

Extensions of the Subjective Expected Utility Model

Robert Nau¹

Fuqua School of Business, Duke University

In *Advances in Decision Analysis: From Foundations to Applications*

Ward Edwards, Ralph Miles, Jr., and Detlof von Winterfeldt (eds.)

Cambridge University Press, 2007

Abstract: The subjective expected utility (SEU) model rests on very strong assumptions about the consistency of decision making across a wide range of situations. The descriptive validity of these assumptions has been extensively challenged by behavioral psychologists during the last few decades, and the normative validity of the assumptions has also been reappraised by many statisticians, philosophers, and economists, motivating the development of more general utility theories and decision models. These generalized models are characterized by features such as imprecise probabilities, nonlinearly weighted probabilities, source-dependent risk attitudes, and state-dependent utilities, permitting the pattern of the decision maker's behavior to change with the decision context and to perhaps satisfy the usual SEU assumptions only locally. Recent research in the emerging field of neuroeconomics sheds light on the physiological basis of decision making, the nature of preferences and beliefs, and interpersonal differences in decision competence. These findings do not necessarily invalidate the use of SEU-based decision analysis tools, but they suggest that care needs to be taken to structure preferences and to assess beliefs and risk attitudes in a manner that is appropriate for the decision and also for the decision maker.

Key words: subjective probability, expected utility, non-expected utility, Savage's axioms, sure-thing principle, Allais' paradox, Ellsberg's paradox, risk aversion, uncertainty aversion, exponential utility function, logarithmic utility function, risk tolerance, ambiguity, incomplete preferences, imprecise probabilities, robust Bayesian analysis, rank-dependent utility, comonotonic acts, probability weighting function, cumulative prospect theory, Choquet expected utility, Knightian uncertainty, maxmin expected utility, second-order utility, state-dependent utility, state-preference theory, risk neutral probabilities, inseparable probabilities and utilities, neuroeconomics, cognitive neuroscience, affective decision making, emotions

1. The SEU model and its assumptions

The subjective expected utility (SEU) model provides the conceptual and computational framework that is most often used to analyze decisions under uncertainty. In the SEU model, uncertainty about the future is represented by a set of *states of the world*, which are mutually exclusive and exhaustive events. Possible outcomes for the decision maker are represented by a set of *consequences*, which could be amounts of money in the bank or more general “states of the person” such as health, happiness, pleasant or unpleasant experiences, and so on. A decision alternative, known as an *act*, is defined by an assignment of consequences to states of the world. In the case where the set of states is a finite set (E_1, \dots, E_n) , an act can be written as a vector $\mathbf{x} = (x_1, \dots, x_n)$, where x_i is the consequence that is received or experienced in state E_i . The decision

¹ I am grateful to Jill Stowe, Jim Smith, and the editors for helpful comments. Any errors or omissions are my own.

maker's beliefs concerning states of the world are represented by a *subjective probability distribution* $\mathbf{p} = (p_1, \dots, p_n)$, where p_i is the probability of E_i , and her values for consequences are represented by a *utility function* $v(x)$, in terms of which the value she assigns to an act \mathbf{x} for decision making purposes is its *subjective expected utility*:

$$SEU(\mathbf{x}) = \mathbf{E}_p[v(\mathbf{x})] = \sum_{i=1}^n p_i v(x_i). \quad (1)$$

This recipe for rational decision making has ancient roots: it was first proposed by Daniel Bernoulli (1738) to explain aversion to risk in problems of gambling and insurance as well as to solve the famous St. Petersburg Paradox. Bernoulli recognized that different individuals might display different risk attitudes, especially if they differ in wealth, and he recommended using the logarithmic utility function $v(x) = \log(x)$ because it implies that “the utility resulting from any small increase in wealth will be inversely proportionate to the quantity of goods previously possessed.” The idea of seeking to maximize the expected value of a utility function – particularly a logarithmic one – was discarded and even ridiculed by later generations of economists, who doubted that utility could ever be measured on a cardinal numerical scale. (See Stigler 1950 for an excellent historical review.) However, it was revived and rehabilitated in dramatic fashion by von Neumann and Morgenstern (1944/1947) and Savage (1954), who showed that the expected-utility model could be derived from simple and seemingly reasonable axioms of consistent preferences under risk and uncertainty, in which a pivotal role is played by an independence condition known as the *sure-thing principle* (Fishburn and Wakker 1995).

Von Neumann and Morgenstern considered the special case in which states of the world have objectively known probabilities (as in games of chance), and Savage extended the model to include situations where probabilities are subjectively determined by the decision maker. The key axioms of Savage are as follows. (P1) Preferences among acts are weakly ordered, that is, *complete* and *transitive*. (P2) Preferences satisfy the *independence* condition (sure-thing principle) which requires that if two acts “agree” (i.e., yield the same consequence) in some state, it does not matter *how* they agree there. This permits a natural definition of conditional preferences, namely that \mathbf{x} is preferred to \mathbf{y} conditional on event E if \mathbf{x} is preferred to \mathbf{y} and they agree in the event not- E . (P3) Preferences among consequences are *state-independent* in the sense that conditional preferences between “constant” acts (those which yield the same consequence in all states) do not depend on the conditioning event. (P4) Events can be unambiguously *ordered by probability* in the following way: if x and y are any two consequences such that x is preferred to y (as a constant act), and if the act that yields x if E and y if not- E is preferred to the act that yields x if F and y if not- F , then E is revealed to be at least as probable as F . These four substantive behavioral postulates, together with a few purely technical assumptions, imply the *SEU* formula. Other systems of axioms also lead to SEU (e.g., Anscombe and Aumann 1963; Wakker 1989).

The SEU model had a revolutionary impact on statistical decision theory and social science in the 1950s and 1960s, providing the mathematical foundation for a broad range of social and economic theories under the general heading of “rational choice,” including the development of Bayesian methods of statistical inference, the emergence of decision analysis as an applied science taught in engineering and business schools, the establishment of game theory as a foundation for microeconomics, and the development of expected-utility-based models of portfolio optimization and competitive equilibria in asset markets by finance theorists. The

logarithmic utility function originally proposed by Bernoulli even came to be hailed as the “premier” utility model for investors in financial markets (Rubinstein 1976).

The SEU model also had its early detractors, most notably Allais (1953) and Ellsberg (1961) who constructed famous paradoxes consisting of thought-experiments in which most individuals willingly violate the independence axiom (Savage’s P2), but for several decades it was widely accepted as both an appropriate normative standard and a useful descriptive model, as if it were self-evident that a thinking person should be an expected-utility maximizer. That consensus began to break down in the late 1970s, however, as an emerging body of behavioral decision research showed that subjects in laboratory experiments display an array of predictable “heuristics and biases” that are inconsistent with SEU theory, even beyond the paradoxical behavior identified by Allais and Ellsberg. The normative status of the SEU model was also questioned, insofar as violations of completeness or independence do not necessarily expose a decision maker to exploitation as long as she respects more fundamental principles such as dominance and transitivity. In response to these developments, decision theorists and economists proceeded to extend the SEU model by weakening various of its axioms, giving rise to a host of theories of “nonexpected utility” (e.g., Kahneman and Tversky 1979; Machina 1982; Fishburn 1982; Quiggin 1982; Loomes and Sugden 1982; Bell 1982, 1985; Chew 1983; Luce and Narens 1985; Yaari 1987; Becker and Sarin 1987; Schmeidler 1989; Gilboa and Schmeidler 1989; Machina and Schmeidler 1992; Tversky and Kahneman 1992; Wakker and Tversky 1993). This chapter provides a nontechnical summary of some extensions of the SEU model which appear most relevant to decision analysis and which can be defended normatively as well as descriptively. For more breadth and technical depth, the recent surveys by Starmer (2000), Sugden (2004), and Schmidt (2004) are highly recommended; a vast online bibliography has been compiled by Wakker (2006).

2. Incomplete preferences, imprecise probabilities and robust decision analysis

Arguably the strongest and most unrealistic assumption of the SEU model is that the decision maker’s preferences are *complete*, meaning that between any two alternatives that might be proposed, no matter how complicated or hypothetical or even counterfactual, the decision maker is always able to say either that she strictly prefers one to the other or else she is exactly indifferent between them: she is never “undecided.” This assumption is somewhat antithetical to the spirit of decision analysis, which provides tools for *constructing* preferences where they may not already exist. The completeness assumption also amplifies the effects of all the other axioms, making it relatively easy to generate examples in which they are violated.

Incompleteness of preferences is implicitly acknowledged whenever ad hoc methods of sensitivity analysis are applied to subjectively-assessed probabilities and utilities, which are often among the most controversial and hard-to-measure parameters in a decision model, especially where decisions must be taken on behalf of a group whose members may disagree. Incompleteness can also be acknowledged more explicitly by carrying out the entire analysis in terms of *imprecise probabilities and utilities*, which are intervals of numbers rather than point values. Thus, for example, each state of the world might be assigned a distinct *lower and upper probability*, each consequence might be assigned a distinct *lower and upper utility*, and each alternative would then be assigned a lower expected utility and an upper expected utility by computing the minimum and maximum values attained by the SEU formula (1) as the probabilities and utilities are varied between their lower and upper bounds in all possible combinations (e.g., Rios Insua 1990, 1992; Nau 1989, 1992, 2006b; Moskowitz et al. 1993; Rios

Insua and Ruggeri 2000). This type of analysis, known as *robust Bayesian analysis*, need not yield a unique optimal decision: there may be several “potentially optimal” decisions whose expected-utility intervals overlap to some extent. Further introspection and analysis may or may not shrink the intervals to the point where a unique optimal decision emerges, hence at the end of the day it may be necessary to acknowledge that the analysis leaves some room for doubt and to base the final decision at least partly on other criteria not explicitly included in the quantitative model.

