Choice under Uncertainty

Advanced article

Mark J Machina, University of California, San Diego, California, USA

CONTENTS

The standard theory of individual choice under uncertainty consists of the joint hypothesis of expected utility risk preferences and probabilistic beliefs. Experimental work by both psychologists and economists has uncovered systematic departures from both hypotheses, and has led to the development of alternative, usually more general, models.

INTRODUCTION

Decisions under uncertainty take place in two types of settings. In settings of 'objective uncertainty', the probabilities attached to the various outcomes are specified in advance, and the objects of choice consist of 'lotteries' of the form $P = (x_1, p_1; \ldots; x_n, p_n)$, which yield outcomes or monetary pay-offs x_i with probability p_i, where $p_1 + \ldots + p_n = 1$. Examples include games of chance involving dice and roulette wheels, as well as ordinary lotteries.

In settings of 'subjective uncertainty', probabilities are not given, and the objects of choice consist of 'bets' or 'acts' $f(\cdot) = [x_1 \text{ on } E_1; \ldots; x_n \text{ on } E_n]$, which yield outcomes or pay-offs x_i in event E_i, for some mutually exclusive and exhaustive collection of events $\{E_1, \ldots, E_n\}$ which can be thought of as a partition of the set S of all possible 'states of nature'. Examples include bets on horse races or the weather, as well as standard insurance contracts.

Under objective uncertainty, choices are determined by an individual's attitudes towards risk. Under subjective uncertainty, they are additionally determined by the individual's subjective beliefs about the likelihoods of the various states and events.

EXPECTED UTILITY THEORY AND EXPERIMENTAL EVIDENCE

Axiomatic and Normative Foundations of Expected Utility Theory

The earliest formal hypothesis of individual attitudes towards risk, proposed by Pascal, Fermat and others in the seventeenth century, was that individuals evaluate monetary lotteries $P = (x_1, p_1; \ldots; x_n, p_n)$ simply on the basis of their mathematical expectation $E[P] = \sum_{i=1}^{n} x_i \cdot p_i$. This hypothesis was dramatically refuted by Daniel Bernoulli's 'St Petersburg paradox'. In this game, a fair coin is repeatedly flipped until it lands heads. If it lands heads on the first flip, the player wins \$1; if it does not land heads until the second flip, the player wins \$2; and in general, if it does not land heads until the i^{th} flip, the player wins $\$2^{i-1}$. Most people would prefer to receive a sure payment of, say, \$50 than a single play of the St Petersburg game, even though the expected pay-off of the game is $\frac{1}{2}\cdot\$1 + \frac{1}{4}\cdot\$2 + \frac{1}{8}\cdot\$4 + \ldots = \$\frac{1}{2} + \$\frac{1}{2} + \$\frac{1}{2} + \ldots = \$\infty$. In the first of what has turned out to be a long series of such developments, Bernoulli weakened the prevailing expected-value hypothesis by positing that individuals instead evaluate lotteries on the basis of their 'expected utility' $\sum_{i=1}^{n} U(x_i)\cdot p_i$, where the utility $U(x)$ of receiving a monetary amount x is probably subproportional to x. Bernoulli himself proposed the form $U(x) = \ln(x)$, which leads to an evaluation of the St Petersburg game consistent with typical actual play.

The expected utility hypothesis came to dominate decision theory on the twin bases of its elegant and highly normative axiomatic development (von Neumann and Morgenstern, 1944; Marschak, 1950) and its analytical power (Arrow, 1965; Pratt, 1964). In the modern approach, risk preferences are denoted by the individual's 'weak preference' relation $\succeq$ over lotteries, where $P^* \succeq P$ reads 'P^* is weakly preferred to P', and its implied 'strict preference' relation $\succ$ (where $P^* \succ P$ iff $P^* \succeq P$ but not $P \succeq P^*$) and 'indifference' relation $\sim$ (where $P^* \sim P$ iff $P^* \succeq P$ and $P \succeq P^*$). The preference relation $\succeq$ is said to be 'represented' by an expected utility preference function $V_{EU}(P) = \sum_{i=1}^{n} U(x_i)\cdot p_i$ if $P^* \geq P \Leftrightarrow V_{EU}(P^*) \geq V_{EU}(P)$. $U(\cdot)$ is called the 'von Neumann–Morgenstern utility function'.

The axiomatic and normative underpinnings of expected utility theory are based on the notion of a 'probability mixture' $\alpha\cdot P + (1-\alpha)\cdot P^*$ of two lotteries $P = (x_1, p_1; \ldots; x_n, p_n)$ and $P^* = (x_1^*, p_1^*; \ldots; x_{n^*}^*, p_{n^*}^*)$, which is the lottery that would be generated by a coin flip yielding the lotteries P and P^* as prizes with respective probabilities α and $1-\alpha$, and where both stages of uncertainty (the coin flip and the resulting lottery) are realized simultaneously, so that we can write $\alpha\cdot P + (1-\alpha)\cdot P^* = (x_1, \alpha\cdot p_1; \ldots; x_n, \alpha\cdot p_n; x_1^*, (1-\alpha)\cdot p_1^*; \ldots; x_{n^*}^*, (1-\alpha)\cdot p_{n^*}^*)$. A preference relation $\succeq$ will then be represented by an expected utility preference function $V_{EU}(\cdot)$ for some utility function $U(\cdot)$ if and only it satisfies the following axioms:

- *Completeness*. For all lotteries P and P^*, either $P \succeq P^*$ or $P^* \succeq P$, or both.
- *Transitivity*. For all lotteries P, P^* and P^{**}, if $P \succeq P^*$ and $P^* \succeq P^{**}$ then $P \succeq P^{**}$.
- *Mixture Continuity*. For all lotteries P, P^* and P^{**}, if $P \succ P^*$ and $P^* \succ P^{**}$ then $P^* \sim \alpha\cdot P + (1-\alpha)\cdot P^{**}$ for some $\alpha \in (0, 1)$.
- *Independence Axiom*. For all lotteries P, P^* and P^{**} and all $\alpha \in (0, 1)$, if $P \succeq P^*$ then $\alpha\cdot P + (1-\alpha)\cdot P^{**} \succeq \alpha\cdot P^* + (1-\alpha)\cdot P^{**}$.

Completeness and Transitivity are standard axioms in preference theory, and Mixture Continuity serves as the standard Archimedean property in the context of choice over lotteries. The key normative and behavioral axiom of the theory is the Independence axiom. Behaviorally, it corresponds to the property of separability across mutually exclusive events. Normatively, it corresponds to the following argument: 'Say you weakly prefer P to P^*, and have to choose between an α:$(1-\alpha)$ coin flip yielding P if heads and P^{**} if tails, or an α:$(1-\alpha)$ coin flip yielding P^* if heads and P^{**} if tails. Now, either the coin will land tails, in which case your choice won't have mattered, or it will land heads, in which case you are back to a choice between P and P^*, so you should weakly prefer the first coin flip to the second.'

