Synthese (2010) 172:1–6
DOI 10.1007/s11229-009-9468-1

Introduction

Horacio Arló-Costa · Jeffrey Helzner

Published online: 26 February 2009

We are pleased to present the first of two special issues of Synthese that are dedicated to the foundations of the decision sciences. The inspiration for this project came from a variety of sources. Among the most significant of these were our own experiences in teaching courses on rational choice, and related areas, to philosophy students at Carnegie Mellon and Columbia. The literature on these topics is vast. Like logic, decision theory is a significant example of a topic that is of interest to people working in a variety of fields: economics, philosophy, psychology, and statistics. Yet despite this variety and the large amount of research that it has generated, we found it difficult to compile appropriate reading lists for our philosophy students. As one would expect, and again resembling the analogous situation in logic, much of the literature is highly specialized. There is no doubt that some of this highly specialized work is both interesting and important, but for the purposes of teaching our philosophy students about these areas of research we think that it is best to assign readings of a more foundational nature.

As noted in the previous paragraph, our intentions in preparing these special issues were focused on the possibility of adding to the existing body of philosophical literature concerning the decision sciences. The reader may have already surmised that we have taken *decision science* as an umbrella term that is meant to include fields such as decision theory, game theory, and social choice. The purpose of this introduction is to provide a brief survey of some of the variety that exists in the foundations of the decision sciences.

H. Arló-Costa (✉)
Department of Philosophy, Carnegie Mellon University, Baker Hall 135, Pittsburgh, PA 15213, USA
e-mail: hcosta@andrew.cmu.edu

J. Helzner
Department of Philosophy, Columbia University, 708 Philosophy Hall, 1150 Amsterdam Ave,
MC: 4971, New York, NY 10027, USA

As Peter Gardenfors and Nils-Eric Sahlin note in their introduction Gardenfors and Sahlin (1988) to *Decision, Probability, and Utility*, itself an excellent collection of older essays on the foundations of decision theory, a significant amount of work in these areas has focused on the integration of two key factors: desires and beliefs. Traditionally, work on the desires side has been intimately connected with the concept of utility. Many are familiar with Daniel Bernoulli's now famous discussions of the law of diminishing marginal utility, a concept that was extended in various ways that were famously discredited by economists such as Vilfredo Pareto.[1]

The reason for the discrediting of these extensions stems largely from the fact that the utility functions at issue preserve only ordinal properties of the underlying value judgements. Thus, such real-valued utility measures were superfluous at best, since one could just as well work with the underlying qualitative judgements, and they were misleading at worst, since one might be tempted to subject these utilities to forms of mathematical analysis that had no justification in light of the fact that only the ordinal properties of these measures were meaningful. The modern conception of utility, due largely to von Neumann and Morgenstern (1947), is immune to these particular criticisms. In particular, the "cardinal" utilities of von Neumann and Morganstern are, like the measurement of temperature, unique up to a positive linear transformation. Although none of the essays in this volume deal exclusively with the concept of utility, the importance of this concept continues to be evident throughout foundational work in the decision sciences (Krantz et al. 1971; Fishburn 1979; Keeney and Raiffa 1993).

Relevant foundational work on the other central factor, belief, is likely to be more familiar to philosophers. Much of the early work on expected utility theory was concerned with situations, such as gambling, where the decision maker has access to an "objective" probability distribution over the relevant states. It is not difficult to convince oneself that objective probabilities are not available to the decision maker in many, if not most, situations. The applicability of the recommendation to maximize expected utility is increased considerably by allowing for "subjective" probabilities, where these subjective probabilities are usually interpreted as a representation of the decision maker's degrees of belief. Following a tradition that dates at least from the seminal works of Ramsey (1931) and de Finetti (1964) to the work of Savage (1972), the received account of decision making under uncertainty requires the rational agent to select an alternative that maximizes his or her subjective expected utility. Insofar as the decision maker has a subjective probability over the relevant states, along with a cardinal utility on the set of outcomes, the relevant expected utility calculations are well-defined: of course, the practical difficulties involved with such calculations might be substantial, but typically these difficulties are acknowledged as a separate issue.

Although the use of subjective probabilities extends the applicability of the expected utility hypothesis, there are difficulties associated with the use of these probabilities in normative accounts of decision making. A central difficulty of this sort derives from the fact that several agents, similarly positioned with respect to the evidence that is available to them, may have vastly different subjective probabilities. Similarly, nothing in this minimal standard demands that the agent be responsive to information

[1] See Sect. 5.6 of Savage (1972) for a nice discussion concerning the evolution of utility theory.

concerning objective probabilities. Quite a bit of work has been done on the problem of how to further constrain this standard beyond the modest requirements of coherence that are demanded by the probability axioms. The relevant literature is wide and deep, and this is not the place for a survey of these developments.[2]

