Rational expectations models thus often sacrifice descriptive accuracy in the hope that the models would exhibit stability in the presence of interventions of the kind envisioned by makers of government macroeconomic policy. The models may not be generally well-suited to forecasting when the policy regime is unchanged. They are most appropriately considered policy analysis tools.

ROBERT J. SCHILLER

See also ADAPTIVE EXPECTATIONS; INTERTEMPORAL EQUILIBRIUM AND EFFICIENCY; RATIONAL EXPECTATIONS; UNCERTAINTY; UNCERTAINTY AND GENERAL EQUILIBRIUM.

BIBLIOGRAPHY

Cagan, P. 1956. The monetary dynamics of hyperinflation. In *Studies in the Quantity Theory of Money*, ed. M. Friedman, Chicago: University of Chicago Press.

Fisher, I. 1930. *The Theory of Interest*. New York: Macmillan.

Friedman, B.M. 1979. Optimal expectations and the extreme information assumptions of 'rational' expectations models. *Journal of Monetary Economics* 5, 23–41.

Katona, G. 1975. *Psychological Economics*. New York: Elsevier.

Lucas, R.E., Jr. 1973. Some international evidence on output–inflation tradeoffs. *American Economic Review* 63, 326–34.

Lucas, R.E., Jr. and Sargent, T.J. 1981. After Keynesian macroeconomics. In *Rational Expectations and Econometric Practice*, ed. Lucas and Sargent, Minneapolis: University of Minnesota Press.

Meiselman, D. 1962. *The Term Structure of Interest Rates*. Englewood Cliffs, New Jersey: Prentice-Hall.

Modigliani, F. and Sutch, R. 1966. Innovations in interest rate policy. *American Economic Review, Papers and Proceedings* 56, 178–97.

Muth, J.F. 1961. Rational expectations and the theory of price movements. *Econometrica* 29, 315–35.

Phillips, A.W. 1958. The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica* 25, 283–99.

Sargent, T.J. 1976. A classical macroeconomic model for the United States. *Journal of Political Economy* 84, 207–37.

Sargent, T.J. and Wallace, N. 1973. Rational expectations and the dynamics of hyperinflation. *International Economic Review* 14, 328–50.

Sargent, T.J. and Wallace, N. 1976. Rational expectations and the theory of economic policy. *Journal of Monetary Economics* 2, 169–84.

Sargent, T.J. and Wallace, N. 1981. 'Rational' expectations, the optimal monetary instrument, and the optimal money supply rule. In *Rational Expectations and Econometric Practice*, ed. Robert E. Lucas and Thomas J. Sargent, Minneapolis: University of Minnesota Press.

Shiller, R.J. 1978. Rational expectations and the dynamic structure of rational expectations models: a critical review. *Journal of Monetary Economics* 4, 1–44.

Simon, H.A. 1956. Dynamic programming under uncertainty with a quadratic objective function. *Econometrica* 24, 74–81.

Tobin, J. 1980. *Asset Accumulation and Economic Activity*. Yrjö Jahnsson Lectures, Oxford: Basil Blackwell.

expected utility and mathematical expectation. 1. Expected utility theory deals with choosing among acts where the decision-maker does not know for sure which consequence will result from a chosen act. When faced with several acts, the decision-maker will choose the one with the highest 'expected utility', where the expected utility of an act is the sum of the products of probability and utility over all possible consequences.

The introduction of the concept of expected utility is usually attributed to Daniel Bernoulli (1738). He arrived at this concept as a resolution of the so-called St Petersburg paradox. It involves the following gamble: A 'fair' coin is flipped until the first time heads up. If this is at the kth flip, then the gambler receives $\$2^k$. The question arose how much to pay for participation in this gamble. Since the probability that heads will occur for the first time in the kth flip is 2^{-k} (assuming independence of the flips), and the gain then is $\$2^k$, the 'expected value' (i.e. the *mathematical expectation* of the gain) of the gamble is infinite. It has been observed though that gamblers were not willing to pay more then \$2 to \$4 to participate in such a gamble. Hence the 'paradox' between the mathematical expectation of the gain, and the observed willingness to pay.

Bernoulli suggested that the gambler's goal is not to maximize his expected gain, but to maximize the expectation of the logarithm of the gain which is $\Sigma_{j=1}^{\infty} 2^{-j} \log 2^j$, i.e. 2 log 2(= log 4). Then the gambler is willing to pay \$4 for the gamble. The idea that *homo economicus* considers the expected utility of the gamble, and not the expected value, is a cornerstone of expected theory.

