DECISION ANALYSIS

INTRODUCTION

Decision analysis is engineering applied to decision making. The term decision analysis is typically used to refer to a set of analytical tools applied by decision analysts to arrive at recommendations for action. While these tools are important, identifying them with decision analysis is like identifying carpentry with hammers and saws. More important than the tools is the structured process of decomposition, analysis and synthesis that decision analysts apply to decision problems. Fundamentally, decision analysis is a way of thinking. Like all engineering disciplines, decision analysis is based on application of the scientific method and rational analysis. But decision analysis makes explicit what other engineering disciplines often leave implicit, that rationality and the scientific process are tools for helping to devise policies and construct solutions that serve our values. Decision analysis forms a bridge between the rational world of structured, analytic thinking and the esthetic world of values and feelings. Properly applied, the decision analysis process improves individual, organizational and societal decision making by helping to construct decision policies that serve our core values.

The decision analysis process decomposes a decision problem into components, analyzes the components, synthesizes them into a decision model, and uses the model to evaluate options and recommend actions. Decomposition focuses on three basic questions:

1. *What do I value?* The decision maker considers the fundamental objectives served by the decision, develops ways to measure how well these objectives are achieved, and organizes the information into a preference model for evaluating policy options.
2. *What can I do?* The decision maker identifies a set of policy options under consideration.
3. *What might happen?* The decision maker considers the consequences of the options under consideration and evaluates how well the consequences meet the fundamental objectives. When consequences are uncertain and the uncertainty has an impact on which option is preferred, an uncertainty model is constructed. Options for gathering information to reduce uncertainty are identified and considered.

Decision analysis follows the standard phases of the engineering process: problem definition, analysis, design of a solution, criticism and refinement of the proposed solution, and implementation. Problem definition involves setting the decision context: who are the actors, what are their roles, whose objectives (individual or organizational) are to be served by the decision, how broad or narrow is the scope of options to be considered, what constraints must be taken into account, what is the time frame for decision making and policy implementation, what sources of information are available. The analysis phase constructs a preference model and/or an uncertainty model. In the solution design phase the model is applied to evaluate candidate policy options and select a preferred option. Solution criticism and refinement involves examining the results of the model, performing sensitivity analysis on key inputs, evaluating how well the model captures the essentials of the problem for the purpose at hand, and making a final policy selection. The final step of a decision analysis is implementation of the chosen option. Although often omitted from texts on decision analysis, strategies for implementation and execution monitoring are often as important to achieving the decision maker's objectives as identifying the best policy.

PREFERENCE MODELING

In decision analysis, preferences are modeled by a *utility function* that measures the decision maker's relative degrees of preference for different consequences. Preference modeling is the process of constructing a utility function that captures the decision maker's values.

The first step in preference modeling is identifying the fundamental objectives for the decision context. It is important to distinguish fundamental objectives, which are intrinsically important to the decision maker, from means objectives, which are important only to the degree to which they support fundamental objectives. Fundamental objectives are decomposed hierarchically. A good final set of objectives is complete, concise, non-redundant, separable, measurable, operational, and controllable.

Next, attributes of value are defined to measure how well outcomes satisfy each of the fundamental objectives. The set of attributes should cover the fundamental objectives completely but without redundancy. Once attributes have been defined, a *single attribute utility function* is constructed for each attribute, and these are aggregated into a *multiattribute utility function*. The model is constructed according to by following a structured elicitation process. The preference elicitation process is based on constraints imposed by the mathematics of utility theory, insights and methods from experimental psychology and psychometrics, and the distilled wisdom of years of decision analysis practice. The decision maker provides judgments that enable the modeler to select an appropriate functional form for the utility function and determine parameters of the function (most notably, the shape of the single-attribute utility curves and the relative weight to be given to different attributes). The most commonly applied functional form for the multiattribute utility function is a linear weighted average of single-attribute utility functions:

$$u(x) = \sum w_i u_i(x_i) \tag{1}$$

In this expression, $u_i(x_i)$ is the single-attribute utility function for the ith attribute of value. It is commonly specified to range on a scale between 0 (the utility of a "reasonable worst" option) and 1 (the utility for a "reasonable best" option). The weights are positive numbers that sum to 1. The weights measure of the relative value to the decision maker of "swings" from reasonable worst to reasonable best on the respective attributes. When the problem involves uncertainty, the utility function reflects not just ordinal preferences but also attitude toward risk. A concave utility function on a numerical attribute reflects aversion to risk, i.e., the decision maker prefers a certain option to a gamble with the same expected value. A linear utility function is risk neutral; a convex utility function is risk-seeking.

OPTION GENERATION

Some decision problems involve selecting from among a discrete, denumerable set of options. For such problems, option generation means developing the list from which an option is to be selected. In a portfolio problem, the options under consideration are subsets of a set of elementary options. In other problems, the option space is defined implicitly by constraints defining a feasible region of solutions in some solution space, or by operators applied to options to generate other options.

Option generation and preference modeling support each other. An explicit focus on values and fundamental objectives helps to spur creative generation of new options for meeting fundamental objectives, avoiding a narrow focus on a few salient options. Attention to fundamental objectives helps to mitigate the tendency to underweight or ignore options that meet fundamental objectives but score poorly on salient means objectives. Comparing options to see how they differ can help to identify missed objectives. Examining objectives on which an otherwise good option scores poorly can help to generate ideas for modifying the option to address its shortcomings.

UNCERTAINTY MODELING

Uncertainty over consequences is measured by a probability distribution that quantifies the decision maker's degree of belief in different consequences. Decision analysis embraces the subjectivist view of probability, in which probability measures degrees of belief in propositions about which the decision maker is uncertain. Uncertainty modeling is useful because it provides a structured and theoretically sound approach to sort through the implications of different contingencies, organize information about their impact on the decision, and summarize them to arrive at an overall evaluation of an option.

To build an uncertainty model, the decision maker first identifies the key uncertain contingencies that affect the decision. A qualitative assessment is made about how much the uncertainties affect the choice of which option is preferred. Uncertainties that matter are selected for further modeling. Uncertain contingencies are modeled as random variables. A joint probability distribution is defined to express the decision maker's beliefs about the uncertain contingencies. As with preference modeling, a structured elicitation process is used to specify the uncertainty model. Subjectivist Bayesian theory provides a sound methodology for integrating empirical data with informed expert judgment to form a model that accurately reflects available knowledge about the uncertain contingencies.

The uncertainty model may include contingencies whose outcome is not known at the time the model is constructed but will be known at the time the choice is made. The model is specified so that the recommended decision policy may be contingent on these outcomes. Suppose X is an uncertain contingency affecting value to the decision maker and Y is a related observable contingency whose outcome depends probabilistically on X. If Y becomes known before the decision is made, then the optimal decision is based on the conditional distribution of X given Y, which is computed from the prior distribution using *Bayes rule*:

$$P(X|Y) = \frac{P(Y|X)P(X)}{P(Y)} \tag{2}$$

The optimal policy may depend on which outcome occurs for Y.

MAXIMIZATION OF SUBJECTIVE EXPECTED UTILITY

Once preference and uncertainty models have been specified, the model is solved for the optimal policy, which is the policy that maximizes the expected value of the utility function, where the expectation is taken over the uncertain contingencies. The principle of maximizing subjective expected utility can be derived mathematically from various systems of axioms reflecting principles of rational choice under uncertainty. Axiomatic justifications of expected utility maximization are compelling because of the guarantee that a decision theoretically sound model is internally consistent and the chosen policy is optimal given the modeling assumptions. From an engineering perspective, theoretical soundness is attractive but not sufficient. More important is the judgment of experienced practitioners that the structured decomposition, analysis, and synthesis process improves decision making by providing a framework for organizing, analyzing, and integrating the many factors involved in complex decisions.