The tension between the compelling nature of the Independence axiom and its systematic violations by experimental subjects has led to a sustained debate over the validity of the expected utility model, with some researchers continuing to posit expected utility maximization, and others developing and testing alternative models of risk preferences.

Analytics of Expected Utility Theory

Analytically, the expected utility hypothesis is characterized by the simplicity of its representation (involving the standard concepts of utility and mathematical expectation) as well as by the elegance of the correspondence between standard features of risk preferences and mathematical properties of $U(\cdot)$. The most basic of these properties is 'first-order stochastic dominance preference', which states that raising the level of some pay-off x_i in a lottery $P = (x_1, p_1; \ldots; x_n, p_n)$ – or alternatively, increasing its probability p_i at the expense of a reduction in the probability p_j of some smaller pay-off x_j – will lead to a preferred lottery. An expected utility maximizer's preferences will exhibit first-order stochastic dominance preference if and only if $U(\cdot)$ is an increasing function of x.

A second property is 'risk aversion'. Originally, this was defined as the property whereby the individual would always prefer receiving the expected value of a given lottery with certainty, rather than bearing the risk of the lottery itself. This is equivalent to the condition that the individual's 'certainty equivalent' CE(P) of a nondegenerate lottery $P = (x_1, p_1; \ldots; x_n, p_n)$ – that is, the value that satisfies $U(\mathrm{CE}(P)) = \sum_{i=1}^{n} U(x_i)\cdot p_i$ – is always less than the expected value of P. In modern treatments, risk aversion is defined as an aversion to all 'mean-preserving spreads' from any (degenerate or nondegenerate) lottery, where a mean-preserving spread consists of a decrease in the probability of a pay-off x_i by some amount Δp, and increases in the probabilities of some higher and lower pay-offs $x_i + \alpha$ and $x_i - \beta$ by the respective amounts $\Delta p\cdot\beta/(\alpha+\beta)$ and $\Delta p\cdot\alpha/(\alpha+\beta)$. This 'spreads' the probability mass of the lottery in a manner that does not change its expected value, so it can be thought of as a 'pure increase in risk'. An expected utility maximizer will be risk-averse in both the original and the modern senses if and only if $U(\cdot)$ is a strictly concave function of x. If $U(\cdot)$ is twice continuously differentiable, strict concavity is equivalent to a negative second derivative $U''(\cdot)$. Although the widespread purchase of actuarially unfair state lottery tickets is evidence of the opposite property of 'risk preference', the even more widespread purchase of insurance and the prevalence of other risk-reducing instruments has led researchers to hypothesize that individuals are for the most part risk-averse.

After these basic characterizations, the most important analytical result in expected utility theory is the Arrow–Pratt characterization of 'comparative risk aversion', which states that the following four conditions on a pair of risk-averse von Neumann–Morgenstern utility functions $U_A(\cdot)$ and $U_B(\cdot)$ are equivalent:

- *Comparative Concavity.* $U_A(\cdot)$ is an increasing concave transformation of $U_B(\cdot)$, that is, $U_A(x) \equiv \rho(U_B(x))$ for some increasing concave function $\rho(\cdot)$.
- *Comparative Arrow–Pratt Measures.* $-U''_A(x)/U'_A(x) \geq -U''_B(x)/U'_B(x)$ for all x.
- *Comparative Certainty Equivalents.* For any lottery $P = (x_1, p_1; \ldots; x_n, p_n)$, if $\mathrm{CE}_A(P)$ and $\mathrm{CE}_B(P)$ satisfy $U_A(\mathrm{CE}_A(P)) = \sum_{i=1}^{n} U_A(x_i)\cdot p_i$ and $U_B(\mathrm{CE}_B(P)) = \sum_{i=1}^{n} U_B(x_i)\cdot p_i$, then $\mathrm{CE}_A(P) \leq \mathrm{CE}_B(P)$.
- *Comparative Demand for Risky Assets.* For any initial wealth W, constant $r > 0$, and random variable $\tilde{x}$ such that $E[\tilde{x}] > r$ but $P(\tilde{x} < r) > 0$, if γ_A^* and γ_B^* respectively maximize $E[U_A(\gamma\cdot\tilde{x} + (W - \gamma)\cdot r)]$ and $E[U_B(\gamma\cdot\tilde{x} + (W - \gamma)\cdot r)]$, then $\gamma_A^* \leq \gamma_B^*$.

(Note: here and elsewhere we write $P(\cdot)$ for the probability of an event. This should not be confused with the use of P to stand for a lottery.)

Each of these conditions can be interpreted as saying that $U_A(\cdot)$ is at least as risk-averse as $U_B(\cdot)$. The first condition extends the above characterization of risk aversion by the concavity of $U(\cdot)$ to its comparative version across individuals, and the second shows that this can be expressed in terms of a numerical index $-U''(x)/U'(x)$, known as the 'Arrow–Pratt index of absolute risk aversion'. The third condition extends the original notion of risk aversion as low certainty equivalents (lower than the mean) to its comparative form.

The fourth condition involves comparative optimization behavior. Consider an individual with initial wealth W to be divided between a riskless asset yielding gross return r, and a risky asset whose gross return $\tilde{x}$ has a higher expected value, but offers some risk of doing worse than the riskless asset. This condition states that the less risk-averse utility function $U_B(\cdot)$ will always choose to invest at least as much in the risky asset as will the more risk-averse $U_A(\cdot)$.

The equivalence of the above four conditions, the first two mathematical and the second two behavioral, and their numerous additional behavioral equivalencies and implications, has made the Arrow–Pratt characterization one of the central theorems in the analytics of expected utility theory, with applications in insurance, financial markets, auctions, the demand for information, bargaining, and game theory.