In the next section the approach of Savage to decisions under uncertainty is presented. In section 3 the von Neumann–Mörgenstern characterization of expected utility maximization for the context of decisions under risk is given. Section 4 briefly mentions some related approaches. Section 5, the Appendix, defines (mathematical) expectation.

2. EXPECTED UTILITY WHEN APPLIED TO DECISIONS UNDER UNCERTAINTY; SAVAGE'S APPROACH

2.1 The main ingredients of a decision problem under uncertainty are acts consequences and states of nature. Suppose that a decision-maker has to choose one of three feasible acts f, g, h. Act f leads to one (only) of the two consequences a and b. Act g leads to a or c, act h to b or d. Thus the set of consequences, C, is in this example {a, b, c, d}.

The matching of feasible acts to consequences is expressed by the concept of 'state of nature', or 'state' for short. More precisely, a given state of nature indicates for each feasible act what the resulting consequence will be. In the above example, there are three feasible acts f, g, h, each leading to one of two possible consequences. See Table 2.1.

TABLE 2.1 The eight logically possible matchings of feasible acts to consequences

	States							
Acts	s_1	s_2	s_3	s_4	s_5	s_6	s_7	s_8
f	a	a	a	a	b	b	b	b
g	a	a	c	c	a	a	c	c
h	b	d	b	d	b	d	b	d

A state of nature completely resolves the uncertainty relating acts to consequences. If the decision-maker would know for sure which state of nature is the true one, then he would choose an act which results in a most desirable consequence. The desirability of a consequence neither depends on the act nor on the state of nature leading to it.

In constructing a table like Table 2.1 some of the states of nature may be deleted if the decision-maker is certain that they cannot occur.

Schmeidler, D. & P.P. Wakker (1987), 'Expected Utility and Mathematical Expectation', In J. Eatwell & M. Milgate, & P. Newman (Eds.), The New Palgrave: A Dictionary of Economics, 229–232. Macmillan, Cambridge.

The next step in the process of selecting the best act is to construct 'conceivable' acts, which are not feasible. Thus the set of acts, F, in Savage's set-up consists of all functions from the set of states of nature, S, to the set C of consequences. In our example there are 4^8 acts. Of these, three acts f, g and h are actually feasible: The additional 65533 acts are only conceivable. The construction of the conceivable acts and the possibility of ranking all acts of F is a basic assumption of the present approach. For the sake of presentation we will in the next subsection assume the validity of the expected utility theory and then we will return to the rationale of our construction.

2.2 Suppose for the present that the decision-maker, in choosing between acts, indeed computes the expected utility of each act, and selects a feasible act with the highest expected utility. Thus we are assuming that he has assigned probability $P(s)$ to every state of nature s in S, and the utility $U(c)$ to every consequence c in C. So, given an act f in F, the expected utility $EU(f)$ of f equals $\Sigma_{s \in S} P(s)U[f(s)]$. More generally, if the set S is infinite, then P is a finitely additive probability measure defined on all events (i.e. subsets of S), and $EU(f)$ equals $\int U[f(s)]\,\mathrm{d}P(s)$ (assuming the integral to exist; say U is bounded; see the Appendix, on Mathematical Expectation, section 5). So in fact in this case the decision-maker has a well-defined 'preference relation' (i.e. binary relation) $\gtrsim$ on the set of acts F, with, for all f, g in F:

$$f \gtrsim g \quad \text{iff } EU(f) \geqslant EU(g) \tag{2.3}$$

It is easily seen that the preference relation, defined in (2.3), is not affected when the utility function $U: C \mapsto R$ is replaced by any positive linear transformation of it (say $\bar{U}: c \mapsto \alpha U(c) + \beta$, for some real β and positive α).

2.4 If a preference relation, $\gtrsim$, over acts is derived from comparisons of expected utility as in (2.3), then it must satisfy several properties. We follow the terminology and order of Savage (1954). He listed seven postulates, five of which (P1 up to P4, and P7) are implied by (2.3). Postulate P1 says that the preference relation is complete ($f \gtrsim g$ or $g \gtrsim f$ for all acts f, g) and transitive. Postulate P2 is referred to as the sure-thing principle. It says that, when comparing two acts, only those states of nature matter, on which these acts differ. In other words, for the comparison between two acts, if they coincide on an event A, it really does not matter what actually the consequence is for each state in A. Thus P2 makes it possible to derive a preference relation over acts, conditioned on the event A^c; this for any event A.

