

Myerson, Roger B.

Working Paper

Axiomatic Foundations of Bayesian Decision Theory

Discussion Paper, No. 671

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Myerson, Roger B. (1986) : Axiomatic Foundations of Bayesian Decision Theory, Discussion Paper, No. 671, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221030>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Center for Mathematical Studies in Economics and Management Science
Northwestern University, Evanston, Illinois 60201

Discussion Paper No. 671

AXIOMATIC FOUNDATIONS OF BAYESIAN DECISION THEORY

by

Roger B. Myerson

January 1986

Abstract. This paper offers a unified axiomatic development of the fundamental theorems of decision theory: the maximization of subjective expected utility, and the use of Bayes' formula for updating subjective probabilities. The main pedagogical innovation is an integrated treatment of Bayesian updating after all events, including events that have zero subjective probability. Finite sets of prizes and states are assumed, together with extraneous events having objective probabilities such as could be provided by a roulette wheel. Methods for assessing subjective probability and utility functions are discussed. Decision-theoretically equivalent representations are characterized.

Acknowledgements. This paper was written with support of a fellowship from the Alfred P. Sloan Foundation.

AXIOMATIC FOUNDATIONS OF BAYESIAN DECISION THEORY

by Roger B. Myerson

1. Basic definitions.

At some point, anyone who is interested in the mathematical social sciences should ask the question, why should we expect that any simple quantitative model can give a reasonable description of people's behavior? The fundamental results of decision theory directly address this question, by showing that any decision-maker who satisfies certain intuitive axioms should always behave so as to maximize the mathematical expected value of some utility function, with respect to some subjective probability distribution. That is, any rational decision-maker's behavior should be describable by a utility function, which gives a quantitative characterization of his preferences for outcomes or prizes, and a subjective probability distribution, which characterizes his beliefs about all relevant unknown factors. This result may be called the subjective expected utility theorem. A further result is that, when new information becomes available to such a decision-maker, his subjective probabilities should be revised in accordance with a mathematical equation known as Bayes' formula.

There is a vast literature on axiomatic derivations of the subjective expected utility theorem, beginning with von Neumann and Morgenstern [1947] and Savage [1954]; for an overview, see Fishburn [1968] and [1970]. The goal of this paper is to fill the need for a simple axiomatic derivation that treats utility, subjective probability, and Bayesian updating together in a unified way, so as to provide a short self-contained introduction to the basic results

of decision theory. There are relatively few new ideas or techniques introduced here. The main pedagogical innovation may be the development of Bayes' formula together with utility and subjective-probability theory, in a way that may offer a clearer conceptual foundation for the study of sequential rationality in game theory (see Kreps and Wilson [1982]).

Decisions under uncertainty are commonly described in two ways: using a probability model or a state-variable 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 of Luce and Raiffa [1957]). In a state-variable model, lotteries are functions from a set of possible states of nature into a set of prizes (for example, see Chapter 13 of Luce and Raiffa).

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 prizes will depend on events which have obvious objective probabilities; we shall refer to such events as objective unknowns. These gambles are the "roulette lotteries" of Anscombe and Aumann [1963], or the "risks" of Knight [1921]. 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 of an urn containing a known population of identically sized but differently colored balls, all could be adequately described in a probability model. An important assumption being used here is that two objective unknowns with the same probability are completely equivalent for 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 determined by tossing 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 result 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 Anscombe and Aumann [1963] or the "uncertainties" of Knight [1921]. 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.

In this paper, we define our lotteries so as to include both the probability and the state-variable models as special cases. That is, we study lotteries in which the prize may depend on both objective unknowns (which may be directly described by probabilities) and subjective unknowns (which must be described by a state-of-nature variable). (In the terminology of Fishburn [1970], we are allowing extraneous probabilities in our model.)