Experimental Evidence on the Independence Axiom

Experimental testing of the expected utility hypothesis has centered on the Independence axiom, either directly or via its implication that the expected utility preference function $V_{\mathrm{EU}}(P) = \sum_{i=1}^{n} U(x_i)\cdot p_i$ is linear in the probabilities p_i. One of the best-known tests is the 'Allais paradox' (Allais, 1953). An individual is asked to rank each of the following two pairs of lotteries (where $1M = $1,000,000):

- $a_1 = \{1.00$ chance of \$1M versus

$$a_2 = \begin{cases} .10 \text{ chance of } \$5\text{M} \\ .89 \text{ chance of } \$1\text{M} \\ .01 \text{ chance of } \$0 \end{cases}$$

- $a_3 = \begin{cases} .10 \text{ chance of } \$5\text{M} \\ .90 \text{ chance of } \$0 \end{cases}$ versus

$$a_4 = \begin{cases} .11 \text{ chance of } \$1\text{M} \\ .89 \text{ chance of } \$0 \end{cases}$$

The expected utility hypothesis implies that the individual's choices from these two pairs must either be a_1 and a_4 (whenever $.11\cdot U(\$1\text{M}) > .10\cdot U(\$5\text{M}) + .01\cdot U(\$0)$), or else a_2 and a_3 (whenever $.11\cdot U(\$1\text{M}) < .10\cdot U(\$5\text{M}) + .01\cdot U(\$0)$). However, when presented with these choices, most subjects choose a_1 from the first pair and a_3 from the second, which violates the hypothesis. Only a small number violate the hypothesis in the opposite direction, by choosing a_2 and a_4.

Although the Allais paradox was originally dismissed as an 'isolated example', subsequent experimental work by psychologists, economists and others has uncovered a similar pattern of violation over a range of probability and pay-off values, and the Allais paradox is now seen to be just one example of a type of systematic violation of the Independence axiom known as the 'common-consequence effect'. It is observed that for lotteries P, P^* and P^{**}, pay-off c, and mixture probability $\alpha \in (0,1)$, such that P^{**} first-order-stochastically dominates P^* and c lies between the highest and lowest pay-offs in P, preferences depart from the Independence axiom in the direction of exhibiting $\alpha\cdot P + (1-\alpha)\cdot P^* \succ \alpha\cdot c + (1-\alpha)\cdot P^*$ yet $\alpha\cdot P + (1-\alpha)\cdot P^{**} \prec \alpha\cdot c + (1-\alpha)\cdot P^{**}$. (In the Allais paradox, these constructs are $P = (\$5\text{M}, 10/11; \$0, 1/11)$, $P^* = (\$0, 1), P^{**} = (\$1\text{M}, 1), c = \$1\text{M}$ and $\alpha = .11$.)

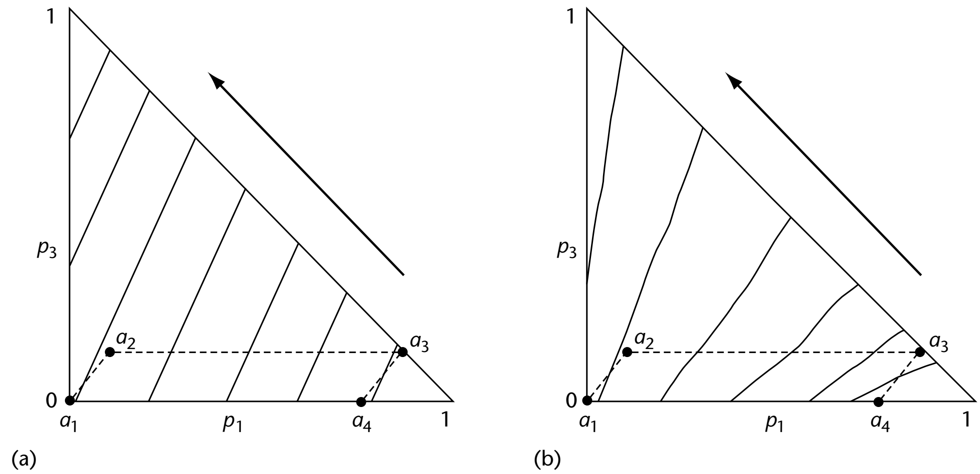

Figure 1. Indifference curves in the probability triangle. (a) Expected utility indifference curves, which are parallel straight lines. (b) Non-expected utility indifference curves, which 'fan out', illustrating the common-consequence effect.

Both the implications of the Independence axiom and the nature of this violation can be illustrated in the special case of all lotteries $P = (\bar{x}_1, p_1; \bar{x}_2, p_2; \bar{x}_3, p_3)$ over a triple of fixed pay-off values $\bar{x}_1 < \bar{x}_2 < \bar{x}_3$. Since we can write $P = (\bar{x}_1, p_1; \bar{x}_2, 1 - p_1 - p_3; \bar{x}_3, p_3)$, each such lottery is uniquely associated with a point in the (p_1, p_3) triangles of Figures 1(a) and 1(b). Since we can write $V_{\mathrm{EU}}(P) = U(\bar{x}_1)\cdot p_1 + U(\bar{x}_2)\cdot(1 - p_1 - p_3) + U(\bar{x}_3)\cdot p_3$, the loci of constant expected utility ('expected utility indifference curves') consist of parallel straight lines as in Figure 1(a). Since upward shifts in the triangle represent increases in p_3 at the expense of p_2, and leftward shifts represent reductions in p_1 to the benefit of p_2, first-order stochastic dominance preference implies that these indifference curves will be upward-sloping, with increasing levels of preference in the direction indicated by the arrows.

Fixing the pay-offs at $\bar{x}_1 = \$0$, $\bar{x}_2 = \$1\mathrm{M}$ and $\bar{x}_3 = \$5\mathrm{M}$, the Allais paradox lotteries a_1, a_2, a_3 and a_4 are seen to form a parallelogram when plotted in the probability triangle, which explains why parallel straight-line expected utility indifference curves must either prefer a_1 and a_4 (as illustrated for the relatively steep indifference curves of Figure 1(a)) or else prefer a_2 and a_3 (for relatively flat expected utility indifference curves). Figure 1(b) illustrates 'non-expected utility indifference curves' that 'fan out', and are seen to exhibit the typical Allais paradox rankings of $a_1 \succ a_2$ and $a_3 \succ a_4$.

Another type of systematic experimental violation of the Independence axiom that has been uncovered is known as the 'common-ratio effect'. For pay-offs $x^* > x > 0$, probabilities $p^* < p$ and $r \in (0, 1)$, preferences depart from the Independence axiom in the direction of exhibiting $(x^*, p^*; 0, 1 - p^*) \prec (x, p; 0, 1 - p)$ yet $(x^*, r\cdot p^*; 0, 1 - r\cdot p^*) \succ (x, r\cdot p; 0, 1 - r\cdot p)$. For losses $0 > -x > -x^*$, with $p^* < p$ and $r \in (0, 1)$, preferences depart in the reflected direction of $(-x^*, p^*; 0, 1 - p^*) \succ (-x, p; 0, 1 - p)$ yet $(-x^*, r\cdot p^*; 0, 1 - r\cdot p^*) \prec (-x, r\cdot p; 0, 1 - r\cdot p)$.

