Discussion Paper No. 338

AN AXIOMATIC DERIVATION OF SUBJECTIVE PROBABILITY, UTILITY, AND EVALUATION FUNCTIONS

by

Roger B. Myerson*

August, 1978

Abstract. Subjective expected utility maximization is derived from four axioms, using an argument based on the separating hyperplane theorem. It is also shown that the first three of these axioms imply a more general maximization formula, involving an *evaluation function*, which can still serve as a basis for decision analysis.

*Graduate School of Management, Northwestern University, Evanston, IL 60201.

1. Introduction

Decisions under uncertainty are commonly described in two ways: using a <u>probability</u> model or a <u>state-variable</u> model. In each case, we speak of the decision-maker as choosing among lotteries, but the two models differ in how a lottery is defined. In a probability model, lotteries are probability distributions over a set of prizes (for example, see Section 2.4 in [7]). In a state-variable model, lotteries are functions from a set of possible <u>states of nature</u> into a set of prizes (for example, see Chapter 13 in [7]).

The distinction between a probability model and a state-variable model is not simply a matter of mathematical style. A probability model is appropriate to describe gambles in which the prize will depend on events which have obvious objective probabilities; we shall refer to such events as <u>objective unknowns</u>. These gambles correspond to the "roulette lotteries" of [1], or the "risks" of [6]. For example, gambles which depend on the toss of a fair coin, the spin of a roulette wheel, or the blind draw of a ball out an urn containing a known population of identically-shaped but different-colored balls, all could be adequately described in a probability model. (There is an implicit assumption here that two objective unknowns with the same probability are completely equivalent for all decision-making purposes. For example, if we describe a lottery by saying that it "offers a prize of $100 or $0, each with probability 1/2", we are assuming that it does not matter whether the prize is to be determined by flipping a fair coin, or by drawing a ball from an urn which contains 50 white and 50 black balls.)

On the other hand many events do not have obvious probabilities; the outcome of a future sports event or the future course of the stock market are good examples. We shall refer to such events as subjective unknowns. Gambles which depend on subjective unknowns correspond to the "horse lotteries" of [1], or the "uncertainties" of [6]. They are more readily described in a state-variable model, because these models allow us to describe how the prize will be determined by the unpredictable events, without our having to specify any probabilities for these events.

The fundamental result in decision theory is the subjective expected utility theorem. This theorem asserts that if a decision-maker satisfies certain basic assumptions then his preferences over lotteries will always be consistent with an expected-utility-maximization model. In the context of a probability model, this means that there must exist some real-valued utility function, defined on the set of prizes, such that a lottery giving a higher expected utility will always be preferred to one giving a lower expected utility. In the context of a state-variable model, this means that there must exist a utility function on the set of prizes and a personal subjective probability distribution over the set of possible states of nature such that, again, the lottery giving higher expected utility is always preferred.

In this paper, we derive the subjective expected utility theorem. We define our lotteries so as to include both the probability and state-variable models as special cases. That is, we study lotteries in which the prize may depend on both objective unknowns (which may be modeled by probabilities) and subjective unknowns (which must be modeled by a state-of-nature variable). We also show that a weaker set of assumptions implies a more general (but still useful) maximization formula, in which a single evaluation function replaces the utility and subjective probability functions.

There is already a vast literature on axiomatic derivations of the subjective expected utility theorem, beginning with [9] and [10]; see [3] for a summary through 1968. Nevertheless, we hope that small further contributions might still be welcomed in this area, since it is basic to so much of operations research and economics. This paper builds particularly on the approach set forth in two papers by Fishburn: our definition of a lottery follows [2], and our axioms and method of proof are in the spirit of [4]. However, we have tried to keep the discussion self-contained in this paper, so that it might also serve as an independent introduction to subjectivist decision theory.

2. Basic Definitions

Let X be a set of prizes, and let S be a set of states of nature. To simplify the mathematics, we assume that X and S are both finite sets. We define a lottery to be any function f which specifies a nonnegative real number $f(x|s) \geq 0$ for every $x \in X$ and every $s \in S$, such that

$$\sum_{x \in X} f(x|s) = 1$$

for every $s \in S$.