Let us now develop some basic notation. For any finite set Z , we let $\Delta(Z)$ denote the set of probability distributions over the set Z . That is,

$$\Delta(Z) = \{q:Z \rightarrow \mathbb{R} \mid \sum_{y \in Z} q(y) = 1 \text{ and } q(z) \geq 0, \forall z \in Z\}.$$

Let X denote the set of possible prizes, and let Ω denote the set of possible states of nature. To simplify the mathematics, we assume that X and Ω are both finite sets. We define a lottery to be any function f which specifies a nonnegative real number $f(x|t)$, for every prize x in X and every state t in Ω , such that $\sum_{x \in X} f(x|t) = 1$ for every t in Ω . Let L denote the set of all such lotteries. That is,

$$L = \{f:\Omega \rightarrow \Delta(X)\}.$$

For any state t in Ω and any lottery f in L , $f(\cdot|t)$ denotes the probability distribution over X designated by f in state t . That is,

$$f(\cdot|t) = (f(x|t))_{x \in X} \in \Delta(X).$$

Each number $f(x|t)$ here is to be interpreted as the (objective) probability of getting prize x in lottery f if t is the true 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 has been specified, only objective probabilities will remain, and an objective probability distribution over the possible prizes can be calculated for any well-defined gamble. So our formal definition of a lottery allows us to represent any gamble in which the prize may depend on both objective and subjective unknowns.

A "prize" in our sense could be any commodity bundle or resource allocation. We are assuming that the prizes in X have been defined so that they are mutually exclusive and exhaust the possible consequences of the decision-maker's decisions. Furthermore, we assume that each prize in X represents a complete specification of all aspects that the decision-maker cares about in the situation that results from his decisions. Thus, the decision-maker should be able to assess a preference ordering over the set of lotteries, given any information that he might have about the state of nature.

The information that the decision-maker might have about the true state of nature can be described by an event, which is a nonempty subset of Ω . We let F denote the set of all such events, so that

$$F = \{S \mid S \subseteq \Omega \text{ and } S \neq \emptyset\}.$$

For any two lotteries f and g in L , and any event S in F , we write $f \succeq_S g$ if and only if the lottery f would be at least as desirable as g , in the opinion

of the decision-maker, if he learned that the true state of nature was in the set S . That is, $f \succeq_S g$ iff the decision-maker would be willing to choose the lottery f when he has to choose between f and g and he knows only that the event S had occurred. Given this relation (\succeq_S), we define relations ($>_S$) and (\sim_S) so that

$$f \sim_S g \text{ iff } f \succeq_S g \text{ and } g \succeq_S f;$$

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

That is, $f \sim_S g$ means that the decision-maker would be indifferent between f and g , if he had to choose between them after learning S ; and $f >_S g$ means that he would strictly prefer f over g in this situation.

For any number α such that $0 \leq \alpha \leq 1$, and for any two lotteries f and g in L , $\alpha f + (1 - \alpha)g$ denotes the lottery in L such that

$$(\alpha f + (1 - \alpha)g)(x|t) = \alpha f(x|t) + (1 - \alpha)g(x|t), \quad \forall x \in X, \quad \forall t \in \Omega.$$

To interpret this definition, suppose that a ball is going to be drawn from an urn in which α is the proportion of black balls and $1 - \alpha$ is the proportion of white balls. Suppose that, if the ball is black then the decision-maker will get to play lottery f , and if the ball is white then the decision-maker will get to play lottery g . Then the decision-maker's ultimate probability of getting prize x if t is the true state is $\alpha f(x|t) + (1 - \alpha)g(x|t)$. Thus, $\alpha f + (1 - \alpha)g$ represents the compound lottery which is built up from f and g by this random lottery-selection process.

For any prize x , we let $[x]$ denote the lottery that always gives prize x for sure. That is, for every state t ,

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

2. Axioms.

We now list some basic properties that a rational decision-maker's preferences may be expected to satisfy. Unless otherwise stated, these axioms are to hold for all lotteries $e, f, g,$ and h in L , for all events S and T in F , and for all numbers α and β between 0 and 1.

Axioms 1A and 1B assert that preferences should always form a complete transitive order over the set of lotteries.

Axiom 1A (completeness). $f \succeq_S g$ or $g \succeq_S f$.

Axiom 1B (transitivity). If $f \succeq_S g$ and $g \succeq_S h$ then $f \succeq_S h$.

