The Handbook of Economic Methodology

Edited by

John B. Davis

Department of Economics, Marquette University, USA

D. Wade Hands

Department of Economics, University of Puget Sound, USA

Uskali Mäki

Department of Philosophy, Erasmus University, Rotterdam, The Netherlands

Edward Elgar

Cheltenham, UK • Northampton, MA, USA

creates an immediate tension. Standard traditional accounts of the emergence and growth of scientific knowledge see science as a progressive enterprise which, under the appropriate conditions of rational and free inquiry, generates a body of knowledge which progressively converges on the truth. Selectionist models of biological evolution, on the other hand, are generally construed to be non-progressive or, at most, locally so. Rather than generating convergence, biological evolution produces diversity. Popper's evolutionary epistemology attempts to embrace both, but does so uneasily. Kuhn's 'scientific revolutions' account draws tentatively upon a Darwinian model but, when criticized, Kuhn has retreated (compare Kuhn, 1962: 172f with Lakatos and Musgrave, 1970: 264). Toulmin (1972) is a noteworthy exception. On his account, concepts of rationality are purely 'local' and subject themselves to evolution. The net result is the need to abandon any sense of 'goal directedness' in scientific inquiry. This is a radical consequence. Pursuing the evolutionary approach to its logical conclusion raises fundamental questions about the concepts of knowledge, truth, realism, justification and rationality.

Disciplines often borrow each other's metaphors. For example, economics, evolutionary theory and epistemology trade concepts of competition back and forth. The 'struggle for existence', which Darwin borrowed from Malthus, was reappropriated by the Social Darwinists in the latter part of the nineteenth century and used as a justification for laissez-faire capitalism. In the selectionist epistemological models of Popper (1968), Toulmin (1972) and Hull (1988), there is a struggle for existence in the 'market-place of ideas'. Economic considerations of costs and benefits are appropriate here (cf. Rescher, 1978; 1989; Wible, 1994). This has obvious implications for the advancement of knowledge, and more subtle implications for the concept of knowledge.

The gathering and dissemination of information involves costs. Who controls the purse strings controls the direction of research and the flow of information. If there is a *market-place* of ideas, should it be regulated? If so, how and to what end? Here both epistemic and economic considerations come into play. If the growth of knowledge is truly evolutionary, there is no guarantee that convergence on Truth or a Final Theory will occur. Indeed, if that evolution is governed by market forces that include economic factors intermixed with epistemic factors, a fundamental reassessment of the concepts of truth and knowledge seems in order.

Taking Darwin seriously demands a fundamental re-examination of the human epistemological condition in the light of our understanding of biological evolution. This task is not one that can or should be consigned to philosophers or biologists alone. Rather, it is one that invites and demands interdisciplinary efforts from a range of specialists in philosophy, biology and the social sciences.

MICHAEL BRADIE

Bibliography

Bradie, M. (1986), 'Assessing Evolutionary Epistemology', *Biology & Philosophy*, **1**, 401–59.

Campbell, D.T. (1974), 'Evolutionary Epistemology', in P.A. Schilpp (ed.), *The Philosophy of Karl Popper. I*, La Salle, IL: Open Court.

Changeux, Jean-Pierre (1985), *Neuronal Man*, New York: Pantheon.

Edelman, G.M. (1987), *Neural Darwinism: The Theory of Neuronal Group Selection*, New York: Basic Books.

Hull, D. (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*, Chicago: University of Chicago Press.

Kuhn, T. (1962), *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

Lakatos, I. and A. Musgrave (eds) (1970), *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press.

Lorenz, Konrad (1977), *Behind the Mirror*, London: Methuen.

Munz, Peter (1993), *Philosophical Darwinism: On the Origin of Knowledge by Means of Natural Selection*, London: Routledge.

Plotkin, H.C. (ed.) (1982), *Learning, Development, and Culture: Essays in Evolutionary Epistemology*, New York: John Wiley & Sons.

Popper, K.R. (1968), *The Logic of Scientific Discovery*, New York: Harper.

Popper, K.R. (1972), *Objective Knowledge: An Evolutionary Approach*, Oxford: The Clarendon Press.