Each number $f(x|s)$ is to be interpreted as the (objective) probability of getting prize x in lottery f if s is the state of nature. For this to make sense, the state must be defined broadly enough to summarize all subjective unknowns which might influence the prize to be received. Then, once a state s has been determined, only objective unknowns will remain, and an objective probability distribution over the prizes can be calculated for any well-defined gamble. So our formal model of a lottery allows us to represent any gamble, in which the prize may depend on both objective and subjective unknowns.

Let L be the set of all lotteries, so that:

$$(1) \qquad L = \left\{ f \in \mathbb{R}^{X \times S} \middle| \sum_{x' \in X} f(x'|s) = 1 \text{ and } f(x|s) \geq 0,\ \forall x \in X,\ \ \forall s \in S \right\}.$$

As vectors in $\mathbb{R}^{X \times S}$, lotteries may be added together or multiplier by a scalar. This result will be a vector in $\mathbb{R}^{X \times S}$, but not necessarily a lottery in L. However if α is a scalar satisfying $0 \leq \alpha \leq 1$, then $\alpha f + (1-\alpha)g$ will be in L if $f \in L$ and $g \in L$. So L is a convex set (or a mixture set, in the language of [5]).

Our decision-maker's feelings about gambles define a preference relation on the set of lotteries. Formally, for any $f \in L$ and $g \in L$, we write

$$f \gtrsim g$$

if and only if the lottery f is at least as desirable as the lottery g, in the opinion of the decision-maker. Given this relation ($\gtrsim$), we define the relations ($\sim$) and ($>$) so that:

$$f \sim g \iff f \gtrsim g \text{ and } g \gtrsim f;$$

$$f > g \iff f \gtrsim g \text{ and } g \not\gtrsim f .$$

3. The First Three Axioms

To build a theory of decisions under uncertainty, we must make some basic assumptions about this preference relation ($\gtrsim$) over lotteries.

Our first assumption is that the decision-maker can assess his preferences over any pair of lotteries in L.

<u>Axiom 1</u>. For any $f \in L$ and $g \in L$, $f \gtrsim g$ or $g \gtrsim f$.

Suppose the decision-maker prefers lottery f to lottery g, and thinks that lottery f' is at least as good as lottery g'. Then consider two compound gambles: one giving a ticket to lottery f if a fair coin toss comes out Heads, and a ticket to f' if Tails; and the other giving g if Heads, and g' if Tails. If Heads were the known outcome of the toss, then the first gamble would be strictly better than the second, and if Tails were the outcome then the first would be at least as good as the second. So it is reasonable to assume that the decision-maker should prefer the first gamble to the second, before the coin is tossed. But the first gamble is equivalent to the lottery $\frac{1}{2}f + \frac{1}{2}f'$ (since the objective probability of getting prize x given than s is the state of the world would be $\frac{1}{2}f(x|s) + \frac{1}{2}f'(x|s)$, before the coin is tossed), and the second gamble is equivalent to the lottery $\frac{1}{2}g + \frac{1}{2}g'$. So we conclude that $f > g$ and $f' \gtrsim g'$ should imply $\frac{1}{2}f + \frac{1}{2}f' > \frac{1}{2}g + \frac{1}{2}g'$. Our second axiom is a generalization of this principle.

> Axiom 2. For any $f \in L$, $f' \in L$, $g \in L$, and $g' \in L$, and for any number α such that $0 < \alpha < 1$, if $f > g$ and $f' \gtrsim g'$, then $\alpha f + (1-\alpha)f' > \alpha g + (1-\alpha)g'$.

To motivate our third assumption, consider the following situation. Suppose that we offer the decision-maker a choice between two lotteries, f and g, and he reports that he strictly prefers f to g ($f > g$). Now suppose we learn that there may be some small error in the objective conditional probabilities $f(x|s)$ which we reported to the decision-maker. So we ask the decision-maker how accurately does he need to know these probabilities. It would make matters very difficult if he replied that he must know all the numbers exactly, that even the slightest error could reverse his preference. To avoid this situation, we make the following assumption.

> Axiom 3. For any $f \in L$ and $g \in L$, if $f > g$ then there exists some $\epsilon > 0$ such that, for any $f' \in L$, if $|f'(x|s) - f(x|s)| \leq \epsilon$ for all $x \in X$ and $s \in S$ then $f' \gtrsim g$.