In financial decisions where the outcomes are measured in monetary terms, the most critical imprecision is that of the *probabilities* rather than the utilities because it often suffices to assume linear utility or else to use a simple parametric utility function such as the exponential function ($v(x) = -\exp(-\lambda x)$), which has a single parameter (the risk aversion coefficient λ) that can be manipulated. Axiomatic models of imprecise probability have a long history in the literature of statistics and philosophy, including the work of Koopman (1940), Smith (1961), Hacking (1967), Kyburg (1974), and Levi (1980). Even de Finetti’s “fundamental theorem of probability” (1974) is stated in terms of lower and upper bounds that can be inferred for probabilities of events, given knowledge of probabilities of other events. However, the publication of Walley’s (1991) book *Statistical Reasoning with Imprecise Probabilities* sparked an upsurge of interest in this subject, and over the last 15 years a large literature on imprecise-probability models has emerged, as well as a related professional society, the Society for Imprecise Probabilities and Their Applications (SIPTA), which has held biannual meetings since 1999.

3. Allais’ paradox, transformed probabilities, and rank-dependent utility

One of the most compelling objections against Savage’s SEU model was raised at its inception by Allais (1953), who constructed the following paradox. Suppose there are three states of the world (E_1, E_2, E_3) whose probabilities — whether objective or subjective — are approximately 0.89, 0.10, and 0.01. (The exact values are unimportant: the states could just be regarded as “very likely,” “rather unlikely,” and “very unlikely,” respectively.) Now consider the following two pairs of alternative gambles over these states:

| | $E_1 (p \approx .89)$ | $E_2 (p \approx .10)$ | $E_3 (p \approx .01)$ |
|------|-----------------------|-----------------------|-----------------------|
| x | \$1M | \$1M | \$1M |
| y | \$1M | \$5M | \$0 |
| x' | \$0 | \$1M | \$1M |
| y' | \$0 | \$5M | \$0 |

Most individuals strictly prefer x over y and also strictly prefer y' over x' . The intuition for this pattern is clear: most would prefer to get \$1M for sure rather than accept a very small (1 percent) risk of getting nothing at all in order to obtain a small (10 percent) chance of getting \$5M instead of \$1M, but if they are most likely going to get nothing anyway, they would prefer to have a small (10 percent) chance at \$5M rather than a slightly larger (11 percent) chance at \$1M. (The effect is even stronger if the second pair of alternatives is simply presented as “an 11 percent chance of \$1M versus a 10 percent chance of \$5M” rather than a specific match-up of events and payoffs.) This preference pattern strikes most people as rational, yet it cannot be rationalized by

the SEU model because it violates the independence axiom. Here, \mathbf{x} and \mathbf{y} agree in state E_1 , both yielding \$1M, so according to the independence axiom it should not matter if some other agreeing payoff is substituted there. But replacing \$1M with \$0 in state E_1 yields \mathbf{x}' and \mathbf{y}' , reversing the direction of preference for most individuals, a phenomenon known as the “common consequence effect.”

The Allais paradox was once regarded as a curiosity, a cognitive illusion that would probably disappear under closer scrutiny. Savage (1954) himself admitted that he was taken in at first, but after further reflection he “corrected an error” and reversed his instinctive preference for \mathbf{y}' over \mathbf{x}' . However, after the frontier of SEU theory had been well explored in the 1950s and 1960s, and after violations of the independence axiom began to emerge as a very robust finding in behavioral experiments in the 1970s, many decision theorists began to explore the possibilities for a theory that would relax the independence axiom in some way.

The remainder of this section will focus on the general *rank-dependent utility* model, an extension of the SEU model that has emerged as the most widely-studied alternative in connection with the Allais paradox and other violations of the independence axiom that occur in situations where probabilities of events are assumed to be known. Variants of this model were developed independently by Quiggin (1982), Schmeidler (1989), Luce and Narens (1985), and Yaari (1987), and were later refined by others (e.g., Chew 1985; Chew et al. 1987; Luce and Manders 1988; Segal 1989; Wakker 1991, 1996; Chew and Wakker 1996); its intuition is nicely discussed by Diecidue and Wakker (2001).

The rank-dependent utility model is motivated by two key observations concerning violations of the independence axiom. One observation is that violations often occur when comparing acts whose relative riskiness is dramatically changed by replacing an agreeing payoff with a different agreeing payoff. In the Allais paradox as shown above, replacing \$1M with \$0 in state E_1 changes a comparison of a safe alternative against a risky alternative to a comparison of two almost-equally-risky alternatives. Perhaps the independence axiom would be easier to obey if it applied only to comparisons among alternatives with qualitatively similar risk profiles.

The second observation is that attitudes toward risk ought to be explained at least in part as a response to risk *per se*, rather than “as if” they are due to diminishing marginal utility for money. Intuitively, the decision maker’s perception of risk is rooted in her *beliefs*, which are represented by probabilities in the SEU model, but the independence axiom permits her utility to depend only linearly on those probabilities. If the independence axiom were relaxed in some way, perhaps the decision maker’s evaluation of an act could depend *nonlinearly* on her probabilities, analogously to the way it is allowed to depend nonlinearly on payoffs.

In the rank-dependent utility model, the decision maker’s utility for an act \mathbf{x} is computed according to a formula which, at first glance, looks very much like the SEU formula:

$$RDU(\mathbf{x}) = \sum_{i=1}^n \pi_i v(x_i), \quad (2)$$

The coefficients $\{\pi_i\}$, which are called *decision weights*, are positive and sum to 1, exactly as if they were subjective probabilities of states. The key difference is that the decision weights are *not* necessarily subjective probabilities, and the decision weight attached to a particular state is not necessarily the same for all acts. In particular, the decision weight π_i that is applied to state E_i

when evaluating the act x may depend on how that state is *ranked* relative to other states in terms of the goodness or badness of its payoff, as well as on the decision maker's beliefs.

The most general forms of the RDU model include the *Choquet expected utility* model (Schmeidler 1989), the *cumulative prospect theory* model (Tversky and Kahneman 1992; axiomatized by Wakker and Tversky 1993), and the *rank-and-sign-dependent utility* model (Luce and Fishburn 1991, Luce 2000). In Choquet expected utility and cumulative prospect theory the decision maker's beliefs are represented by nonadditive probability measures known as "capacities," and in cumulative prospect theory and rank-and-sign-dependent utility there are also reference-point effects: the utility function may be kinked at the status quo wealth position so that the decision maker is risk averse even for very small gambles, a behavioral phenomenon known as "loss aversion." However, in the simplest version of RDU, beliefs are represented by subjective or objective probabilities, and for decision making purposes the cumulative distribution of payoffs is merely distorted by an increasing function $w(p)$ that satisfies $w(0)=0$ and $w(1)=1$, the so-called *probability weighting function*. For a given act x , let the states be labeled in *decreasing* order of payoff, so that $x_1 \geq \dots \geq x_n$, and let (p_1, \dots, p_n) denote the corresponding probabilities. Then the decision weights in the RDU formula (2) are given by:

$$\pi_i = w(p_1) \quad \text{and} \quad \pi_i = w(p_1 + \dots + p_i) - w(p_1 + \dots + p_{i-1}) \quad \text{for } i=2, \dots, n.$$

It follows that $\pi_1 + \dots + \pi_i = w(p_1 + \dots + p_i)$, that is, the cumulative decision weight attached to the top i payoffs of x is equal to the transformed cumulative probability of those states.

If the probability weighting function w is linear, then the RDU model reduces to the SEU model with $\pi_i = p_i$ for all i . However, if w is nonlinear, the decision maker may exhibit pessimism or optimism in her attitude toward risk, even though she has probabilistic beliefs and even though she may have linear utility for money. In particular, if w is a *convex* (upward curving) function, then payoffs near the top of the ranking tend to be underweighted relative to their probabilities, whereas those near the bottom of the ranking tend to be overweighted. Thus, an RDU decision maker with a convex probability weighting function will pessimistically behave in a risk averse fashion (provided her utility function $v(x)$ is also either linear or concave) because she gives disproportionate "attention" to the very worst outcomes. For example, in the Allais paradox, an RDU decision maker with convex w would strongly overweight state E_3 (the zero-payoff event that receives the last 1 percent of probability) when evaluating act y in comparison to act x , whereas the effect of probability weighting would be much less pronounced in evaluating x' and y' (where the zero-payoff event receives either the last 10 percent or 11 percent of probability).

The key underlying assumption that distinguishes the RDU model from the SEU model is the replacement of the independence axiom by a weaker axiom of *comonotonic independence* (a.k.a. the "comonotonic sure-thing principle"). Two acts x and y are said to be comonotonic if they order the states in the same way, so that the same state yields the highest payoff under both acts, the same state yields the second-highest payoff, and so on down the line. The comonotonic independence axiom requires (only) that whenever two *comonotonic* acts agree in some states, then it does not matter how they agree there. Because comonotonic pairs of acts have the same qualitative risk profile, any change of an agreeing payoff in a particular state cannot hedge the risk of one of the acts without similarly hedging the other. The Allais paradox does not violate

comonotonic independence because the pairs of acts are not comonotonic: all states are equally good under x whereas they are strictly ranked by y in the order $E_2 > E_1 > E_3$.

