Decision Theory

Itzhak Gilboa

May 2009

Historical background

People have been consciously coping with decision problems for at least as long as recorded history exists. Some of the best advice modern theory can offer appears to have been known to the ancients. For example, the biblical Jacob, fearing his brother Esau's might, divides his camp into two bands, saying, "If Esau come to the one company and smite it, then the other company which is left shall escape" (Genesis, 32:8) – thereby holding a claim to having invented diversification.

Yet, decision theory, and, in particular, decision under uncertainty, has not been explicitly studied until the mid-17th century, when probability theory has been ushered. Interestingly, the person who is most associated with the concepts of probability and expectation, Blaise Pascal, also introduced decision theory in his famous "wager" (see Hacking, 1975). In this argument, designed to convince non-believers that they will be better off becoming believers, Pascal introduced several basic notions of decision theory: (i) the decision matrix, in which one's acts are independent of Nature's choices, or the "states of the world"; (ii) domination between acts, where one act is better than another no matter which state obtains; (iii) expected utility maximization, according to which the choice between undominated acts should be according to the mathematical expectation of the utility of the outcomes they yield; (iv) subjective probability over the states, which is an application of the mathematical probability model as a way to capture one's beliefs; and (v) non-unique probabilities, where one's beliefs are too vague to be captured by a single probability vector. But even after this dramatic inauguration, decision theory has been largely neglected for the next three centuries, with the exception of Daniel Bernoulli's (1738) explicit introduction of the expected utility hypothesis applied to monetary payoff and given probabilities.

By contrast, mathematicians and philosophers were interested in the notion and the mathematical theory of probability throughout the past centuries. Jacob Bernoulli (1713) discovered the law of large numbers, and also discussed different types of probability. (See Shafer, 1986). Thomas Bayes (1763) introduced the idea of Bayesian updating of "prior" probabilities to "posterior" ones.

Axiomatic foundations

Whether all uncertainty can be quantified probabilistically has remained a topic of dispute from the very early writings on probability to this very day. Frank Knight (1921) famously argued that this is not the case, and he distinguished between situations of "risk", where probabilities can be assumed given, and situations of "uncertainty", where probabilities are neither given, nor can they be inferred from past statistical data. A major proponent of the opposite view was Frank Ramsey (1931), who, in the spirit of logical positivism suggested defining and measuring one's subjective probability by one's willingness to bet. He suggested that a reasonable decision maker will behave *as if* she had a subjective probability that guided her decisions, even if objective probabilities are not part of the description of the decision problem. Bruno de Finetti (1937) offered an "axiomatization" of subjective probabilities in the context of maximization of expected monetary value. That is, he provided a set of conditions on presumably observable choices and showed that they are equivalent to the claim that the decision maker maximizes expected value relative to *some* probability vector, which is taken to be that person's subjective probability. The conditions are called "axioms" partly because they are presented as intuitive, if not compelling, and partly because they are in line with the logical positivist dictum of defining theoretical concepts ("subjective probability") with observations (pairwise choices between bets).

A similar axiomatic derivation of the concept of "utility" was offered by John von Neumann and Oskar Morgenstern, as a by-product to their introduction of game theory in "Games and Economic Behavior" (1944). They considered a presumably-observable preference relation between pairs of "lotteries", namely random variables with known distributions, and they showed that a set of axioms on the relation is equivalent to the

claim that this relation can be represented by a utility function, such that, confronted with any two choices, the decision maker would opt for the one that has a higher expected utility.

von Neumann and Morgenstern provided a definition of utility, coupled with the expected utility paradigm, based on a primitive notion of probability. de Finetti did the opposite: defined subjective probability based on a primitive notion of "utility". However, when neither utility nor probability is well defined, it is not obvious that the theory of expected utility maximization relative to a subjective probability is very convincing, whether interpreted descriptively or normatively. This problem was rectified by Savage (1954), who showed that both utility and subjectively probability can be derived, with the expected utility maximization rule, from basic axioms on "acts", which are not defined numerically, and pre-suppose neither probabilities nor utilities.