Postulate P3 entails that the desirability of a consequence does not depend on the combination of state and act that lead to it; hence the possibility to express the desirability of consequences by a utility function on C.

P4 guarantees that the preference relation over acts induces a qualitative probability relation ('at least as probable as') over events, which is transitive and complete. P7 is a technical monotonicity condition.

P5 and P6 are Savage's only postulates which are not a necessary implication of (2.3). P5 simply serves to exclude the trivial case where the decision-maker is indifferent between any two acts. P6 implies some sort of continuity of the preference relation, and non-atomicity of the probability measure; the last term means that any non-impossible event can be partitioned into two non-impossible events. Hence there must be an infinite number of states.

Savage's great achievement was not to *assume* (2.3), but to show that his list of postulates P1–P7 *implies* that the preference relation over acts has an expected utility representation as in (2.3). Savage argued compellingly for the appropriateness of his postulates. Furthermore, Savage showed that the probability measure in (2.3) is uniquely determined by the preference relation $\geqslant$, and that the utility function is unique up to a positive linear transformation.

2.5. The significance of Savage's achievement is that it gives the first, and until today most complete, conceptual foundation to expected utility. Savage's conclusion, to use expected utility for the selection of optimal acts, can be used even if we do not have the structure and the seven postulates of Savage. Indeed, the assumption needed on consequences, states, acts, and preferences, is that they can be extended so as to satisfy all requirements of Savage's model. Also other models, as mentioned in section 4, can be used to obtain expected utility representations.

Given a decision problem under uncertainty, if we assume that it can be embedded in Savage's framework, then it is not necessary to actually carry out this embedding. In other words, if the decision-maker is convinced that in principle it is possible to construct the conceivable acts as in subsection 2.1 and the ranking of all acts in accordance with the postulates, then this construction does not have to be made. Instead one can directly try to assess probabilities and utilities, and apply the expected utility criterion. As an example, suppose a market-vendor has to decide whether to order 50 portions of ice-cream (f), or not (g). One portion costs \$1, and is sold for \$2. If the weather will be nice the next day, the school nearby will allow the children to go to the market, and all 50 portions, if ordered, will be sold, yielding a profit of \$50. If the weather is not nice, no portion will be sold. We assume that the ice-cream cannot be kept in stock and hence bad weather will yield a 'gain' of \$−50 if the portions have been ordered.

Instead of embedding the above example into Savage's framework, the market salesman may immediately assess P_1 (or $1 - P_1$), the probability for good (bad) weather; next assess the utilities of gaining \$50, \$0 and −\$50; finally order the 50 portions if $P_1 U(\$50) + (1 - P_1)U(-\$50) > U(\$0)$.

Theoretical conclusions can be derived from the mere assumption of expected utility maximization, without an actual assessment of the probabilities and utilities. Examples are the theories of attitudes towards risk, with applications to insurance, portfolio choice, etc. The validity of these applications depends on expected utility theory, which in turn depends on the plausibility of Savage's model (or other derivations of expected utility).

Another important theoretical application of Savage's model is to neo-Bayesian statistics. For applied statistics, in this vein, the availability of a 'prior distribution', as proved by Savage's approach, is essential.

3. EXPECTED UTILITY WHEN APPLIED TO DECISIONS UNDER RISKS; THE VON NEUMANN–MORGENSTERN APPROACH

Special and extreme cases of decisions under uncertainty are decisions in 'risky' situations. In decisions under uncertainty, as exposited in the previous section, the decision-maker who follows the dictum of expected utility has to assign utilities to the consequences and probabilities to the states. He can do it by mimicking the proof of Savage's theorem, or more directly by organizing his information, as the case may be.

Decision-making under risk considers the special case where the formulation of the problem for the decision-maker includes probabilities for the events, so that he only has to derive the utilities of consequences. As an example, consider a gambler in a casino who assumes that the roulette is really unbiased, so that each number has probability 1/37 (or 1/38). Another example is the St Petersburg paradox, described in subsection 1.2.