With the pay-offs $\bar{x}_1 = 0$, $\bar{x}_2 = x$ and $\bar{x}_3 = x^*$, the line segment between the lotteries $b_1 = (x^*, p^*; 0, 1 - p^*)$ and $b_2 = (x, p; 0, 1 - p)$ in the probability triangle of Figure 2(a) is seen to be parallel to that between $b_3 = (x^*, r\cdot p^*; 0, 1 - r\cdot p^*)$ and $b_4 = (x, r\cdot p; 0, 1 - r\cdot p)$, and the common-ratio-effect rankings of $b_1 \prec b_2$ and $b_3 \succ b_4$ again suggests that indifference curves depart from expected utility by fanning out. For losses, with $\bar{x}_1 = -x^*$, $\bar{x}_2 = -x$ and $\bar{x}_3 = 0$ (to maintain the ordering $\bar{x}_1 < \bar{x}_2 < \bar{x}_3$), the reflected rankings of $-b_1 \succ -b_2$ and $-b_3 \prec -b_4$ again suggest fanning out, as in Figure 2(b). Fanning out is consistent with other observed forms of departure from the Independence axiom, although it is not universal across subjects, and seems to be more pronounced near the edges of the triangle than in its central region.

GENERALIZATIONS OF EXPECTED UTILITY THEORY

Non-Expected Utility Functional Forms

The above phenomena, as well as other systematic departures from linearity in the probabilities, have prompted researchers to develop more general models of preferences over lotteries, primarily by generalizing the functional form of the lottery preference function $V(P) = V(x_1, p_1; \ldots; x_n, p_n)$. The earliest of these attempts, which used the form $V(P) = \sum_{i=1}^{n} U(x_i)\cdot\pi(p_i)$, was largely abandoned

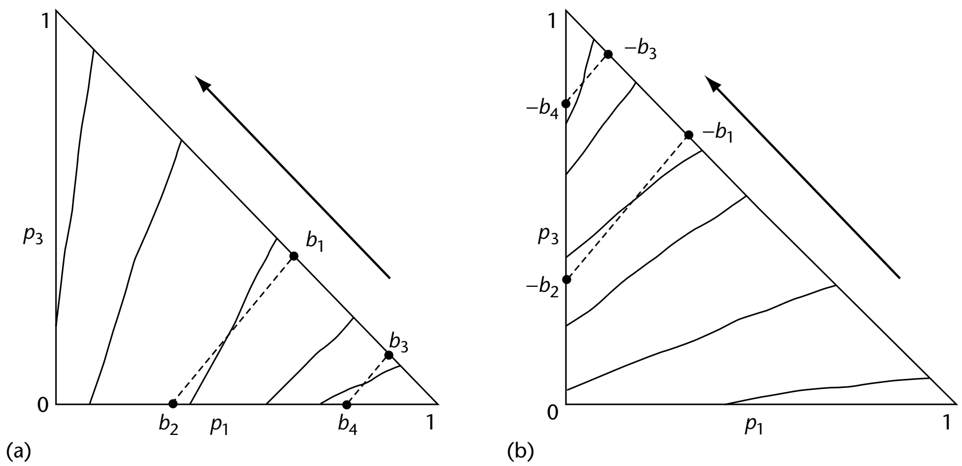

Figure 2. Probability triangles illustrating the common-ratio effect. (a) Positive pay-offs. (b) Negative pay-offs (losses).

when it was realized that, except for the case $\pi(p) \equiv p$ when it reduced to expected utility, it was inconsistent with the property of first-order stochastic dominance preference. Current models include the following:

- *Weighted Utility.* $V(P) = \sum_{i=1}^{n} U(x_i)\cdot\pi(p_i)/\sum_{i=1}^{n} S(x_i)\cdot\pi(p_i)$
- *Moments of Utility.* $V(P) = F(\sum_{i=1}^{n} U(x_i)\cdot p_i, \sum_{i=1}^{n} U(x_i)^2\cdot p_i, \sum_{i=1}^{n} U(x_i)^3\cdot p_i)$
- *Rank-Dependent Expected Utility.* $V(P) = \sum_{i=1}^{n} U(x_i)\cdot(G(\sum_{j=1}^{i} p_j) - G(\sum_{j=1}^{i-1} p_j))$for $x_1 \preceq \ldots \preceq x_n$
- *Quadratic in the Probabilities.* $V(P) = \sum_{i=1}^{n}\sum_{j=1}^{n} T(x_i, x_j)\cdot p_i\cdot p_j$

Under the appropriate monotonicity or curvature assumptions on their constituent functions $U(\cdot)$, $\pi(\cdot)$, $G(\cdot)$, etc., each of these forms is capable of exhibiting first-order stochastic dominance preference, risk aversion and comparative risk aversion, as well as many of the types of observed systematic violations of the Independence axiom. Researchers have also used these forms to revisit many of the applications previously modeled by expected utility theory (e.g. insurance, financial markets, auctions), to determine which of the earlier expected utility-based results are crucially dependent on preferences exhibiting the expected utility functional form, and which are robust to departures from expected utility.

Generalized Expected Utility Analysis

An alternative branch of research on non-expected utility preferences does not rely on any specific functional form, but links properties of attitudes towards risk directly with the probability derivatives of a general (i.e. not necessarily expected utility) preference function $V(P) = V(x_1, p_1; \ldots; x_n, p_n)$ over lotteries. Such analysis reveals that the basic analytics of the expected utility model as outlined above are in fact quite robust to general smooth departures from linearity in the probabilities. It proceeds from the correspondence between the properties of a linear function as determined by its coefficients and the properties of a nonlinear function as determined by its partial derivatives – in this case, between the 'probability coefficients' $U(x_1), \ldots, U(x_n)$ of the expected utility form $\sum_{i=1}^{n} U(x_i)\cdot p_i$ and the 'probability derivatives' $\partial V(x_1, p_1; \ldots; x_n, p_n)/\partial p_1, \ldots, \partial V(x_1, p_1; \ldots; x_n, p_n)/\partial p_n$ of a general smooth preference function $V(x_1, p_1; \ldots; x_n, p_n)$. Under such a correspondence, most of the fundamental analytical results of expected utility theory pass through directly (Machina, 1982). For example:

