E

Expected Utility Hypothesis

Mark J. Machina

Abstract

The expected utility hypothesis – that is, the hypothesis that individuals evaluate uncertain prospects according to their expected level of 'satisfaction' or 'utility' – is the predominant descriptive and normative model of choice under uncertainty in economics. It provides the analytical underpinnings for the economic theory of risk-bearing, including its applications to insurance and financial decisions, and has been formally axiomatized under conditions of both objective (probabilistic) and subjective (event-based) uncertainty. In spite of evidence that individuals may systematically depart from its predictions, and the development of alternative models, expected utility remains the leading model of economic choice under uncertainty.

Keywords

Arrow–Pratt index of absolute risk aversion; Bernoulli, D.; Bernoulli, N.; Cobb–Douglas functions; Comparative likelihood; Consumer theory; Cramer, G.; Environmental economics; Expected utility hypothesis; First-order stochastic dominance preference; Increasing risk; Independence axiom; Inequality (measurement); International trade; Lotteries; Malinvaud, E.; Marschak, J.; Menger, K.; Objective vs. subjective uncertainty; Ordinal revolution; Preference functions; Preference orderings; Probability; Risk; Risk aversion; St Petersburg paradox; Stochastic dominance; Subjective probability; Sure-thing principle; Transitivity; Uncertainty; von Neumann–Morgenstern utility function

JEL Classifications

D8

The expected utility hypothesis is the predominant descriptive and prescriptive theory of individual choice under conditions of risk or uncertainty.

The expected utility hypothesis of behaviour towards risk is the hypothesis that the individual possesses (or acts as if possessing) a 'von Neumann–Morgenstern utility function' $U(\cdot)$ or 'von Neumann–Morgenstern utility index' $\{U_i\}$ defined over some set $\mathscr{X}$ of alternative possible outcomes, and when faced with alternative risky prospects or 'lotteries' over these outcomes, will choose the prospect that maximizes the expected

This chapter was originally published in *The New Palgrave Dictionary of Economics*, 2nd edition, 2008. Edited by Steven N. Durlauf and Lawrence E. Blume

Palgrave Macmillan (ed.), *The New Palgrave Dictionary of Economics*,
DOI 10.1057/978-1-349-95121-5_127-2

value of $U(\cdot)$ or $\{U_i\}$. Since the outcomes could be alternative wealth levels, multidimensional commodity bundles, time streams of consumption, or even non-numerical consequences (such as a trip to Paris), this approach can be applied to a tremendous variety of situations, and most theoretical research in the economics of uncertainty, as well as virtually all applied work in the field (for example, insurance or investment decisions) is undertaken in the expected utility framework.

As a branch of modern consumer theory (for example, Debreu 1959, ch. 4), the expected utility model proceeds by specifying a set of objects of choice and assuming that the individual possesses a preference ordering over these objects which may be represented by a real-valued maximand or 'preference function' $V(\cdot)$, in the sense that one object is preferred to another if and only if it is assigned a higher value by this preference function. However, the expected utility model differs from the theory of choice over non-stochastic commodity bundles in two important respects. The first is that, since it is a theory of choice under uncertainty, the objects of choice are not deterministic outcomes but rather uncertain prospects. The second difference is that, unlike in the non-stochastic case, the expected utility model imposes a very specific restriction on the functional form of the preference function $V(\cdot)$.

The formal representation of the objects of choice, and hence of the expected utility preference function, depends upon the set of possible outcomes. When the outcome set $\mathscr{X}=\{x_1,\ldots, x_n\}$ is finite, we can represent any probability distribution over this set by its vector of probabilities $\boldsymbol{P} = (p_1,\ldots,p_n)$ (where $p_i = \text{prob}(x_i)$) and the expected utility preference function takes the form

$$V(\boldsymbol{P}) = V(p_1,\ldots,p_n) \equiv \sum U_i p_i.$$

When the outcome set consists of the real line or some interval subset of it, probability distributions can be represented by their cumulative distribution functions $F(\cdot)$ (where $F(x) = \text{prob}(\tilde{x} \leq x)$), and the expected utility preference function takes the form $V(F) \equiv \int U(x)\mathrm{d}F(x)$ (or $\int U(x)f(x)\mathrm{d}x$ when $F(\cdot)$ possesses a density function $f(\cdot)$). When the outcomes are commodity bundles of the form $(z_1, \ldots, z_m)$, cumulative distribution functions are multivariate, and the preference function takes the form $\int \ldots \int U(z_1, \ldots, z_m)\, dF(z_1, \ldots, z_m)$. The expected utility model derives its name from the fact that in each case the preference function consists of the mathematical expectation of the von Neumann–Morgenstern utility function $U(\cdot)$, $U(\cdot, \ldots, \cdot)$ or utility index $\{U_i\}$ with respect to the probability distribution $F(\cdot)$, $F(\cdot, \ldots, \cdot)$ or $\boldsymbol{P}$.

Mathematically, the hypothesis that the preference function $V(\cdot)$ takes the form of a statistical expectation is equivalent to the condition that it be 'linear in the probabilities', that is, either a weighted sum of the components of $\boldsymbol{P}$ (i.e. $\sum U_i p_i$) or else a weighted integral of the functions $F(\cdot)$ or $f(\cdot)$ ($\int U(x)\mathrm{d}F(x)$ or $\int U(x)f(x)\mathrm{d}x$). Although this still allows for a wide variety of attitudes towards risk depending upon the shape of the utility function $U(\cdot)$ or utility index $\{U_i\}$, the restriction that $V(\cdot)$ be linear in the probabilities is the primary empirical feature of the expected utility model, and provides the basis for many of its observable implications and predictions.