The idea that decision makers use nonlinearly transformed cumulative probabilities in their evaluations of acts has been very widely studied in the last 20 years, and empirical estimation of the probability weighting function has been an active industry among behavioral decision theorists (e.g., Tversky and Fox 1995; Wu and Gonzalez 1996; Prelec 1998; Bleichrodt and Pinto 2000). A common finding is that many individuals behave as though their probability weighting function is inverse S-shaped, that is, concave for very small cumulative probabilities and convex for moderate or large cumulative probabilities. Such a weighting function potentially explains why otherwise-risk-averse individuals will often gladly pay more than expected value for state lottery tickets offering very small probabilities of very large gains. It is also consistent with the folk wisdom that most people do not discriminate very well among probabilities that are anything other than 0 or 1, tending to treat them as closer to one-half than they really are.

The relevance of this body of work for decision analysis is somewhat ambiguous. The fact that unaided experimental subjects exhibit nonlinear weighting of cumulative probabilities could be interpreted either as proof of the need for SEU-based methods of decision analysis or else as proof that SEU-based methods do not capture the preferences of many decision makers. An interesting synthetic approach has been proposed by Bleichrodt et al. (2001), who argue that nonlinear probability weighting (as well as loss aversion, that is, overweighting of small losses relative to small gains) should be anticipated during preference elicitation and should be corrected in order to estimate the individual's "true" underlying probabilities and utilities. A recent experiment by van de Kuilen and Wakker (2006) also shows that nonlinear probability weighting in Allais-type experiments tends to disappear if subjects are given the opportunity to learn through both experience and thought, although not through thought alone.

4. Ellsberg's paradox, Knightian decision theory, maxmin expected utility, and second-order utility

In recent years, interest among decision theorists and economic theorists has shifted somewhat from away from Allais-type paradoxes and toward Ellsberg-type paradoxes involving "ambiguous" probabilities, which are perceived as a more serious challenge to SEU theory insofar as they raise deep questions about the existence of personal probabilities (not merely the linearity of the decision maker's response to probabilities) and the distinction (if any) between risk and uncertainty. The claim that uncertainty is something different from risk is commonly traced to the following passage by Frank Knight (1921):

"There is a fundamental distinction between the reward for taking a known risk and that for assuming a risk whose value itself is not known. It is so fundamental, indeed, that as we shall see, a known risk will not lead to any reward or special payment."

This was written, however, at a time before expected utility and subjective expected utility had been given their modern form by von Neumann and Morgenstern and Savage (in the 1940s and 1950s) and before the principles of portfolio optimization and the distinction between diversifiable and nondiversifiable risks had been worked out by finance theorists (in the 1950s and 1960s). It is, by now, well-understood that an agent (e.g., an individual or firm) should be rewarded for bearing *nondiversifiable* risk according to the market price of risk and that subjective probabilities can, in principle, be used where objective ones are lacking, and that

interagent differences in risk aversion and subjective probabilities present opportunities for profit in the face of uncertainty. However, there are situations in which the uncertainty is of a more fundamental nature and the conventional decision-analytic and finance-theoretic models do not seem to apply, which was emphatically demonstrated by Ellsberg (1961).

In one of Ellsberg’s paradoxes, a subject is presented with two urns, one of which is said to contain exactly 50 red balls and 50 black balls whereas the other contains 100 balls that are red and black in unknown proportions. (In Ellsberg’s own informal experiments, he not only portrays the composition of the second urn as unknown, but even invites the suspicion that it is somehow “rigged.”) A single ball is to be drawn randomly from each urn, and the subject is presented with the following pairs of acts whose payoffs are pegged to the four possible results:

| | | | | |
|--|------------|--------------|--------------|--------------|
| <i>Ball from urn 1 (known 50% red)</i> | <i>red</i> | <i>red</i> | <i>black</i> | <i>black</i> |
| <i>Ball from urn 2 (unknown % red)</i> | <i>red</i> | <i>black</i> | <i>red</i> | <i>black</i> |
| <i>x</i> | \$100 | \$100 | \$0 | \$0 |
| <i>y</i> | \$100 | \$0 | \$100 | \$0 |
| <i>x'</i> | \$0 | \$0 | \$100 | \$100 |
| <i>y'</i> | \$0 | \$100 | \$0 | \$100 |

Most subjects strictly prefer x over y and also strictly prefer x' over y' , whereas they are indifferent between x and x' and also indifferent between y and y' . In other words, they would prefer to bet on the color of a ball drawn from the “unambiguous” urn rather than the “ambiguous” one, regardless of whether the winning color is red or black. This is another direct violation of the independence axiom because swapping the agreeing payoffs in the first and last column converts x to y' and y to x' , and, moreover, it is a violation that cannot be explained by any model such as RDU in which states are distinguished *only* by the probabilities and payoffs assigned to them, even if probabilities are evaluated in a nonlinear way. Even a model of state-dependent utility (to be discussed in the next section) cannot explain this pattern.

Two different interpretations can be placed on this phenomenon. One is that the subject is unable to assign a probability to the event of drawing a red ball from urn 2, and she is averse to betting on an event whose probability is undetermined. (The SEU model of course would require the decision maker to subjectively assign *some* probability and then proceed exactly as if it were objective.) This interpretation suggests the need for a preference model in which beliefs are represented by something more general and “fuzzier” than probabilities, perhaps *sets* of probabilities or *functions* defined on sets of probabilities. The other possible interpretation is that the individual thinks that red and black are equally likely to be drawn from the second urn, exactly as they are from the first urn, but she is nevertheless *more risk averse toward bets on the second urn* in a way that is incompatible with the independence axiom. The independence axiom permits the decision maker’s risk attitude to be *state*-dependent but does not permit it to be *source*-dependent, that is, dependent on how states are grouped together to form events on which a given winning or losing payoff is received. In the two-urn problem, the four states ought to be regarded symmetrically by the decision maker, each being the conjunction of a color from urn 1 and a color from urn 2, hence it should not matter which two states are combined to form the event on which the winning payoff of \$100 is received.

Regardless of how Ellsberg’s phenomenon is viewed, whether as an aversion to undetermined probabilities or as a source-dependent attitude toward risk, it strikes at the heart of a key principle of applied decision analysis, namely that the decision maker’s attitude toward risk can be assessed by contemplating simple reference gambles with objective probabilities, and the same risk attitude can be safely assumed to apply to decisions with respect to events whose probabilities are highly subjective. From a normative perspective it is hard to condemn violations of this principle as irrational because (unlike some examples of nonlinear probability weighting) they cannot be portrayed as misperceptions of an objective reality.

A variety of models have been proposed for explaining the typical pattern of behavior in Ellsberg’s paradox. One such model has already been introduced, namely the incomplete-preference model in which beliefs may be represented by imprecise probabilities. Bewley (1986) has used this model as the basis for a “Knightian” decision theory by adding to it a behavioral assumption of *inertia*, namely that when presented with a new alternative, the decision maker will choose it only if its minimum possible expected value or expected utility exceeds that of the status quo. Under this theory, the decision maker’s beliefs with respect to the draw of a red ball from urn 2 can be represented by an interval of probabilities, perhaps the entire interval [0, 1]. The minimum expectation of acts y and y' is therefore something less than \$50 (perhaps as low as \$0), and their maximum expectation is greater than \$50 (perhaps as great as \$100), whereas acts x and x' have precise expectations of \$50. Such a decision maker is technically undecided between the two alternatives in each pair, that is, she regards them as noncomparable. However, if a status quo position is specified and inertia is assumed, the decision maker’s choice can be predicted. For example, if the status quo is to “do nothing” or perhaps receive some fixed payment of less than \$50, then acts x and x' are preferable to the status quo whereas y and y' are not. Bewley’s inertia assumption has interesting implications for intertemporal decisions, where it gives rise to hysteresis and path-dependence, but its use of the status quo as a reference point seems arbitrary in some applications and has been criticized as an ad hoc assumption rather than a compelling normative postulate. (Its descriptive validity has also been questioned, for example, by Eisenberger and Weber 1995.) To return to Ellsberg’s example, if the status quo is possession of y or y' , the inertial decision maker will never wish to trade it for x or x' , contrary to intuition.

As an alternative to the inertia assumption, it can be assumed that, no matter what alternatives are offered, the decision maker always chooses the one that *maximizes the minimum possible expected utility* over a set P of probability distributions that represents her imprecise beliefs. This is the implication of Gilboa and Schmeidler’s (1989) *maxmin expected utility model*, also commonly known as the “multiple priors” model, in which the utility of act x is:

$$MEU(x) = \min_{p \in P} \mathbf{E}_p[v(x)] = \min_{p \in P} \sum_{i=1}^n p_i v(x_i).$$

This model rationalizes the typical pattern of responses in Ellsberg’s paradox as long as the set P of prior probabilities for drawing a red ball from urn 2 includes values on both sides of $\frac{1}{2}$. It is derived from a modification of Anscombe and Aumann’s version of the SEU axioms in which the independence condition is replaced by a weaker condition of “certainty independence” and an explicit axiom of uncertainty aversion is adopted. The uncertainty aversion axiom states that whenever a decision maker is indifferent between two acts, she prefers to objectively randomize between them. This axiom jibes nicely with another observation about Ellsberg’s two-urn problem, namely that when subjects are presented with a head-to-head choice between the two

ambiguous acts y and y' , they prefer to randomize between them by flipping a coin, which resolves the ambiguity by rendering the probability of a red ball irrelevant.

