for π-stability of a set of priors. It will be frequently used in Section 4.

4. Cores of Convex Distortions

Important sets of priors are cores of convex distortions of probability measures. Let $f : [0, 1] \to [0, 1]$ be an increasing and convex function satisfying $f(0) = 0$ and $f(1) = 1$. Set function $f \circ \pi$ is the distortion of probability π by f. The *core* of $f \circ \pi$ is

$$\text{core}(f \circ \pi) = \{P \in \mathcal{M} : P(A) \geq f(\pi(A)), \forall A\} \tag{17}$$

Multiple-prior expected utility (1) with the set of priors $\text{core}(f \circ \pi)$ for convex distortion f can be written as

$$\sum_{i=1}^{S} v(x_{(i)})[f(\pi\{s : x_s \geq x_{(i)}\}) - f(\pi\{s : x_s \geq x_{(i-1)}\})]. \tag{18}$$

where we used $x_{(i)}$ to denote the i-th highest value from among all x_s, so that $x_{(1)} \geq x_{(2)} \geq \cdots \geq x_{(S)}$. Utility function (18) is the rank-dependent expected utility axiomatized by Quiggin (1982) and, in the case of linear v, by Yaari (1987).

Proposition 1: *If f is convex, then* $\text{core}(f \circ \pi)$ *is π-stable.*

Proof: Every contingent claim x when treated as random variable on probability space $(S, 2^S, \pi)$ dominates $E_\pi(x|F)$ in the sense of second order stochastic dominance for every partition F. It is well known (see Yaari (1987), or Lemma 2.2 in Dana (2005)) that rank-dependent expected utility (18) with linear v and convex f is monotone decreasing with respect to the second order stochastic dominance. Therefore

$$\min_{P \in \text{core}(f \circ \pi)} E_P(x) \leq \min_{P \in \text{core}(f \circ \pi)} E_P(E_\pi(x|F)). \tag{19}$$

We apply (19) to $x = \chi_A$ for $A \subset S$. The left-hand side equals $f(\pi(A))$ while the right-hand side is $\min P_F^\pi(A)$ over all P in $\text{core}(f \circ \pi)$. Thus

$$f(\pi(A)) \leq P_F^\pi(A) \tag{20}$$

for every $P \in \text{core}(f \circ \pi)$ and every A and F. This shows that $\text{core}(f \circ \pi)$ is π-stable. □

The sets of priors $\mathcal{P}_l$ and $\mathcal{P}_u$ of Section 3 are examples of cores. It is easy to see that $\mathcal{P}_l$ is the core of $f_l \circ \pi$ for any convex function f_l such that $f_l(t) = \lambda t$ for every $t \leq \max_{A \subset S, A \neq S} \pi(A)$ and $f_l(1) = 1$. $\mathcal{P}_u$ is the core of $f_u \circ \pi$ for a convex function f_u given by $f_u(t) = \max\{\lambda(t-1)+1, 0\}$ for every $t \in [0,1]$. Also the interal $[\gamma\pi, \lambda\pi]$ is the core of $f \circ \pi$ for convex function $f = \max\{f_l, f_u\}$.

Proposition 1 and Theorem 1 imply that every rank-dependent expected utility with concave utility and convex distortion is averse to mean-independent risk. Yaari (1987) proved (using an inequality of Hardy, Litlewood and Polya) that rank-dependent expected utility (18) with linear utility v and convex distortion f is decreasing with respect to the relation of Rothschild-Stiglitz more risky. Therefore it is also averse to mean-independent risk. Chew, Karni and Safra (1987) proved that rank-dependent expected utility on the set of all distributions on a real interval is decreasing with respect to the relation of R-S more risky if and only if utility v is concave and distortion f is convex. The proof in Chew, Karni and Safra (1987) (see also Chew and Mao (1995)) relies on Gateaux differentiability of RDEU utility on the space of distribution.

5. Neighborhoods in Statistical Measures of Distance

For a convex function $\Phi : \mathcal{R}_+ \to \mathcal{R}$ such that $\Phi(1) = 0$ the statistical Φ-measure of distance between probability measures $P \in \mathcal{M}$ and $\pi \in \mathring{\mathcal{M}}$ is

$$d_\Phi(P, \pi) = \sum_{s=1}^{S} \pi(s) \Phi\Big(\frac{P(s)}{\pi(s)}\Big). \tag{21}$$

It can be shown that the Φ-measure d_Φ is non-negative and a convex function of P. If Φ is strictly convex, then $d_\Phi(P, \pi)$ equals zero if and only if $P = \pi$. In general d_Φ is not a metric for it is asymmetric and violates triangle inequality.