- *First-Order Stochastic Dominance Preference.* Under expected utility, this is equivalent to $U(x)$ (the coefficient of $P(x)$) being an increasing function of x. For a general smooth $V(\cdot)$, if $\partial V(P)/\partial P(x)$ is an increasing function of x at every lottery P, then for any pay-offs $x_i > x_j$ we will have $\partial V(P)/\partial p_i > \partial V(P)/\partial p_j$ at each P, so any (small or large) rise in p_i and matching fall in p_j will lead to an increase in $V(P)$ and hence will be preferred.
- *Risk Aversion.* Under expected utility, this is equivalent to $U(x)$ being a strictly concave function of x. For a general smooth $V(\cdot)$, if $\partial V(P)/\partial P(x)$ is a strictly concave function of x at each P, then for any pay-offs $x_i - \beta < x_i < x_i + \alpha$ we will have $[\partial V(P)/\partial P(x_i) - \partial V(P)/\partial P(x_i - \beta)]/\beta > [\partial V(P)/\partial P(x_i + \alpha) - \partial V(P)/\partial P(x_i)]/\alpha$ at each P, which implies that each mean-preserving spread over the pay-offs $x_i - \beta < x_i < x_i + \alpha$ will lead to a reduction in $V(P)$ and hence will be dispreferred.
- *Comparative Risk Aversion.* Under expected utility, this is equivalent to $U_A(\cdot)$ being an increasing concave transformation of $U_B(\cdot)$. For general smooth $V_A(\cdot)$ and

$V_B(\cdot)$, if at each P the function $\partial V_A(P)/\partial P(x)$ is some increasing concave transformation of $\partial V_B(P)/\partial P(x)$, then $V_A(\cdot)$ and $V_B(\cdot)$ will exhibit the above conditions for comparative certainty equivalence and comparative demand for risky assets.

In addition to the above theoretical results, this approach also allows for a direct characterization of the fanning-out property in terms of how the probability derivative $\partial V(P)/\partial P(x)$, treated as a function of x, varies with the lottery P. Namely, the indifference curves of a preference function $V(\cdot)$ will fan out for all pay-offs $\bar{x}_1 < \bar{x}_2 < \bar{x}_3$ if and only if $\partial V(P^*)/\partial P(x)$ is a concave transformation of $\partial V(P)/\partial P(x)$ whenever P^* first-order-stochastically dominates P.

Regret Theory

Another type of non-expected utility model dispenses with the assumption of an underlying preference order $\succeq$ over lotteries, and instead derives choice behavior from the underlying psychological notion of 'regret' – that is, the reaction to receiving an outcome x when an alternative decision would have led to a preferred outcome x^* (Loomes and Sugden, 1982). The opposite experience, namely of receiving an outcome that is preferred to what the alternative decision would have yielded, is termed 'rejoice'. The primitive for this model is a 'rejoice function' $R(x, x^*)$ which is positive if x is preferred to x^*, negative if x^* is preferred to x, and zero if they are indifferent, and satisfies the skew-symmetry condition $R(x, x^*) \equiv -R(x^*, x)$.

In the simplest case of pairwise choice over two lotteries $P = (x_1, p_1; \ldots; x_n, p_n)$ and $P^* = (x_1^*, p_1^*; \ldots; x_{n^*}^*, p_{n^*}^*)$ that are realized independently, the individual's expected rejoice from choosing the lottery P over the alternative lottery P^* is given by the formula $\sum_{i=1}^{n}\sum_{j=1}^{n^*} R(x_i, x_j^*)\cdot p_i \cdot p_j^*$, and the individual is predicted to choose P if this value is positive, to choose P^* if it is negative, and to be indifferent if it is zero. Various proposals for extending this approach beyond pairwise choice have been made, including a formal result that shows that for any finite collection of lotteries, one of these lotteries or some randomization over them will exhibit nonnegative expected rejoice with respect to every other lottery or randomization.

As with the non-expected utility functional forms listed above, various monotonicity and curvature assumptions on the rejoice function $R(\cdot,\cdot)$ can be shown to correspond to various properties of risk preferences, such as risk aversion and comparative risk aversion, as well as to the general fanning-out property. Since this model derives from the pairwise comparison of lotteries rather than from their individual evaluation by some preference function, it allows pairwise choice to be intransitive, so that an individual could choose P over P^*, P^* over P^{**}, and P^{**} over P. Although some have argued that such cyclic choice allows for the phenomenon of 'money pumps', it also allows the model to solve the problem of 'preference reversal' described below.

SUBJECTIVE EXPECTED UTILITY AND AMBIGUITY

Axiomatic and Normative Foundations of Subjective Expected Utility

The expected utility model of choice under subjective uncertainty – sometimes called the 'subjective expected utility' model – hypothesizes that the individual's preference relation $\succeq$ over subjective acts $f(\cdot) = [x_1 \text{ on } E_1; \ldots; x_n \text{ on } E_n]$ can be represented by a preference function of the form $W_{SEU}(f(\cdot)) = \sum_{i=1}^{n} U(x_i)\cdot\mu(E_i)$ for some von Neumann–Morgenstern utility function $U(\cdot)$ and 'subjective probability measure' $\mu(\cdot)$ over events. Thus, both attitudes towards risk and subjective beliefs are specific to the individual, and the values $\mu(E_1), \ldots, \mu(E_n)$ are sometimes called 'personal probabilities'. By virtue of its independent representation of risk attitudes by the utility function $U(\cdot)$, and beliefs by the subjective probability measure $\mu(\cdot)$, the subjective expected utility model is sometimes described as achieving a 'separation of preferences and beliefs'.

By analogy with the probability mixture of two objective lotteries, the axiomatic and normative underpinnings of the subjective expected utility model are based on the notion of a 'subjective mixture' $[f(\cdot) \text{ on } E; f^*(\cdot) \text{ on } \sim E]$ of two acts $f(\cdot) = [x_i \text{ on } E_1; \ldots; x_n \text{ on } E_n]$ and $f^*(\cdot) = [x_1^* \text{ on } E_1^*; \ldots; x_{n^*}^* \text{ on } E_{n^*}^*]$, which is the act that would yield the same outcome as $f(\cdot)$ should the event E occur, and the same outcome as $f^*(\cdot)$ should the event $\sim E$ occur, so that we can write $[f(\cdot) \text{ on } E; f^*(\cdot) \text{ on } \sim E] = [x_1 \text{ on } E\cap E_1; \ldots; x_n \text{ on } E\cap E_n; x_1^* \text{ on } \sim E\cap E_1^*; \ldots; x_{n^*}^* \text{ on } \sim E\cap E_{n^*}^*]$. An event E is said to be 'null' for the individual if $[x^* \text{ on } E; f(\cdot) \text{ on } \sim E] \sim [x \text{ on } E; f(\cdot) \text{ on } \sim E]$ for all outcomes x and x^* and all acts $f(\cdot)$, so that the individual effectively treats E as if it had zero likelihood. Since we can identify each outcome x with the 'constant act' $[x \text{ on S}]$, we can write $x^* \succeq x$ if and only if $[x^* \text{ on S}] \succeq [x \text{ on S}]$. The individual's preferences over subjective acts can then be represented by a subjective expected utility preference function

