CHAPTER 3

DECISION MAKING

INTRODUCTION

by
Glenn Shafer

This chapter is concerned with making decisions under uncertainty. The emphasis is on Bayesian decision theory, which recommends maximizing subjective expected utility, and on decision analysis, which uses decision trees, utility curves, multiattribute utility techniques, and influence diagrams to implement Bayesian decision theory.

Decision analysis has not been widely implemented in artificial intelligence, and many workers in AI consider its computer implementation impractical. Yet Bayesian decision theory has come to define what is meant by a general theory of decision making under uncertainty. Its elegance retains a hold on the imagination of scholars, and its familiarity makes it a reference point to which other theories must relate. Any general account of decision making must define itself relative to Bayesian decision theory. Any such account must explain why it does not do things that Bayesian decision theory does, or why it does some things differently.

Our first two articles are concerned with the justification of Bayesian decision theory and with the basics of decision analysis. The first article, by D. Warner North, reviews the basic ideas of subjective probability, utility curves, and decision trees. The second article, by Ross Shachter, is concerned with a more recent elaboration, influence diagrams.

Our third and fourth articles discuss the inadequacy of Bayesian decision theory as a theory of human behavior. The article by Amos Tversky and Daniel Kahneman reviews some of the ways in which people diverge from the theory when they are given simple choices. The article by Benjamin Kuipers and his colleagues discusses how even experts diverge from the prescription that decision trees or influence diagrams should be analyzed in detail before any decisions are made.

In our fifth article, I review Savage's postulates for subjective expected utility and challenge the claim that it is normative for people to follow these postulates. In our last article, Edward Shortliffe reviews the area of medical decision support systems.

A huge literature is available to readers who want to learn more about decision analysis. Among books by advocates, the best is still Howard Raiffa's classic *Decision Analysis* (1968). A more recent introduction, with a wider variety of examples, is Behn and Vaupel's *Quick Analysis for Busy Decision Makers* (1982). The readings edited by Elster (1986) and by Gärdenfors and Sahlin (1988) provide a critical philosophical perspective.

The remainder of this introduction will be concerned with the historical development of Bayesian decision theory and decision analysis, their relevance to AI, their descriptive and normative claims, and their place in the context of other decision methods.

1. The Historical Development of Bayesian Decision Theory and Decision Analysis

Bayesian decision theory and decision analysis have come into being since the 1950s, but the most basic ideas, including decision trees and utility elicitation, derive from work by the frequentist school of probability in the 1930s and 1940s.

These frequentist roots lie primarily in John von Neumann and Oskar Morgenstern's work on game theory and Abraham Wald's work on statistical decision theory. The decision tree is a special case of von Neumann and Morgenstern's game tree, and the methodology for constructing utility curves derives from von Neumann and Morgenstern's axiomatization of utility. The Bayesian treatment of the decision tree is essentially due to Wald.

Wald had set out in the late 1930s to extend Neyman and Pearson's frequentist theories of hypothesis testing and confidence intervals to a unified theory of statistical decision. The task of mathematical statistics, in Wald's mind, was to choose among tests, estimates, or confidence interval procedures. This required that costs be made explicit. Neyman and Pearson had evaluated procedures for accepting or rejecting hypotheses in terms of the frequency with which these procedures would make errors of different types. Wald emphasized that we can actually choose between procedures only when the relative costs of the different errors are also specified.

When von Neumann and Morgenstern's *Theory of Games and Economic Behavior* appeared in 1944, Wald recognized the relevance of game theory to his own work. Statistical decision problems, Wald realized, were games between a statistician and nature, and the sequential nature of statistical observation was a special case of the sequential nature of two-person games.