Popper, K.R. (1984), 'Evolutionary Epistemology', in J.W. Polland (ed.), *Evolutionary Theory: Paths into the Future*, London: John Wiley.

Radnitzky, G. and W.W. Bartley (1987), *Evolutionary Epistemology, Theory of Rationality and the Sociology of Knowledge*, LaSalle, IL: Open Court.

Rescher, N. (1978), *Scientific Progress: A Philosophical Essay on the Economics of Research in Natural Science*, Oxford: Basil Blackwell.

Rescher, N. (1989), *Cognitive Economy: The Economic Dimension of the Theory of Knowledge*, Pittsburgh: University of Pittsburgh Press.

Rescher, Nicholas (1990), *A Useful Inheritance: Evolutionary Aspects of the Theory of Knowledge*, Lanham, MD: Rowman.

Riedl, Rupert (1984), *Biology of Knowledge: The Evolutionary Basis of Reason*, Chichester: John Wiley.

Ruse, M. (1986), *Taking Darwin Seriously: A Naturalistic Approach to Philosophy*, Oxford: Blackwell.

Skinner, B.F. (1981), 'Selection by Consequences', *Science*, **213**, 501–4.

Toulmin, Stephen (1967), 'The Evolutionary Development of Natural Science', *American Scientist*, **55**, 4.

Toulmin, Stephen (1972), *Human Understanding: The Collective Use and Evolution of Concepts*, Princeton: Princeton University Press.

Wible, J.R. (1994), 'Rescher's Economic Philosophy of Science: A Review of Nicholas Rescher's *Cognitive Economy*, *Scientific Progress*, and Peirce's *Philosophy of Science*', *The Journal of Economic Methodology*, **1/2**, 314–29.

Expected Utility Theory

Expected Utility Theory (EUT) states that the decision maker (DM) chooses between risky or uncertain prospects by comparing their expected utility values, that is, the weighted sums obtained by adding the utility values of outcomes multiplied by their respective probabilities. This elementary and seemingly commonsensical decision rule raises at once some of the most important questions in contemporary decision theory. We will focus here on two of these questions. First, what do the utility numbers in the formula refer to, and in particular do they belong to the same value scale as do the utility numbers that represent the DM's choices under certainty? Second, is the weighted sum procedure of combining probability and utility values the only one to be considered, and if there are indeed alternative modellings, how will the theorist choose? Most of this entry is concerned with risk as opposed to uncertainty, that is, with choice contexts in which probabilities are *given*, be they objective or not. Corresponding to the standard distinction between risk and uncertainty, there are two received versions of the theory, Von Neumann–Morgenstern theory (VNMT) and subjective expected utility theory (SEUT), respectively. We will only touch on the latter. Before examining the two questions in turn, we restate the sources and basic axiomatic structure of EUT. It has been used as both a positive and normative (or prescriptive) theory. Methodological discussions are primarily concerned with the theory in its positive role, and with the second question, which in effect concerns the relaxation of the 'VNM independence axiom'.

The history of EUT is often construed in terms of the following smooth generalization process: the principle of maximizing expected monetary values antedates EUT, which is now in the process of being generalized in two directions, by either non-additive or non-probabilistic decision theories. The highlights in this sequence are Bernoulli's (1738) resolution of the St Petersburg paradox, and Allais' (1953) invention of a thought-provoking problem widely referred to as the Allais paradox. In the St Petersburg game, people were asked how much they would pay for the following prospect: if tails comes out of the first toss of a fair coin, to receive nothing and stop the game and, in the complementary case, to receive two guilders and stay in the game; if tails comes out of the second toss of the coin, to receive nothing and stop the game and in the complementary case, to receive four guilders and stay in the game; and so on ad infinitum. The expected monetary value of this prospect is $\Sigma_n(2^n \times 1/2^n)$ = infinite. Since the people always set a definite, possibly quite small upper value on the St Petersburg prospect, it follows that they do not price it in terms of its expected monetary value. Bernoulli argued in effect that they estimate it in terms of the utility of money outcomes, and defended the log function as a plausible idealization, given its property of quickly decreasing marginal utilities. Because the resulting series, $\Sigma_n(\log 2^n \times 1/2^n)$, is convergent, Bernoulli's hypothesis was supposed to deliver a solution to the paradox; more on the history of this problem in Todhunter (1865).