$W_{\mathrm{SEU}}(f(\cdot)) = \sum_{i=1}^{n} U(x_i)\cdot\mu(E_i)$ for some $U(\cdot)$ and $\mu(\cdot)$ if and only they satisfy the following axioms (Savage, 1954):

- *Completeness*. For all acts $f(\cdot)$ and $f^*(\cdot)$, either $f(\cdot)\succeq f^*(\cdot)$ or $f^*(\cdot)\succeq f(\cdot)$, or both.
- *Transitivity*. For all acts $f(\cdot)$, $f^*(\cdot)$ and $f^{**}(\cdot)$, if $f(\cdot)\succeq f^*(\cdot)$ and $f^*(\cdot)\succeq f^{**}(\cdot)$ then $f(\cdot)\succeq f^{**}(\cdot)$.
- *Eventwise Monotonicity*. For all outcomes x^* and x, non-null events E and acts $f(\cdot)$, $[x^*$ on E; $f(\cdot)$ on $\sim E] \succeq [x$ on E; $f(\cdot)$ on $\sim E]$ if and only if $x^*\succeq x$.
- *Weak Comparative Probability*. For all events A and B and outcomes $x^*\succ x$ and $y^*\succ y$, $[x^*$ on A; x on $\sim A] \succeq [x^*$ on B; x on $\sim B]$ implies $[y^*$ on A; y on $\sim A] \succeq [y^*$ on B; y on $\sim B]$.
- *Small-Event Continuity*. For all acts $f(\cdot)\succ g(\cdot)$ and outcomes x, there exists a partition $\{E_1,\ldots,E_n\}$ such that $f(\cdot)\succ[x$ on E_i; $g(\cdot)$ on $\sim E_i]$ and $[x$ on E_i; $f(\cdot)$ on $\sim E_i]\succ g(\cdot)$ for each $i=1,\ldots,n$.
- *Sure-Thing Principle*. For all events E and acts $f(\cdot)$, $f^*(\cdot)$, $g(\cdot)$ and $h(\cdot)$, $[f^*(\cdot)$ on E; $g(\cdot)$ on $\sim E] \succeq [f(\cdot)$ on E; $g(\cdot)$ on $\sim E]$ implies $[f^*(\cdot)$ on E; $h(\cdot)$ on $\sim E] \succeq [f(\cdot)$ on E; $h(\cdot)$ on $\sim E]$.

Completeness and Transitivity are as before, and Eventwise Monotonicity is the subjective analogue of first-order stochastic dominance preference. Weak Comparative Probability essentially ensures that the individual's 'revealed likelihood ranking' of a pair of events A and B, as given by their preference for staking the more preferred of a pair of prizes on A versus staking it on B, is stable in the sense that it does not depend on the particular prizes involved. Small-Event Continuity serves as the standard Archimedean property in the context of choice over subjective acts. The key normative and behavioral axiom of subjective expected utility theory is the Sure-Thing Principle. Behaviorally, it once again corresponds to the property of separability across mutually exclusive events. Normatively, it corresponds to the same argument as for the Independence axiom, with the objective randomization by the $\alpha{:}(1-\alpha)$ coin replaced by the 'subjective randomization' via the events E and $\sim E$.

State-dependent Utility

In some subjective settings, the individual's valuation of outcomes may depend on the source of uncertainty itself. Thus, for the mutually exclusive and exhaustive events ('rain','shine') and prizes ('umbrella', 'sun lotion'), each of which is preferred to \$0, the individual may well exhibit the preferences [umbrella on rain; \$0 on shine] $\succ$ [umbrella on shine; \$0 on rain] and [sun lotion on rain; \$0 on shine] $\prec$ [sun lotion on shine; \$0 on rain], which violates the Weak Comparative Probability axiom for $x^*=$ umbrella, $y^*=$ lotion, $x=y=\$0$, $A=$ rain and $B=$ shine, and hence is inconsistent with the subjective expected utility preference function $W_{\mathrm{SEU}}(\cdot)$. This phenomenon, known as 'state dependence', can occur even when the outcomes are monetary pay-offs: if the state of nature is the individual's health, the utility of a \$50,000 prize may be very high in states where the individual requires a \$50,000 operation to survive, much lower in states where the individual requires much more than that for the operation, and somewhere in between in states of good health.

The subjective expected utility model can be easily adapted to accommodate the phenomenon of state dependence, by allowing the utility function $U(\cdot|E)$ to depend upon the event or state of nature, so that the preference function over acts takes the 'state-dependent expected utility' form $W_{\mathrm{SDEU}}(x_1 \text{ on } E_1; \ldots x_n \text{ on } E_n) = \sum_{i=1}^{n} U(x_i|E_i)\cdot\mu(E_i)$. Most of the analytics of the standard (i.e. 'state-independent') form $W_{\mathrm{SEU}}(\cdot)$ extend to the state-dependent case (Karni, 1985). However, under state dependence, subjective probabilities cannot be uniquely inferred from preferences over acts: for any state-dependent preference function $W_{\mathrm{SDEU}}(f(\cdot)) = \sum_{i=1}^{n} U(x_i|E_i)\cdot\mu(E_i)$, and any distinct subjective probability measure $\mu^*(\cdot)$ that satisfies $\mu^*(E) > 0 \Leftrightarrow \mu(E) > 0$, $W_{\mathrm{SDEU}}(\cdot)$ is indistinguishable from the preference function $W^*_{\mathrm{SDEU}}(f(\cdot)) = \sum_{i=1}^{n} U^*(x_i|E_i)\cdot\mu^*(E_i)$ with $U^*(\cdot|\cdot)$ defined by $U^*(x|E) = U(x|E)\cdot[\mu(E)/\mu^*(E)]$.

Ambiguity and Nonprobabilistic Beliefs

A more serious departure from the notion of well-defined probabilistic beliefs arises from the phenomenon of 'ambiguity', which is distinct from the phenomenon of state dependence and much more difficult to model. The best-known example is the 'Ellsberg paradox' (Ellsberg, 1961). An individual must draw a ball from an opaque urn that contains 30 red balls and 60 black or yellow balls in an unknown proportion, and is offered four possible bets on the color of the drawn ball, as shown in Figure 3.