It is important to distinguish between the preference function $V(\cdot)$ and the von Neumann–Morgenstern utility function $U(\cdot)$ (or index $\{U_i\}$) of an expected utility maximizer, in particular with regard to the prevalent though mistaken belief that expected utility preferences are somehow 'cardinal' in a sense not exhibited by preferences over non-stochastic commodity bundles. As with any real-valued representation of a preference ordering, an expected utility *preference function* $V(\cdot)$ is 'ordinal' in that it may be subject to any increasing transformation without affecting the validity of the representation – thus, the preference functions $\int U(x)\mathrm{d}F(x)$ and $[\int U(x)dF(x)]^3$ represent identical risk preferences. On the other hand, the von Neumann–Morgenstern utility function $U(\cdot)$ is 'cardinal' in the sense that a different utility function $U^*(\cdot)$ will generate an ordinally equivalent preference function $V^*(F) \equiv \int U^*(x)\mathrm{d}F(x)$ if and only if it satisfies the cardinal relationship $U^*(x) \equiv a \cdot U(x) + b$ for some $a > 0$ (in which case $V^*(F) \equiv a \cdot V(F) + b$. However, the same distinction holds in the theory of preferences over non-stochastic commodity bundles: the Cobb–Douglas preference function $\alpha \cdot \ln(z_1) + \beta \cdot \ln(z_2) + \gamma \cdot \ln(z_3)$ (written here in its additive form) can be subject to any increasing transformation and is clearly ordinal, even though a vector of parameters $(\alpha^*, \beta^*, \gamma^*)$ will generate an ordinally equivalent additive form $\alpha^* \cdot \ln(z_1) + \beta^* \cdot \ln(z_2) + \gamma^* \cdot \ln(z_3)$ if and only if it satisfies the cardinal relationship $(\alpha^*, \beta^*, \gamma^*) = \lambda \cdot (\alpha, \beta, \gamma)$ for some $\lambda > 0$.

In the case of a simple outcome set of the form $\{x_1, x_2, x_3\}$, it is possible to graphically illustrate the 'linearity in the probabilities' property of expected utility preferences. Since every probability distribution (p_1, p_2, p_3) over these outcomes must satisfy $p_1 + p_2 + p_3 = 1$, we may represent such distributions by points in the unit triangle in the (p_1, p_3) plane, with p_2 given by $p_2 = 1 - p_1 - p_3$ (Figs. 1 and 2). Since they represent the loci of solutions to the equations

$$\begin{aligned} U_1p_1 + U_2p_2 + U_3p_3 &= U_2 - [U_2 - U_1] \cdot p_1 \\ &\quad + [U_3 - U_2] \cdot p_3 \\ &= \text{constant} \end{aligned}$$

for the fixed utility indices $\{U_1, U_2, U_3\}$, the indifference curves of an expected utility maximizer consist of parallel straight lines in the triangle, with slope $[U_2 - U_1]/[U_3 - U_2]$, as illustrated by the solid lines in Fig. 1. Indifference curves which do *not* satisfy the expected utility hypothesis (that is, are not linear in the probabilities) are illustrated by the solid curves in Fig. 2.

When the outcomes consist of different wealth levels $x_1 < x_2 < x_3$, this diagram can be used to illustrate other possible features of an expected utility maximizer's attitudes towards risk. On the principle that more wealth is better, it is typically postulated that any change in a distribution (p_1, p_2, p_3) which increases p_3 at the expense of p_2, increases p_2 at the expense of p_1, or both, will be preferred: this property is known as 'first-order stochastic dominance preference'. Since such shifts of probability mass are represented by north, west, or north-west movements in the diagram, first-order stochastic dominance preference is equivalent to the condition that indifference curves are upward sloping, with more preferred indifference curves lying to the north-west. Algebraically, this is equivalent to the condition $U_1 < U_2 < U_3$.

Another widely (though not universally) hypothesized aspect of attitudes towards risk is that of 'risk aversion' (for example, Arrow 1974, ch. 3; Pratt 1964). To illustrate this property, consider the dashed lines in Fig. 1, which represent loci of solutions to the equations

$$\begin{aligned} x_1p_1 + x_2p_2 + x_3p_3 &= x_2 - [x_2 - x_1] \cdot p_1 \\ &\quad + [x_3 - x_2] \cdot p_3 \\ &= \text{constant} \end{aligned}$$

and hence may be termed 'iso-expected *value* loci'. Since north-east movements along any of these loci consist of increasing the tail probabilities p_1 and p_3 at the expense of the middle probability p_2 in a manner which preserves the mean of the distribution, they correspond to what are termed 'mean-preserving increases in risk' (Rothschild and Stiglitz 1970, 1971). An individual is said to be 'risk averse' if such increases in risk always lead to less preferred indifference curves, which is equivalent to the graphical condition that the indifference curves be steeper than the iso-expected value loci. Since the slope of the latter is given by $[x_2 - x_1]/[x_3 - x_2]$, this is equivalent to the algebraic condition that $[U_2 - U_1]/[x_2 - x_1] > [U_3 - U_2]/[x_3 - x_2]$. Conversely, individuals who *prefer* mean-preserving increases in risk are termed 'risk loving': such individuals' indifference curves will be flatter than the iso-expected value loci, and their utility indices will satisfy $[U_2 - U_1]/[x_2 - x_1] < [U_3 - U_2]/[x_3 - x_2]$.

Note finally that the indifference map in Fig. 1 indicates that the lottery $\boldsymbol{P}$ is indifferent to the origin, which represents the degenerate lottery yielding x_2 with certainty. In such a case the amount x_2 is said to be the 'certainty equivalent' of the lottery $\boldsymbol{P}$. The fact that the origin lies on a lower iso-expected value locus than $\boldsymbol{P}$ reflects a general property of risk-averse preferences, namely, that the certainty equivalent of any lottery will always be less than its mean. (For risk lovers, the opposite is the case.)