The axiomatic results of von Neumann and Morgenstern and of Savage had a tremendous impact on research in decision theory, game theory, and in their applications in the social sciences. The mainstream view among theorists is that expected utility maximization, with respect to a subjective probability, is the only rational way of behavior. Moreover, many also believe that it is the only reasonable model to be used in applications where a formal model attempts to describe reality. However, the theory has been challenged, mostly, but not solely, from a descriptive viewpoint.

Challenges

One of the earliest, and perhaps the most radical objection to the theory was raised by Herbert Simon (1957). He coined the term "bounded rationality", and argued that people do not optimize; rather, they "satisfice": as long as their performance is above a certain "aspiration level", they stick to their previous choice. Only when their performance is below that threshold do they experiment with other choices. Simon thus challenged the very paradigm of optimization. While his theory is relatively seldom incorporated into formal decision models, it has had a remarkable impact on the thinking of many decision theorists, who developed models that are classified as "bounded rationality" even if their departure from the basic paradigm is much less dramatic than that of satisficing behavior.

Expected utility maximization was also attacked based on concrete examples in which it turned out to provide a poor prediction of people's choices. Maurice Allais (1953) provided a "paradox" in the context of decision under risk (with known probabilities). In this example many people violate a key axiom of von Neumann and Morgenstern (the "Independence Axiom)", and therefore behave in a way that cannot be captured by expected utility maximization (for *any* utility function). Allais's example showed that people tend to put more weight on certainty than the standard theory predicted. In other words, people tend to behave in a way that is non-linear in probabilities. Daniel Ellsberg (1961) proposed examples (also dubbed "paradoxes"), where many people violate one of Savage's basic axioms (the "Sure Thing Principle"). In Ellsberg's examples many people behave in a way that cannot be described by subjective probability at all. Specifically, people tend to prefer situations with known probabilities to situations with unknown probabilities. This phenomenon is referred to as "uncertainty aversion", or "ambiguity aversion" (following Knight's and Ellsberg's terms, respectively).

Starting in the late 1960s, Daniel Kahneman and Amos Tversky launched a systematic experimental study of decision theoretic axioms. In carefully designed experiments, they have shown that practically any axiom of decision theory is violated in some examples. Importantly, Kahneman and Tversky also uncovered several implicit assumptions of the decision theory, which were also too idealized to describe actual choices. For example, they documented the "framing effect" (Tversky and Kahneman, 1974), which shows that different representations of the same problem may result in different choices. Kahneman and Tversky (1979) suggested "Prospect Theory" as an alternative to expected utility maximization.

One key idea in Prospect Theory is that people respond to given probabilities in a way that is non-linear in the probability, especially near the extreme values of 0 and 1. Another idea, with potentially far-reaching implications to research in political science, is that people react differently to gains as compared to losses. That is, the (monetary) bottom-line is not all that matters to the decision maker: it also matters whether this bottom-line is perceived as a gain or a loss relative to a "reference point" that the decision maker has in her mind.

Main directions

There are many phenomena that the classical theory of expected utility maximization relative to a subjective probability cannot explain. Which are the most important? Specifically, if we take the viewpoint that formal models in political science should provide basic insights to help us think about real-world phenomena, which assumptions of the classical theory should we relax first? The answer to this question is inevitably subjective. I point here to two issues that I find fundamental, one having to do with the utility function, and the other – with the notion of probability.