The multiple-priors model does not really distinguish between the decision maker's *perception* of ambiguous probabilities and her *attitude* toward the ambiguity: once the set of priors has been specified, the only permissible ambiguity attitude is the maximally-pessimistic one that is implicit in the maxmin-EU decision rule. A more nuanced approach, which permits a range of attitudes toward ambiguity, is provided by the concept of a *second-order utility function*, which appears in recent papers by Klibanoff et al. (2005), Ergin and Gul (2004), Chew and Sagi (2006), and Nau (2001b, 2006a). The second-order utility function is applied to the expected utilities computed at an intermediate stage of solving the decision tree, while the usual first-order utility function is applied to the terminal payoffs, so that the overall evaluation of an act is based on an expected-utility-of-an-expected-utility. There are two versions of the second-order utility model, one of which involves a unique prior distribution (the "source-dependent risk attitude" model), and one of which involves multiple priors (the "uncertain prior" model).

To illustrate the source-dependent risk aversion attitude model, suppose that the decision problem involves two distinct sources of uncertainty, represented by logically independent sets of events (A_1, \dots, A_J) and (B_1, \dots, B_K) , so that the state space is the Cartesian product $(A_1, \dots, A_J) \times (B_1, \dots, B_K)$, where the A -events are a priori "more uncertain" (such as the color of the ball from Ellsberg's second urn) and the B -events are a priori "less uncertain" (such as the color of the ball from the first urn). Suppose that the decision maker has probabilistic beliefs in which her unconditional distribution on the A -events is $\mathbf{p} = (p_1, \dots, p_J)$ and given event A_j her conditional distribution on the B -events is $\mathbf{q}_j = (q_{j1}, \dots, q_{jK})$. Thus, her probability for state $A_j B_k$ is exactly $p_j q_{jk}$. Let acts be represented by doubly-subscripted payoff vectors, so that x_{jk} is the payoff of act \mathbf{x} in state $A_j B_k$. If the decision maker is an "uncertainty neutral" SEU maximizer, her risk attitudes are described by a first-order utility function $v(x)$, and her evaluation of \mathbf{x} is:

$$SEU(\mathbf{x}) = \sum_{j=1}^J p_j \sum_{k=1}^K q_{jk} v(x_{jk}).$$

When the decision tree is solved by dynamic programming (backward induction), the conditional expected utility given event A_j is computed first, and the expected value of the conditional expected utility is computed second, although the order of events is unimportant: the decision tree could be "flipped" by an application of Bayes' rule. If she has no prior stakes in events, such a decision maker will exhibit the same degree of risk aversion toward bets on A -events or B -events, as determined by the concavity of $v(x)$. Suppose, however, that she also has a second-order utility function $u(v)$, which is applied to the conditional expected utility given a particular A -event, so that her evaluation of \mathbf{x} is based on the two-stage expected-utility calculation:

$$SOU(\mathbf{x}) = \sum_{j=1}^J p_j u \left(\sum_{k=1}^K q_{jk} v(x_{jk}) \right). \quad (3)$$

Under this preference model, if there are no prior stakes, the decision maker will behave toward bets on B -events as if her utility function were $v(x)$, and she will behave toward bets on A -events as if her utility function were $u(v(x))$. If u is a concave function, then despite the fact that she assigns unique probabilities to all events, the decision maker will be uniformly more risk averse toward bets on A -events than toward bets on B -events, as though she is averse to uncertainty.

Tree-flipping can no longer be performed in this case: the decision tree is solvable by dynamic programming only if it is drawn so that A -events are resolved first.

To see how this model potentially explains Ellsberg's paradox, suppose that the first-order utility function is simply the linear function $v(x) = x$, that is, the individual is risk neutral, and the second-order utility function is any concave function, such as the exponential utility function $u(v) = -\exp(-\alpha v)$, where the parameter α is now a measure of (constant) uncertainty aversion. For simplicity, assume $\alpha = \ln(4)/100$, so that $u(100) = -0.25$, $u(50) = -0.5$, and $u(0) = -1$. Then the Ellsberg acts \mathbf{x} and \mathbf{y} are evaluated as follows:

$$SOU(\mathbf{x}) = \frac{1}{2} u(\frac{1}{2}(100) + \frac{1}{2}(0)) + \frac{1}{2} u(\frac{1}{2}(100) + \frac{1}{2}(0)) = \frac{1}{2} u(50) + \frac{1}{2} u(50) = -0.5 \quad (4a)$$

$$SOU(\mathbf{y}) = \frac{1}{2} u(\frac{1}{2}(100) + \frac{1}{2}(100)) + \frac{1}{2} u(\frac{1}{2}(0) + \frac{1}{2} v(0)) = \frac{1}{2} u(100) + \frac{1}{2} u(0) = -0.625 \quad (4b)$$

The corresponding certainty equivalents for \mathbf{x} and \mathbf{y} are $u^{-1}(-0.5) = \$50$ and $u^{-1}(-0.625) = \$33.90$, respectively, and exactly the same values would be obtained for \mathbf{x}' and \mathbf{y}' , rationalizing the usual strict preferences for \mathbf{x} over \mathbf{y} and for \mathbf{x}' over \mathbf{y}' .

The source-dependent risk attitude model can be derived from a two-stage application of the independence axiom, in which it is first assumed to apply only to acts that agree on entire A -events, which permits conditional preferences to be defined on A -events, and then the A -conditional preferences are assumed to satisfy independence with respect to agreements on B -events. In general, the first- and second-order utility functions may be state-dependent (Model I in Nau 2006a), although state-independence can be forced by replacing the nested independence axioms with stronger nested axioms of tradeoff consistency (Model II in Nau, 2006a, analogous to Theorem 3 of Ergin and Gul 2004).

Alternatively, a second-order utility function can be embedded in a model of uncertain priors, which is somewhat more general than the source-dependent risk attitude model (3) and is easier to relate to the multiple-priors model as well as to the literature of hierarchical Bayesian models. Suppose that the decision maker's ambiguous beliefs are represented by a *second-order probability distribution* over a set of possible first-order priors, as in a hierarchical Bayesian model, except with the added twist that the decision maker is averse to the uncertainty about which prior is "correct," or more precisely she is *averse to the uncertainty in her first-order expected utility* that is induced by her uncertainty about which prior is correct. The second-order utility function models her attitude toward the second-order uncertainty. Suppose that the set of priors is a finite set with I elements, $(\mathbf{q}_1, \dots, \mathbf{q}_I)$, and to maintain continuity with the previous model, assume that the set of payoff-relevant events is the set $(A_1, \dots, A_J) \times (B_1, \dots, B_K)$, although this Cartesian product structure is no longer essential. Then \mathbf{q}_i is a vector with elements indexed by jk , where q_{ijk} denotes the probability of event $A_j B_k$ under prior i . For simplicity, assume that under every prior the decision maker has the same risk attitude, represented by a first-order utility function $v(x)$, although more generally her first-order risk attitude could also vary with the index i . (For example, the decision maker could be acting on behalf of a group of individuals with different beliefs *and* risk attitudes, and she might be averse to any lack of consensus among them.) Finally, let $\mathbf{p} = (p_1, \dots, p_I)$ denote the decision maker's second-order probability distribution over the first-order priors, and let $u(v)$ denote her second-order utility function that

represents aversion to second-order uncertainty, so that her overall evaluation of an act x is given by the two-stage expected-utility calculation:

$$SOU(x) = \sum_{i=1}^I p_i u \left(\sum_{j=1}^J \sum_{k=1}^K q_{ijk} v(x_{jk}) \right). \quad (5)$$

This version of the second-order-utility model is discussed by Nau (2001b) and axiomatized for general (possibly infinite) sets of states and consequences by Klibanoff et al. (2005).

The uncertain-prior model is strictly more general than the source-dependent risk attitude model and the maxmin EU model. The source-dependent-risk attitude model can be viewed as the special case in which $I=J$ and the j^{th} prior simply assigns probability 1 to event A_j . Then \mathbf{p} is effectively the probability distribution over the A -events and q_{ijk} is the conditional probability of B_j given A_j if $i=j$, and $q_{ijk} = 0$ otherwise, whence (5) reduces to (3). For example, the specific model (4ab) that was used above to explain Ellsberg's two-urn paradox can be reinterpreted to mean that the decision maker thinks urn 2 is rigged, but she does not know which way. Specifically, she thinks that it either contains 100 percent red balls or 100 percent black balls, she regards these possibilities as equally likely, and her aversion to this second-order uncertainty is modeled by the second-order utility function $u(v) = -\exp(-\alpha v)$. The maxmin EU model, meanwhile, is a limiting case of the uncertain-prior model in which the decision maker has a uniform distribution over some set of possible priors but is pathologically averse to second-order uncertainty (i.e., $u(v)$ is a radically concave function, like the exponential function with a large risk aversion coefficient α), so her second-order expected utility for an act is simply its worst-case first-order expected utility.

To sum up this section, there is fairly compelling evidence that individuals do not – and perhaps should not – treat all risks the same. Thus, a decision maker might display different risk attitudes in betting on the spin of a roulette wheel, betting on the outcome of a football game, purchasing various kinds of insurance, and managing a retirement portfolio. A variety of formal extensions of the SEU model have been developed to model this phenomenon. Whether or not one of these models is adopted, the bottom line is that assessments of the decision maker's risk attitude ideally should be carried out in terms of hypothetical bets on real events that are similar to the ones involved in the decision at hand, not arbitrary, artificial events.