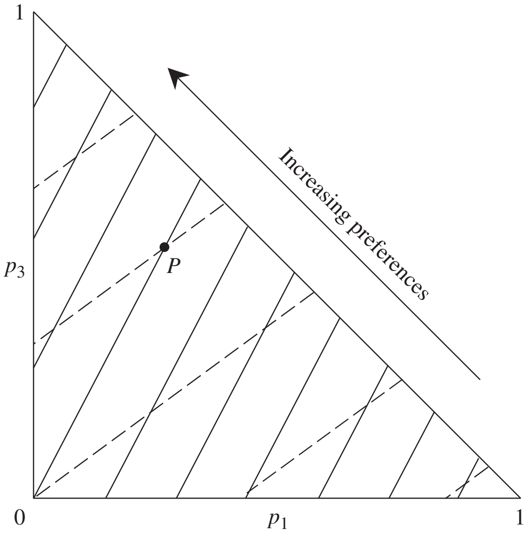

Expected Utility Hypothesis, Fig. 1 Expected utility indifference curves

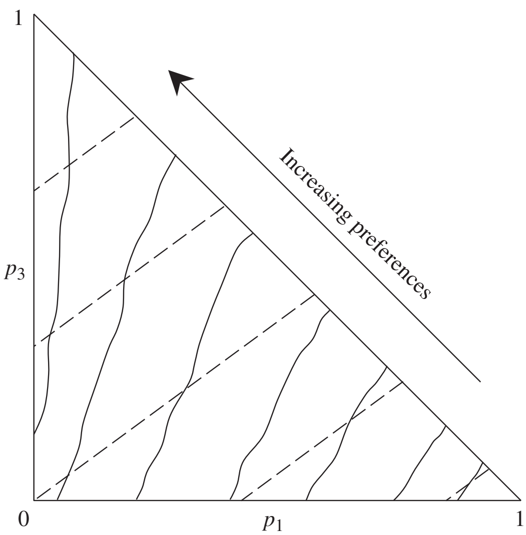

Expected Utility Hypothesis, Fig. 2 Non-expected utility indifference curves

When the outcomes are elements of the real line, it is possible to represent the above (as well as other) aspects of preferences in terms of the shape of the von Neumann–Morgenstern utility function $U(\cdot)$, as seen in Figs. 3 and 4. In each figure, consider the lottery which assigns the probabilities 2/3:1/3 to the outcome levels x': x''. The

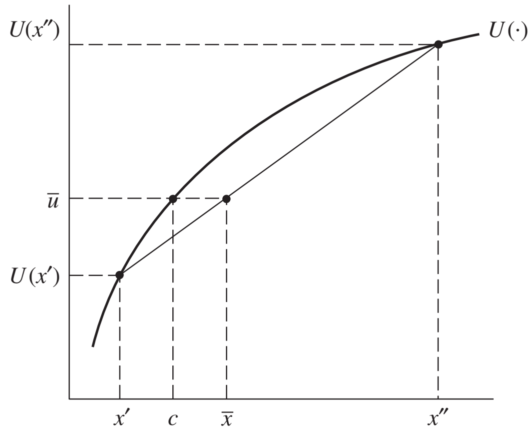

Expected Utility Hypothesis, Fig. 3 Von Neumann-Morgenstern utility function of a risk averse individual

expected value of this lottery, $\bar{x} = 2/3 \cdot x' + 1/3 \cdot x''$, lies between these two values, two-thirds of the way towards x'. The expected *utility* of this lottery, $\bar{u} = 2/3 \cdot U(x') + 1/3 \cdot U(x'')$ lies between $U(x')$ and $U(x'')$ on the vertical axis, two-thirds of the way towards $U(x')$. The point $(\bar{x},\bar{u})$ thus lies on the line segment connecting the points $(x', U(x'))$ and $(x'', U(x''))$, two-thirds of the way towards the former. In each figure, the certainty equivalent of this lottery is given by the sure outcome c that also yields a utility level of $\bar{u}$.

The property of first-order stochastic dominance preference can be extended to the case of distributions over the real line (Quirk and Saposnick 1962), and it is equivalent to the condition that $U(x)$ be an increasing function of x, as in Figs. 3 and 4. It is also possible to generalize the notion of a mean-preserving increase in risk to density functions or cumulative distribution functions (Rothschild and Stiglitz 1970, 1971), and the earlier algebraic condition for risk aversion generalizes to the condition that $U''(x) < 0$ for all x, that is, that the von Neumann–Morgenstern utility function $U(\cdot)$ be concave, as in Fig. 3. As before, risk aversion implies that the certainty equivalent of any lottery will lie below its mean, as seen in Fig. 3; the opposite is true for the convex utility function of a risk lover, as in Fig. 4. Two of the earliest and most important analyses of risk attitudes in terms of the shape of the von

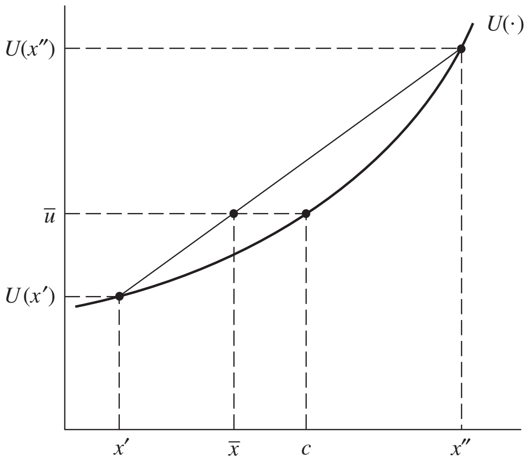

Expected Utility Hypothesis, Fig. 4 von Neumann–Morgenstern utility function of a risk loving individual

Neumann–Morgenstern utility function are those of Friedman and Savage (1948) and Markowitz (1952).

Analytics

The tremendous analytic capabilities of the expected utility model derive largely from the work of Arrow (1974) and Pratt (1964), who showed that the 'degree' of concavity of the utility function provides a measure of an expected utility maximizer's 'degree' of risk aversion. Formally, the Arrow–Pratt characterization of comparative risk aversion is the result that the following conditions on a pair of (increasing, twice differentiable) von Neumann–Morgenstern utility functions $U_a(\cdot)$ and $U_b(\cdot)$ are equivalent:

- $U_a(\cdot)$ is a concave transformation of $U_b(\cdot)$ (that is, $U_a(x) \equiv \rho(U_b(x))$ for some increasing concave function $\rho(\cdot)$),
- $-U''_a(x)/U'_a(x) \geq -U''_b(x)/U'_b(x)$ for each x,
- if c_a and c_b solve $U_a(c_a) = \int U_a(x)\mathrm{d}F(x)$ and $U_b(c_b) = \int U_b(x)\mathrm{d}F(x)$ for some distribution $F(\cdot)$, then $c_a \leq c_b$,

and if $U_a(\cdot)$ and $U_b(\cdot)$ are both concave, these conditions are in turn equivalent to:

- if $r > 0$, $\mathrm{E}[\tilde{z}] > r$, $\mathrm{prob}(\tilde{z} < r) > 0$, and α_a and α_b maximize $\int U_a((I-\alpha)\cdot r + \alpha\cdot z)\mathrm{d}F(z)$ and $\int U_b((I-\alpha)\cdot r + \alpha\cdot z)\mathrm{d}F(z)$ respectively, then $\alpha_a \leq \alpha_b$.