It is straightforward to check that Axiom 1B implies a number of other transitivity results, such as: if $f \sim_S g$ and $g \sim_S h$ then $f \sim_S h$; and if $f \succ_S g$ and $g \succeq_S h$ then $f \succ_S h$.

Axiom 2 asserts that only possible states are relevant to the decision-maker, so that, given an event S , he would be indifferent between two lotteries that differ only in states outside of S .

Axiom 2 (relevance). If $f(\bullet|t) = g(\bullet|t) \quad \forall t \in S$, then $f \sim_S g$.

Axiom 3 asserts that a higher probability of getting a better lottery is always better.

Axiom 3 (monotonicity). If $f \succ_S h$ and $0 \leq \beta < \alpha \leq 1$, then $\alpha f + (1 - \alpha)h \succ_S \beta f + (1 - \beta)h$.

Building on Axiom 3, Axiom 4 asserts that $\gamma f + (1 - \gamma)h$ gets better in a continuous manner as γ increases, so that any lottery that is ranked between

f and h is just as good as some randomization between f and h.

Axiom 4 (continuity). If $f \succeq_S g$ and $g \succeq_S h$ then there exists some number γ such that $0 \leq \gamma \leq 1$ and $g \sim_S \gamma f + (1 - \gamma)h$.

The substitution axioms (also known as independence or sure-thing axioms) are probably the most important axioms in our system, in the sense that they generate strong restrictions on what the decision-maker's preferences must look like even without the other axioms. They should also be very intuitive axioms. The idea that they express is that, if the decision-maker must choose between two alternatives, and if there are two mutually exclusive events, one of which must occur, such that in each event he would prefer the first alternative, then he must prefer the first alternative before he learns which event occurs. (Otherwise, he would be expressing a preference that he would be sure to want to reverse after learning which of these events was true!) In Axioms 5A and 5B, these events are objective randomizations in a random lottery-selection process, as discussed in the preceding section. In Axioms 6A and 6B, these events are subjective unknowns, subsets of Ω .

Axiom 5A (objective substitution). If $e \succeq_S f$ and $g \succeq_S h$ and $0 \leq \alpha \leq 1$, then $\alpha e + (1 - \alpha)g \succeq_S \alpha f + (1 - \alpha)h$.

Axiom 5B (strict objective substitution). If $e \succ_S f$ and $g \succeq_S h$ and $0 < \alpha \leq 1$, then $\alpha e + (1 - \alpha)g \succ_S \alpha f + (1 - \alpha)h$.

Axiom 6A (subjective substitution). If $f \succeq_S g$ and $f \succeq_T g$ and $S \cap T = \emptyset$, then $f \succeq_{S \cup T} g$.

Axiom 6B (strict subjective substitution). If $f \succ_S g$ and $f \succ_T g$ and $S \cap T = \emptyset$, then $f \succ_{S \cup T} g$.

Axiom 7 asserts that the decision-maker is never indifferent between all prizes. This is just a regularity condition, to make sure that there is something of interest that could happen in each state.

Axiom 7 (interest). For every state t in Ω , there exist prizes y and z in X such that $[y] \succ_{\{t\}} [z]$.

Axiom 8 is optional in our analysis, in the sense that we can state a version of our main result with or without this axiom. It asserts that the decision-maker has the same preference ordering over objective gambles in all states of nature. If this axiom fails, it is because the same prize might be valued differently in different states.

Axiom 8 (state-neutrality). For any two states r and t in Ω , if $f(\cdot|r) = f(\cdot|t)$ and $g(\cdot|r) = g(\cdot|t)$ and $f \succ_{\{r\}} g$, then $f \succ_{\{t\}} g$.

3. The main representation theorem.

A conditional-probability function on Ω is any function $p:F \rightarrow \Delta(\Omega)$, that specifies conditional probabilities $p(t|S)$, for every state t in Ω and every event S , such that

$$\sum_{t \in S} p(t|S) = 1, \quad \forall S \in F.$$

Given any such conditional-probability function, we may write

$$p(R|S) = \sum_{r \in R} p(r|S), \quad \forall R \subseteq \Omega, \quad \forall S \in F.$$

A utility function can be any function from $X \times \Omega$ into the real numbers \mathbb{R} . We say that a utility function $u: X \times \Omega \rightarrow \mathbb{R}$ is state-independent iff it does not actually depend on the state, so that there exists some function $U: X \rightarrow \mathbb{R}$ such that $u(x,t) = U(x)$ for all x and t .