5. State-preference theory, state-dependent utility, and decision analysis with risk-neutral probabilities

Much of the controversy over SEU theory in the last few decades has focused on the normative or descriptive validity of particular preference axioms such as independence or completeness, taking as given the underlying analytic framework introduced by Savage, in which acts consist of arbitrary mappings from a set of states of the world to a set of consequences, or the more tractable framework of Anscombe and Aumann that also includes objective randomization. However, the primitive concepts of those analytic frameworks are also open to question, particularly the concept of a "consequence" that yields the same value to the decision maker no matter what state of the world it occurs in, and the concept of an "act" in which states are mapped to consequences in arbitrary and often counterfactual ways. Savage introduced these concepts in order to build a theory in which the decision maker's beliefs would turn out to be representable by subjective probabilities, which in turn would be uniquely determined by her

preferences. The use of subjective probabilities to represent beliefs was controversial then, and to some extent it is still controversial today, partly due to the persistence of alternative views of probability among many economists and philosophers, but also partly because of an emerging realization that the preferences among counterfactual acts which are required by Savage's definition of subjective probability are not really observable because they cannot be instantiated by feasible choices. But more importantly, even if they *were* observable, they *still* would not enable the decision maker's subjective probabilities to be uniquely separated from her utilities for consequences.

The problem with trying to define probabilities in terms of preferences is that no matter how consequences are defined, it is impossible to verify whether their utilities are really state-independent. The axiom that Savage uses to force state-independence of preferences for consequences (P3) does not actually imply that *utilities* must be state-independent: it requires the decision maker's utility function to rank the consequences in the same order in every state of the world, but their utilities could still have *state-dependent scale factors* whose effects would be impossible to separate from the effects of subjective probabilities. Such an entanglement of probability and utility is quite likely to arise in decisions that involve life-changing events, such as insurance, health care, savings for retirement, or choosing a career or a spouse. In a famous exchange of letters with Savage, Aumann (1971) raised the example of a man who must decide whether his critically ill wife should undergo a risky operation. If the man's enjoyment of utterly everything would be severely diminished by the event of his wife's death, then there may be no "consequence" whose utility is state-independent for him, and the contemplation of bets on her survival may not help him to think more clearly about the probability of that event.

A related problem is that the decision maker could have *unobservable prior stakes in events*, which would give rise to the appearance of state-dependent utilities and/or distorted probabilities in her observed preferences. For example, suppose that the decision maker is an SEU-maximizer with subjective probability distribution \mathbf{p} and state-independent exponential utility function $v(x) = -\exp(-\alpha x)$, but meanwhile she has large prior financial stakes in events – that is, "background risk" – represented by a wealth vector $\mathbf{w} = (w_1, \dots, w_n)$, known only to herself. Then her expected utility of an act whose observable payoff vector is \mathbf{x} is actually:

$$SEU(\mathbf{w}+\mathbf{x}) = -\sum_{i=1}^n p_i \exp(-\alpha(w_i + x_i)) = -\sum_{i=1}^n (p_i \exp(-\alpha(w_i))) \exp(-\alpha x_i),$$

which, up to a scale factor, is also the expected utility of \mathbf{x} for someone with identical risk aversion who has *no* prior stakes but whose probability for state i is $p_i' \propto p_i \exp(-\alpha w_i)$. Hence, even under the assumptions of the SEU model, it is impossible for an outside observer to infer the decision maker's true probabilities from her preferences without independently knowing her prior stakes in events, and for the same reason it may be hard for a decision maker with large vested interests to think clearly about her own beliefs by introspectively examining her own preferences. This is likely to be especially problematic when contemplating events that strongly affect everyone, such as natural or economic disasters. Many individuals now have their retirement portfolios invested in mutual funds, even if they are not professional investors, and moreover everyone's economic future is likely to be affected in some complex way by a "boom" or "bust" in the market, so it is hard to conceive of a realistic way to model attitudes toward uncertainty in financial markets that does not admit the presence of background risk.

Various authors have studied the difficulties of subjective probability measurement that are raised by state-dependent utility and prior stakes (e.g., Fishburn 1970; Karni et al. 1983; Karni 1985; Shafer 1986; Rubin 1987; Drèze 1987; Kreps 1988; Kadane and Winkler 1988; Schervish et al. 1990; Karni and Schmeidler 1993; Karni 1996; Wakker and Zank 1999; Karni and Mongin 2000; Nau 1995, 2001a), and the results can be described as a mix of good and bad news for decision analysis. The *bad* news is that there does not appear to be any foolproof way to measure “true” subjective probabilities by eliciting preferences among acts that are actually available or even remotely possible. Rather, it generally seems to be necessary to elicit preferences among acts that are by definition impossible, such as acts in which arbitrary objective probabilities are assigned to states of world (as in Karni et al. 1983), or else to abandon the preference-based approach and instead treat subjective probability as an undefined psychological primitive that can be measured verbally without reference to specific acts (as in Degroot 1970). The conspicuous exceptions to this rule are situations where events really do have objective probabilities (or “almost-objective” probabilities as defined by Machina 2005) and/or the decision maker has no intrinsic interest in the events. But the *good* news is that *it does not matter*: it is possible to carry out decision analysis without measuring the “true” subjective probabilities of the decision maker, particularly in problems involving financial markets.

An alternative framework for analyzing choice under uncertainty, which sidesteps the measurement of subjective probabilities, was proposed by Arrow (1951, 1953) at around the same time that Savage unveiled his SEU model and Allais responded with his paradox. In fact, Savage, Allais, and Arrow all gave presentations of their nascent theories at a legendary econometrics colloquium held in Gif-sur-Yvette, near Paris, in May 1952 (CNRS 1953). Arrow’s approach, which has come to be known as *state-preference theory*, was refined and extended by Debreu (1959), Hirshleifer (1965), Yaari (1969), and Drèze (1970, 1987), among others. It has been very widely adopted in financial economics, especially in the theories of insurance, general equilibrium under uncertainty, and asset pricing by arbitrage. Economists have been receptive to this approach in part because it is agnostic concerning the existence of subjective probabilities, and in part because it uses old familiar concepts of utility measurement that date back to the work of Walras, Pareto, and Edgeworth in the nineteenth century. State-preference theory is merely a straightforward adaptation of neoclassical consumer theory to an environment in which goods may be distinguished by the times and states in which they are delivered or consumed. Thus, objects of preference are not just apples, bananas, and money, but also *state-contingent* and/or *time-contingent claims* to apples, bananas, and money.

In the simplest case, which resembles the framework of the SEU and non-SEU models discussed earlier, there is a single date, a single consumption good (money), and n states of the world. The states could be possible values for a stock index or an interest rate, or various hazards that might be insured against, or any other verifiable events on which financial contracts can be written. In this setting, an act is a vector \mathbf{x} whose i th component is an amount of money x_i to be paid or received in state i . These are not consequences in Savage’s sense because receiving a given amount of money need not count as the same experience in all states of the world, and it need not represent the decision maker’s entire wealth in that state – she could have prior stakes. If preferences among such acts are merely assumed to be continuous, reflexive, transitive, and monotonic, it follows that there is some continuous utility function $U(\mathbf{x})$, increasing in x_i for each i , that represents the decision maker’s preferences under uncertainty.

If no stronger assumptions are imposed, then $U(\mathbf{x})$ represents general nonexpected utility preferences, and it is merely an *ordinal* utility function in the tradition of neoclassical consumer theory: any monotonic transformation of U represents the same preferences, hence, it is not meaningful to compare utility differences between different pairs of acts. However, if U is a *differentiable* function – which can be assumed if preferences are sufficiently smooth – it is meaningful to compare the *marginal utilities* of different state-contingent commodities, and a great deal of analysis can be done in these terms. Assuming that more money is strictly preferred to less, the partial derivatives of U are positive and can be normalized to yield a local probability distribution associated with act \mathbf{x} , which will be denoted here by $\boldsymbol{\pi}(\mathbf{x})$. That is:

$$\pi_i(\mathbf{x}) \equiv \frac{\partial U(\mathbf{w}) / \partial x_i}{\sum_{j=1}^n \partial U(\mathbf{w}) / \partial x_j} .$$

The distribution $\boldsymbol{\pi}(\mathbf{x})$ need not represent the decision maker’s true probabilities, but it is very useful information because it completely characterizes the decision maker’s local gambling behavior in the vicinity of \mathbf{x} . In particular, one who already possesses \mathbf{x} would strictly prefer to accept a small gamble \mathbf{z} (i.e., would strictly prefer $\mathbf{x}+\mathbf{z}$ over \mathbf{x}) if and only if \mathbf{z} has positive expected value under $\boldsymbol{\pi}(\mathbf{x})$, that is, $\mathbf{E}_{\boldsymbol{\pi}(\mathbf{x})}[\mathbf{z}] > 0$, provided that \mathbf{z} is small enough so that $\boldsymbol{\pi}(\mathbf{x}+\mathbf{z})$ is not significantly different from $\boldsymbol{\pi}(\mathbf{x})$. Thus, every decision maker with sufficiently smooth preferences behaves locally like a risk-neutral SEU-maximizer, and it is therefore appropriate to refer to the local distribution $\boldsymbol{\pi}(\mathbf{x})$ as her *risk-neutral probability distribution* evaluated at \mathbf{x} .