The first two of these conditions provide equivalent formulations of the notion that $U_a(\cdot)$ is a more concave function than $U_b(\cdot)$. The curvature measure $R(x) \equiv -U''(x)/U'(x)$ is known as the 'Arrow–Pratt index of (absolute) risk aversion', and plays a key role in the analytics of the expected utility model. The third condition states that the more risk averse utility function $U_a(\cdot)$ will never assign a higher certainty equivalent to any lottery $F(\cdot)$ than will $U_b(\cdot)$. The final condition pertains to the individuals' respective demands for risky assets. Specifically, assume that each must allocate $\$I$ between two assets, one yielding a riskless (gross) return of r per dollar, and the other yielding a risky return $\tilde{z}$ with a higher expected value but with some chance of doing worse than r. This condition says that the less risk-averse utility function $U_b(\cdot)$ will generate at least as great a demand for the risky asset as the more risk-averse utility function $U_a(\cdot)$. It is important to note that it is the *equivalence* of the above concavity, certainty equivalent and asset demand conditions which makes the Arrow-Pratt characterization such an important result in expected utility theory. (Ross 1981, provides an alternative, stronger, characterization of comparative risk aversion.)

Although the applications of the expected utility model extend to virtually all branches of economic theory (for example, Hey 1979), much of the flavour of these analyses can be gleaned from Arrow's (1974, ch. 3) analysis of the portfolio problem of the previous paragraph. If we rewrite $(I-\alpha)\cdot r + \alpha\cdot z$ as $I\cdot r + \alpha\cdot(z-r)$, the first-order condition for this problem can be expressed as:

$$\int z\cdot U'(I\cdot r + \alpha\cdot(z-r))\mathrm{d}F(z) - \int r\cdot U'(I\cdot r + \alpha\cdot(z-r))\mathrm{d}F(z) = 0,$$

that is, the marginal *expected* utility of the last dollar allocated to each asset is the same. The second-order condition can be written as:

$$\int (z-r)^2 \cdot U''' (I \cdot r + \alpha \cdot (z-r)) \ \mathrm{d}F(z) < 0$$

and is ensured by the property of risk aversion (i.e. $U''(\cdot) < 0$).

As usual, we may differentiate the first-order condition to obtain the effect of a change in some parameter, say initial wealth I, on the optimal level of investment in the risky asset (the optimal value of α). Differentiating the first-order condition (including α) with respect to I, solving for $d\alpha/dI$, and invoking the second-order condition and the positivity of r yields that this derivative possesses the same sign as:

$$\int (z-r) \cdot U''(I \cdot r + \alpha \cdot (z-r)) \ \mathrm{d}F(z).$$

Substituting $U''(x) \equiv -R(x) \cdot U'(x)$ and subtracting $R(I \cdot r)$ times the first-order condition yields that this expression is equal to:

$$-\int (z-r) \cdot [R(I \cdot r + \alpha \cdot (z-r)) - R(I \cdot r)] \cdot U'(I \cdot r + \alpha \cdot (z-r))\mathrm{d}F(z).$$

On the assumption that α is positive and $R(\cdot)$ is monotonic, the expression $(z-r) \cdot [R(I \cdot r + \alpha \cdot (z-r)) - R(I \cdot r)]$ will possess the same sign as $R'(\cdot)$. This implies that the derivative $\mathrm{d}\alpha/\mathrm{d}I$ will be positive (negative) whenever the Arrow–Pratt index $R(x)$ is a decreasing (increasing) function of the individual's wealth level x. In other words, an increase in initial wealth will always increase (decrease) demand for the risky asset if and only if $U(\cdot)$ exhibits decreasing (increasing) absolute risk aversion in wealth. Further examples of the analytics of the expected utility model may be found in the above references, as well as the surveys of Hirshleifer and Riley (1979), Lippman and McCall (1981), Machina (1983) and Karni and Schmeidler (1991).

Axiomatic Development

Although there exist dozens of formal axiomatizations of the expected utility model, most proceed by specifying an outcome space and postulating that the individual's preferences over probability distributions on this outcome space satisfy the following four axioms: completeness, transitivity, continuity and the Independence Axiom. Although it is beyond the scope of this entry to provide a rigorous derivation of the expected utility model in its most general setting, it is possible to illustrate the meaning of the axioms and sketch a proof of the expected utility representation theorem in the simple case of a finite outcome set $\{x_1, \ldots, x_n\}$.

Recall that in such a case the objects of choice consist of probability distributions $\boldsymbol{P} = (p_1, \ldots, p_n)$ over $\{x_1, \ldots, x_n\}$, so that the following axioms refer to the individuals' weak preference relation $\succcurlyeq$ over these prospects, where $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$ is read '$\boldsymbol{P}^*$ is weakly preferred (that is, preferred or indifferent) to $\boldsymbol{P}$' (the associated strict preference relation $\succ$ and indifference relation $\sim$ are defined in the usual manner):

- *Completeness*: For any two distributions $\boldsymbol{P}$ and $\boldsymbol{P}^*$, either $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$, $\boldsymbol{P} \succcurlyeq \boldsymbol{P}^*$, or both.
- *Transitivity*: If $\boldsymbol{P}^{**} \succcurlyeq \boldsymbol{P}^*$ and $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$, then $\boldsymbol{P}^{**} \succcurlyeq \boldsymbol{P}$.
- *Mixture continuity*: If $\boldsymbol{P}^{**} \succcurlyeq \boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$, then there exists some $\lambda \in [0, 1]$ such that $\boldsymbol{P}^* \sim \lambda \cdot \boldsymbol{P}^{**} + (1-\lambda) \cdot \boldsymbol{P}$.
- *Independence*: For any two distributions $\boldsymbol{P}$ and $\boldsymbol{P}^*$, $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$ if and only if $\lambda \cdot \boldsymbol{P}^* + (1-\lambda) \cdot \boldsymbol{P}^{**} \succcurlyeq \lambda \cdot \boldsymbol{P} + (1-\lambda) \cdot \boldsymbol{P}^{**}$ for all $\lambda \in [0, 1]$ and all $\boldsymbol{P}^{**}$

where $\lambda \cdot \boldsymbol{P} + (1-\lambda) \cdot \boldsymbol{P}^{**}$ denotes the $\lambda : (1-\lambda)$ 'probability mixture' of $\boldsymbol{P}$ and $\boldsymbol{P}^{**}$, that is, the lottery with probabilities $(\lambda \cdot p_1 + (1-\lambda) \cdot p_1^{**}, \ldots, \lambda \cdot p_n + (1-\lambda) \cdot p_n^{**})$.