4. The Evaluation Function

Our first result characterizes the preferences which satisfy Axioms 1 through 3.

Theorem 1. If the preference relation ($\succsim$) satisfies Axioms 1, 2, and 3, then there exists a function $w: X \times S \to \mathbb{R}$ such that:

$f \succsim g$ if and only if

$$\sum_{s \in S} \sum_{x \in X} f(x|s) \cdot w(x,s) \geq \sum_{s \in S} \sum_{x \in X} g(x|s) \cdot w(x,s),$$

for any lotteries $f \in L$ and $g \in L$.

Following [11], we call a function $w(\cdot)$ an evaluation function if it satisfies the property described in Theorem 1. So if a decision-maker's preferences satisfy Axioms 1 through 3, then a decision analyst only needs to know $|X| \cdot |S|$ evaluation numbers ($w(x,s)$) to compute the most-perferred alternative for the decision-maker, whenever he has a choice among lotteries.

We defer the proof of Theorem 1 to Section 6. However, since the proof will be nonconstructive, we now discuss how to find the evaluation function for a given decision-maker, assuming Theorem 1.

For the evaluation function w, we define $W: L \to \mathbb{R}$ so that:

$$(2) \qquad W(f) = \sum_{s \in S} \sum_{x \in X} f(x|s) w(x,s) .$$

So, assuming that $w(\cdot)$ is the evaluation function from Theorem 1, we have:

$$f \succsim g \iff W(f) \geq W(g).$$

This immediately implies that the preference relation ($\succsim$) is a transitive order. Also, since $W(\cdot)$ is a linear functional, we know that if $f \succsim g \succsim h$ (so $W(f) \geq W(g) \geq W(h)$) then there exists some λ, $0 \leq \lambda \leq 1$, such that $W(g) = W(\lambda f + (1-\lambda)h)$ and $g \sim \lambda f + (1-\lambda)h$. (Furthermore, if $f \succ h$ then this λ must be unique.)

For each state s, let $z_0(s)$ and $z_1(s)$ be prizes in X such that

$$(3) \qquad w(z_0(s),s) = \min_{x \in X} w(x,s)$$

$$w(z_1(s),s) = \max_{x \in X} w(x,s)$$

Without knowing $w(\cdot)$, we could still find $z_1(s)$ (or $z_0(s)$) by describing any lottery f to the decision-maker and then asking him the following question: "Consider all lotteries in the set $\{f_{x,s} | x \in X\}$, where

$$f_{x,s}(y|t) = \begin{cases} 1 & \text{if} \quad y = x \text{ and } t = s, \\ 0 & \text{if} \quad y \neq x \text{ and } t = s, \\ f(z|t) & \text{if} \quad t \neq s. \end{cases}$$

Which of these lotteries seems best (or worst) to you?" If Theorem 1 holds for the decision-maker, his answer will be $f_{z_1(s),s}$ (or $f_{z_0(s),s}$).

Let h_0^* and h_1^* be lotteries such that (for i = 0,1)

$$(4) \qquad h_i^*(x|s) = \begin{cases} 1 & \text{if} \quad x = z_i(s), \\ 0 & \text{if} \quad x \neq z_i(s) . \end{cases}$$

That is, h_0^* and h_1^* are the lotteries which give respectively the worst and best prizes in every state.

For every $x \in X$ and $s \in S$, let $h_{x,s}$ be the lottery such that:

$$(5) \qquad h_{x,s}(y|t) = \begin{cases} 1 & \text{if} \quad y = x \text{ and } t = s, \\ 0 & \text{if} \quad y \neq x \text{ and } t = s, \\ h_0^*(y|t) & \text{if} \quad t \neq s. \end{cases}$$

That is, $h_{x,s}$ is the lottery which differs from h_0^* only in that $h_{x,s}$ gives x instead of $z_0(s)$ in state s.