Within the framework of expected utility theory, for the evaluation of an act, only its probability distribution over the consequences has to be taken into account. Thus, for decision-making under risk, with probabilities known in advance, one may just as well describe acts as probability distributions over consequences instead of as functions from the states to the consequences.

3.1. Let us denote by L the set of probability distributions over C with finite support. We refer to them as lotteries. Von Neumann and Morgenstern (1947, Appendix) suggested conditions on a preference relation $\geqslant$ between lotteries, necessary and sufficient for the existence of a real-valued utility function U on C, such that for any two lotteries P and Q in L:

$$P \geqslant Q \qquad \text{iff} \sum_{c \in C} P(c)U(c) \geqslant \sum_{c \in C} Q(c)U(c) \tag{3.2}$$

It is easy to see that the utility function, U, is unique up to positive linear transformations. Before we present a version of von Neumann–Morgenstern's theorem, recall that for any $0 \leqslant \alpha \leqslant 1$, and for any two lotteries P and W, $R := \alpha P + (1-\alpha)Q$ is again a lottery, assigning probability $R(c) = \alpha P(c) + (1-\alpha)Q(c)$ to any c in C. Also note that the assumption that all lotteries are given, is sometimes as heroic as Savage's assumption that all functions from S to C are conceivable acts.

The first axiom of von Neumann–Morgenstern, NM1, says that the preference relation over the lotteries is complete and transitive. NM2, the continuity axiom, says that, if $P \succ Q \succ R$, then there are α, β in $]0, 1[$, such that $\alpha R + (1-\alpha)P \succ Q \succ \beta P + (1-\beta)R$. Here the strict preference relation $\succ$ is derived from $\succsim$ in the usual way: $P \succ Q$ if $P \succsim Q$ and not $Q \succsim P$.

The third axiom NM3 is the independence axiom. It says that for α in $]0,1]$, P is preferred to Q iff $\alpha P + (1-\alpha)R$ is preferred to $\alpha Q + (1-\alpha)R$. This condition is the antecedent of Savage's sure-thing principle, and is the most important innovation of the above axioms.

3.3 Von Neumann and Morgenstern originally stated their theorem for more general sets than L. They did it for so-called mixture spaces, i.e. spaces endowed with some sort of convex combination operation. This has been done more precisely by Herstein and Milnor (1953).

Von Neumann and Morgenstern introduced their theory of decision-making under risk as a normative tool for playing zero-sum games in strategic form. There the 'player' (i.e. decision-maker) can actually construct any lottery he wishes over his pure-strategies (but not over his consequences).

The theorem of von Neumann and Morgenstern, stated above, is a major step in the proof of Savage's theorem.

Recently there has been much research on decision making under risk for its own end. Some of this research is experimental, subjects are asked to express their preferences between lotteries. These experiments, or polls, reveal violations of most of the axioms. They lead to representations different from expected utility.

4. OTHER APPROACHES AND BIBLIOGRAPHICAL REMARKS

The first suggestion for expected utility theory in decision-making under uncertainty in the vein of Savage was Ramsey's (1931). His model was not completely formalized. The work of Savage was influenced by de Finetti's approach to probabilities, as in de Finetti (1931, 1937). The decision theoretic framework to which Savage's expected utility model owes much is that of Wald (1951), who regards a statistician as a decision-maker.

A model which can be considered intermediate between those of Savage and von Neumann and Morgenstern is that considered by Anscombe and Aumann (1963). Formally it is a special case of a mixture set, but like Savage it introduces states of nature, and gives a simultaneous derivation of probabilities for the states, and of utilities for the consequences. A consequence in this model consists of a lottery over deterministic outcomes; this involves probabilities known in advance, as in the approach of von Neumann and Morgenstern. The Anscombe and Aumann theory, as well as most of the technical results up to 1970, are presented in detail in Fishburn (1970).

In the expected utility theory, described above, the desirability (utility) of consequences does not depend on acts or states of nature. This is a restriction in many applications. For example the desirability of family income may depend on whether the state of nature is 'head of family alive' or 'head of family deceased'. Karni (1985) summarized and developed the expected utility theory without the restrictive assumption of state-independent preferences over consequences.