However, at least three of the papers in this volume deal directly with issues of this sort. In 'Bayesian Probability,' Patrick Maher argues that the usual interpretation of subjective probabilities, i.e. as a representations of the agent's degrees of belief, is not adequate in the context of rational choice and, moreover, that there is an alternative interpretation, one informed by the notion of inductive probability, that is well-suited to this context. In 'On the Definition of Objective Probabilities by Empirical Similarity,' the trio of Itzhak Gilboa, Offer Lieberman and David Schmeidler use similarity-weighted empirical frequencies as the basis for an account of objective probabilities that extends the traditional frequentist approach in ways that are congenial to decision making. As noted above, the subjectivist position provides for the possibility that distinct agents who are supplied with the same evidence may differ with respect to their degrees of belief in light of this common evidence, and, in particular, the subjectivist position does not discriminate according to these differences so long as each of the agents satisfies the requirements that are expressed in the probability axioms. In contrast, theorists such as Jeffreys, Keynes, Carnap, and Kyburg have argued that probability should be determined by the agent's evidence and certain logical principles. In his essay 'Probability Logic, Logical Probability, and Inductive Support,' Isaac Levi calls into question some of the central claims that have been advanced on behalf of the logicist position.

As we have noted, the introduction of subjective probabilities extends the domain of potential applications of the expected utility hypothesis. Following the well-known taxonomy that is presented by Luce and Raiffa (1989), situations in which the decision maker has access to objective probabilities are instances of decision making under risk while decision situations in which such access is lacking are instances of decision making under uncertainty. In light of the extended applicability of expected utility theory that is afforded by the introduction of subjective probabilities, it is sometimes maintained that uncertainty is reducible to risk. The insights of theorists such as Keynes (1921), Knight (1921), and Ellsberg (1961, 2001) have called this reduction into question. Empirical evidence concerning simple, yet compelling, examples such as Ellsberg's three-color problem suggest that a significant number of people deviate from the norms of Savage's subjective expected utility theory (SEU). Generally, this evidence has been taken as motivation for constructing theories of decision making under uncertainty that are more descriptively adequate than SEU. However, since among those who deviate from these norms are sophisticated decision makers who are well aware of the fact that their behavior is incompatible with at least one of the axioms of SEU, others have interpreted some of these violations as evidence of SEU's inadequacy as a normative theory of decision making under uncertainty: this was the interpretation that was taken by Ellsberg in response to some of the violations of SEU that are associated with his now-famous examples. A variety of normative

[2] The interested reader might want to begin by looking at some of the work on principles of direct inference (Kyburg 1974; Levi 1977).

alternatives to SEU have been proposed. Among the most distinctive of these is the one that has been suggested by Levi (1974, 1986). This theory is unusual in that it drops the assumption that the rational decision maker has a complete preference ranking over the set of alternatives. In 'Coherent Choice Functions under Uncertainty,' the trio of Teddy Seidenfeld, Mark Schervish and Joseph Kadane axiomatize a special case Levi's decision theory.

Thus far we have been focusing on the study of single-agent decision making. However, the study of groups of interacting decision makers constitutes another major component of the decision sciences. Like single-agent decision theory, the study of groups of interacting agents, commonly known as game theory, traces a significant amount of its contemporary inspiration to the seminal work of von Neumann and Morgenstern. Despite this common ancestry, contemporary game theory draws upon ideas that are in many ways quite distinct from those that are central to contemporary work on single-agent decision making. The concept of common knowledge provides an outstanding example of such an idea (Lewis 1969; Aumann 1976). More generally, a significant amount of contemporary game theory overlaps with traditional concerns in epistemology. The relational semantics for epistemic logic, familiar to philosophers through the work of Hintikka (1962), are closely related to the "event-based" models that are familiar to economists through the work of Aumann (1976). In 'S5 Knowledge without Partitions,' Dov Samet examines the familiar S5 conditions for knowledge in the context of set algebras that are equipped with an additional operator.

While Samet's paper has clear connections to the "interactive epistemology" tradition in contemporary game theory (Aumann 1999a, b), 'Evolutionary Dynamics of Lewis Signaling Games: Signaling Systems vs. Partial Pooling' by Simon Huttegger, Brian Skyrms, Rory Smead, and Kevin Zollman is a contribution to the subfield known as evolutionary game theory. This subfield, which has its origins in attempts to apply game theory to biological systems (Lewontin 1961; Maynard-Smith 1982), has developed into an area that is of interest to researchers in a variety of fields including philosophy. In their contribution to this volume, Huttegger, Skyrms, Smead, and Zollman analyze the evolutionary aspects of the "signaling games" that were discussed by Lewis in the context of his well-known study of conventions (Lewis 1969). The third paper in this volume that concerns groups of decision makers is 'The Puzzle of the Hats' by Luc Bovens and Wlodek Rabinowicz. To be sure, several of the papers in this volume touch on multiple themes that are relevant to the foundations of the decision sciences. This multiplicity is perhaps most clear in this work by Bovens and Rabinowicz, which presents an analysis of the Puzzle of the Hats that promises to have implications for areas ranging from social choice to recent work on the Sleeping Beauty problem.