Bernoulli's hypothesis counts as the first systematic occurrence of EUT theory. Two centuries later, Allais questioned the naturalness of EU-based choices by devising the following questionnaire.

Question 1: which prospect would you choose of x_1 = to receive 100 million FF with probability 1, and y_1 = to receive 500 million FF with probability 0.10, 100 million FF with probability 0.89, and nothing with probability 0.01?
Question 2: which prospect would you choose of x_2 = to receive 100 million FF with probability 0.11, and nothing with probability 0.89, and y_2 = to receive 500 million FF with probability 0.10, and nothing with probability 0.90?

Allais found that the majority answers were x_1 to question 1 and y_2 to question 2, and argued that this pair of prospects could indeed be chosen for good reasons. But it violates EUT, since there is no function U that would satisfy both:

$$U(100) > 10/100\ U(500) + 89/100\ U(100) + 1/100\ U(0)$$

and

$$11/100\ U(100) + 89/100\ U(0) < 10/100\ U(500) + 90/100\ U(0).$$

Although the word 'paradox' is frequently used in the history of EUT, it should be clear from this and the previous examples that it does not refer to deeply ingrained conceptual difficulties, such as Russell's paradox in set theory, or the EPR paradox in physics, but rather just to problems or anomalies for the theory that is currently taken for granted: expected monetary value theory in the St Petersburg case and EUT in the Allais case.

There are few explicit EU calculations in economics before von Neumann and Morgenstern (1944), who chose to determine the utility value of a randomized strategy in this mathematically convenient way. Their theoretical choice has proved to be of long-lasting influence. Not only is current game theory (including its branch specializing in incomplete information games, which dates back to Harsanyi's work in the late 1960s) still heavily dependent on EU calculations, but the same can be said, although to a lesser degree, of today's microeconomics of imperfect information, as textbooks will confirm. Like Bernoulli, von Neumann and Morgenstern are concerned with the case in which the probabilities are part of the decision problem. Their work did not yet amount to an axiomatization in the sense decision theorists and economists have become accustomed to. It is only with Marschak and with Herstein and Milnor, in the early 1950s, that the EU formula was derived as the numerical counterpart of a qualitatively defined preference structure subjected to various axiomatic constraints. In honour of the founders, this derivation was named the VNM theorem, and the crucial axiom in the construction VNM independence. This terminology might hide the historical fact that the very axiom which was to arouse innumerable discussions is not even stated in von Neumann and Morgenstern's account! Fishburn (1989) and Fishburn and Wakker (1995) discuss this fact while surveying the work of the formative years.

All available axiomatizations assume that there is a binary relation $\leq$ on the set X of all risky prospects, called also lotteries, and subject this relation to the *preordering* (that is, transitivity and completeness), *continuity* and *independence* properties. Little will be said here about the first axiom, not because it lacks empirical content, but because it is not specific to the theory of risky or uncertain choices. (However, the transitivity condition has come to be discussed widely in the EUT context, in particular because of the 'preference reversal' phenomenon, on which the reader is referred to Hausman, 1992.) The second axiom typically says that if, x is strictly preferred to z, which is strictly preferred to y, then a suitable mixture of x and y will be strictly preferred to z, and z will be strictly preferred to another suitable mixture of x and y. In the presence of the first, this axiom makes it possible to 'represent' the qualitative datum $\leq$ by some, yet unspecific, numerical function $u(x)$. It has some empirical content but plays mostly a technical role. The third axiom (VNM independence) can be stated in the following easy form, due to Samuelson: for all x, y and z in X, and any number α such that $0 < \alpha < 1$, $x \leq y$ if and only if $\alpha x + (1 - \alpha)z \leq \alpha y + (1 - \alpha)z$.

Expressed in words, preference inequalities are preserved when the initial two lotteries are mixed in a given proportion with a third lottery. This axiom is responsible for the specific, expectational form of the function $u(x)$ provided by the VNM theorem. A few constructions also involve a compound lottery axiom, which says in effect that any lottery having further lotteries as its outcomes can be reduced to a one-stage lottery. This further axiom has a definite empirical content and is now regarded as being responsible for some cases of violation of EUT. Notice that it is automatically satisfied by the standard formalization of lotteries in terms of probability functions (since a mixture, that is, convex combination, of probabilities is again a probability).