The completeness and transitivity axioms are analogous to their counterparts in standard consumer theory. Mixture continuity states that if the lottery P^{**} is weakly preferred to P^* and P^* is weakly preferred to P, then exists some probability mixture of the most and least preferred lotteries which is indifferent to the intermediate one.

As in standard consumer theory, completeness, transitivity and continuity serve to establish the existence of a real-valued preference function $V(p_1, \ldots, p_n)$ which represents the relation $\succcurlyeq$, in the sense that $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$ if and only if $V(p_1^*, \ldots, p_n^*) \geq V(p_1, \ldots, p_n)$. It is the Independence Axiom which gives the theory its primary empirical content by implying that $\succcurlyeq$ can be represented by a linear preference function of the form $V(p_1, \ldots, p_n) \equiv \sum U_i p_i$. To see the meaning of this axiom, assume that individuals are always indifferent between a two-stage compound lottery and its probabilistically equivalent single-stage lottery, and that $\boldsymbol{P}^*$ happens to be weakly preferred to $\boldsymbol{P}$. In that case, the choice between the mixtures $\lambda \cdot \boldsymbol{P}^* + (1 - \lambda) \cdot \boldsymbol{P}^{**}$ and $\lambda \cdot \boldsymbol{P} + (1 - \lambda) \cdot \boldsymbol{P}^{**}$ is equivalent to being presented with a coin that has a $(1 - \lambda)$ chance of landing tails (in which case the prize will be $\boldsymbol{P}^{**}$) and being asked *before the flip* whether one would rather win $\boldsymbol{P}^*$ or $\boldsymbol{P}$ in the event of a head. The normative argument for the Independence Axiom is that either the coin will land tails, in which case the choice won't have mattered, or it will land heads, in which case one is 'in effect' facing a choice between $\boldsymbol{P}^*$ and $\boldsymbol{P}$ and one 'ought' to have the same preferences as before. Note finally that the above statement of the axiom in terms of the weak preference relation $\succcurlyeq$ also implies its counterparts in terms of strict preference and indifference.

In the following sketch of the expected utility representation theorem, expressions such as '$x_i \succcurlyeq x_j$' should be read as saying that the individual weakly prefers the degenerate lottery yielding x_i with certainty to that yielding x_j with certainty, and '$\lambda \cdot x_i + (1 - \lambda) \cdot x_j$' will be used to denote the $\lambda : (1 - \lambda)$ probability mixture of these two degenerate lotteries.

The first step in the proof is to define the von Neumann–Morgenstern utility index $\{U_i\}$ and the expected utility preference function $V(\cdot)$. Without loss of generality, we may order the outcomes so that $x_n \succcurlyeq x_{n-1} \succcurlyeq \ldots \succcurlyeq x_2 \succcurlyeq x_1$. Since $x_n \succcurlyeq x_i \succcurlyeq x_1$ for each outcome x_i, mixture continuity implies that there exist scalars $\{Ui\} \subset [0, 1]$ such that $x_i \sim U_i \cdot x_n + (1 - U_i) \cdot x_1$ for each i (which implies $U_1 = 0$ and $U_n = 1$). Given this, define $V(\boldsymbol{P}) = \sum U_i p_i$ for each $\boldsymbol{P}$.

The second step is to show that each lottery $\boldsymbol{P} = (\boldsymbol{p}_1, \ldots, \boldsymbol{p}_n)$ is indifferent to the mixture $\lambda \cdot x_n + (1 - \lambda) \cdot x_1$ where $\lambda = \sum U_i p_i$. Since $(p_1, \ldots, p_n)$ can be written as the n-component probability mixture $p_1 \cdot x_1 + p_2 \cdot x_2 + \ldots + p_n \cdot x_n$, and each outcome x_i is indifferent to the mixture $U_i \cdot x_n + (1 - U_i) \cdot x_1$, an n-fold application of the Independence Axiom yields that $\boldsymbol{P} = (p_1, \ldots, p_n)$ is indifferent to the mixture

$$p_1 \cdot [U_1 \cdot x_n + (1 - U_1) \cdot x_1] + p_2 \cdot [U_2 \cdot x_n + (1 - U_2) \cdot x_1] + \ldots + p_n \cdot [U_n \cdot x_n + (1 - U_n) \cdot x_1],$$

which is equivalent to $\left(\sum_{i=1}^n U_i p_i\right) \cdot x_n + \left(1 - \sum_{i=1}^n U_i p_i\right) \cdot x_1$.

The third step is to demonstrate that a mixture $\lambda^* \cdot x_n + (1 - \lambda^*) \cdot x_1$ is weakly preferred to a mixture $\lambda \cdot x_n + (1 - \lambda) \cdot x_1$ if and only if $\lambda^* \geq \lambda$. This follows immediately from the Independence Axiom and the fact that $\lambda^* \geq \lambda$ implies that these two lotteries may be expressed as the respective mixtures $(\lambda^* - \lambda) \cdot x_n + (1 - \lambda^* + \lambda) \cdot \boldsymbol{Q}$ and $(\lambda^* - \lambda) \cdot x_1 + (1 - \lambda^* + \lambda) \cdot \boldsymbol{Q}$, where $\boldsymbol{Q}$ is defined as the lottery $(\lambda/(1 - \lambda^* + \lambda)) \cdot x_n + ((1 - \lambda^*)/(1 - \lambda^* + \lambda)) \cdot x_1$.

The completion of the proof is now simple. For any two distributions $\boldsymbol{P}^* = (p_1^*, \ldots, p_n^*)$ and $\boldsymbol{P} = (p_1, \ldots, p_n)$, transitivity and the second step imply that $\boldsymbol{P}^* \succcurlyeq \boldsymbol{P}$ if and only if

$$\left(\sum_{i=1}^n U_i p_i^*\right) \cdot x_n + \left(1 - \sum_{i=1}^n U_i p_i^*\right) \cdot x_1 \succcurlyeq \left(\sum_{i=1}^n U_i p_i\right) \cdot x_n + \left(1 - \sum_{i=1}^n U_i p_i\right) \cdot x_1,$$

which by the third step is equivalent to the condition $\sum U_i p_i^* \geq \sum U_i p_i$, or in other words, that $V(\boldsymbol{P}^*) \geq V(\boldsymbol{P})$.