By the time of his death in an airplane accident in 1950, Wald's influence had pervaded mathematical statistics. He had convinced the majority of mathematical statisticians that their enterprise was decision theory, not inference. This decision theory was frequentist, but it gave a technical role to Bayesian ideas. One of the most striking theorems in Wald's last book, *Statistical Decision Functions* (1950), was the complete class theorem, which says that any procedure in a statistical decision problem can be beaten or at least matched in performance by a Bayesian procedure, a procedure based on the adoption of some set of prior probabilities for hypotheses (Ferguson 1967). But for Wald and his colleagues, such prior probabilities were simply a technical tool for the generation of procedures. They were not taken seriously as expressions of belief.

Wald's work was very mathematical, but after his death his ideas were developed and popularized by many able writers, including Blackwell and Girshick (1954), Luce and Raiffa (1957), and Raiffa and Schlaifer (1961). After the publication of L.J. Savage's *Foundations of Statistics* in 1954, many of these writers became increasingly sympathetic to subjectivist ideas and began to interpret Wald's Bayesian sequential decision procedures in subjectivist terms.

The next step in this evolution was to interpret not only prior probabilities, but also statistical models in subjectivist terms. In the statistical context, debate is usually about the legitimacy of subjective probabilities for hypotheses; both Bayesians and non-Bayesians tend to accept frequentist interpretations for the probabilities in the statistical models themselves. By giving a subjective interpretation to these probabilities as well, authors such as Raiffa (1968) and North were able to claim a range of application for their ideas far broader than the range of application for statistical methods.

In the late 1960s, when North published the article reprinted here, Bayesian decision theory was only beginning to break away from its roots in statistical decision theory. Since then, it has been enriched by a number of new techniques, including multiattribute utility assessment and influence diagrams. The term “decision analysis” is generally used to refer to the methodology that uses decision trees along with these newer techniques.

Multiattribute utility assessment, developed by Ralph Keeney and Howard Raiffa on the basis of traditional economic ideas about utility, is intended to help people make tradeoffs between different objectives. The tradeoffs in the decision problems discussed by North are mainly between value and certainty, but many problems involve tradeoffs between competing values, and their solution requires techniques for getting people to specify in simple ways how they want to make these tradeoffs. Readers interested in multiattribute utility may consult Keeney and Raiffa (1976) or von Winterfeldt and Edwards (1986).

Influence diagrams were developed by Ronald Howard, James Matheson, and their colleagues at Stanford in the 1970s, as part of an effort to automate the construction of decision trees. Since it provides an economical representation of part of the information that might be contained in a very large decision tree, an influence diagram can serve as a starting point for a person constructing a decision tree. Howard and Matheson envisaged a person using an influence diagram as a way to tell a computer how to construct a decision tree.

In fact, the influence diagram can replace the decision tree altogether. When a decision tree is too complex for us to organize directly in our heads, we may not want to look at it even if a computer draws it for us. It will be too large and bushy to be illuminating. The influence diagram itself is a more succinct visual representation, and as Shachter explains in his article in this chapter, the decision tree can be bypassed in the process of providing the further information for the analysis and making the expected utility computations. (For a more mathematical treatment of the methods discussed by Shachter, see Smith (1989).)

In the simple examples discussed by North and Raiffa, decision trees can be drawn in full, and the visual power of this representation contributes significantly to the persuasiveness of the analysis. In more complex examples, however, the tree representation can be impractical and wasteful. If the probabilities for outcomes at a certain step do not actually depend on all the decisions and observations made earlier, then the branching that displays these probabilities will be repeated over and over in the tree, once for each way the irrelevant decisions and observations can come out.

As we saw in Chapter 2, the irrelevance of some variables to others is a relation of conditional independence. An influence diagram represents variables and decisions by nodes, and it represents these relations of conditional independence by omitting arcs between variables. Thus it is closely related to the diagrams discussed by Pearl, Geiger, and Verma in Chapter 2. It differs from these diagrams only in the special treatment given to decision nodes. Moreover, Shachter's algorithm for evaluating an influence diagram is related to the algorithms discussed by Pearl and by Lauritzen and Spiegelhalter in Chapter 6. See Pearl (1988, pp. 306-311).