Ellsberg (1961) argued against the expected utility approach of Savage by proposing an example, inconsistent with it. A way of resolving the inconsistency is to relax the additivity property of the involved probability measures. Schmeidler (1984) formulated expected utility theory with non-additive probabilities for the framework of Anscombe and Aumann (1963). Gilboa (1985) did the same for the original framework of Savage. Wakker (1986) obtained expected utility representation, including the non-additive case, for a finite number of states of nature and non-linear utility.

5. APPENDIX: MATHEMATICAL EXPECTATION

5.1 *Expectation with respect to finitely additive probability.* A non-empty collection Σ of subsets (called events) of a non-empty set S is said to be an algebra if it contains the complement of each set belonging to it, and it contains the union of any two sets belonging to it. A (finitely additive) probability P on Σ assigns to every event in Σ a number between 0 and 1 such that $P(S) = 1$ and for any two *disjoint* events A and B, $P(A \cup B) = P(A) + P(B)$.

A random variable X is a real-valued function on S such that, for any open or closed (bounded or unbounded) interval I, $\{s \in S \mid X(s) \in I\}$ (or $[X \in I]$ for short) is an event i.e., in Σ. Given such a random variable X, its (mathematical) expectation is:

$$E(X) = \int_0^{\infty} P[X \geqslant \alpha]\,\mathrm{d}\alpha - \int_{-\infty}^{0} (1 - P[X \geqslant \alpha])\,\mathrm{d}\alpha, \tag{5.2}$$

where the integration above is Riemann-integration and it is assumed that the integral exist. The integrands in (5.2) are monotonic, so $E(X)$ exist if X is bounded. If the random variable X has finitely many values, say $x_1, \ldots, x_n$ then (5.2) reduces to

$$E(X) = \sum_{i=1}^{n} P(X = x_i)x_i, \tag{5.3}$$

However, an equation like that above may not hold if the random variable obtains countably many different values. An example will be provided in subsection 5.7.

5.4 *σ-additive probability.* Kolmogorov (1933) imposed an additional continuity assumption on probability P on Σ: To simplify presentation he first assumed that Σ is a σ-algebra, i.e., an algebra such that for every sequence of events $(A_i)_{i=1}^{\infty}$ it contains its union $\cup_{i=1}^{\infty} A_i$. He then required that $P(\cup_{i=1}^{\infty} A_i) = \Sigma_{i=1}^{\infty} P(A_i)$ if the A_i's are pairwise disjoint.

This last property is referred to as σ-additivity of the probability P. In this way Kolmogorov transformed large parts of probability theory into (a special case of) measure theory. Thus an expectation of a random variable X is

$$(5.5) \qquad E(X) = \int_S X(s)\,\mathrm{d}P(s)$$

where the right side is a Lebesgue integral (if it exists...), defined as a limit of integrals of random variables with countably many values. Let Y be such a random variable with values $(y_i)_{i=1}^{\infty}$, then

$$(5.6) \qquad E(Y) = \sum_{i=1}^{\infty} P(Y = y_i)y_i$$

if the right side is absolutely convergent.

5.7 *An example* will now be introduced of a finitely additive probability, i.e. a probability for which (5.3) holds but (5.6) does not hold. Let S be the set of rational numbers in the interval [0, 1] and let Σ be the algebra of all subsets of S. (It is in fact a σ-algebra.) For $0 \leqslant \alpha \leqslant \beta \leqslant 1$ define $P(S \cap [\alpha, \beta]) = \beta - \alpha$ and extend P to all subsets of Σ. For each s in S, $P(s) = 0$. Since S is countable we can write $S = \{s_1, s_2, \ldots\}$ and $1 = P(S) > \Sigma_{i=1}^{\infty} P(s_i) = 0$. Defining $Y(s_i) = 1/i$ for all i, we get a contradiction to (5.6). The finitely additive probability P has also the property implied by Savage's P6 (see 2.4): If $P(A) > 0$ then there is an event $B \subset A$ such that $0 < P(B) < P(A)$.