As mentioned, the expected utility model has been axiomatized many times and in many contexts. The most comprehensive accounts of the axiomatics of the model are undoubtedly Fishburn (1982) and Kreps (1988).

Subjective Expected Utility

In addition to the above setting of 'objective' (that is, probabilistic) uncertainty, it is possible to define expected utility preferences under conditions of 'subjective' uncertainty. In this case, uncertainty is represented by a set $\mathscr{S}$ of mutually exclusive and exhaustive 'states of nature,' which can be a finite set $\{s_1,\ldots, s_n\}$ (as with a horse race), a real interval $[\underline{s}, \overline{s}] \subseteq R^1$ (as with tomorrow's temperature), or a more abstract space. The objects of choice are then 'acts' a(·): $\mathscr{S} \to \mathscr{X}$ which map states to outcomes. In the case of a finite state space, acts are usually expressed in the form $\{x_1 \text{ if } s_1;\ldots; x_n \text{ if } s_n\}$. When the state space is infinite, finite-outcome acts can be expressed in the form $a(\cdot) = [x_1 \text{ on } E_1;\ldots; x_m \text{ on } E_m]$ for some partition of $\mathscr{S}$ into a family of mutually exclusive and exhaustive 'events' $\{E_1,\ldots, E_m\}$. Except for casino games and state lotteries, virtually all real-world uncertain decisions (including all investment or insurance decisions) are made under conditions of subjective uncertainty.

In such a setting, the 'subjective expected utility hypothesis' consists of the *joint* hypothesis that the individual possesses *probabilistic beliefs*, as represented by a 'personal' or 'subjective' probability measure $\mu(\cdot)$ over the state space, and *expected utility risk preferences*, as represented by a von Neumann–Morgenstern utility function $U(\cdot)$ over outcomes, and evaluates acts according a preference function of the form $W(x_1 \text{ if } s_1;\ldots;x_n \text{ if } s_n) \equiv \sum_{i=1}^{n} U(x_i)\cdot\mu(s_i)$, $W(x_1 \text{ on } E_1;\ldots;x_m \text{ on } E_m) \equiv \sum_{i=1}^{m} U(x_i)\cdot\mu(E_i)$, or more generally, $W(a(\cdot)) \equiv \int U(a(s))\mathrm{d}\mu(s)$. Whereas all individuals facing a given *objective* prospect $\boldsymbol{P} = (x_1, p_1;\ldots; x_n, p_n)$ are assumed to 'see' the same probabilities $(p_1,\ldots, p_n)$ (though they may have different utility functions), individuals facing a given *subjective* prospect $\{x_1 \text{ if } s_1;\ldots; x_n \text{ if } s_n\}$ or $[x_1 \text{ on } E_1;\ldots; x_m \text{ on } E_m]$ will generally possess differing subjective probabilities over these states or events, reflecting their different beliefs, past experiences, and so on.

Researchers such as Arrow (1974), Debreu (1959, ch. 7) and Hirshleifer (1965, 1966) have shown how the analytics of the objective expected utility model can be extended to both the positive and normative analysis of decisions under subjective uncertainty. As a simple example, consider an individual deciding whether to purchase earthquake insurance, and if so, how much. A simple specification of this decision involves the state space $\mathscr{S} = \{s_1, s_2\} = \{$earthquake; no earthquake$\}$, the individual's von Neumann–Morgenstern utility of wealth function $U(\cdot)$, their subjective probabilities $\{\mu(s_1), \mu(s_2)\}$ (which sum to unity), and the price γ of each dollar of insurance coverage. An individual with initial wealth w would then purchase q dollars' worth of coverage, where q was the solution to

$$\max_q \left[U(w-\gamma q+q)\cdot\mu(s_1)+U(w-\gamma q)\cdot\mu(s_2)\right]$$

Note that this formulation does not require that the individual and the insurance company agree on the likelihood of an earthquake.

As in the objective case, subjective expected utility can be derived from axiomatic foundations. Completeness and transitivity carry over in a straightforward way, and continuity with respect to mixture probabilities is replaced by continuity with respect to small changes in the events. The existence of additive personal probabilities is obtained by the following axiom:

Comparative likelihood: For all events A, B and outcomes $x^* \succ x$ and $y^* \succ y$, $[x^* \text{ on } A; x \text{ on} \sim A] \succcurlyeq [x^* \text{ on } B; x \text{ on} \sim B]$ implies $[y^* \text{ on } A; y \text{ on} \sim A] \succcurlyeq [y^* \text{ on } B; y \text{ on} \sim B]$.

This axiom states that if the individual 'reveals' event A to be at least as likely as event B by their preference for staking the preferred outcome x^* on A rather than on B, then this likelihood ranking will hold for all other pairs of ranked outcomes $y^* \succ y$. Finally, under subjective uncertainty the Independence Axiom is replaced by its subjective analogue, first proposed by Savage (1954):

Sure-Thing Principle: For all events E and acts $a(\cdot)$, $a^*(\cdot)$, $b(\cdot)$ and $c(\cdot)$, $[a^*(\cdot) \text{ on } E; b(\cdot) \text{ on} \sim E]$

$\succcurlyeq [a(\cdot) \text{ on } E; b(\cdot) \text{ on} \sim E]$ implies $[a^*(\cdot) \text{ on } E; c(\cdot) \text{ on} \sim E] \succcurlyeq [a(\cdot) \text{ on } E; c(\cdot) \text{ on} \sim E]$.

where $[a(\cdot)$ on E; $b(\cdot)$ on $\sim E]$ denotes the act yielding outcome $a(s)$ for all $s \in E$ and $b(s)$ for all $s \in \sim E$.

Under subjective uncertainty, an individual's utility of outcomes might sometimes depend upon the particular state of nature. Given a health insurance decision with a state space of $\mathscr{S} = \{s_1, s_2\} = \{\text{cancer; no cancer}\}$, an individual may feel a greater need for \$100,000 in state s_1 than in state s_2. This can be modelled by means of a 'state-dependent' utility function $\{U(\cdot|s)|\, s \in \mathscr{S}\}$ and a'state-dependent expected utility' preference function $\hat{W}(x_1 \text{ if } s_1; \ldots; x_n \text{ if } s_n) = \sum_{i=1}^{n} U(x_i|\, s_i) \cdot \mu(s_i)$or $\hat{W}(a(\cdot)) = \int U(a(s)|\, s) d\mu(s)$. The analytics of state-dependent expected utility preferences have been extensively developed by Karni (1985).