2. The Social Context of Decision Analysis

Since the 1960s, decision trees have become a standard part of the curriculum in United States business schools, and they have penetrated into many other disciplines as well. Many business consultants make their living applying decision analysis. Yet only a tiny fraction of the major decisions made by businesses and other organizations actually use formal decision analysis. Why is it used for some problems and not for others?

Accounts of successful decision analyses by Keeney and Raiffa (1976) and von Winterfeldt and Edwards (1986) make it clear that the success of decision analysis as a technique is closely tied to the social role of the decision analyst. Organizations usually appeal to consultants when their own decision processes are not working—i.e., when decision makers are divided or are unwilling to act because they fear negative reactions to any choice they make. The decision analyst, in the most successful cases, is able to divert attention from potential or real conflicts to a more technical setting, where everyone can participate in constructing a solution.

In such situations, a decision tree can embody compromises on what values are to be emphasized, what choices are to be considered possible, and what weights are to be given to contingencies feared or hoped for by different constituencies. Once it is drawn, such a tree is also a tool of persuasion. Because it represents visually a large amount of information about tradeoffs and uncertainties, it can help persuade adversaries that all their competing concerns are being addressed.

Not all situations permit mediation by a decision analyst. Many compromises are based on ambiguity and diffusion of responsibility rather than on clarity about alternatives. This may be one reason for the relatively rare use of decision analysis.

3. Computer Implementation

To what extent have Bayesian decision theory and decision analysis been implemented by computer? This question is difficult to answer because it raises many questions of definition. If we use a broad definition of Bayesian decision theory, we will find many implementations, most more frequentist than subjectivist in spirit. On the other hand, we will find few complete implementations of the art of decision analysis, the art of helping people construct decision trees and eliciting utilities and probabilities for those trees.

Among implementations of Bayesian decision theory, there are many systems that implement decisions automatically (e.g., in navigation, chemical engineering, manufacturing, and pattern recognition), as well as systems that give advice. Most of these systems use estimated or observed frequencies rather than subjective beliefs. As one example, we can cite the Kalman filter, and its use in aerial navigation. Since its development in the 1960s, this Bayesian technique has been incorporated into the navigational and guidance systems of almost all transoceanic aircraft (Spall 1988). The probabilities that it uses are usually based on direct experience, however, and in this sense the system is frequentist.

Since most computer systems that use statistical information have been developed outside the AI community, they are often not called expert systems in the computer science literature. The distinction is breaking down, however, as more and more of these systems are coupled with rule-based systems, planning systems, explanation systems, or knowledge-acquisition systems that have recognizable AI pedigrees.

The art of decision analysis has not lent itself so readily to computer implementation. Samuel Holtzman (1989) has discussed some of the difficulties involved. Of particular importance is the difficulty a computer has in playing the social role of the decision analyst. In most situations, the point of constructing a decision tree is to build consensus, and the technical aspects of the analysis are secondary to the analyst's ability to craft this consensus. It accomplishes nothing for a computer to build a decision tree all by itself.

Decision analysts use a number of computer decision aids, however. They have long used computers to carry out expected utility computations and to perform sensitivity analyses. More recently, they have developed computer aids to help decision makers construct preferences (Humphreys and McFadden 1980), generate explanations (Langlotz, Shortliffe, and Fagan 1988), and carry out other general or problem-specific tasks (von Winterfeldt and Edwards 1986).

Many of the older computer decision aids show little influence from AI, but this is changing. The article in the next chapter by Curtis Langlotz and his colleagues provides a remarkable example of a coupling of AI and decision analysis. The authors of this article advocate using planning methods to construct decision trees, which are then evaluated using a combination of statistical information and clinical judgment.

Decision theory is also beginning to play a more fundamental role in other areas of AI. In the late 1970s and early 1980s, there was considerable work on constructing decision trees using game searching (Leal and Pearl 1977; Pearl, Leal, and Saleh 1982; Kim 1977). Now, in contrast, we are seeing use of Bayesian decision theory to improve game searching (Russel and Wefald 1989).