The distinction between the study of single-agent decision making and the study of multi-agent decision making is just one of many ways to classify work in the decision sciences. Another equally important classification distinguishes between normative and descriptive work. Although the division between these classes is not as sharp as is sometimes supposed, at least one would hope that normative theories are somewhat descriptive of our best practices, the papers considered above are mainly focused on normative matters. In contrast, the three remaining papers in this volume address, at least in part, several issues that concern descriptive accounts of decision making. The

"heuristics and biases" program, which is most clearly associated with the work of Daniel Kahneman and Amos Tversky, has informed much of the contemporary work of this sort (Tversky and Kahneman 1974). Roughly, the basic thesis of this program is that human decision makers rely on certain heuristics and that this results in systematic biases that are realized as predictable deviations from what is required by prevailing normative accounts. In 'Cognitive Biases in Moral Judgments that Affect Political Behavior,' Jonathan Baron explores this sort of program in the context of political decision making. While work in the heuristics and biases tradition is often interpreted as painting a bleak picture of human rationality, some of the central examples of non-normative behavior that are cited in this tradition can also be taken as evidence that the relevant normative theories need to be reexamined with respect to their normative status. In 'Ambiguity Aversion: the Explanatory Power of Indeterminate Probabilities,' Horacio Arlo-Costa and Jeffrey Helzner examine the possibility of explaining *ambiguity aversion*—which is perhaps the most well-known psychological effect that is invoked to explain deviations from subjective expected utility theory in Ellsberg problems—in terms of other normative accounts that are based on imprecise probabilities. Finally, in response to the non-normative behavior that has been documented in psychological studies of decision making, some theorists have sought to locate a safe-haven for Normative Man within a "dual process" theory. In 'Decision Science: From Ramsey to Dual Process Theories,' the trio of Nils-Eric Sahlin, Annika Wallin, and Johannes Persson present a critical analysis of these dual process theories.

References

Aumann, R. J. (1976). Agreeing to disagree. *Annals of Statistics, 4*(6), 1236–1239.

Aumann, R. J. (1999a). Interactive epistemology i: Knowledge. *International Journal of Game Theory, 28*(3), 263–300.

Aumann, R. J. (1999b). Interactive epistemology ii: Probability. *International Journal of Game Theory, 28*(3), 301–314.

de Finetti, B. (1964). Foresight: Its logical laws, its subjective sources. In: H.E. Kyburg & H. Smokler (Eds.), *Studies in subjective probability*. London: Wiley. (Translation of the 1937 paper).

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

Ellsberg, D. (2001). *Risk, ambiguity and decision*. New York: Garland Publishing.

Fishburn, P. (1979). *Utility theory for decision making*. NY: Krieger.

Gardenfors, P. & Sahlin, N. (1988). *Introduction: Bayesian decision theory—foundations and problems*. In: P. Gardenfors & N. Sahlin (Eds.), *Decision, probability, and utility*. London: Cambridge University Press.

Hintikka, J. (1962). *Knowledge and belief*. New York: Cornell University Press.

Keeney, R. L. & Raiffa, H. (1993). *Decisions with multiple objectives*. Cambridge University. (The Cambridge edition is a republication of the 1976 work).

Keynes, J. M. (1921). *A treatise on probability*. NY: MacMillan.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Boston: Houghton-Mifflin.

Krantz, D. H., Luce, R. D. Suppes, P. & Tversky, A. (1971). *Foundations of measurment: Additive and polynomial representations* (Vol. 1). New York: Academic Press.

Kyburg, H. E. (1974). *The logical foundations of statistical inference*. Dordrecht: Reidel.

Levi, I. (1974). On indeterminate probabilities. *Journal of Philosophy, 71*, 391–418.

Levi, I. (1977). Direct inference. *The Journal of Philosophy, 74*, 5–29.

Levi, I. (1986). The paradoxes of allais and ellsberg. *Economics and Philosophy, 2*, 23–53.

Lewis, D. (1969). *Convention: A philosophical study*. Cambridge: Harvard University Press.

Lewontin, R. C. (1961). Evolution and the theory of games. *Journal of Theoretical Biology, 1*, 382–403.

Luce, R. D. & Raiffa, H. (1989). *Games and decisions: Introduction and critical survey*. NY: Dover Publications. (The Dover edition is a republication of the 1957 work).

Maynard-Smith, J. (1982). *Evolution and the theory of games*. London: Cambridge University Press.

Ramsey, F. P. (1931). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). London: Routledge and Kegan Paul.

Savage, L. J. (1972). *The foundations of statistics*. NY: Dover Publications. (The Dover edition is a republication of the 1954 work).

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*, 1124–1131.

von Neumann, J. & Morgenstern, O. (1947). *Theory of games and economic behavior*. NJ: Princeton University Press.