From the way h_0^* and h_1^* were constructed, we know that $W(h_1^*) \geq W(h_{x,s}) \geq W(h_0^*)$, for any x and s. So there exists a number $\lambda_{x,s}$ such that

$$(6) \qquad W(h_{x,s}) = \lambda_{x,s} W(h_1^*) + (1 - \lambda_{x,s}) W(h_0^*) .$$

This number can be obtained by asking the decision-maker: "For what value of λ would you be indifferent between $h_{x,s}$ and $\lambda h_1^* + (1-\lambda)h_0^*$?" His answer will be $\lambda_{x,s}$.

Equations (5) and (6) imply that:

(7) $$w(x,s) - w(z_0(s),s) = \lambda_{x,s}(W(h_1^*) - W(h_0^*)).$$

Thus, to determine the evaluation function w, it seems that we need to know $w(z_0(s),s))$ for each state s and $(W(h_1^*) - W(h_0^*))$. However, assuming that $h_1^* > h_0^*$, we can show that if $w(\cdot)$ satisfies Theorem 1, then $w'(\cdot)$ will also satisfy Theorem 1, where

$$w'(x,s) = \frac{1}{W(h_1^*) - W(h_0^*)}(w(x,s) - w(z_0(s),s))$$

for any $x \in X$ and $s \in S$. To check this, observe that $W'(f) = \dfrac{W(f) - W(h_0^*)}{W(h_1^*) - W(h_0^*)}$, so $W'(f) \geq W'(g)$ if and only if $W(f) \geq W(g)$.) So unless $h_1^* \sim h_0^*$ (in which case $w \equiv 0$ would satisfy Theorem 1), we may assume that the evaluation function in Theorem 1 satisfies $w(z_0(s),s) = 0$ for all s and $W(h_1^*) - W(h_0^*) = 1$. Then equation (7) becomes:

$$w(x,s) = \lambda_{x,s}, \quad \text{for all } x \in X \text{ and } s \in S.$$

So we can determine the evaluation function by asking the decision-maker a finite number of questions about his preferences over specific lotteries.

5. Probabilities and Utilities

The evaluation numbers w(x,s) measure how much the prospect of receiving prize x if state s occurs would influence the happiness of the decision-maker now (before he knows the true state of nature). As such, w(x,s) is influenced by two factors: how much would the decision-maker enjoy prize x, and how likely does he think state s is to occur. To distinguish these factors into two separate functions, we need one additional assumption.

Axiom 4. Let $s \in S$ and $s' \in S$ be any two states. Let f, f', and g be any lotteries in L which satisfy:

$$g(x|s) = g(x|s'),$$

$$f(x|t) = g(x|t) \text{ if } t \neq s,$$

$$f'(x|t) = g(x|t) \text{ if } t \neq s', \text{ and}$$

$$f(x|s) = f'(x|s')$$

for all $x \in X$ and all $t \in S$. Then $f \gtrsim g$ if and only if $f' \gtrsim g$.

This axiom expresses the idea that, if a particular change in the lottery were considered an improvement in state s, then the same change should also be considered an improvement in state s'. That is, Axiom 4 will hold if the decision-maker's preferences over prizes and objective gambles do not depend on the state of nature.

With Axiom 4, the evaluation function becomes the product of a vonNeumann-Morgenstern utility function, $u(\cdot)$ on X and a Savage subjective probability function, $p(\cdot)$ on S.

Theorem 2. If the preference relation ($\gtrsim$) satisfies Axioms 1, 2, 3, and 4, then there exists a function $u: X \to \mathbb{R}$ and a function $p: S \to \mathbb{R}$ such that:

(a) $\sum_{s \in S} p(s) = 1$, and $p(s) > 0 \quad \forall s \in S$; and

(b) $f \gtrsim g$ if and only if

$$\sum_{s \in S} \sum_{x \in X} f(x|s)p(s)u(x) \geq \sum_{s \in S} \sum_{x \in X} g(x|s)p(s)u(x),$$

for any lotteries $f \in L$ and $g \in L$.

This theorem considerably simplifies the problem of assessing the evaluation function described in the last section. Now, instead of requiring $|X|\cdot|S|$ evaluation numbers, we only need $|X|$ utility numbers and $|S|$ probability numbers. These numbers can be assessed as follows.