A comprehensive survey of recent efforts to use decision analysis in AI is provided by Horvitz, Breese, and Henrion (1988).

4. Decision Support Systems

Since the mid-1970s, management scholars have discussed *decision support systems*—computer systems that would allow executives to access, organize, and analyze the scattered information needed to improve decisions (Keen and Scott-Morton 1978). In recent years, with the explosion in the power of personal computers, a number of such systems have been implemented; Sprague and Watson (1986) provide a relatively up-to-date bibliography.

Some scholars have seen a facility for Bayesian decision theory as an integral part of a decision support system. Most decision support systems put much greater emphasis, however, on the management of information. This is clearly justified. Making a decision once one has the right information is often the easy part. Even when computation is required, the appropriate computational technique may be an optimization technique, such as linear programming, rather than expected utility.

In his article in this chapter, Edward Shortliffe looks at the problem of computer decision support in medicine. Here, as in the case of business, the management of information seems to be most important. In medicine, however, there is more repetitive decision making, and more cooperation between organizations faced with similar decisions. So the management of statistics and their combination through Bayesian decision theory may be an important aspect of decision support in medicine.

5. The Descriptive Inadequacy of Subjective Expected Utility

As we learned in Chapter 2, psychologists have looked in detail over the past few decades at the ways in which people deviate from the rules of probability theory in their judgments of probability and from Bayesian decision theory in their choices and decisions. It has become clear from this research that the deviations are substantial. People do not use subjective expected utility to make most of their everyday decisions, and often their choices cannot be reconciled with any ranking by subjective expected utility.

How do everyday choices and decisions deviate from subjective expected utility? Among the many deviations that can be identified, one of the most fundamental is the occurrence of framing effects. As Tversky and Kahneman explain in their article in this chapter, choices and preferences are strongly influenced by the way in which a problem is framed. If a problem is described in terms of possible gains, for example, people tend to be risk averse, but if the same problem is described in terms of losses relative to a possible maximum gain, they tend to be risk seeking. If people used pre-defined subjective probabilities and utilities to make choices, they would not be affected in this way by the description of the problem.

These deviations from Bayesian decision theory seem unavoidable. The choices that we can be called on to make are too numerous and diverse for us to settle on relevant utilities and probabilities in advance. We often lack appropriate evidence to justify probabilities, and we often lack the experience that might inspire utilities. Even if we want to compute subjective expected utilities, and even if we have the computational means to do so, the fact that we do not have probabilities and utilities stored in advance means that our results will be sensitive to the way the problem is framed. Since our search for evidence or analogies on which to base probabilities and utilities must begin somewhere, we will surely exploit the clues provided by the framing of the problem.

This chapter's article by Benjamin J. Kuipers and his colleagues illuminates further how even experts deviate from Bayesian decision theory. As these authors show, physicians, even when they have reasonably good ingredients for a subjective expected utility computation, prefer to make a decision step-by-step, in a way that assigns distinct reasons to each salient aspect of the decision. Their decision process looks more like successive refinement of an abstract plan, combined with opportunistic insertion of plan steps, than like the formulation of a decision tree followed by computation. Whereas decision-tree computations must be made working backward from the details, physicians prefer to settle larger questions before moving on to details.

Decision trees can show us how to improve our decisions in particular cases, but there is a sense in which the strategy of successive refinement, in which one makes broader choices first and then refines these choices, is unavoidable. Any analysis, whether it uses decision trees or some other technique, must be made at a limited level of detail. There will always be further choices to make later. In the problem considered by Kuipers, for example, there will be decisions about how to perform open lung biopsy, about how much amphotericin to administer, and so on. In theory, the decision tree could and should be enlarged to consider these subtleties. In practice, we must leave some bridges to cross when we get to them.

6. The Normative Claims of Subjective Expected Utility

Most decision analysts readily acknowledge that people's unaided decisions usually do not conform to subjective expected utility. Following Savage, they argue that subjective expected utility is an improvement over everyday human deliberation. It leads to more rational and generally better decisions. Its use is normative.