Historically, SEUT can be said to result from two distinct traditions, one the Bernoulli–VNM tradition of decision theory, the other the mathematical and philosophical tradition of subjective probability, which can be traced back to the British empiricists and Bayes, and which was revived in the 1930s by Ramsey and de Finetti. (The contrast between 'objective' and 'subjective' schools of probability is surveyed in Fine, 1973; see also Fishburn's 1986 introduction to subjective probability.) De Finetti was particularly emphatic in claiming that probability does not exist in any substantial sense. As he conceived of it, probability does not even necessarily exist in the subject's mind; it might just be the numerical expression, as defined by an outside observer, of the property that the subject behaves coherently when choosing between uncertain prospects. This interpretation elaborates on the Dutch book theorem, which was first sketched by Ramsey: a Dutch book is a list of bets on all possible events which leads to a net loss of money whichever state of the world is realized; the theorem shows that to avoid Dutch books is equivalent to choosing among prospects according to the expectation of their monetary values, where the expectation is taken with respect to some well-defined probability. Leaving aside the strong anti-realist stand taken by de Finetti, as well as (though to a lesser degree) Ramsey, these authors were the first to bridge the analysis of probability with a (rudimentary) decision theory. Savage (1954) consolidated and enlarged the bridge by showing that to satisfy certain behavioural requirements in the style of, but more abstract and general than, the no-Dutch-book assumption is equivalent to choosing among prospects according to the expectation of their utility values, where the expectation is taken with respect to some well-defined pair of probability and utility function. The non-trivial step in Savage's contribution is to reveal these two items simultaneously from the axiomatically constrained preference behaviour. To do so, he made implicit use of VNM theory. While two of his axioms (P3 and P4) are reminiscent of the Dutch book scheme, his postulate P2, or 'sure-thing principle', is the counterpart of VNM independence in the subjective probability framework. Unsurprisingly, VNM independence and the 'sure-thing principle' have been criticized in broadly similar ways, and have led to parallel generalizations. Savage's axiomatization has also induced a specific 'paradox', Ellsberg's (1954), which is usually understood as contradicting the existence of subjective probabilities and constitutes the starting point of another generalization trend.

To make precise the first question raised at the outset, a formal statement of the VNM theorem is needed. Denoting by $c_1, \ldots, c_k$ the outcomes of lottery x, and by $p_1, \ldots, p_k$ the attached probability values, the theorem says that if the three axioms of preordering, continuity and independence hold, there is a representation of the preference relation in terms of the expectation of some utility function U on the outcomes, that is, $\Sigma_i p_i\, U(c_i)$ and that the U function in this representation is 'unique up to a positive linear transformation'. Now, the question is, does U refer to the same quantities as does the utility function of non-stochastic theories, such as consumer theory in microeconomics? The 'measurability controversy' in which Baumol, Friedman and Savage, Ellsberg, Luce and Raiffa were involved in the 1950s, was mostly concerned with this problem. The participants agreed that, for a given individual, the *ordinal* properties of $U(c)$ should be the same as those of any alternative index $V(c)$ provided by the non-stochastic theory of preferences among outcomes. However, there were severe disagreements on what added properties the VNM theorem delivered. A view shared by some prominent economists, including perhaps von Neumann and Morgenstern themselves, was that the $U(c)$ index succeeded where earlier indexes from non-stochastic theory had failed; that is, it had *cardinal* properties in the sense of measuring the individual's preference differences or intensities over the outcome set. This conclusion was supported by the following argument: take three outcomes c, c' and c'' such that $U(c) > U(c') > U(c'')$, and suppose that: $1/2U(c) + 1/2U(c'') > U(c')$ holds. Then $U(c) - U(c') > U(c') - U(c'')$ trivially follows; because the last inequality does not depend on the particular representation U (see the uniqueness part of the VNM theorem), it would seem natural to interpret it as implying that the individual's intensity of preference of c over c' is stronger than the intensity of his preference of c' over c''. Hence the conclusion that VNM had gone beyond the purely ordinalist stand at which the Paretian school had stopped. By and large, the 'measurability controversy' resulted in the rejection of this optimistic interpretation. It was shown to rely on a superficial understanding of the uniqueness part of the VNM theorem. To make sense of a cardinal index in the desired sense, one should first of all impose a special axiom on the preference relation on outcomes, to the effect that preference differences or intensities are meaningful. In the presence of suitably strengthened versions of the preordering and continuity axioms, this added axiom will have the effect of determining another utility function W on outcomes, which is itself 'unique up to a linear transformation'. But unless this is explicitly required by adding still another axiom,