When Theorem 2 holds, the worst-prize and best-prize function $z_0(s)$ and $z_1(s)$ defined in (3) become constants. That is, there are two prizes z_0 and z_1 in X such that

$$u(z_0) = \min_{x\in X} u(x) \quad \text{and} \quad u(z_1) = \max_{x\in X} u(x).$$

In all interesting cases we will have $u(z_1) > u(z_0)$ (otherwise the decision-maker will be indifferent between all lotteries). If $u(\cdot)$ satisfies Theorem 2, then $u'(\cdot) = a \cdot u(\cdot) + b$ will also satisfy Theorem 2, for any numbers $a > 0$ and b. So without loss of generality we may assume that $u(z_0) = 0$ and $u(z_1) = 1$.

For any prize x, let e_x be the lottery giving x for sure; that is:

$$e_x(y|s) = \begin{cases} 1 & \text{if } y = x, \\ 0 & \text{if } y \neq x. \end{cases}$$

Now, if $u(z_0) = 0$ and $u(z_1) = 1$, then Theorem 2 implies:

(8) $$e_x \sim u(x)e_{z_1} + (1-u(x))e_{z_0}$$

since both sides give expected utility $u(x)$. So, when the decision-maker is asked what value of λ would make $\lambda e_{z_1} + (1-\lambda)e_{z_0}$ as good as e_x, his answer will be $\lambda = u(x)$. Thus we can assess the utility function $u(\cdot)$ with a finite number of questions.

Similarly, to assess the subjective probability $p(s)$ for any state s, let g_s be the lottery giving z_1 if s happens and giving z_0 otherwise; that is:

$$g_s(x|t) = \begin{cases} e_{z_1}(x|t) & \text{if } t = s, \\ e_{z_0}(x|t) & \text{if } t \neq s. \end{cases}$$

Observe that

(9) $$g_s \sim p(s)e_{z_1} + (1-p(s))e_{z_0},$$

since both sides give expected utility $p(s)$. So, when the decision-maker is asked what value of λ would make $\lambda e_{z_1} + (1-\lambda)e_{z_0}$ as good as g_s, his answer will be $\lambda = p(s)$.

6. Proofs

PROOF OF THEOREM 1.

We assume that the preference relation ($\gtrsim$) satisfies Axioms 1, 2, and 3.

Let $M \subseteq \mathbb{R}^{X \times S}$ be defined by:

$$M = \{d \in \mathbb{R}^{X \times S} \mid \sum_{x \in X} d(x,s) = 0, \quad \forall_s \in S\}$$

Observe that M is a finite dimensional subspace of $\mathbb{R}^{X \times S}$. Notice also that M satisfies

$$M = \{\alpha \cdot (f-g) \mid f \in L,\ g \in L,\ \alpha \in \mathbb{R}\}.$$

That is, any scalar multiple of the difference between two lotteries (thought of as vectors in $\mathbb{R}^{X \times S}$) will be a vector in M.

Using Axioms 1 and 2, we can show that, if $\alpha \cdot (f-g) = \alpha' \cdot (f'-g')$ (for any lotteries $f \in L$, $g \in L$, $f' \in L$, $g' \in L$) where $\alpha > 0$ and $\alpha' > 0$, then $f > g$ if and only if $f' > g'$. To prove this, suppose to the contrary that $f > g$ and $g' \gtrsim f'$. Then by Axiom 2 we get $\frac{\alpha}{\alpha+\alpha'} f + \frac{\alpha'}{\alpha+\alpha'} g' > \frac{\alpha}{\alpha+\alpha'} g + \frac{\alpha}{\alpha+\alpha'} f'$. But this is impossible, since $\alpha f + \alpha' g' = \alpha g + \alpha' f'$. Similarly, $f' > g'$ and $g \gtrsim f$ cannot hold at the same time, so $f > g$ if and only if $f' > g'$, using Axiom 1.

We can now define $H \subseteq M$ by

$$H = \{\alpha \cdot (f-g) \mid \alpha > 0 \text{ and } f > g\}.$$

By the above paragraph, we know that, if $d = \alpha \cdot (f-g)$ for some $\alpha > 0$, $f \in L$, and $g \in L$, then $d \in H$ if and only if $f > g$.