A neighborhood of probability measure $\pi \in \mathring{\mathcal{M}}$ in Φ-measure of distance is the set

$$\mathcal{N}_\Phi = \{P \in \mathcal{M} : d_\Phi(P, \pi) \leq \epsilon\} \tag{22}$$

for $\epsilon > 0$. We restrict our attention of neighborhoods that are contained in the interior of the probability simplex.

Proposition 2: *If Φ is strictly convex and differentiable, then the neighorhood $\mathcal{N}_\Phi$ of π is π-stable.*

Proof: We verify that condition (10) of Theorem 1 holds. For non-deterministic x the set of minimizing probabilities $\mathcal{N}_\Phi(x)$ consists of unique probability measure P_x^*. The first-order conditions for P_x^* as a solution to the minimization in (5) imply that

$$x_s - x_{s'} = \lambda[\Phi'\big(\frac{P_x^*(s)}{\pi(s)}\big) - \Phi'\big(\frac{P_x^*(s')}{\pi(s')}\big)], \tag{23}$$

for every $s, s' \in S$, and some $\lambda > 0$. From this we obtain that if $x_s = x_{s'}$, then $\frac{P_x^*(s)}{\pi(s)} = \frac{P_x^*(s')}{\pi(s')}$, so that condition (10) holds. If x is deterministic, then $\mathcal{N}_\Phi(x) = \mathcal{N}_\Phi$. In particular, $\pi \in \mathcal{N}_\Phi(x)$, and hence (10) holds. □

Example 5.1 (Kullback-Leibler relative entropy.) For $\Phi(t) = t\ln(t)$, the Φ-measure of distance is the *relative entropy*

$$d_\Phi(P, \pi) = \sum_{s=1}^{S} P(s) \ln\big(\frac{P(s)}{\pi(s)}\big). \tag{24}$$

Neighborhoods in the relative entropy distance have been used as sets of priors in applications of multiple-prior expected utility by several authors, see Hansen et al (2002) and Kogan and Wang (2002) and Cao, Wang and Zhang (2003).

Example 5.2 (Gini χ^2-index.) For $\Phi(t) = (t-1)^2$, the Φ-measure of distance is the *Gini χ^2-index*

$$\sum_{s=1}^{S} \frac{(P(s) - \pi(s))^2}{\pi(s)}. \tag{25}$$

Example 5.3 (total variation.) For $\Phi(t) = |t-1|$, the Φ-measure of distance is the *total variation*

$$\sum_{s=1}^{S} |P(s) - \pi(s)|. \tag{26}$$

Function $\Phi(t) = |t-1|$ is neither strictly convex nor differentiable. Nevertheless, it is easy to show that total variation (26) has the distance properties common to strictly convex statistical measures. It is non-negative, convex and equal to zero only if $P = \pi$. Further, total variation neighborhoods of π are π-stable.

5. Risk Aversion and Unambiguous Events.

Under multiple-prior expected utility, the agent's probabilistic beliefs about events are described by the set of probability measures $\mathcal{P}$. A natural definition (see Nehring (1999)) of an *unambiguous event* is as such event $A \subset S$ that $P(A) = P'(A)$ for all $P, P' \in \mathcal{P}$. Of course, the *trivial* events, $\emptyset$ and S, are always unambiguous.

It turns out that, if a set of priors - other than a singleton set - permits non-trivial unambiguous events, then it cannot be π-stable for any π. Thus, the existence of a non-trivial unambiguous event precludes mean-independent risk aversion.

Theorem 2: *If $\mathcal{P}$ is π-stable and there exists a non-trivial unambiguous event $A \subset S$, then $\mathcal{P} = \{\pi\}$.*

PROOF: Suppose by contradiction that $\mathcal{P}$ is π-stable, has a non-trivial unambiguous event A, and there exists $P \in \mathcal{P}$ such that $P \neq \pi$. Let s be such that $\pi(s) \neq P(s)$. Suppose first that $s \notin A$. Consider a partition F of S into two sets: $A \cup \{s\}$, and its complement. Note that the complement of $A \cup \{s\}$ is non-empty, for it cannot be that $P(A) = \pi(A)$, $\pi(s) \neq P(s)$, and $A \cup \{s\} = S$. For the probability measure P_F^π defined by (6), we have

$$P_F^\pi(A) = \pi(A) \frac{\pi(A) + P(s)}{\pi(A) + \pi(s)} \neq \pi(A). \tag{27}$$