there is no reason why W should be a linear transformation of the VNM index, U. In other words, the uniqueness part of the theorem does provide a formal method for comparing utility differences but the numbers derived in this way might be unrelated to the measurement of preference differences on the outcome set. The refutation sketched here is in accord with Fishburn's (1989) and further clarified by Bouyssou and Vansnick (1990).

Although the negative point just made might now count as standard doctrine, it is not universally recognized. Harsanyi is prominent among those who interpret the VNM index as measuring the individual's true preference differences. He actually needs this interpretation in order to tighten up and clarify the important work he pursued in the 1950s to connect EU theory with utilitarianism. Because of Harsanyi's influence on current welfare economics, the 'measurability controversy' is not yet closed; for further details, see Mongin and d'Aspremont (forthcoming). Allais' notion of cardinality is also at variance with the standard doctrine, although for reasons of his own.

The second of the two questions mentioned at the outset is concerned with the VNM independence axiom and has led to a host of competing answers. Among the many available generalizations of this axiom, two will be singled out here. Machina's proposed theory directly generalizes the EU representation without investigating the corresponding properties of the preference relation. The first step is to replace the linearity-in-the-probabilities property of the EU representation by the weaker property of differentiability with respect to the probabilities. It can be checked that EUT then holds as a local approximation of the new theory: there are as many EU representations as there are lotteries (formally, probability distributions) and these many representations can be used to approximate relevant properties of the global, differentiable representation in the neighbourhood of the lottery x they are associated with. Each of them gives rise to a function on outcomes which will be denoted by $U_x(c)$. (Note the difference from VNM theory, which deduced a function $U(c)$ independent of the particular lottery x.) Clearly, the U_x may exhibit widely differing curvature properties. The second step in Machina's theory precisely consists in determining how the curvature of U_x varies with x: following the 'fanning out' hypothesis, it should satisfy a condition of increasing risk aversion (as measured in terms of concavity) when x varies in the direction of increasing stochastic dominance (roughly speaking, in the direction of lotteries that move probabilistic weight towards the better outcomes). The whole of this technical construction can be illustrated elegantly in terms of indifference curves: step 1 then means that the agent's indifference curves are of any smooth shape, instead of being parallel straight lines, as in the VNM particular case; and step 2 simply means that they become steeper when one looks in the direction defined by the best outcome. As Machina (1983) explains, it becomes possible to account for the common, systematic pattern underlying a number of well-established anomalies of EUT: not only the 'common consequence effect' (which is Allais' paradox in more abstract form), but also the 'common ratio effect' (a related anomaly, which was investigated by Kahneman and Tversky), the 'utility evaluation effect' (an anomaly which emerged from the attempts at numerically estimating the $U(c)$ function), and a lesser known effect, called 'oversensitivity to changes in small probability-outlying events'.

Another generalization of EUT which has perhaps become more popular than any other among decision theorists is the so-called rank-dependent or anticipated utility theory (AUT). It is now endowed with an axiomatic version which clarifies the sense in which VNM independence is weakened, but is easier to discuss in terms of functional representations. Assuming that the outcomes $c_1, \ldots, c_k$ of lottery x are ranked in increasing preference order, AUT evaluates x as follows:

$$f(p_1)\, U(c_1) + \{f(p_1 + p_2) - f(p_1)\}\, U(c_2) + \ldots + \{f(p_1 + p_2 + \ldots + p_k) - f(p_1 + p_2 + \ldots + p_{k-1})\}\, U(c_k)$$

where f: $[0, 1] \to [0, 1]$ and $f(0) = 0, f(1) = 1$ and f is weakly increasing.