Most individuals exhibit the preference rankings $f_1(\cdot)\succ f_2(\cdot)$ and $f_4(\cdot)\succ f_3(\cdot)$. When asked why, they explain that the probability of winning under $f_2(\cdot)$ could be anywhere from 0 to $\frac{2}{3}$ whereas the probability of winning under $f_1(\cdot)$ is known to be exactly $\frac{1}{3}$, and they prefer the act that offers the known probability. Similarly, the probability of winning under $f_3(\cdot)$ could be anywhere from $\frac{1}{3}$ to 1 whereas the probability of winning under $f_4(\cdot)$ is known to be exactly $\frac{2}{3}$, so it is preferred. However, these preferences are inconsistent with any assignment of

	30 balls	60 balls	
	Red	Black	Yellow
$f_1(\cdot)$	\$100	\$0	\$0
$f_2(\cdot)$	\$0	\$100	\$0
$f_3(\cdot)$	\$100	\$0	\$100
$f_4(\cdot)$	\$0	\$100	\$100

Figure 3. Four possible bets on the color of the drawn ball in the 'Ellsberg paradox'. The proportion of black to yellow balls is unknown. Most individuals prefer $f_1(\cdot)$ to $f_2(\cdot)$ and $f_4(\cdot)$ to $f_3(\cdot)$.

numerical subjective probabilities μ(red), μ(black), μ(yellow) to the three events: if the individual were choosing on the basis of such probabilistic beliefs, the ranking $f_1(\cdot) \succ f_2(\cdot)$ would reveal that $\mu(\text{red}) > \mu(\text{black})$, but the ranking $f_4(\cdot) \succ f_3(\cdot)$ would reveal that $\mu(\text{red}) < \mu(\text{black})$.

This phenomenon be cannot be accommodated by simply allowing the event to enter the utility function and working with the state-dependent form $\sum_{i=1}^{n} U(x_i|E_i)\cdot\mu(E_i)$, since this form still satisfies the Sure-Thing Principle, whereas the preferences $f_1(\cdot) \succ f_2(\cdot)$ and $f_4(\cdot) \succ f_3(\cdot)$ violate this axiom (for $E = \text{red} \cup \text{black}$ and $\sim E = \text{yellow}$). The Ellsberg paradox and related examples are attributed to the phenomenon of 'ambiguity aversion', whereby individuals exhibit a general preference for bets based on probabilistic partitions such as {red, black ∪ yellow} rather than on ambiguous partitions such as {black, red ∪ yellow}.

Just as the Allais paradox and related violations of the Independence axiom led to the development of non-expected utility models of risk preferences, the Ellsberg paradox and related examples have led to the development of nonprobabilistic models of beliefs. The most notable of these involves replacing the additive subjective probability measure $\mu(\cdot)$ over events by a 'capacity' $C(\cdot)$, which is similar to $\mu(\cdot)$ in that it satisfies the properties $C(\emptyset) = 0$, $C(\mathrm{S}) = 1$, and $E \subseteq E^* \Rightarrow C(E) \leq C(E^*)$, but differs from $\mu(\cdot)$ in that it is not necessarily additive. By labeling the outcomes in any act so that $x_1 \preceq \ldots \preceq x_n$ and writing the subjective expected utility formula as $W_{\mathrm{SEU}}(f(\cdot)) = \sum_{i=1}^{n} U(x_i)\cdot\mu(E_i) = \sum_{i=1}^{n} U(x_i)\cdot(\mu(\cup_{j=1}^{i}E_j) - \mu(\cup_{j=1}^{i-1}E_j))$, we can generalize from an additive $\mu(\cdot)$ to a non-additive $C(\cdot)$ to obtain the 'Choquet expected utility' preference function $W_{\mathrm{Choquet}}(f(\cdot)) = \sum_{i=1}^{n} U(x_i)\cdot(C(\cup_{j=1}^{i}E_j) - C(\cup_{j=1}^{i-1}E_j))$ over subjective acts (Schmeidler, 1989). Selecting $U(\$100) = 1$, $U(\$0) = 0$, $C(\text{red}) = \frac{1}{3}$, $C(\text{black} \cup \text{yellow}) = \frac{2}{3}$, $C(\text{black}) = \frac{1}{2}$, $C(\text{red} \cup \text{yellow}) = \frac{3}{4}$ yields the values $W_{\mathrm{Choquet}}(f_1(\cdot)) = \frac{1}{3}$, $W_{\mathrm{Choquet}}(f_2(\cdot)) = \frac{1}{4}$, $W_{\mathrm{Choquet}}(f_3(\cdot)) = \frac{1}{2}$, $W_{\mathrm{Choquet}}(f_4(\cdot)) = \frac{2}{3}$, which correspond to the typical Ellsberg rankings.

Another alternative to the subjective expected utility model of act preferences, also capable of exhibiting ambiguity aversion, is the 'maxmin expected utility' form, which involves a family $\{\mu_\tau(\cdot)\,|\,\tau \in T\}$ of additive probability measures over the events, and the preference function $W_{\max\min}(f(\cdot)) = \min_{\tau\in T}\sum_{i=1}^{n} U(x_i)\cdot\mu_\tau(E_i)$.

DESCRIPTION AND PROCEDURE INVARIANCE

Although the alternative models described above drop or weaken many of the axioms of standard objective and subjective expected utility theory, they typically retain the primary implicit assumptions of the standard theory, namely that: the objects of choice (objective lotteries or subjective acts) can be unambiguously described; net changes in wealth are combined with any initial endowment and evaluated in terms of the final wealth levels they imply; and situations that imply the same set of final opportunities (the same set of objective lotteries or same set of subjective acts over final wealth levels) will lead to the same choice. They also assume that the individual is able to perform the mathematical operations necessary to determine this opportunity set, e.g. to calculate the probabilities of compound or conditional events and add net changes to initial endowments. However, psychologists have uncovered several systematic violations of these assumptions.

Framing Effects

Effects whereby alternative descriptions of the same decision problem lead to systematically different responses are called 'framing' effects. Some framing effects in choice under uncertainty involve alternative representations of the same likelihood. For example, contingency of a gain or loss on the joint occurrence of four independent events, each with probability p, is found to elicit a different response from contingency on the occurrence of a single event with probability p^4. In comparison with the single-event case, making a gain contingent on the joint occurrence of events seems to make it more attractive, and making a loss contingent on the joint occurrence of events seems to make it more unattractive (Slovic, 1969).

Other framing effects in choice under uncertainty involve alternative representations of the same final wealth levels. Consider the following two proposals.

- 'In addition to whatever you own, you have been given 1,000 (Israeli pounds). You are now asked to choose between a ½:½ chance of a gain of 1,000 or 0 or a sure chance of a gain of 500.'
- 'In addition to whatever you own, you have been given 2,000. You are now asked to choose between a ½:½ chance of a loss of 1,000 or 0 or a sure loss of 500.'

These two problems involve identical distributions over final wealth. However, when put to two different groups of subjects, 84% chose the sure gain in the first problem but 69% chose the ½:½ gamble in the second (Kahneman and Tversky, 1979).