There are at least three distinct ideas in psychology that suggest that decisions might be better understood if we take into account the special role of a certain point on the utility scale. As mentioned above, Simon's "aspiration level" was such a special point, serving as the threshold between experimentation and "business as usual". Harry Helson (1947) suggested that people perceive sensory inputs relative to a given "adaptation level", which, as the name suggests, adapts as a result of experiences. This idea has implications to the relationship between utility and happiness, as well as to the measurement of well-being (see Philip Brickman and Donald Campbell, 1971). Finally, we also mentioned Kahneman and Tversky's "reference point", which was introduced to explain the distinction between gains and losses. The notion that a certain level of utility, whether interpreted as an "aspiration level", "adaptation level", or a "reference point", plays a special role in one's conduct is eminently intuitive. An everyday account of negotiation and conflict often involves people's expectations, goals, and hopes. While the standard theory may attempt to capture these concepts by one's subjective beliefs about the possible paths in which future interaction might unfold, it seems to be more parsimonious and more intuitive to summarize these aspirations by an "aspiration level". Aspiration levels are defined for each agent on their utility scale, which might implicitly reflect different subjective beliefs without devoiding equilibrium analysis of its content. Moreover, it appears very intuitive that aspiration levels adjust. For example, aspirations are lowered as a result of failures, and raised following successes (whether by the same agent or by others).

The notion that one can generate probabilistic beliefs over political events is doubtful. One hardly has sufficient information to "figure out" the probability of war, of a particular candidate winning the election two years from now, and so forth. Political events seem to be unique and intertwined: they are unique in the sense that no two events are the same, and relative frequencies can hardly be used to suggest a definition of "probability". They are causally intertwined in the sense that the very fact that a certain event has occurred is likely to change the probability of its recurrence. Hence, one cannot use past events to calculate probabilities of future ones as one might do when causal independence holds.

Decision theory suggests several alternative models to the classical one. Non-Bayesian models were proposed by Schmeidler (1989), Gilboa and Schmeidler (1989), Bewley (2002), and others. A model that is simple to explain is the maxmin expected utility model of Gilboa and Schmeidler. It suggests that agents entertain a set of probabilities, rather than a single one, and that when they have to evaluate an act, they compute the worse-case expected utility it can have, as long as the probability vectors is in the set. Within economics there are several applications of models of this type to finance, macroeconomics, labor economics, and so forth. (See Dow and Werlang, 1994, Epstein and Tang, 1994, 1995, Mukerji, 1998, Epstein and Miao, 2003, Nishimura and Ozaki, 2004, Hansen and Sargent, 2007.) In the context of international relations, Greenberg (2000) interprets a comment by Henry Kissinger on the strategic advantage of ambiguity as a Nash equilibrium in a three-player game, that cannot be supported as an equilibrium with Bayesian beliefs, but can be sustained with uncertainty averse players.

While it might appear that uncertainty averse players will always be less prone to conflict, one finds out that this is the case only in a single-period model. If one considers a second period, in which two players have uncertainty about each other's capacity (say, the ability to develop nuclear weapons), one may find that conflict may arise in the first period, where it could not arise in equilibrium if the players were Bayesian. Thus, models in which uncertainty averse players share the same uncertain (non-Bayesian) beliefs may lead to fundamentally different conclusions than models with Bayesian players who share their beliefs.

A more dramatic departure from the classical model has to do with models that are based on analogical reasoning. Gilboa and Schmeidler (1995) offer such a model, according to which acts are evaluated by their performance in similar past cases, rather than by an exhaustive list of all potential scenarios that might evolve.

Conclusion

Formal decision theory has been extremely powerful in providing important insights into the behavior of agents in social and political environments. Formal models help us analyzing real life situations and seeing analogies that would otherwise might be difficult to identify. At the same time, formal models have been justifiably criticized on various grounds. Some have to do with assumptions of decision theory per se, such as the existence of probabilistic beliefs, and some have to do with assumptions of related fields, such as the concept of equilibrium in games, which is not always compelling. It is important not to discard the powerful insights that formal analysis might generate on account of some assumptions that need to be refined or replaced. Future research will hopefully improve our understanding of political phenomena using formal models, while taking each assumption thereof with a grain of salt.

References

Allais, Maurice (1953), "Le Comportement de L'Homme Rationnel devant le Risque: critique des Postulats et Axiomes de l'Ecole Americaine", Econometrica, 21: 503-546.

Bayes, Thomas (1763), "An Essay towards solving a Problem in the Doctrine of Chances." Communicated by Mr. Price, Philosophical Transactions of the Royal Society of London, 53: 370-418.