This added function is intended to capture the agent's distortion of probability values. When f is the identity function, the formula collapses into EUT. If f satisfies $f(p) \geq p$ for all p, and in particular is concave throughout, it expresses a psychological attitude akin to risk aversion: think of a lottery with two distant monetary outcomes and compare it with the lottery giving its expected monetary value for sure. Conversely, $f(p) \leq p$ for all p expresses a kind of risk-seeking attitude. The case of a S-shaped f has also been explored; it involves a strong tendency to overweight small probability values. Various shapes can be invoked to account for the Allais paradox and the other related effects.

Conceptually, there are two important connections between AUT and earlier work. First, the idea of enriching EUT with a probability distortion element is very natural, and indeed emerged at an early stage of the empirical psychologists' work. For instance, Kahneman and Tversky's 'prospect theory' involved generalizing EUT in terms of the following formula:

$$g(p_1)\, U(c_1) + g(p_2)\, U(c_2) + \ldots + g(p_k)\, U(c_k).$$

Natural though it seems, this formula leads to an unwelcome consequence. Decision theorists were willing to give up the linearity-in-the-probabilities property of EUT, but not this further, much weaker implication of VNM independence: if lottery x stochastically dominates lottery y, the agent prefers x to y. The stochastic dominance property can be violated in Kahneman and Tversky's generalization of EUT. Importantly, it is always satisfied by the AUT formula, in which cumulative distribution values, rather than probability values irrespective of the order of prospects, are assumed to be distorted. Following Quiggin's (1993) interpretation, this is indeed the decisive contribution of AUT: it salvaged the psychologists' intuition in the only way compatible with a hard core postulate of decision theory. Second, AUT connects with the 'measurability controversy'. As explained above, one of the results of this controversy was that the VNM cardinalization should not be confused with the cardinalization relevant to the measurement of preference differences in the certainty case. A closely related conclusion is that the property of risk aversion, as measured by the concavity of the VNM index, should not be confused with the property of diminishing marginal utility of money, which belongs to the certainty context. A richer theory than EUT is needed to express this conceptual distinction formally. According to many decision theorists, AUT is the theory needed: the added function f has the role of expressing risk attitudes, so that U can be reserved for the altogether different use of conveying certainty-related properties, such as diminishing marginal utility.

Perhaps the most important methodological question raised by EUT is whether the process of successful generalization sketched above can be interpreted as evidence of scientific progress. Mongin (1988) gives some perspective on this issue by investigating EUT in the context of Duhem's problem: the problem of choosing which part of a compound hypothesis should be sacrificed when that hypothesis is faced with unfavourable evidence. This paper emphasizes the Duhem problem raised by the axiomatic structure of VNM theory and, using standard

philosophy of science arguments, claims that this problem was solved satisfactorily when – after many hesitations – decision theorists agreed to follow Allais' suggestion of sacrificing VNM independence while retaining the weaker property of stochastic dominance and the other axioms. Another, perhaps equally important, Duhem problem of EUT has to do with the structure of evidence. Even leaving aside thought experiments, introspection and stylized facts, which play a role in decision-theoretic discussions as they do elsewhere in economics, truly experimental evidence can be interpreted as telling only against the background of delicate auxiliary hypotheses. The crucial ones relate to the subject's cognitive abilities (typically, his understanding of probability and expectation), his financial motivations (rewards should be neither too small nor too large) and to the subject's consistency from one set of answers to another (for example, when data on different gambles are matched against each other, or gambling data are matched against market data). The only existing discussions of auxiliary hypotheses are by experimenters themselves, who have invented further experiments to assess their hidden role: see the continuing work around the widely used 'Becker–De Groot–Marschak procedure', as reported in, for example, Camerer (1995). The experimenters' assessments of their own methods rarely lead to clear-cut conclusions. However, there is some reason for believing that the Duhem problem, in the second of the two senses distinguished here, was at least sometimes resolved satisfactorily: it would seem as if the four effects listed by Machina were correctly interpreted as telling evidence against EUT, rather than against auxiliary hypotheses, because they follow a highly systematic pattern of mutually supporting evidence.