Given any such conditional-probability function p and any utility function u , and given any lottery f in L and any event S in F , we let $E_p(u(f)|S)$ denote the expected utility value of the prize determined by f , when $p(\cdot|S)$ is the probability distribution for the true state of nature. That is,

$$E_p(u(f)|S) = \sum_{t \in S} p(t|S) \sum_{x \in X} u(x,t) f(x|t).$$

Theorem 1. Axioms 1AB, 2, 3, 4, 5AB, 6AB, and 7 are jointly satisfied if and only if there exists a utility function $u: X \times \Omega \rightarrow \mathbb{R}$ and a conditional-probability function $p: F \rightarrow \Delta(\Omega)$ such that

- (1) $\max_{x \in X} u(x,t) = 1$ and $\min_{x \in X} u(x,t) = 0$, $\forall t \in \Omega$;
- (2) $p(R|T) = p(R|S) p(S|T)$, $\forall R, \forall S$, and $\forall T$ such that
 $R \subseteq S \subseteq T \subseteq \Omega$ and $S \neq \emptyset$;
- (3) $f \succeq_S g$ if and only if $E_p(u(f)|S) \geq E_p(u(g)|S)$, $\forall f, g \in L$,
 $\forall S \in F$.

Furthermore, given these Axioms 1AB - 7, Axiom 8 is also satisfied if and only if conditions (1)-(3) here can be satisfied with a state-independent utility function.

In this theorem, condition (1) is a normalization condition, asserting that we can choose our utility functions to range between 0 and 1 in every state. (Recall that X and Ω are assumed to be finite.) Condition (2) is a version of Bayes' formula, which establishes how the conditional probabilities

assessed in one event must be related to conditional probabilities assessed in another. The most important part of the theorem is condition (3), however, which asserts that the decision-maker always prefers lotteries with higher expected utility. By condition (3), once we have assessed u and p , we can predict the decision-maker's optimal choice in any decision-making situation. (He will choose the lottery with the highest expected utility among those available to him, using his subjective probabilities conditioned on whatever event in Ω he has observed.) Furthermore, with X and Ω finite, there are only finitely many utility and probability numbers to assess. Thus, the decision-maker's preferences over all of the infinitely many lotteries in L can be completely characterized by finitely many numbers.

To apply this result in practice, we need a procedure for assessing the utilities $u(x,t)$ and the probabilities $p(t|S)$, for all x , t , and S . As Raiffa [1968] has emphasized, such procedures do exist, and form the basis of practical decision analysis. To define one such assessment procedure, and to prove Theorem 1, we begin by defining some special lotteries, using the assumption that the decision-maker's preferences satisfy axioms 1A - 7.

Let a_1 be a lottery that gives the decision-maker one of the best prizes in every state; and let a_0 be a lottery that gives him one of the worst prizes in every state. That is, for every state t , $a_1(y|t) = 1 = a_0(z|t)$ for some prizes y and z such that, for every x in X , $y \succ_{\{t\}} x \succ_{\{t\}} z$. Such best and worst prizes can be found in every state because the preference relation $(\succ_{\{t\}})$ forms a transitive ordering over the finite set X .

For any event S in F , let b_S denote the lottery such that

$$b_S(\cdot|t) = \begin{cases} a_1(\cdot|t) & \text{if } t \in S, \\ a_0(\cdot|t) & \text{if } t \notin S. \end{cases}$$

That is, b_S is a "bet on S" which gives the best possible prize if S occurs and gives the worst possible prize otherwise.

For any prize x and any state t , let $c_{x,t}$ be the lottery such that

$$c_{x,t}(\cdot|r) = \begin{cases} [x](\cdot|t) & \text{if } r = t. \\ a_0(\cdot|r) & \text{if } r \neq t. \end{cases}$$

That is, $c_{x,t}$ is the lottery that always gives the worst prize, except in state t , when it gives prize x .