Since $P_F^\pi \in \mathcal{P}$, this contradicts the assumption that A is unambiguous. If $s \in A$, then we consider the complement event A^c instead of A. Event A^c is unambiguous and $s \notin A^c$, so that the above arguments apply. This concludes the proof. □

For the set of probability priors in our discussion of the Ellsberg paradox in Section 1, the event of red ball drawn from the urn is unambiguous. It has probability 1/3. Theorem 2 implies that there is no measure π in the set of priors such that π-stability holds. Thus, the form of ambiguity of beliefs in the Ellsberg paradox precludes risk aversion.

7. Remarks.

Euclidean neighborhoods of π are in general *not* π-stable. An exception is the case of the uniform probability measure on S for then the Euclidean neighborhood coincides with the Gini-index neighborhood which is π-stable by Proposition 2. If π is not uniform, then one can check that condition (10) does not hold, and hence Euclidean neighborhoods of π are not π-stable.

For the sets of priors $\mathcal{P}_l$ and $\mathcal{P}_u$ of Section 3 (see (7) and (8)), one can show that the probability measure π is the unique measure with respect to which the sets are π-stable. This is not always so for cores of convex distortions. For instance, if the probability measure is such that different contingent claims have different probability distributions and the distortion function is strictly convex, then the core is π-stable for all probability measures in a small neighborhood around the reference probability measure.

References

Arrow, K. 1965. *Aspect of the Theory of Risk Bearing,* Yrjö Jahnsson Foundation, Helsinki.

Aubin, J-P, *Optima and Equilibria. An Introduction to Nonlinear Analysis*, Springer, 1998.

Cao, H., T. Wang and H. Zhang (2003), "Model Uncertainty, Limited Market Participation and Asset Prices," University of North Carolina working paper.

Chew, S., E. Karni and Z. Safra (1987), "Risk aversion in the theory of expected utility with rank dependent preferences," *Journal of Economic Theory*, pp. 370-381.

Chew, S. and M.H. Mao (1995), "A Schur Concave Characterization of Risk Aversion for Non-expected Utility Preferences," *Journal of Economic Theory*, 67 (2): 402-435.

Chateauneuf A., R-A Dana, and J.M. Tallon (2000) "Optimal risk-sharing rules and equilibria with Choquet expected utility." *Journal of Mathematical Economics*, 34(2): 191-214.

Epstein, L. and T. Wang (1995), "Uncertainty, Risk-Neutral Measures and Security Price Booms and Crashes," *Journal of Economic Theory*, 67 (1), pp. 40-83.

Föllmer, H. and A. Schied (2002), *Stochastic Finance - An Introduction in Discrete Time*, de Gruyeter, Berlin.

Gilboa, I. and D. Schmeidler (1989), "Maxmin expected utility with nonunique prior," *Journal of Mathematical Economics*, 18, pp. 141-153.

Hansen, L., T. Sargent, G. Turmuhambetova, and N. Williams (2002), "Robustness and Uncertainty Aversion," University of Chicago.

Kogan, L. and T. Wang (2002), "A Simple Theory of Asset Pricing under Model Uncertainty," MIT working paper.

Machina, M. and D. Schmeidler (1992), "A more robust definition of subjective probability", *Econometrica*, 60, No. 4 (July), pp. 745-780.

Marinacci, M. (2002), "Probabilistic sophistication and multiple priors," *Econometrica*, 70, pp. 755–764.

Maccheroni, F, Marinacci M. and A. Rustichini (2005), "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," Discussion Paper.

Nehring, K. (1999), "Capacities and Probabilistic Beliefs: A Precarious Coexistence," *Mathematical Social Sciences,* 38, pp 197-213.

Pratt, J. 1964. "Risk Aversion in the Small and in the Large" *Econometrica*, 32: 132-136.

Quiggin, J., (1982) "A theory of anticipated utility," *Journal of Economic Behavior and Organization*, 3, pp. 323-343.

Schmeidler, D. (1989), "Subjective Probability and Expected Utility Without Additivity," *Econometrica*, 57: 571-587.

Yaari, M. (1987), "The dual theory of choice under risk," *Econometrica*, 55: 95-115.

Werner, J. (2005) "Risk and Risk Aversion when States of Nature Matter," mimeo, revised May 2005.