H is a convex set, by Axiom 2. To check this, observe that, if $d = \alpha \cdot (f-g)$ and $d' = \alpha'(f'-g')$ (where $\alpha > 0$, $\alpha' > 0$, $f > g$, and $f' > g'$) then:

$$f^{\lambda} = \frac{(1-\lambda)\alpha f + \lambda\alpha' f'}{(1-\lambda)\alpha + \lambda\alpha'} > \frac{(1-\lambda)\alpha g + \lambda\alpha' g'}{(1-\lambda)\alpha + \lambda\alpha'} = g^{\lambda}$$

by Axiom 2, for any λ between 0 and 1. So $(1-\lambda)d + \lambda d' = ((1-\lambda)\alpha + \lambda\alpha')(f^{\lambda}-g^{\lambda}) \in H$.

Axiom 3 implies that H is relatively open as a subset of M. To prove this, let $|\cdot|$ be the sup-norm on M, so that

$$|c| = \underset{(x,s)\in X\times S}{\text{maximum}} |c(x,s)|, \text{ for any } c \in M.$$

Also, let $g^* \in L$ be the lottery such that

$$g^*(x|s) = \frac{1}{|X|} \text{ for all } (x,s).$$

($|X|$ is the number of prizes in the set X.)

Now, suppose $d \in H$. Let $f = g^* + \frac{1}{2|X|\,|d|}\cdot d$. Then $f \in L$, because $d \in M$ and $f(x|s) \geq \frac{1}{2|X|}$ for all (x,s). Furthermore, for any vector $c \in M$, if $|c| < \frac{1}{2|X|}$, then $f + c \in L$ also. Now $f \succ g^*$, because $(2|X|\,|d|)(f-g^*) = d \in H$. Using Axiom 3, choose ϵ so that $\frac{1}{2|X|} > \epsilon > 0$ and so that $|f'-f| < \epsilon$ implies $f' \gtrsim g^*$. So for any $c \in M$ such that $|c| < \frac{\epsilon}{2}$, we have $f + 2c \in L$ and $f + 2c \gtrsim g^*$; and then, by Axiom 2, $f + c = \frac{1}{2}(f+2c) + \frac{1}{2}(f) \succ \frac{1}{2}g^* + \frac{1}{2}g^* = g^*$. So $d + (2|X|\,|d|)c = (2|X|\,|d|)(f + c - g^*) \in H$ for any $c \in M$ such that $|c| < \epsilon/2$. So we have an open ball in H around any $d \in H$. Thus H is relatively open in M.

Thus H is convex and open in M. Furthermore, $\underset{\sim}{0} \notin H$. By the separating hyperplane theorem (Theorem 11.2 in [8]), there exists some vector $w \in M$ such that $w\cdot d = \sum_{(x,s)} w(x,s)d(x,s) > 0$ for all $d \in H$. Therefore, for any $f \in L$ and $g \in L$, if $f \succ g$ then $\sum_{x\in S}\sum_{s\in S} w(x,s)(f(x|s) - g(x|s)) > 0$.

Conversely, suppose $\sum_{x\in X}\sum_{s\in S} w(x,s)(f(x|s) - g(x|s)) > 0$. Is it possible that $g \gtrsim f$? We must consider two cases. First, suppose that there exists at least one pair of lotteries such that $g' \succ f'$. Then let $f^\epsilon = (1-\epsilon)f + \epsilon f'$ and $g^\epsilon = (1-\epsilon)g + \epsilon g'$. For $\epsilon > 0$ sufficiently small we have

$$\sum_{x\in X}\sum_{s\in S} w(x,s)(g^\epsilon(x|s) - f^\epsilon(x|s) < 0,$$

but $g \succsim f$ would imply $g^{\epsilon} > f^{\epsilon}$, a contradiction of the way w was constructed. So $w \cdot (f-g) > 0$ if and only if $f > g$. Thus $w \cdot f \geq w \cdot g$ if and only if $f \succsim g$, which proves the theorem.

If, on the other hand, there does not exist any pair of lotteries such that $f' > g'$, then $w = \underset{\sim}{0} \in M$ would satisfy the theorem.

PROOF OF THEOREM 2.