We can now define a procedure to assess the utilities and probabilities that satisfy the theorem, given preferences that satisfy the axioms. For each x and t , first ask the decision-maker "for what number β would you be indifferent between $[x]$ and $\beta a_1 + (1 - \beta)a_0$, if you knew that t were the true state of nature?" By the continuity axiom, such a number must exist. Then let $u(x,t)$ equal the number that he specifies, so that

$$[x] \sim_{\{t\}} u(x,t) a_1 + (1 - u(x,t)) a_0.$$

For each t and S , ask the decision-maker "for what number γ would you be indifferent between $b_{\{t\}}$ and $\gamma a_1 + (1 - \gamma)a_0$ if you knew that the true state was in S ?" Again, such a number must exist, by the continuity axiom. (The subjective substitution axiom guarantees that $a_1 \succeq_S b_{\{t\}} \succeq_S a_0$.) Then let $p(t|S)$ equal the number that he specifies, so that

$$b_{\{t\}} \sim_S p(t|S) a_1 + (1 - p(t|S)) a_0.$$

Thus, finitely many questions suffice to assess the probabilities and utilities which completely characterize the decision-maker's preferences. We must now show that defining u and p in this way does satisfy the conditions of the theorem.

Derivation of condition (3) from the axioms. The relevance axiom and the definition of $u(x,t)$ implies that, for every state r ,

$$c_{x,t} \sim_{\{r\}} u(x,t) b_{\{t\}} + (1 - u(x,t)) a_0.$$

Then subjective substitution implies that, for every event S ,

$$c_{x,t} \sim_S u(x,t) b_{\{t\}} + (1 - u(x,t)) a_0.$$

Axioms 5A and 5B together imply that $f \succeq_S g$ if and only if

$$(1/|\Omega|) f + (1 - (1/|\Omega|)) a_0 \succeq_S (1/|\Omega|) g + (1 - (1/|\Omega|)) a_0.$$

(Here, $|\Omega|$ denotes the number of states in the set Ω .) Notice that

$$(1/|\Omega|) f + (1 - (1/|\Omega|)) a_0 = (1/|\Omega|) \sum_{t \in \Omega} \sum_{x \in X} f(x|t) c_{x,t}.$$

But, by repeated application of the objective substitution axiom,

$$\begin{aligned} & (1/|\Omega|) \sum_{t \in \Omega} \sum_{x \in X} f(x|t) c_{x,t} \\ & \sim_S (1/|\Omega|) \sum_{t \in \Omega} \sum_{x \in X} f(x|t) (u(x,t) b_{\{t\}} + (1 - u(x,t)) a_0) \\ & \sim_S \frac{1}{|\Omega|} \sum_{t \in \Omega} \sum_{x \in X} f(x|t) (u(x,t) (p(t|S) a_1 + (1 - p(t|S)) a_0) + (1 - u(x,t)) a_0) \\ & = \frac{1}{|\Omega|} \sum_{t \in \Omega} \sum_{x \in X} f(x|t) u(x,t) p(t|S) a_1 + \\ & \quad (1 - \frac{1}{|\Omega|} \sum_{t \in \Omega} \sum_{x \in X} p(t|S) u(x,t) p(t|S)) a_0 \\ & = (1/|\Omega|) E_p(u(f)|S) a_1 + (1 - (1/|\Omega|) E_p(u(f)|S)) a_0. \end{aligned}$$

Similarly,

$$\begin{aligned} & (1/|\Omega|) g + (1 - (1/|\Omega|)) a_0 \\ & \sim_S (1/|\Omega|) E_p(u(g)|S) a_1 + (1 - (1/|\Omega|) E_p(u(g)|S)) a_0. \end{aligned}$$

Thus, by transitivity, $f \succeq_S g$ if and only if

$$\begin{aligned} & (1/|\Omega|) E_p(u(f)|S) a_1 + (1 - (1/|\Omega|) E_p(u(f)|S)) a_0 \\ & \succeq_S (1/|\Omega|) E_p(u(g)|S) a_1 + (1 - (1/|\Omega|) E_p(u(g)|S)) a_0. \end{aligned}$$

But by monotonicity, this final relation holds if and only if

$$E_p(u(f)|S) \geq E_p(u(g)|S).$$

because interest and strict subjective substitution guarantee that $a_1 \succ_S a_0$.