If the independence axiom is also assumed, it follows that the utility function $U(\mathbf{x})$ has the additively separable *cardinal* form:

$$U(\mathbf{x}) = \sum_{i=1}^n u_i(x_i), \tag{6}$$

which is a *state-dependent SEU model without separation of probabilities from utilities*. The state-dependent utility functions $\{u_i\}$ simply lump together the effects of beliefs, risk attitudes, and prior stakes in events, and the decision maker’s risk-neutral probabilities are their normalized derivatives: $\pi_i(\mathbf{x}) \propto u'_i(x_i)$. The standard SEU model is obtained if $u_i(x) = p_i v(x)$ for some “true” probability distribution \mathbf{p} and state-independent utility function $v(x)$, in which case risk-neutral probabilities are the product of true probabilities and state-dependent marginal utilities: $\pi_i(\mathbf{x}) \propto p_i v'(x_i)$. But risk-neutral probabilities exist for decision makers with more general smooth preferences, and they are useful terms of analysis even for decision makers with SEU preferences, especially in multiagent settings such as games and markets.

It is obvious that low-stakes financial decisions can be analyzed in terms of the decision maker’s risk-neutral probabilities because the risk-neutral local approximation applies to small transactions by definition. However, it is also possible, in principle, to analyze *any* decision entirely in terms of risk-neutral probabilities, regardless of the stakes, because the decision maker’s indifference curves are completely determined by $\boldsymbol{\pi}(\mathbf{x})$ – it is merely necessary to take into account the functional dependence of $\boldsymbol{\pi}$ on \mathbf{x} . Decisions that take place in the context of a *complete market for contingent claims* are especially easy to analyze in these terms. In a complete and arbitrage-free market, the Arrow-Debreu security for each state (which pays \$1 if that state occurs and \$0 otherwise) has a unique market price, and normalization of these state

prices yields a probability distribution that is called the *risk neutral distribution of the market*, which will be denoted here by π^* . Suppose the decision maker faces a choice among some set X of alternative capital-investment projects that can be financed through investments in the market, and suppose that she is a “small player” in the sense that her choices are not expected to influence market prices. In the simplest case, there is single future time period and the projects can be described by vectors representing the net present (or future) value of their cash flows in different states of the world. Therefore, let $\mathbf{x} \in X$ denote the vector whose i th element x_i is the NPV of a project in state E_i . Then the decision maker’s utility-maximizing strategy is simply to choose the project with the highest market risk-neutral valuation, that is, choose project \mathbf{x}^* where $\mathbf{x}^* = \arg \max \{E_{\pi^*}[\mathbf{x}], \mathbf{x} \in X\}$, assuming this is positive, and meanwhile execute trades in the market to yield additional cash flows \mathbf{z}^* such that $E_{\pi^*}[\mathbf{z}^*] = 0$ (i.e., the trades are self-financing at relative prices π^*) and such that $\pi(\mathbf{x}^* + \mathbf{z}^*) = \pi^*$ (i.e., her “posterior” risk-neutral probabilities equal those of the market, thus optimally financing the project \mathbf{x}^*). Effectively the decision maker reaps an arbitrage profit equal to $E_{\pi^*}[\mathbf{x}^*]$ and then optimally invests the money in the market. The *only* role for her own preferences is to determine the optimal financing scheme \mathbf{z}^* , not the optimal project \mathbf{x}^* , and the problem of finding \mathbf{z}^* can be parameterized in terms of the decision maker’s risk-neutral probabilities rather than her true probabilities. The same approach generalizes to multiperiod problems, although it becomes necessary to explicitly model the decision maker’s preferences for consumption in different periods as well as different states of the world, and the analysis is greatly simplified (i.e., decision tree can be solved by dynamic programming) if the independence axiom is applied so that utility is additively separable across *both* periods and states of the world.

The presence of a complete market thus allows a financial decision problem to be neatly separated into two parts, one of which is solved by the market and the other of which is solved by the decision maker, and both of which can be solved in terms of risk-neutral probabilities. In the case where X consists of a single project and the only question is how to optimally finance it, this result is known as the “fundamental theorem of risk bearing”: at the risk bearing optimum, the individual’s relative marginal expected utilities must equal relative market prices. If the market is only “partially complete,” then its risk-neutral probabilities for some events will be imprecise and there will be a larger role for the decision maker’s own preferences, but under suitable restrictions it is possible to decompose the dynamic-programming solution of the decision tree into an alternation between steps in which the risk-neutral probabilities of the market are used and steps in which the risk-neutral probabilities of the decision maker are used (Smith and Nau 1995). More details and examples of decision analysis in terms of risk-neutral probabilities are given by Nau and McCardle (1991) and Nau (2001a, 2003).

To sum up this section, Savage’s analytic framework of states, consequences, and acts is not the only possible framework for modeling choice under uncertainty. The alternative framework of state-preference theory is more appropriate for many applications, particularly those involving financial markets, and it does not unduly emphasize the state-independence of utility or the measurement of “true” subjective probabilities. Using this framework, it is possible to recast many of the tools of decision analysis (risk aversion measures, risk premia, decision tree solution algorithms) in terms of risk-neutral probabilities, which are often the most natural parameters in settings where risks can be hedged by purchasing financial assets.

6. Neuroeconomics: the next (and final?) frontier

The SEU model and its extensions (including game theory) follow a centuries-old tradition of importing mathematical methods from physics into economics, descending from Bernoulli to von Neumann. Individuals are imagined to rationally pursue a single goal, namely the maximization of some all-inclusive measure of utility, which enables the analytic tools of physical scientists (calculus, probability theory, equilibrium, etc.) to be applied to the study of human behavior. The brain is conceived as a calculating engine that operates on two distinct sources of data, namely *beliefs* (represented by probabilities of events that might occur) and *values* (represented by utilities of outcomes that might be experienced), which it uses to evaluate every new opportunity. The same types of calculations are performed regardless of whether the situation is simple or complicated, familiar or unfamiliar, certain or uncertain, financial or nonfinancial, competitive or noncompetitive. The great innovation of von Neumann and Morgenstern and Savage was to show that this model of rational decision making is implied, in an “as if” sense, by simple axioms of consistency among preferences. Thus, SEU theory is at bottom a theory of *consistency* rather than a theory with any particular empirical content.

The utilitarian, consistency-focused view of decision making has always been to some extent controversial. During the last 50 years, the bulk of the criticism has come from behavioral economists and psychologists in the tradition of Herbert Simon, Vernon Smith, Richard Thaler, Amos Tversky, and Daniel Kahneman, who have imported experimental methods from psychology into economics. The latest frontier has been opened during the last decade or so by developments in cognitive neuroscience, giving rise to an emerging field of “neuroeconomics,” which studies the root causes of economic behavior at a physiological level, yielding empirical insights into decision making that are potentially explanatory rather than merely descriptive. The most distinctive tools of neuroeconomics are real-time measurements of the neural activity of experimental subjects, which have revealed that different brain areas may be activated by different attributes of decisions – for example, certainty versus uncertainty, gains versus losses, immediate versus delayed rewards. Such measurements include functional MRI, positron emission topography, and magnetoencephalography in human-subject experiments; more invasive studies of single-neuron activity in animal experiments; and recordings of other psychophysical variables such as galvanic skin response and pupil dilation. Other research tools include “depletion studies” that explore the effects on mood and emotion of amino acids such as tryptophan and serotonin; administration of hormones such as oxytocin, which raises the level of “trust” among subjects in experimental games; electrical stimulation of different brain regions to determine the sensations or feelings that are processed there; case studies of decision making by patients with who have suffered lesions in specific brain regions; and behavioral experiments designed to test hypotheses about decision processes that are otherwise suggested by evolutionary and/or neuroscientific arguments. Some excellent recent surveys of the field have been given in review articles by Camerer et al. (2004ab) and Glimcher and Rustichini (2004), as well as a book by Glimcher (2003); philosophical and social implications of research in cognitive neuroscience have also been discussed in popular books by Damasio (1994, 2003) and Pinker (1997, 2002).

There is insufficient space here to do justice to this vibrant area, but some of the stylized facts are as follows. First, the human brain turns out not to be a calculating engine with a knowledge base of coherent beliefs and values. Rather, it is composed of distinct modules and

layers, which are: (1) responsible for different aspects of cognition and behavior; (2) partly “programmable” by experience but also partly “hard-wired” by evolution; (3) sometimes influenced by transient body chemistry; and (4) not always in perfect harmony with each other. Like other anatomical structures, they have evolved by a series of “exaptations” from structures that may have originally served other purposes in other environments. Some brain activities are under conscious control, but others proceed automatically, and “learning” consists not only of acquiring information but also of shifting repetitive tasks from conscious to automatic control.