Adversaries of subjective expected utility have repeatedly challenged Savage's claims for the normativeness of his axioms. My article in this chapter reviews and amplifies these challenges. I emphasize Savage's first axiom, the requirement that a person should have definite preferences among many risky actions. A close look at Savage's framework reveals that these actions include many that are obviously undesirable and many that are not feasible. A person might reasonably decline to make the effort required to rank such uninteresting and imaginary actions.

Savage's argument also depends on choosing a level of detail—a small world, as Savage called it. But as we saw in connection with Kuipers' article, we can always repeat the analysis with more detail, with possibly different results. Presumably a finer analysis is better, but we must stop somewhere. Savage's claim that it is normative to rank all imaginable alternatives at a given level of detail is in conflict with our intuition that more might be gained by probing at a finer level in some promising directions. It seems analogous to a claim, which no one would make, that it is normative for searches to always be breadth-first.

An alternative to Savage's claims is provided by the constructive interpretation advanced by Tversky and me in Chapter 2. According to this interpretation, assessing subjective probabilities and utilities amounts to drawing an analogy between our situation and the situation of a gambler in a game of chance. This analogy may or may not be persuasive.

One important element of the persuasiveness of the analogy is the availability of statistical information on which to base the probabilities. The most successful uses of decision analysis turn out to be those in which such information is available. In this case, of course, the distinction between Bayesian decision theory and frequentist statistical methods disappears; we are simply talking about the use of statistics.

7. Alternatives to Subjective Expected Utility

If it is not always normative to use subjective expected utility when making decisions under uncertainty, then what alternatives should we consider?

Several alternatives are mentioned in the articles in these readings. The article by Kuipers and his colleagues in this chapter advances AI planning as one alternative. The next chapter includes several articles that emphasize the role of the larger architecture of a system in decision making under uncertainty. In some cases, such an architecture may handle numerical judgments like probabilities and numerical scores akin to utilities. But in other cases, as in Rodney Brooks' architecture for a mobile robot, the explicit representation of uncertainty and value may be replaced by a hierarchical system of control, with more reliable information or more important goals being given priority by their position in the control system rather than by numerical scoring.

Outside of AI, within the decision science literature itself, there are also many proposals for alternatives to subjective expected utility. Some scholars have proposed generalizing subjective expected utility by weakening some of the axioms (Fishburn 1988). Others (e.g., Margolis 1982) have proposed alternatives to multiattribute utility functions for managing tradeoffs.

References

Behn, R.D., and J.W. Vaupel (1982). *Quick Analysis for Busy Decision Makers*. Basic Books.

Blackwell, David, and M.A. Girshick (1954). *Theory of Games and Statistical Decisions*. Wiley.

Elster, Jon, ed. (1986). *Rational Choice*. Basil Blackwell, Oxford.

Ferguson, Thomas S. (1967). *Mathematical Statistics: A Decision-Theoretic Approach*. Academic Press.

Ferguson, Thomas S. (1976). Development of the decision model. Pp. 335-346 of *On the History of Statistics and Probability*, Donald B. Owen, ed., Marcel Dekker.

Fishburn, Peter C. (1988). Normative theories of decision making under risk and under uncertainty. Pp. 78-98 of *Decision Making*, David E. Bell, Howard Raiffa, and Amos Tversky, eds., Cambridge University Press.

Gärdenfors, Peter, and Nils-Eric Sahlin, eds. (1988). *Decision, Probability, and Utility*. Cambridge University Press.

Holtzman, Samuel (1989). *Intelligent Decision Systems*. Addison-Wesley.

Horvitz, Eric J., John S. Breese, and Max Henrion (1988). Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning* **2** 247-302.

Howard, Ronald A., and James E. Matheson (1984). *The Principles and Applications of Decision Analysis*. Two volumes. Strategic Decisions Group, Menlo Park, California.

Humphreys, P.C., and W. McFadden (1980). Experiences with MAUD: Aiding decision structuring versus bootstrapping the decision maker. *Acta Psychologica* **45** 51-69.