Thus, condition (3) is satisfied.

Derivation of condition (2) from the axioms. For any events R and S,

$$\frac{1}{|R|} b_R + \left(1 - \frac{1}{|R|}\right) a_0 = \frac{1}{|R|} \sum_{r \in R} b_{\{r\}}$$

$$\sim_S \frac{1}{|R|} \sum_{r \in R} (p(r|S) a_1 + (1 - p(r|S)) a_0)$$

$$= (1/|R|)(p(R|S) a_1 + (1 - p(R|S)) a_0) + (1 - (1/|R|)) a_0.$$

by objective substitution. ($|R|$ is the number of states in the set R.) Then, using Axioms 5A and 5B, we get

$$b_R \sim_S p(R|S) a_1 + (1 - p(R|S)) a_0.$$

By the relevance axiom, $b_S \sim_S a_1$ and, for any r not in S, $b_{\{r\}} \sim_S a_0$. So the above formula implies (using monotonicity and interest) that $p(r|S) = 0$ if $r \notin S$, and $p(S|S) = 1$. Thus, p is a conditional-probability function, as defined above.

Now, suppose that $R \subseteq S \subseteq T$. Using $b_S \sim_S a_1$ again, we get

$$b_R \sim_S p(R|S) b_S + (1 - p(R|S)) a_0.$$

Furthermore, since b_R , b_S , and a_0 all give the same worst prize outside S, relevance also implies

$$b_R \sim_{T \setminus S} p(R|S) b_S + (1 - p(R|S)) a_0.$$

So, by subjective and objective substitution,

$$b_R \sim_T p(R|S) b_S + (1 - p(R|S)) a_0$$

$$\sim_T p(R|S)(p(S|T) a_1 + (1 - p(S|T)) a_0) + (1 - p(R|S)) a_0.$$

$$= p(R|S) p(S|T) a_1 + (1 - p(R|S) p(S|T)) a_0.$$

But $b_R \sim_T p(R|T) b_S + (1 - p(R|T)) a_0$. If $b_S \succ_T a_0$ then monotonicity implies that $p(R|T) = p(R|S) p(S|T)$. On the other hand, if $b_S \sim_T a_0$ then $b_R \sim_T a_0$ also (using subjective substitution and transitivity), so that

$p(R|T) = p(S|T) = 0$ and $p(R|T) = p(R|S)p(S|T)$ again. Thus, Bayes' formula (2) follows from the axioms.

If y is the best prize and z is the worst prize in state t , then $[y] \sim_{\{t\}} a_1$ and $[z] \sim_{\{t\}} a_0$, so that $u(y,t) = 1$ and $u(z,t) = 0$ by monotonicity. So the range condition (1) is also satisfied by the utility function that we have constructed.

If state-neutrality is also given, then the decision-maker will give us the same answer when we assess $u(x,t)$ as when we assess $u(x,r)$ for any other state r (since $[x] \sim_{\{t\}} \beta a_1 + (1 - \beta)a_0$ implies $[x] \sim_{\{r\}} \beta a_1 + (1 - \beta)a_0$, and monotonicity and interest guarantee that his answer is unique). So Axiom 8 implies that u is state-independent.

To complete the proof of the theorem, it remains to show that the existence of functions u and p that satisfy conditions (1) - (3) in the theorem is sufficient to imply all the axioms (using state-independence only for axiom 8). Using the basic mathematical properties of the expected-utility formula, the axioms are all straightforward to verify. To illustrate, we show the proof of one axiom, subjective substitution, and leave the rest as an exercise for the reader.