5.8 *Distributions.* A non-decreasing right continuous function on the extended real line is called a distribution function if $F(-\infty) = 0$ and $F(\infty) = 1$. Given a random variable X, its distribution function F_X is defined by $F_X(\alpha) = P(X \leqslant \alpha)$ for all real α. Then

$$(5.9) \qquad E(X) = \int_0^{\infty} [1 - F_X(\alpha)]\,\mathrm{d}\alpha - \int_{-\infty}^{0} F(\alpha)\,\mathrm{d}\alpha$$

which is the dual of formula (5.2). If the distribution F_X is smooth we say that the random variable X has a density $f_X: R \to R$, which is the derivative of F_X. In this case

$$(5.10) \qquad E(X) = \int_{-\infty}^{\infty} \alpha f(\alpha)\,\mathrm{d}\alpha$$

5.11 *Non-additive probability.* A function $P: \Sigma \to [0, 1]$ is said to be *non-additive* probability (or capacity) if $P(S) = 1$, $P(\phi) = 0$ and for $A \subset B$, $P(A) \leqslant P(B)$. Choquet (1954) suggested to integrate a random variable with respect to non-additive probability by formula (5.2).

David Schmeidler and Peter Wakker

See also ALLAIS PARADOX; MEAN VALUES; RISK; SUBJECTIVE PROBABILITY; UNCERTAINTY; UTILITY THEORY AND DECISION-MAKING

BIBLIOGRAPHY

Anscombe, F.J. and Aumann, R.J. 1963. A definition of subjective probability. *Annals of Mathematical Statistics* 34, 199–205.

Bernoulli, D. 1738. Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5, 175–92. Translated into English by L. Sommer (1954) as: Exposition of a new theory on the measurement of risk, *Econometrica* 12, 23–36; or in *Utility Theory: A Book of Readings*, ed. A.N. Page, New York: Wiley, 1986.

Choquet, G. 1953–54. Theory of capacities. *Annales de l'Institut Fourier* (Grenoble), 131–295.

de Finetti, B. 1931. Sul significato soggettivo della probabilita. *Fundamenta Mathematicae* 17, 298–329.

de Finetti, B. 1937. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré* 7, 1–68. Translated into English in *Studies in Subjective Probability*, ed. H.E. Kyburg and H.E. Smokler, 1964, New York: Wiley.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75, 643–69.

Fishburn, P.C. 1970. *Utility Theory for Decision Making*. New York: Wiley.

Gilboa, I. 1986. Non-additive probability measures and their applications in expected utility theory. PhD Thesis submitted to Tel Aviv University.

Herstein, I.N. and Milnor, J. 1953. An axiomatic approach to measurable utility. *Econometrica* 21, 291–7.

Karni, E. 1985. *Decision-Making under Uncertainty: The Case of State-Dependent Preferences*. Cambridge, Mass.: Harvard University Press.

Kolmogorov, A.N. 1933. *Grundbegriffe der Wahrscheinlichkeitrechnung*. Berlin. Translated into English by Nathan Morrison (1950, 2nd edn, 1956), New York: Chelsea Publishing Company.

Loeve, M. 1963. *Probability Theory*. 3rd edn, Princeton: Van Nostrand.

von Neumann, J. and Morgenstern, O. 1947. *Theory of Games and Economic Behavior*. 2nd edn, Princeton: Princeton University Press.

Ramsey, F.P. 1931. Truth and probability. In *The Foundations of Mathematics and Other Logical Essays*, ed. R.B. Braithwaite, New York: Harcourt, Brace.

Savage, L.J. 1954. *The Foundations of Statistics*. New York: Wiley, 2nd edn, 1972.

Schmeidler, D. 1984. Subjective probability and expected utility without additivity. CARESS, University of Pennsylvania and IMA University of Minnesota, mimeo.

Wald, A. 1951. *Statistical Decision Functions*. New York: Wiley.

Wakker, P.P. 1986. Representations of choice situations. PhD thesis, University of Tilburg, Department of Economics.

expected utility hypothesis. The expected utility hypothesis of behaviour towards risk is essentially the hypothesis that the individual decision-maker possesses (or acts as if possessing) a 'von Neumann–Morgenstern utility function' $U(\cdot)$ or 'von Neumann–Morgenstern utility index' $\{U_i\}$ defined over some set of outcomes, and when faced with alternative risky prospects or 'lotteries' over these outcomes, will choose that prospect which maximizes the expected value of $U(\cdot)$ or $\{U_i\}$. Since the outcomes could represent alternative wealth levels, multidimensional commodity bundles, time streams of consumption, or even non-numerical consequences (e.g. a trip to Paris), this approach can be applied to a tremendous variety of situations, and most theoretical research in the economics of uncertainty, as well as virtually all applied work in the field (e.g. optimal trade, investment or search under uncertainty) is undertaken in the expected utility framework.

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