Keen, Peter G. W., and Michael S. Scott-Morton (1978). *Decision Support Systems: An Organizational Perspective*. Addison-Wesley, Reading, Massachusetts.

Keeney, Ralph L., and Howard Raiffa (1976). *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York.

Kim, J. H. (1987). CONVINCE: A CONVersational INference Consolidation Engine. *IEEE Transactions on Systems, Man, and Cybernetics* **17(2)** 120-132.

Langlotz, Curtis P., Edward H. Shortliffe, and Lawrence M. Fagan (1988). A methodology for generating computer-based explanations of decision-theoretic advice. *Medical Decision Making* **8** 290-303.

Leal, A., and J. Pearl (1977). An interactive program for conversational elicitation of decision structures. *IEEE Transactions on Systems, Man and Cybernetics* **SMC-7(5)** 368-767.

Luce, R. Duncan, and Howard Raiffa (1957). *Games and Decisions*. Wiley, New York.

Margolis, Howard (1982). *Selfishness, Altruism, and Rationality: A Theory of Social Choice*. University of Chicago Press.

Pearl, Judea (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, California.

Pearl, J., A. Leal, and J. Saleh (1982). GODDESS: A goal-directed decision structuring system. *IEEE Transactions on Pattern Recognition and Machine Intelligence* **4(3)** 250-262.

Raiffa, Howard (1968). *Decision Analysis*. Addison-Wesley.

Raiffa, Howard, and Robert Schlaifer (1961). *Applied Statistical Decision Theory*. Harvard University Press.

Russel, S., and E. Wefald (1989). On optimal game-tree search using rational meta-reasoning. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, pp. 334-340.

Savage, Leonard J. (1954). *The Foundations of Statistics*. Wiley, New York. Second revised edition published in 1972 by Dover.

Smith, James Q. (1989). Influence diagrams for Bayesian decision analysis. *European Journal of Operational Research* **40** 363-376.

Spall, James C., ed. (1988). *Bayesian Analysis of Time Series and Dynamic Models*. Marcel Dekker.

Sprague, Ralph H., Jr., and Hugh J. Watson, eds. (1986). *Decision Support Systems: Putting Theory into Practice*. Prentice-Hall, Englewood Cliffs, New Jersey.

von Neumann, John, and Oskar Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press. Second edition, 1947; third edition, 1953.

von Winterfeldt, Detlof, and Ward Edwards (1986). *Decision Analysis and Behavioral Research*. Cambridge University Press.

Wald, Abraham (1950). *Statistical Decision Functions*. Wiley, New York.

Articles for Chapter 3

1. A Tutorial Introduction to Decision Theory, by D. Warner North. (1968) *IEEE Transactions on Systems Science and Cybernetics*, Vol. SSC-4, No. 3, September 1968, pp. 200-210.
2. Evaluating Influence Diagrams, by Ross D. Shachter. (1986) *Operations Research*, Vol. 34, No. 6, November-December 1986, pp. 871-882.
3. Rational Choice and the Framing of Decisions, by Amos Tversky and Daniel Kahneman. (1986) *Journal of Business*, Vol. 59, pp. S251-S278. This special volume of the *Journal of Business* was reprinted by the University of Chicago Press in 1987 under the title *Rational Choice*, edited by Robin M. Hogarth and Melvin W. Reder. Tversky and Kahneman's article appears on pages 67-94 of this version.
4. Critical Decisions under Uncertainty: Representation and Structure, by Benjamin J. Kuipers, Alan J. Moskowitz, and Jerome P. Kassirer. (1988) *Cognitive Science*, Vol. 12, pp. 177-210.
5. Savage Revisited, by Glenn Shafer. (1986) *Statistical Science*, Vol. 1, No. 4, pp. 463-485.
6. Computer Programs to Support Clinical Decision Making, by Edward H. Shortliffe. (1987) *Journal of the American Medical Association*, 3 July 1987, Vol. 258, No. 1, pp. 61-66.