Suppose that $f \succeq_S g$ and $f \succeq_T g$ and $S \cap T = \emptyset$. By (3), $E_p(u(f)|S) \geq E_p(u(g)|S)$ and $E_p(u(f)|T) \geq E_p(u(g)|T)$. But Bayes formula (2) implies that

$$\begin{aligned} E_p(u(f)|SUT) &= \sum_{t \in SUT} \sum_{x \in X} p(t|SUT) f(x|t) u(x,t) \\ &= \sum_{t \in S} \sum_{x \in X} p(t|S)p(S|SUT)f(x|t)u(x,t) + \sum_{t \in T} \sum_{x \in X} p(t|T)p(T|SUT)f(x|t)u(x,t) \\ &= p(S|SUT) E_p(u(f)|S) + p(T|SUT) E_p(u(f)|S) \end{aligned}$$

and $E_p(u(g)|SUT) = p(S|SUT) E_p(u(g)|S) + p(T|SUT) E_p(u(g)|S)$.

So $E_p(u(f)|SUT) \geq E_p(u(g)|SUT)$ and $f \succeq_{SUT} g$. Q.E.D.

4. Equivalent representations.

When we drop the range condition (1), there can be more than one pair of utility and conditional-probability functions that represent the same decision-maker's preferences, in the sense of condition (3). Such equivalent representations are completely indistinguishable in terms of their decision-theoretic properties, and so we should be suspicious of any theory of economic behavior that requires distinguishing between such equivalent representations. Thus, it may be theoretically important to be able to recognize such equivalent representations.

Given any subjective event S , when we say that a utility function v and a conditional-probability function q represent the preference ordering \succeq_S , we mean that, for every pair of lotteries f and g , $E_q(v(f)|S) \geq E_q(v(g)|S)$ if and only if $f \succeq_S g$.

Theorem 2. Let S in F be any given subjective event. Suppose that the decision-maker's preferences satisfy Axioms 1AB through 7, and let u and p be utility and conditional-probability functions satisfying (1) - (3) in Theorem 1. Then v and q represent the preference ordering \succeq_S if and only if there exists a positive number A and a function $B: S \rightarrow \mathbb{R}$ such that

$$q(t|S) v(x,t) = A p(t|S) u(x,t) + B(t), \quad \forall t \in S, \quad \forall x \in X.$$

Proof. Suppose first that A and $B(\cdot)$ exist as described in the theorem. Then, for any lottery f ,

$$E_q(v(f)|S) = \sum_{t \in S} \sum_{x \in X} f(x|t) q(t|S) v(x,t)$$

$$\begin{aligned}
 &= \sum_{t \in S} \sum_{x \in X} f(x|t) (A p(t|S) u(x,t) + B(t)) \\
 &= A \sum_{t \in S} \sum_{x \in X} f(x|t) p(t|S) u(x,t) + \sum_{t \in S} B(t) \sum_{x \in X} f(x|t) \\
 &= A E_p(u(f)|S) + \sum_{t \in S} B(t),
 \end{aligned}$$

because $\sum_{x \in X} f(x|t) = 1$. So expected v -utility with respect to q is an increasing linear function of expected u -utility with respect to p , since $A > 0$. Thus $E_q(v(f)|S) \geq E_q(v(g)|S)$ if and only if $E_p(u(f)|S) \geq E_p(u(g)|S)$, and so v and q together represent the same preference ordering over lotteries as u and p .

Conversely, suppose now that v and q represent the same preference ordering as u and p . Pick any prize x and state t , and let

$$\lambda = (E_q(v(c_{x,t})|S) - E_q(v(a_0)|S)) / (E_q(v(a_1)|S) - E_q(v(a_0)|S)).$$

Then, by the linearity of the expected-value operator,

$$\begin{aligned}
 E_q(v(\lambda a_1 + (1 - \lambda)a_0)|S) &= E_q(v(a_0)|S) + \lambda (E_q(v(a_1)|S) - E_q(v(a_0)|S)) \\
 &= E_q(v(c_{x,t})|S),
 \end{aligned}$$

so $c_{x,t} \sim_S \lambda a_1 + (1 - \lambda)a_0$. In the proof of Theorem 1, we constructed u and p so that

$$\begin{aligned}
 c_{x,t} &\sim_S u(x,t) b_{\{t\}} + (1 - u(x,t)) a_0 \\
 &\sim_S u(x,t)(p(t|S) a_1 + (1 - p(t|S)) a_0) + (1 - u(x,t)) a_0 \\
 &\sim_S p(t|S) u(x,t) a_1 + (1 - p(t|S) u(x,t)) a_0.
 \end{aligned}$$