Response-Mode Effects and Preference Reversal

Effects whereby alternative response formats lead to systematically different inferences about underlying preferences are called 'response-mode' effects. For example, under the expected utility hypothesis, an individual's von Neumann–Morgenstern utility function can be assessed or elicited in a number of different ways, which typically involve a sequence of prespecified lotteries P_1, P_2, P_3,…, and ask for the individual's certainty equivalent $\mathrm{CE}(P_i)$ for each lottery P_i, or else the 'gain equivalent' G_i that would make the lottery $(G_i, \frac{1}{2};\ \$0, \frac{1}{2})$ of equal preference to P_i, or else the 'probability equivalent' $\wp_i$ that would make the lottery $(\$1000, \wp_i;\ \$0, 1-\wp_i)$ of equal preference to P_i. Although such procedures should be expected to generate equivalent assessed utility functions, they have been found to yield systematically different ones (Hershey and Schoemaker, 1985).

In an experiment that demonstrates what is now known as the 'preference reversal phenomenon', subjects were first presented with a number of pairs of bets and asked to choose one bet out of each pair. Each pair of bets took the form of a 'p-bet', which offered a p chance of $\$X^*$ and a $1-p$ chance of \$X, versus a '\$-bet', which offered a q chance of $\$Y^*$ and a $1-q$ chance of \$Y, where $X^* > X, Y^* > Y, p > q$ and $X^* < Y^*$. The names 'p-bet' and '\$-bet' derive from the greater probability of winning in the first bet, and greater possible gain in the second bet (in some cases, X and Y took on small negative values). Subjects were next asked for their certainty equivalents of each of these bets, via a number of standard elicitation techniques.

The expected utility model, and most of the aforementioned alternative models, predict that for each such pair, the bet that was selected in the direct-choice problem would also be the one assigned the higher certainty equivalent. However, subjects exhibit a systematic departure from this prediction in the direction of choosing the p-bet in a direct choice but assigning a higher certainty equivalent to the \$-bet (Lichtenstein and Slovic, 1971). Although this finding initially generated widespread scepticism, especially among economists, it has been widely replicated by both psychologists and economists in a variety of different settings involving real-money gambles, patrons of a Las Vegas casino, group decisions, and experimental market trading. By viewing it as an instance of intransitivity (\$-bet $\sim$ CE(\$-bet) $\succ$ CE(p-bet) $\sim$ p-bet $\succ$ \$-bet), some economists have explained the phenomenon in terms of the regret theory model. However, most psychologists and a growing number of economists regard it as a response-mode effect, specifically, that the psychological processes of valuation (which generates certainty equivalents) and choice are differentially influenced by the probabilities and pay-offs involved in a lottery, and that under certain conditions this can lead to choice and valuation that reveal opposite 'underlying' preference rankings over a pair of gambles.

SUMMARY

Since the work of Bernoulli, the theory of choice under uncertainty has seen both a tension and a scientific interplay between theoretical models of decision making and experimentally observed violations of these models. Current research in the field continues to reflect this tension, while the degree of interplay has increased, with theorists now more willing to model experimentally generated phenomena, and experimenters more willing to provide constructive feedback on these attempts.

References

Allais M (1953) Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica* **21**: 503–546.

Arrow K (1965) *Aspects of the Theory of Risk Bearing*. Helsinki: Yrjö Jahnsson Säätiö.

Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* **75**: 643–669.

Hershey J and Schoemaker P (1985) Probability versus certainty equivalence methods in utility measurement: are they equivalent? *Management Science* **31**: 1213–1231.

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.
Lichtenstein S and Slovic P (1971) Reversals of preferences between bids and choices in gambling decisions. *Journal of Experimental Psychology* **89**: 46–55.
Loomes G and Sugden R (1982) Regret theory: an alternative theory of rational choice under uncertainty. *Economic Journal* **92**: 805–824.
Machina M (1982) 'Expected utility' analysis without the Independence Axiom. *Econometrica* **50**: 277–323.
Marschak J (1950) Rational behavior, uncertain prospects, and measurable utility. *Econometrica* **18**: 111–141. [Errata: *Econometrica* **18**: 312.]
von Neumann J and Morgenstern O (1944) *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press. [Second edition 1947; third edition 1953.]
Pratt J (1964) Risk aversion in the small and in the large. *Econometrica* **32**: 122–136.
Savage L (1954) *The Foundations of Statistics*. New York, NY: Wiley. [Revised and enlarged edition, 1972. New York, NY: Dover.]
Schmeidler D (1989) Subjective probability and expected utility without additivity. *Econometrica* **57**: 571–587.
Slovic P (1969) Manipulating the attractiveness of a gamble without changing its expected value. *Journal of Experimental Psychology* **79**: 139–145.

Further Reading

Camerer C and Weber M (1992) Recent developments in modeling preferences: uncertainty and ambiguity. *Journal of Risk and Uncertainty* **5**: 325–370.
Einhorn H and Hogarth R (1985) Ambiguity and uncertainty in probabilistic inference. *Psychological Review* **92**: 433–461.
Epstein L (1999) A definition of uncertainty aversion. *Review of Economic Studies* **66**: 579–608.
Fishburn P (1982) *The Foundations of Expected Utility*. Dordrecht: Reidel.
Heath C and Tversky A (1991) Preferences and belief: ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty* **4**: 5–28.
Kahneman D, Slovic P and Tversky A (eds) (1982) *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge, UK: Cambridge University Press.
Kelsey D and Quiggin J (1992) Theories of choice under ignorance and uncertainty. *Journal of Economic Surveys* **6**: 133–153.
Starmer C (2000) Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* **38**: 332–382.
Tversky A and Fox C (1995) Weighing risk and uncertainty. *Psychological Review* **102**: 269–283.
Weber M and Camerer C (1987) Recent developments in modeling preferences under risk. *OR Spektrum* **9**: 129–151.

Circadian Rhythms

Introductory article

Ralph E Mistlberger, Simon Fraser University, Burnaby, British Columbia, Canada

CONTENTS

Circadian rhythms are daily (about 24 h) rhythms of behavior, physiology and biochemistry that are controlled by internal clocks. These rhythms are entrained by environmental cues, and modulate cognitive performance.

INTRODUCTION

The rotation of the earth about its axis creates daily cycles of light, temperature, humidity and other geophysical variables that have had a profound impact on the evolution of life. Most living organisms, from single-celled bacteria to fungi, plants and animals, exhibit daily rhythms in their biochemistry, physiology and behavior that mirror the dramatic environmental changes that define the solar day. Some daily rhythms may represent a direct response to environmental stimuli, but most are controlled by one or more internal, 'circadian' clocks (from the Latin *circa*,