History

The hypothesis that individuals might maximize the expectation of 'utility' rather than of monetary value was proposed independently by mathematicians Gabriel Cramer and Daniel Bernoulli, in each case as the solution to a problem posed by Daniel's cousin Nicholas Bernoulli (see Bernoulli 1738). This problem, now known as the 'St Petersburg Paradox', considers the gamble which offers a 1/2 chance of \$1, a 1/4 chance of \$2, a 1/8 chance of \$4, and so on. Although the expected value of this prospect is

$$(1/2) \cdot \$1 + (1/4) \cdot \$2 + (1/8) \cdot \$4 + \ldots\ldots = \$0.50 + \$0.50 + \$0.50 + \ldots = \$\infty,$$

common sense suggests that no one would be willing to forgo a very substantial certain payment in order to play it. Cramer and Bernoulli proposed that, instead of using expected value, individuals might evaluate this and other lotteries by means of their expected 'utility', with utility given by a function such as the natural logarithm or the square root of wealth, in which case the certainty equivalent of the St Petersburg gamble becomes a moderate (and plausible) amount.

Two hundred years later, the St Petersburg paradox was generalized by Karl Menger (1934), who noted that, whenever the utility of wealth function was unbounded (as with the natural logarithm or square root functions), it would be possible to construct similar examples with infinite expected utility and hence infinite certainty equivalents (replace the payoffs \$1, \$2, \$4 ... in the above example by $x_1, x_2, x_3 \ldots$, where $U(x_i) = 2^i$ for each i). In light of this, von Neumann–Morgenstern utility functions are typically (though not universally) postulated to be bounded functions of wealth.

The earliest formal axiomatic treatment of the expected utility hypothesis was developed by Frank Ramsey (1926) as part of his theory of subjective probability, or individuals' 'degrees of belief' in the truth of alternative propositions. Starting from the premise that there exists an 'ethically neutral' proposition whose degree of belief is 1/2, and whose validity or invalidity is of no independent value, Ramsey proposed a set of axioms on how the individual would be willing to stake prizes on its truth or falsity, in a manner which allowed for the derivation of the 'utilities' of these prizes. He then used these utility values and betting preferences to determine the individual's degrees of belief in other propositions. Perhaps because it was intended as a contribution to the philosophy of belief rather than to the theory of risk bearing, Ramsey's analysis did not have the impact among economists that it deserved.

The first axiomatization of the expected utility model to receive widespread attention was that of John von Neumann and Oskar Morgenstern, presented in connection with their formulation of the theory of games (von Neumann and Morgenstern 1944, 1947, 1953). Although this development was recognized as a breakthrough, the mistaken belief that von Neumann and Morgenstern had somehow mathematically overthrown the Hicks–Allen 'ordinal revolution' led to some confusion until the difference between 'utility' in the von Neumann-Morgenstern and the ordinal (that is, non-stochastic) senses was illuminated by writers such as Ellsberg (1954) and Baumol (1958).

Another factor which delayed the acceptance of the theory was the lack of recognition of the role played by the Independence Axiom, which did not explicitly appear in the von Neumann–Morgenstern formulation. In fact, the initial reaction of researchers such as Baumol (1951) and Samuelson (1950) was that there was no reason why preferences over probability distributions must *necessarily* be linear in the probabilities. However, the independent discovery of the Independence Axiom by Marschak (1950), Samuelson (1952) and others, and Malinvaud's (1952) observation that it had been implicitly invoked by von Neumann and Morgenstern, led to an almost universal acceptance of the expected utility hypothesis as both a normative and positive theory of behaviour towards risk. This period also saw the development of the elegant axiomatization of Herstein and Milnor (1953) as well as Savage's (1954) joint axiomatization of utility and subjective probability, which formed the basis of the state-preference approach described above.

While the 1950s essentially saw the completion of foundational work on the expected utility model, subsequent decades saw the flowering of its analytic capabilities and its application to fields such as portfolio selection (Merton 1969), optimal savings (Levhari and Srinivasan 1969; Fleming and Sheu 1999), international trade (Batra 1975; Lusztig and James 2006), environmental economics (Wolfson et al. 1996), medical decision-making (Meltzer 2001) and even the measurement of inequality (Atkinson 1970). This movement was spearheaded by the development of the Arrow–Pratt characterization of risk aversion (see above) and the characterization, by Rothschild–Stiglitz (1970, 1971) and others, of the notion of 'increasing risk'. This latter work in turn led to the development of a general theory of 'stochastic dominance' (for example, Whitmore and Findlay 1978; Levy 1992), which has further expanded the analytical powers of the model.

Although the expected utility model received a small amount of experimental testing by economists in the early 1950s (for example, Mosteller and Nogee 1951; Allais 1953) and continued to be examined by psychologists, economists' interest in the empirical validity of the model waned from the mid-1950s through the mid-1970s, no doubt due to both the normative appeal of the Independence Axiom and model's analytical successes. However, since the late 1970s there has been a revival of interest in the testing of the expected utility model; a growing body of evidence that individuals' preferences *systematically* depart from linearity in the probabilities; and the development, analysis and application of alternative models of choice under objective and subjective uncertainty. It is fair to say that today the debate over the descriptive (and even normative) validity of the expected utility hypothesis is more extensive than it has been in over half a century, and the outcome of this debate will have important implications for the direction of research in the economics of uncertainty.

See Also

▶ Bernoulli, Daniel
▶ Non-Expected Utility Theory
▶ Ramsey, Frank Plumpton
▶ Risk
▶ Risk Aversion
▶ Savage's Subjective Expected Utility Model
▶ Uncertainty
▶ Utility

Bibliography

Allais, M. 1953. Fondements d'une théorie positive des choix comportant un risque et critique des postulats et axiomes de l'école Américaine. *Colloques Internationaux du Centre National de la Recherche Scientifique* 40: 257–332. Trans. as: The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American School. In *Expected utility hypotheses and the Allais paradox*, ed. M. Allais and O. Hagen. Dordrecht: D. Reidel, 1979.

Arrow, K. 1974. *Essays in the theory of risk-bearing*. Amsterdam: North-Holland.