The monotonicity axiom guarantees that only one randomization between a_1 and a_0 can be just as good as $c_{x,t}$, so

$$\lambda = p(t|S) u(x,t).$$

But $c_{x,t}$ differs from a_0 only in state t , where it gives prize x instead of the worst prize, so

$$E_q(v(c_{x,t})|S) - E_q(v(a_0)|S) = q(t|S)(v(x,t) - \min_{z \in X} v(z,t)).$$

Thus, going back to the definition of λ , we get

$$p(t|S) u(x,t) = q(t|S)(v(x,t) - \min_{z \in X} v(z,t)) / (E_q(v(a_1)|S) - E_q(v(a_0)|S)).$$

Now let

$$A = (E_q(v(a_1)|S) - E_q(v(a_0)|S)),$$

and let

$$B(t) = q(t|S) \min_{z \in X} v(z,t).$$

Then

$$A p(t|S) u(x,t) + B(t) = q(t|S) v(x,t).$$

Notice that A is independent of x and t , and $B(t)$ is independent of x .

Furthermore, $A > 0$ because $a_1 \succ_S a_0$ implies $E_q(v(a_1)|S) > E_q(v(a_0)|S)$.

Q.E.D.

It is easy to see from Theorem 2 that more than one probability distribution can represent the decision-maker's beliefs given some event S . In fact, we can make the probability distribution $q(\cdot|S)$ almost anything and still satisfy the equation in Theorem 2, as long as we make reciprocal changes in v , so as to keep the left-hand side of the equation the same. The way to eliminate this indeterminacy is to assume Axiom 8 and require utility functions to be state-independent.

Theorem 3. Let S in F be any given subjective event. Suppose that the decision-maker's preferences satisfy Axioms 1AB through 8, and let u and p be the state-independent utility function and conditional-probability function that satisfy (1) - (3) in Theorem 1. Let v be a state-independent utility function, let q be a conditional-probability function, and suppose that v and q represent the preference ordering \succeq_S . Then

$$q(t|S) = p(t|S), \quad \forall t \in S,$$

and there exist numbers A and C such that $A > 0$ and

$$v(x) = A u(x) + C, \quad \forall x \in X.$$

(For simplicity, we can write $v(x)$ and $u(x)$ here, instead of $v(x,t)$ and $u(x,t)$, because both functions are state-independent.)

Proof. Let $A = (E_q(v(a_1)|S) - E_q(v(a_0)|S))$, and let $C = \min_{z \in X} v(z)$. Then, from the proof of Theorem 2,

$$A p(t|S) u(x) + q(t|S) C = q(t|S) v(x), \quad \forall x \in X, \quad \forall t \in S.$$

Summing this equation over all t in S , we get $A u(x) + C = v(x)$. Then, substituting this equation back, and letting x be the best prize, so that $u(x) = 1$, we get

$$A p(t|S) + q(t|S) C = A q(t|S) + q(t|S) C.$$

Since $A > 0$, the theorem follows. Q.E.D.

REFERENCES

- F. J. Anscombe and R. J. Aumann [1963], "A Definition of Subjective Probability," Annals of Mathematical Statistics 34, 199-205.
- P. C. Fishburn [1968], "Utility Theory," Management Science 14, 335-378.
- P. C. Fishburn [1970], Utility Theory for Decision-Making, Wiley, New York.
- F. H. Knight [1921], Risk, Uncertainty and Profit, Houghton Mifflin, Boston.
- D. Kreps and R. Wilson [1982], "Sequential Equilibria," Econometrica 50, 863-894.
- R. D. Luce and H. Raiffa [1957], Games and Decisions, Wiley, New York.
- J. von Neumann and O. Morgenstern [1947], Theory of Games and Economic Behavior, 2nd ed., Princeton University Press, Princeton.
- H. Raiffa [1968], Decision Analysis, Addison-Wesley, Reading, Massachusetts.
- L. J. Savage [1954], The Foundations of Statistics, Wiley, New York.