An important distinction can be drawn between cognitive processes (reasoning) and affective processes (emotions). Cognitive processes are controlled by areas of the higher brain (in the cerebral cortex), and affective processes are controlled by areas of the lower brain (especially the amygdala, which is closely associated with the sense of smell). Affective responses have an intrinsic positive or negative valence (i.e., they are automatically classified as good or bad), they may be accompanied by physiological reactions in other parts of the body (literally “gut feelings”), and they often occur automatically and unconsciously, so that the higher brain is not always aware of the stimuli that may have provoked a particular emotional response (Bechara et al. 1999). One of the most striking findings is that affective processes appear to be critical for decision making: patients who have suffered brain lesions that disengage their cognitive processes from their emotions are sometimes unable to make effective decisions on their own behalf, despite otherwise being able to reason as logically as before (Bechara et al. 1994; Damasio 1994). Risk and ambiguity normally give rise to emotions of fear or discomfort (although persons who are genetically predisposed to compulsive behavior may get emotional thrills from inappropriate risk-taking), and they appear to activate different brain areas (Hsu et al. 2005; Huettel et al. 2006); patients with lesions in these areas have been shown to be less risk-averse or less ambiguity-averse than normal subjects (Shiv et al. 2005), for better or worse.

The brain, due to its modular structure, is not equally good at solving all kinds of inference and decision problems. Rather, some of its parts are specialized for efficiently solving classes of problems that were especially important for the survival and reproduction of early humans, including the use of language and strategic reasoning about reciprocity and exchange, as well as intuition about the physical and biological environment. For example, there appears to be a built-in cheater-detection system: if a logical inference problem can be framed in terms of detecting an instance of cheating on an obligation, it is much more likely to be correctly solved by an experimental subject than if it is presented in abstract terms. There is also some evidence of a utility-for-money effect, that is, money appears to have a direct utility rather than merely an indirect utility for the consumption that it buys later (Prelec and Loewenstein 1998), whereas von Neumann-Morgenstern and Savage deliberately constructed their theories so that money would play no distinguished role. The human brain also appears to have a module that harbors a “theory of mind” concerning the motivation and reasoning of other human agents, which is necessary for social interactions and game-playing but is absent or damaged in autistic individuals. However, as Camerer et al. (2004b) observe, “the modularity hypothesis should not be taken too far. Most complex behaviors of interest to economics require collaboration among more specialized modules and functions. So the brain is like a large company – branch offices specialize in different functions, but also communicate to one another, and communicate more feverishly when an important decision is being made.”

Although much of neuroeconomics has sought to import findings from neuroscience into economics, some of it has aimed the other way, looking for evidence of Bayesian inference or

expected-utility maximization at a neuronal level, on the theory that evolutionary pressure for “rational choice” algorithms may be strongest in more primitive areas of the brain where the simplest decisions are made. In single-neuron studies in which animal subjects receive uncertain rewards for either strategic or nonstrategic choices, where both the probabilities and values of the rewards are manipulated, the firing rate of activated neurons in the parietal area appears to be proportional to the “relative subjective desirability” of the chosen alternative, that is, it depends on the probability as well as the value in relative terms, compared to other available alternatives (Glimcher and Platt 1999; Dorris and Glimcher 2004). This finding tentatively suggests that probability may not be separated from utility at a primitive level as learning occurs under risk and uncertainty.

The neuroeconomic view of a modular, internally-collaborative, emotion-driven and not-fully-transparent brain raises new questions about the extent to which individual decisions can be improved by externalizing beliefs and values and forcing them into a single, coherent structure that is supposed to govern preferences across a wide range of situations. Effective decision making under risk and uncertainty appears to also require drawing on the right part of the brain at the right time, harnessing appropriate emotions, and suppressing inappropriate ones. Individuals may differ widely in their innate and acquired abilities to do this, hence the same prescriptive tools may not work equally well for all decisions and decision makers.

References

- Allais, M. (1953). Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Américaine. *Econometrica*, 21, 503–546.
- Anscombe, F., and Aumann, R. (1963). A Definition of Subjective Probability. *Annals of Mathematical Statistics*, 34, 199–205.
- Arrow, K. (1951). An Extension of the Basic Theorems of Classical Welfare Economics. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley: University of California Press.
- Arrow, K. (1953/1964). The Role of Securities in the Optimal Allocation of Risk-Bearing. *Quarterly Journal of Economics*, 31, 91–96.
- Aumann, R. (1971). Letter to L. J. Savage. Reprinted in Drèze, J. (1987).
- Bechara, A., Damasio, A., Damasio, H., and Anderson, S. (1994). Insensitivity to Future Consequences Following Damage to Human Prefrontal Cortex. *Cognition*, 50, 7–15.
- Bechara, A., Damasio, H., Damasio, A., and Lee, G. (1999). Different Contributions of the Human Amygdala and Ventromedial Prefrontal Cortex to Decision-Making. *J. Neuroscience*, 19, 473–481.
- Becker, J., and Sarin, R. (1987). Lottery Dependent Utility. *Management Science*, 33, 1367–1382.
- Bell, D. (1982). Regret in Decision Making under Uncertainty. *Operations Research*, 30, 961–981.
- Bell, D. (1985). Disappointment in Decision Making under Uncertainty. *Operations Research*, 33, 1–27.
- Bernoulli, D. (1738). Specimen Theoriae Novae de Mensura Sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 5, 175–192. Translation by L. Sommer (1954) Exposition of a New Theory on the Measurement of Risk. *Econometrica*, 22, 23–36.
- Bewley, T. (1986). Knightian Decision Theory Part I. Cowles Foundation Discussion Paper No. 807. Reprinted in (2002). *Decisions in Economics and Finance*, 25, 79–110.
- Bleichrodt, H., and Pinto, J. (2000). A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis. *Management Science*, 46, 1485–1496.
- Bleichrodt, H., Pinto, J., and Wakker, P. (2001). Making Descriptive Use of Prospect Theory to Improve the Prescriptive Use of Expected Utility. *Management Science*, 47, 1498–1514.
- Camerer, C., Loewenstein, G., and Prelec, D. (2004a). Neuroeconomics: How Neuroscience Can Inform Economics. *Journal of Economic Literature*, XLIII, 9–64.
- Camerer, C., Loewenstein, G., and Prelec, D. (2004b). Neuroeconomics: Why Economics Needs Brains. *Scandinavian Journal of Economics*, 106, 555–579.

- Chew, S. H. (1983). A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox. *Econometrica*, 51, 1065–1092.
- Chew, S.H. and Sagi, J. (2006). Small Worlds: Modeling Attitudes Toward Sources of Uncertainty. Working paper, University of California at Berkeley.
- Chew, S. H., and Wakker, P. (1996). The Comonotonic Sure-Thing Principle. *Journal of Risk and Uncertainty*, 12, 5–27.
- Chew, S.H., Karni, E. and Safra, Z. (1987). Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities. *Journal of Economic Theory*, 42, 370–381.
- CNRS. (1953). *Econometrie, Paris, 12–17 Mai 1952*. Colloques Internationaux, 40, Centre National de la Recherche Scientifique, Paris.
- Damasio, A. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: Putnam.
- Damasio, A. (2003). *Looking For Spinoza: Joy, Sorrow, and the Feeling Brain*. Orlando: Harcourt.
- de Finetti, B. (1974). *Theory of Probability, Vol. 1*. New York: Wiley.
- Debreu, G. (1959). *Theory of Value*. New Haven: Cowles Foundation.
- Degroot, M. (1970). *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Diecidue, E., and Wakker, P. (2001). On the Intuition of Rank-Dependent Utility. *Journal of Risk and Uncertainty*, 23, 281–298.
- Dorris, M., and Glimcher, P. (2004) Activity in Posterior Parietal Cortex is Correlated with the Subjective Desirability of an Action. *Neuron*, 44, 365–378.
- Drèze, J. (1970). Market Allocation Under Uncertainty. *European Economic Review*, 2, 133–165.
- Drèze, J. (1987). *Essays on Economic Decision Under Uncertainty*. London: Cambridge University Press.
- Eisenberger, R., and Weber, M. (1995). Willingness-to-Pay and Willingness-to-Accept for Risky and Ambiguous Lotteries. *Journal of Risk and Uncertainty*, 10, 223–233.
- Ellsberg, D. (1961). Risk, Ambiguity and the Savage Axioms. *Quarterly Journal of Economics*, 75, 643–669.
- Ergin, H., and Gul, F. (2004). A Subjective Theory of Compound Lotteries. Working paper, MIT.
- Fishburn, P. (1970). *Utility Theory for Decision Making*. New York: Wiley
- Fishburn, P. (1982). Nontransitive Measurable Utility. *Journal of Mathematical Psychology*, 26, 31–67.
- Fishburn, P. and Wakker, P. (1995). The Invention of the Independence Condition for Preferences. *Management Science*, 41, 1130–1144.