Bernoulli, Daniel (1738), "Exposition of a New Theory on the Measurement of Risk", Econometrica, 22 (1954): 23-36.

Bernoulli, Jacob (1713), Ars Conjectandi. 1713.

Bewley, Truman (2002), "Knightian Decision Theory: Part I", Decisions in Economics and Finance, 25: 79-110. (Working paper, 1986)

Brickman, Philip, and Donald T. Campbell (1971), "Hedonic Relativism and Planning the Good Society". In M. H. Appley (Ed.), Adaptation level theory: A symposium. New York: Academic Press.

de Finetti, Bruno (1937), "La Prevision: Ses Lois Logiques, Ses Sources Subjectives", Annales de l'Institute Henri Poincare, 7: 1-68.

Dow, James, and Sergio R. C. Werlang (1992), "Uncertainty aversion, risk aversion, and the optimal choice of portfolio", Econometrica, 60:, 197-204.

Epstein, Larry G. and Jianjun Miao (2003), "A Two-Person Dynamic Equilibrium under Ambiguity", Journal of Economic Dynamics and Control, 27: 1253-1288.

Epstein, Larry G. and Tan Wang (1994), "Intertemporal asset pricing under Knightian uncertainty", Econometrica, 62: 283-322.

Epstein, Larry G. and Tan Wang (1995), "Uncertainty, risk-neutral measures and security price booms and crashes", Journal of Economic Theory, 67: 40-82.

Ellsberg, Daniel (1961), "Risk, Ambiguity and the Savage Axioms", Quarterly Journal of Economics, 75: 643-669.

Gilboa, Itzhak and David Schmeidler (1989), "Maxmin Expected Utility with a Non-Unique Prior", Journal of Mathematical Economics, 18, 141-153. (Working paper, 1986)

Gilboa, Itzhak and David Schmeidler (1995), "Case-Based Decision Theory", The Quarterly Journal of Economics, 110 605-639.

Greenberg, Joseph (2000), "The Right to Remain Silent", Theory and Decision, 48: 193-204.

Hacking, Ian (1975), The Emergence of Probability. Cambridge: Cambridge University Press.

Hansen, Lars Peter, and Thomas J. Sargent (2006), Robustness. Princeton: Princeton University Press, forthcoming.

Helson, Harry (1947). "Adaptation-Level as Frame of Reference for Prediction of Psychophysical Data", American Journal of Psychology, 60: 1-29.

Kahneman, Daniel and Amos Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," Econometrica, 47: 263-291.

Knight, Frank H. (1921), Risk, Uncertainty, and Profit. Boston, New York: Houghton Mifflin.

Lindley, Dennis (1965), Introduction to Probability and Statistics from a Bayesian Viewpoint. Cambridge: Cambridge University Press.

Mukerji, Sujoy (1998) "Ambiguity Aversion and the Incompleteness of Contractual Form", American Economic Review, 88: 1207-1232.

Nishimura, Kiyohiko and Hiroyuki Ozaki (2004), "Search and Knightian Uncertainty", Journal of Economic Theory, 119: 299-333.

Ramsey, Frank P. (1931), "Truth and Probability", in The Foundation of Mathematics and Other Logical Essays. New York, Harcourt, Brace and Co.

Savage, Leonard J. (1954), The Foundations of Statistics. New York: John Wiley and Sons. (Second addition in 1972, Dover…)

Schmeidler, David (1989), "Subjective Probability and Expected Utility without Additivity", Econometrica, 57: 571-587. (Working paper, 1982)

Shafer, Glenn (1986), "Savage Revisited", Statistical Science, 1: 463-486.

Simon, Herbert (1957), Models of Man. New York: John Wiley and Sons.

Tversky, Amos, & Kahneman, Daniel (1974), "Judgment under uncertainty: Heuristics and biases", Science, 185: 1124-1131.

von Neumann, John and Oskar Morgenstern (1944), Theory of Games and Economic Behavior. Princeton, N.J.: Princeton University Press.