To solve the Duhem problem in the theory of risky choice would be an important negative step, but it would not yet provide a positive guarantee of scientific progress. Mongin noted that Machina's generalized expected utility theory (GEUT) satisfies neither Popper's nor Lakatos criteria of progressivity; the same could probably be said of AUT. Leaving aside the refutationist context in which this conclusion is phrased, the underlying philosophical problem appears to be that both GEUT and AUT are generalizations of EUT in the sense of being logically weaker than EUT. As a result, they account for violations of EUT in a loose way: they are compatible with these violations whereas EUT was not, but they do not *imply* them; hence, on a standard construal, they do not *explain* them. Similarly, neither GEUT nor AUT really explain the occasional empirical success of EUT. When philosophers of physics say that relativity theory is more general than Newtonian mechanics, they do not mean to suggest that the former is a logical weakening of the latter; rather the contrary: they mean to say that, for some values of relevant parameters, the former implies the latter. The theory of risky choices is far from this stage, even if both Machina's discussion of 'fanning out' and the investigation of the f function in AUT can be construed as coarse attempts to identify the parameters whose special values would turn the more general theories into the particular case of EUT. It is at least a reassuring feature of recent experimental work that the logical structure of the test problem is fully taken into account, and that methods are being devised to compare not only EUT with alternatives but also these alternatives among themselves (for example, Hey and Orme, 1994).

The previous assessment was concerned with EUT viewed as a positive theory of individual behaviour, but part of the current discussion of VNM independence is normative in character. Friedman and Savage (1952) had claimed for this axiom the normative force of the stochastic dominance principle; the flaw in their argument is blatant. A more important defence revolves around a restatement of independence in terms of temporal consistency; it has gained wide acquiescence. To explain why Allais' solution might be irrational, some authors reformulate the lotteries in the Allais paradox by using decision trees and assuming that the agent's and nature's choices follow each other in a specific order. Hammond axiomatized the underlying structure of this argument under the name 'consequentialism'; his theory also provides a defence of Savage's sure-thing principle; see McClennen's (1990) critical account.

By and large, the still lively discussion around dynamic rationality does not interact with experimental work. The specialists' wisdom is that these are completely separate areas, since one belongs to positive decision theory and the other to normative decision theory, which allegedly do not communicate with each other. Methodologists should be impressed by this application of Hume's thesis. There are several reasons for believing that the cleavage between two groups of investigations is inappropriate. First, any piece of axiomatic decision theory is *prima facie* open to both positive and normative interpretations, so that the cleavage cannot be a material one but at most a distinction between two vantage points. For instance, Hammond's axioms could be discussed also at the experimental level, even it is not obvious how to test them. Second, it can be argued that the positive vantage point is not entirely self-contained. One way of controlling for experimental data is simply to attempt to reproduce them in different circumstances; another, which was used in the context of the Allais paradox, is to check verbally whether subjects acquiesce in the (normative) principles of choice they spontaneously apply. Third, some decisions theorists are primarily interested in applying EUT and its variants to actual decisions, typically in business and medicine. Their *prescriptive* use of the theory implies that they have to face, and in some sense supersede, the positive–normative dichotomy. For instance, the parameter values in the prescriptive model derive from data obtained from the subject, which data can be obtained only by assuming, that the subject obeys some theory. Incoherences will result if the (normative) theory which underlies the prescriptions is too distant from the (positive) theory assumed in order to collect the data.

Methodologists have hardly begun to explore the developments of EUT. The above shows that they constitute a rich vein of case studies. Actually, a stronger suggestion can be made. Over the years, EUT assumptions have strengthened their grip on economic theorizing: today, they are no longer reserved for a particular department that is, the economics of risk and uncertainty, but in a sense channelled everywhere by game-theoretic reasoning (which we said earlier heavily relies on VNMT and SEUT). We believe that every attempt at constructing a general economic methodology should be subjected to the test of whether or not it delivers a coherent account of EUT. A famous precedent here is Friedman, whose articles with Savage anticipate his methodological themes.