- Gilboa, I., and Schmeidler, D. (1989). Maxmin Expected Utility with Nonunique Prior. *Journal of Mathematical Economics*, 18, 141–153.
- Glimcher, P. (2002). Decisions, Decisions, Decisions: Choosing a Neurobiological Theory of Choice. *Neuron*, 36, 323–332.
- Glimcher, P., and Rustichini, A. (2004). Neuroeconomics: The Concurrence of Brain and Decision. *Science*, 306, 447–452.
- Glimcher, P. (2003). *Decisions, Uncertainty, and the Brain: The Science of NeuroEconomics*. Cambridge, MA and London: MIT Press.
- Hacking, I. (1967). Slightly More Realistic Personal Probability. *Philosophy of Science*, 34, 311–325.
- Hirshleifer, J. (1965). Investment Decision under Uncertainty: Choice-Theoretic Approaches. *Quarterly Journal of Economics*, 74, 509–536.
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., and Camerer, C. (2005). Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making. *Science*, 310, 1680–1683.
- Huettel, S., Stowe, J., Gordon, E., Warner, B., and Platt, M. (2006). Neural Signatures of Economic Preferences for Risk and Ambiguity. *Neuron*, 49, 765–775.
- Kadane, J.B., and Winkler, R. L. (1988). Separating Probability Elicitation from Utilities. *Journal of the American Statistical Association*, 83, 357–363.
- Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47, 263–291.
- Karni, E. (1985). *Decision-Making under Uncertainty: The Case of State-Dependent Preferences*. Cambridge, MA: Harvard University Press.
- Karni, E. (1996). Probabilities and Beliefs. *Journal of Risk and Uncertainty*, 13, 249–262.
- Karni, E. and Mongin, P. (2000) On the Determination of Subjective Probabilities by Choices. *Management Science*, 46, 233–248.
- Karni, E., Schmeidler, D., and Vind, K. (1983). On State Dependent Preferences and Subjective Probabilities. *Econometrica*, 51, 1021–1031.
- Karni, E., and Schmeidler, D. (1993). On the Uniqueness of Subjective Probabilities. *Economic Theory*, 3, 267–277.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A Smooth Model of Decision Making Under Ambiguity. *Econometrica*, 73, 1849–1892.
- Knight, F. (1921). *Risk, Uncertainty, and Profit*. New York: Houghton Mifflin.
- Koopman, B. (1940). Axioms and Algebra of Intuitive Probability. *Annals of Mathematics*, 41, 269–292.
- Kreps, D. (1988). *Notes on the Theory of Choice*. Boulder: Westview Press.
- Kyburg, H. (1974). *The Logical Foundations of Statistical Inference*. Dordrecht: Reidel.

- Levi, I. (1980). *The Enterprise of Knowledge*. Cambridge, MA: MIT Press.
- Loomes, G., and Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice under Uncertainty. *Economic Journal*, 92, 805–824.
- Luce, R. D. (2000). *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. Mahwah, NJ: Erlbaum.
- Luce, R. D., and Fishburn, P. (1991). Rank-and Sign-Dependent Linear Utility Models for Finite First-Order Gambles. *Journal of Risk and Uncertainty*, 4, 29–59.
- Luce, R. D., and Narens, L. (1985). Classification of Concatenation Measurement Structures According to Scale Type. *Journal of Mathematical Psychology*, 29, 1–72.
- Luce, R. D., and Manders, K. (1988). Rank-Dependent Subjective Expected-Utility Representations. *Journal of Risk and Uncertainty*, 1, 305–332.
- Machina, M. (1982). Expected Utility without the Independence Axiom. *Econometrica*, 50, 227–323.
- Machina, M. (2005). ‘Expected Utility/Subjective Probability’ Analysis without the Sure-Thing Principle or Probabilistic Sophistication. *Economic Theory*, 26, 1–62.
- Machina, M. and Schmeidler, D. (1992). A More Robust Definition of Subjective Probability. *Econometrica*, 60, 745–780.
- Moskowitz, H., Preckel, P. and Yang, A. (1993). Decision Analysis with Incomplete Utility and Probability Information. *Operations Research*, 41, 864–879.
- Nau, R. (1989). Decision Analysis with Indeterminate or Incoherent Probabilities. *Annals of Operations Research*, 19, 375–403.
- Nau, R. (1992). Indeterminate Probabilities on Finite Sets. *Annals of Statistics*, 20, 1737–1767.
- Nau, R. (1995). Coherent Decision Analysis with Inseparable Probabilities and Utilities. *Journal of Risk and Uncertainty*, 10, 71–91.
- Nau, R. (2001a). De Finetti Was Right: Probability Does Not Exist. *Theory and Decision*, 59, 89–124.
- Nau, R. (2001b). Uncertainty Aversion with Second-Order Probabilities and Utilities. *Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications*. Web site: <http://www.sipta.org/isipta01/proceedings/063.html>.
- Nau, R. (2002). The Aggregation of Imprecise Probabilities. *Journal of Statistical Planning and Inference*, 105, 265–282.
- Nau, R. (2003). A Generalization of Pratt-Arrow Measure to Non-Expected-Utility Preferences and Inseparable Probability and Utility. *Management Science*, 49, 1089–1104.
- Nau, R. (2006a). Uncertainty Aversion with Second-Order Utilities and Probabilities. *Management Science*, 52, 136–145.
- Nau, R. (2006b). The Shape of Incomplete Preferences. *The Annals of Statistics*, 34, 2430–2448
- Nau, R. and McCardle, K. (1991). Arbitrage, Rationality, and Equilibrium. *Theory and Decision*, 31, 199–240.

- Pinker, S. (1997). *How the Mind Works*. New York: Norton.
- Pinker, S. (2002). *The Blank Slate: the Modern Denial of Human Nature*. New York: Viking.
- Platt, M. and Glimcher, P. (1999). Neural Correlates of Decision Variables in Parietal Cortex. *Nature*, 400, 233–238.
- Prelec, D. (1998). The Probability Weighting Function. *Econometrica*, 66, 497–527.
- Prelec, D., and Loewenstein, G. (1998). The Red and the Black: Mental Accounting of Savings and Debt. *Marketing Science*, 17, 4–28.
- Quiggin, J. (1982). A Theory of Anticipated Utility. *Journal of Economic Behaviour and Organization*, 3, 323–343.
- Rios Insua, D. (1990). *Sensitivity Analysis in Multiobjective Decision Making*. Berlin: Springer-Verlag.
- Rios Insua, D. (1992). On the Foundations of Decision Making Under Partial Information. *Theory and Decision*, 33, 83–100.
- Rios Insua, D., and Ruggeri, F. (2000). *Robust Decision Analysis*. New York: Springer-Verlag.
- Rubin, H. (1987). A Weak System of Axioms for “Rational” Behavior and the Non-Separability of Utility from Prior. *Statistics and Decisions*, 5, 47–58.
- Rubinstein, M. (1976). The Strong Case for the Generalized Logarithmic Utility Model as the Premier Model of Financial Markets. *Journal of Finance*, 31, 551–571.
- Savage, L. (1954). *The Foundations of Statistics*. New York: Wiley.
- Schervish, M., Seidenfeld, T., and Kadane, J. (1990). State-Dependent Utilities. *Journal of the American Statistical Association*, 85, 840–847.
- Schmeidler, D. (1989). Subjective Probability and Expected Utility without Additivity. *Econometrica*, 57, 571–587.
- Schmidt, U. (2004). Alternatives to Expected Utility: Formal Theories. In S. Barberà, P. Hammond, and C. Seidl, *Handbook of Utility Theory, Vol. 2*. Dordrecht: Kluwer.
- Segal, U. (1989). Anticipated Utility: A Measure Representation Approach. *Annals of Operations Research*, 19, 359–373.
- Shafer, G. (1986). Savage Revisited (including comments). *Statistical Science*, 1, 463–501.
- Shiv, B., Loewenstein, G., Bechara, A., Damasio, H., and Damasio, A. (2005). Investment Behavior and the Negative Side of Emotion. *Psychological Science*, 16, 435–439.
- Smith, C. (1961). Consistency in Statistical Inference and Decision. *Journal of the Royal Statistical Society B*, 23, 1–25.
- Smith, J., and Nau, R. (1995). Valuing Risky Projects: Option Pricing Theory and Decision Analysis. *Management Science*, 41, 795–816.
- Starmer, C. (2000). Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk. *Journal of Economic Literature*, XXXVIII, 332–82.

- Stigler, G. (1950). The Development of Utility Theory: I; II. *Journal of Political Economy*, 58, 307–327; 373–396.
- Sugden, R. (2004). Alternatives to Expected Utility: Foundations. In S. Barberà, P. Hammond, and C. Seidl, *Handbook of Utility Theory*, Vol. 2. Dordrecht: Kluwer.
- Tversky, A., and Fox, C. (1995). Weighing Risk and Uncertainty. *Psychological Review*, 102, 269–283.
- Tversky, A., and Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323.
- van de Kuilen, G., and Wakker, P. (2006). Learning in the Allais Paradox. *Journal of Risk and Uncertainty*, forthcoming.
- von Neumann, J., and Morgenstern, O. (1947). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Wakker, P. (1989). *Additive Representations of Preferences: A New Foundation for Decision Analysis*. Dordrecht: Kluwer.
- Wakker, P. (1991). Additive Representations on Rank-Ordered Sets. I. The Algebraic Approach. *Journal of Mathematical Psychology*, 35, 501–531.
- Wakker, P. (1996) The Sure-Thing Principle and the Comonotonic Sure-Thing Principle: An Axiomatic Analysis. *Journal of Mathematical Economics*, 25, 213–227.
- Wakker, P. (2006). Annotated References on Decisions and Uncertainty. Web site: <http://www1.fee.uva.nl/creed/wakker/refs/rfrncs.htm>.
- Wakker, P., and Tversky, A. (1993). An Axiomatization of Cumulative Prospect Theory. *Journal of Risk and Uncertainty*, 7, 147–176.
- Wakker, P., and Zank, H. (1999). State Dependent Expected Utility for Savage's State Space. *Mathematics of Operations Research*, 24, 8–34
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. London: Chapman & Hall.
- Wu, G., and Gonzalez, R. (1996). Curvature of the Probability Weighting Function. *Management Science*, 42, 1676–1690.
- Yaari, M. (1969). Some Remarks on Measures of Risk Aversion and their Uses. *Journal of Economic Theory*, 1, 315–329.
- Yaari, M. (1987). The Dual Theory of Choice under Risk. *Econometrica*, 55, 95–115.