Let $M^* = \{b \in \mathbb{R}^X \mid \sum_{x \in X} b(x) = 0\}$. M^* is a finite dimensional subspace of $\mathbb{R}^X$. Let w be the evaluation function from Theorem 1. For any state $s \in S$, let $w(s)$ be the vector

$$w(s) = (w(x,s))_{x \in X} \in \mathbb{R}^X$$

The evaluation function derived in the proof of Theorem 1 in fact was in M, so we may assume without loss of generality that $w(s) \in M^*$ for every $s \in S$.

Now select any states s and s' in S. Let $C(s) \subseteq M^*$ be

$$C(s) = \{\alpha \cdot w(s) \mid \alpha \in \mathbb{R}, \alpha \geq 0\}.$$

We will show that $w(s') \in C(s)$.

If $w(s') \notin C(s)$, then there exists some $b \in M^*$ such that $w(s') \cdot b = \sum_{x \in X} w(x,s') b(x) > 0$ but $c \cdot b \leq 0$ for all $c \in C(s)$, using the separating hyperplane theorem. Multiplying by a small positive scalar if necessary, we may assume that

$$\underset{x \in X}{\text{minimum}} \; b(x) \geq \frac{-1}{|X|} .$$

Now let g^*, f, and f' be lotteries in L such that:

$$g^*(x|t) = 1/|X|$$

$$f(x|t) = \begin{cases} g^*(x|s) + b(x) & \text{if } t = s \\ g^*(x|t) & \text{if } t \neq s \end{cases}$$

$$f'(x|t) = \begin{cases} g^*(x|s') + b(x) & \text{if } t = s' \\ g^*(x|t) & \text{if } t \neq s' \end{cases}$$

for any $x \in X$ and $t \in S$. Then

$$w \cdot (f - g^*) = w(s) \cdot b > 0, \text{ so } f \succ g^* \text{ and } w \cdot (f' - g^*) = w(s') \cdot b \leq 0,$$

so $g^* \succsim f'$. But this contradicts Axiom 4, which requires that $f \succsim g^*$ if $f' \succsim g^*$.

So for any s and s', there exists some $\alpha \geq 0$ such that $w(s') = \alpha \cdot w(s)$.

If any $w(s) = \underset{\sim}{0}$ then all $w(s') = \underset{\sim}{0}$, for all $s' \in S$. But then letting $u(x) = 0$ for all x and letting $p(\cdot)$ be any probability distribution on S would satisfy the theorem, since the evaluation function is identically zero.

So let us assume that $w(s) \neq \underset{\sim}{0}$ for every state $s \in S$. Fix $t \in S$. For any $s \in S$ we can find $\alpha(s) > 0$ such that $w(s) = \alpha(s)w(t)$. Let $p(s) = \alpha(s)/(\sum_{s' \in S} \alpha(s'))$ for every $s \in S$, and let $u(x) = (\sum_{s' \in S} \alpha(s'))w(x,t)$, for every $x \in X$. Then $w(x,s) = \alpha(s)w(x,t) = p(s)u(x)$ for every $s \in S$ and $x \in X$. Every $p(s) > 0$ and $\sum_{s' \in S} p(s') = 1$. This proves the theorem.

References

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

[2] Fishburn, P.C.: "Preference-Based Definitions of Subjective Probability," Annals of Mathematical Statistics, 38 (1967), 1605-1617.

[3] Fishburn, P.C.: "Utility Theory," Management Science, 14 (1968), 335-378.

[4] Fishburn, P.S.: "Separation Theorems and Expected Utilities," Journal of Economic Theory, 11 (1975), 16-33.

[5] Herstein, I.N. and J. Milnor: "An Axiomatic Approach to Measurable Utility," Econometrica, 21 (1953), 291-297.

[6] Knight, F.H.: Risk, Uncertainty and Profit, Houghton Mifflin, Boston (1921).

[7] Luce, R.D. and H. Raiffa: Games and Decisions, Wiley, New York (1957).

[8] Rockafellar, R.T.: Convex Analysis, Princeton University Press, Princeton (1970).

[9] Savage, L.J.: The Foundations of Statistics, Wiley, New York (1954).

[10] von Neumann, J. and O. Morgenstern: Theory of Games and Economic Behavior, 2nd ed., Princeton University Press, Princeton (1947).

[11] Wilson, R.: "The Theory of Syndicates," Econometrica, 38 (1968), 119-132.