PHILIPPE MONGIN

References

Allais, M. (1953), 'Le comportement de l'homme rationnel devant le risque', *Econometrica*, **21**, 503–46.
Bernoulli, D. (1738), 'Specimen theoriae novae de mensura sortie', *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, **5**, 175–92.
Bouyssou, D. and J.C. Vansnick (1990), 'Utilité cardinale dans le certain et choix dans le risque', *Revue économique*, **6**, 979–1000.
Camerer, C. (1995), 'Individual Decision Making', in J.K. Kagel and A.E. Roth, *The Handbook of Experimental Economics*, Princeton: Princeton University Press.
Ellsberg, D. (1954), 'Classic and Current Notions of "Measurable Utility"', *Economic Journal*, **64**, 528–56.
Fine, T.L. (1973), *Theories of Probabilities*, New York: Academic Press.
Fishburn, P.C. (1986), 'The Axioms of Subjective Probability', *Statistical Science*, **1**, 335–58.
Fishburn, P.C. (1989), 'Retrospective on the Utility Theory of von Neumann and Morgenstern', *Journal of Risk and Uncertainty*, **2**, 127–58.
Fishburn, P.C. and P. Wakker (1995), 'The Invention of the Independence Condition', *Management Science*, **41**, 1130–1144.
Friedman, M. and L. Savage (1952), 'The Expected Utility Hypothesis and the Measurability of Utility', *Journal of Political Economy*, **60**, 463–74.
Hausman, D. (1992), *The Inexact and Separate Science of Economics*, Cambridge: Cambridge University Press.
Hey J. and C. Orme (1994), 'Investigating Generalizations of Expected Utility Theory Using Experimental Data', *Econometrica*, **62**, 1291–1326.
Machina, M. (1983), 'Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom', in B.P. Stigum and F. Wenstop (eds), *Foundations of Utility and Risk Theory with Applications*, Dordrecht: Reidel.
McClennen, E.F. (1990), *Rationality and Dynamic Choice*, Cambridge: Cambridge University Press.
Mongin, P. (1988), 'Problèmes de Duhem en théorie de l'utilité espérée', *Fundamenta Scientiae*, **9**, 299–327.
Mongin, P. and C. D'Aspremont (forthcoming), 'Utility Theory and Ethics', in S. Barbera, P. Hammond and C. Seidl (eds), *Handbook of Utility Theory*, Dordrecht: Reidel.
Quiggin, J. (1993), *Generalized Expected Utility. The Rank-Dependent Model*, Boston: Kluwer.
Savage, L.J. (1954), *The Foundations of Statics*, New York: Dover; 2nd rev. edn, 1972.
Todhunter, I. (1865), *A History of the Mathematical Theory of Probability*, Cambridge; reprinted, Bronx, NY: Chelsea Publishers.
Von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

Experimental Economics

Methodologists have often found social sciences, including economics, problematic because of their relationship to experiment. The worry is that experiment is sometimes thought an integral feature of scientific inquiry, but experiment, especially controlled, repeatable laboratory experiment, has been almost non-existent in the practice of social science. It might be argued that there are subdisciplines of physics that are similarly impoverished; astrophysics and seismology come to mind as fields in which scientists hungry for data must often wait for Nature's cooperation. But this argument involves a gross exaggeration. Many of the phenomena studied in astrophysics and seismology are independently confirmed on a different scale in laboratories. And many celestial objects of interest are constantly emitting radiation, while seismologists have been known to use explosives to generate data. There is no analogy for the social sciences, especially for such theories as historical materialism where the whole thing can happen only once, so to speak.

There is, however, a growing literature that reports on the quite new practice of laboratory experimentation in economics. Some of this work attempts to study the behaviour of individual responses to lotteries, or two-person bargaining situations. Probably the most well developed experimental paradigms deal instead with market behaviour. These experiments, pioneered by Vernon Smith and his associates, all place experimental subjects at computer terminals where they make entries which they are to interpret as buying and selling. Each subject thus interacts only with the terminal and never face-to-face with the other agents in the market.

An enormous obstacle to investigating the market behaviour of economic agents in this way is ensuring that the subjects behave as economic agents. If the hypothesis being experimentally tested is, for example, that a market consisting of economic agents in specified conditions will reach equilibrium according to some piece of theory, it is necessary that the subjects' behaviour conform to the theory's description of them. Only in this way can the theory's implications for such agents be tested. The obstacle is removed by putting in place a schedule of incentives for the subjects that 'induces' an appropriate utility function. In other words, the subject will leave the laboratory well off if he behaves as though he were maximizing the utility function that the experimenter wishes. On the whole, the results of these experiments appear to confirm