Atkinson, A. 1970. On the measurement of inequality. *Journal of Economic Theory* 2: 244–263.
Batra, R. 1975. *The pure theory of international trade under uncertainty*. London: Macmillan.
Baumol, W. 1951. The Neumann–Morgenstern utility index: An ordinalist view. *Journal of Political Economy* 59: 61–66.
Baumol, W. 1958. The cardinal utility which is ordinal. *Economic Journal* 68: 665–672.
Bernoulli, D. 1738. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*. Trans. as Exposition of a new theory on the measurement of risk. *Econometrica* 22 (1954): 23–36.
Debreu, G. 1959. *Theory of value: An axiomatic analysis of economic equilibrium*. New Haven: Yale University Press.
Ellsberg, D. 1954. Classical and current notions of 'measurable utility'. *Economic Journal* 64: 528–556.
Fishburn, P. 1982. *The foundations of expected utility*. Dordrecht: D. Reidel.
Fleming, W., and S.-J. Sheu. 1999. Optimal long term growth rate of expected utility of wealth. *Annals of Applied Probability* 9: 871–903.
Friedman, M., and L. Savage. 1948. The utility analysis of choices involving risk. *Journal of Political Economy* 56: 279–304.
Herstein, I., and J. Milnor. 1953. An axiomatic approach to measurable utility. *Econometrica* 21: 291–297.
Hey, J. 1979. *Uncertainty in microeconomics*. Oxford/New York: Martin Robinson/New York University Press.
Hirshleifer, J. 1965. Investment decision under uncertainty: Choice theoretic approaches. *Quarterly Journal of Economics* 79: 509–536.
Hirshleifer, J. 1966. Investment decision under uncertainty: Applications of the state-preference approach. *Quarterly Journal of Economics* 80: 252–277.
Hirshleifer, J., and J. Riley. 1979. The analytics of uncertainty and information – An expository survey. *Journal of Economic Literature* 17: 1375–1421.
Karni, E. 1985. *Decision making under uncertainty: The case of state-dependent preferences*. Cambridge, MA: Harvard University Press.
Karni, E., and D. Schmeidler. 1991. Utility theory with uncertainty. In *Handbook of mathematical economics*, ed. W. Hildenbrand and H. Sonnenschein, vol. 4. Amsterdam: North-Holland.
Kreps, D. 1988. *Notes on the theory of choice*. Boulder: Westview Press.
Levhari, D., and T.N. Srinivasan. 1969. Optimal savings under uncertainty. *Review of Economic Studies* 36: 153–164.
Levy, H. 1992. Stochastic dominance and expected utility: Survey and analysis. *Management Science* 38: 555–593.
Lippman, S., and J. McCall. 1981. The economics of uncertainty: Selected topics and probabilistic methods. In *Handbook of mathematical economics*, ed. K. Arrow and M. Intriligator, vol. 1. Amsterdam: North-Holland.
Lusztig, M., and P. James. 2006. How does free trade become institutionalised? An expected utility model of the Chrétien era. *World Economy* 29: 491–505.
Machina, M. 1983. *The economic theory of individual behavior toward risk: Theory, evidence and new directions*. Technical report no. 433. Institute for Mathematical Studies in the Social Sciences, Stanford University.
Malinvaud, E. 1952. Note on von Neumann–Morgenstern's strong independence axiom. *Econometrica* 20: 679–680.
Markowitz, H. 1952. The utility of wealth. *Journal of Political Economy* 60: 151–158.
Marschak, J. 1950. Rational behavior, uncertain prospects, and measurable utility. *Econometrica* 18: 111–141.
Meltzer, D. 2001. Addressing uncertainty in medical cost-effectiveness analysis: Implications of expected utility maximization for methods to perform sensitivity analysis and the use of cost-effectiveness analysis to set priorities for medical research. *Journal of Health Economics* 20: 109–129.
Menger, K. 1934. Das Unsicherheitsmoment in der Wertlehre. *Zeitschrift für Nationalökonomie*. Trans. as: The role of uncertainty in economics. In *Essays in mathematical economics in honor of Oskar Morgenstern*, ed. M. Shubik. Princeton: Princeton University Press, 1967.
Merton, R. 1969. Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics* 51: 247–257.
Mosteller, F., and P. Nogee. 1951. An experimental measurement of utility. *Journal of Political Economy* 59: 371–404.
Pratt, J. 1964. Risk aversion in the small and in the large. *Econometrica* 32: 122–136.
Quirk, J., and R. Saposnick. 1962. Admissibility and measurable utility functions. *Review of Economic Studies* 29: 140–146.
Ramsey, F. 1926. Truth and probability. In *The foundations of mathematics and other logical essays*, ed. R. Braithwaite. New York: Harcourt, Brace and Co, 1931. Reprinted in *Foundations: Essays in philosophy, logic, mathematics and economics*, ed. D. Mellor. New Jersey: Humanities Press, 1978.
Ross, S. 1981. Some stronger measures of risk aversion in the small and in the large, with applications. *Econometrica* 49: 621–638.
Rothschild, M., and J. Stiglitz. 1970. Increasing risk I: A definition. *Journal of Economic Theory* 2: 225–243.
Rothschild, M., and J. Stiglitz. 1971. Increasing risk II: Its economic consequences. *Journal of Economic Theory* 3: 66–84.
Samuelson, P. 1950. Probability and attempts to measure utility. *Economic Review* 1: 167–173.
Samuelson, P. 1952. Probability, utility, and the independence axiom. *Econometrica* 20: 670–678.

Savage, L. 1954. *The foundations of statistics*. New York: John Wiley & Sons. Revised edition: New York: Dover, 1972.

von Neumann, J., and O. Morgenstern. 1944. *Theory of games and economic behavior*. Princeton: Princeton University Press.

von Neumann, J., and O. Morgenstern. 1947. *Theory of games and economic behavior*. 2nd ed. Princeton: Princeton University Press.

von Neumann, J., and O. Morgenstern. 1953. *Theory of games and economic behavior*. 3rd ed. Princeton: Princeton University Press.

Whitmore, G., and M. Findlay, eds. 1978. *Stochastic dominance: An approach to decision making under risk*. Lexington: D.C. Heath.

Wolfson, L., J. Kadane, and M. Small. 1996. Expected utility as a policy making tool: An environmental health example. In *Bayesian biostatistics*, ed. D. Berry and D. Stangl. New York: Marcel Dekker.