DECISION TREES AND INFLUENCE DIAGRAMS

Two commonly applied visual and analytic tools for constructing a decision model are the decision tree and the influence diagram. Figures 1 and 2 show an influence dia-

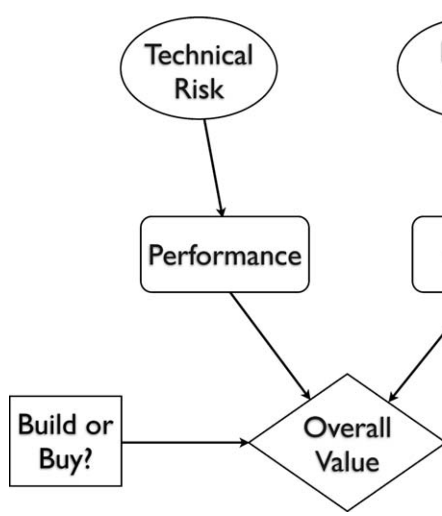

Figure 1. Influence Diagram for Build or Buy Decision.

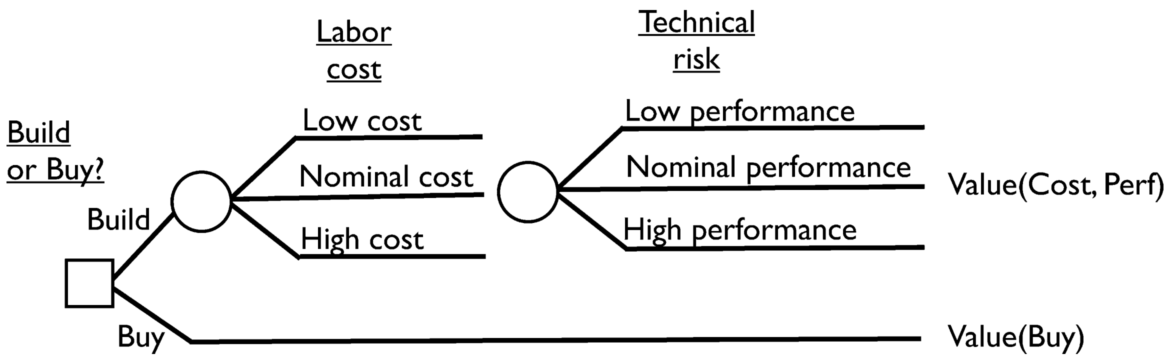

Figure 2. Decision Tree for Build or Buy Decision.

gram and a decision tree for a decision of whether to satisfy a client's requirement with an off-the-shelf solution or to design and build a custom solution. The two tools show complementary views of the same problem. The influence diagram displays independence relationships between uncertain contingencies and the decomposition of utility into attributes of value. For example, Figure 1 shows a decomposition of value into cost and performance. Labor costs are modeled as independent of technical risk. The decision tree shows how the contingencies affect the decision options. For example, Figure 2 shows that labor cost and technical risk affect the decision maker's value for the build option but not for the buy option. Decision trees serve both as graphical aids for specifying a model and as computational architectures for solving the model. The decision tree of Figure 2 is shown in schematic form. A full decision tree would have ten branches: a branch for each combination of values of cost and performance under the build option, and another branch for the buy option. A number of commercial software packages exist for specifying and solving decision models using decision trees and influence diagrams.

VALUE OF INFORMATION

Decision analysis provides a sound basis for evaluating options for collecting information to resolve uncertainty. When the optimal decision differs for different outcomes of an uncertain contingency, resolving the uncertainty before the decision is made may increase expected value to the decision maker. The *expected value of perfect information* (*EVPI*) is the increase in expected utility if costless information is provided prior to the decision that completely resolves the uncertainty in question. The EVPI determines an upper bound on the price the decision maker would be willing to pay for information. The *expected value of sample information* (*EVSI*) is the difference in utility if a realistically achievable information collection option is implemented.

SENSITIVITY ANALYSIS AND MODEL CRITICISM

It is a key tenet of decision analysis that the value of modeling derives as much from the insight the decision maker gains into the problem than from the answer provided by the model. The modeling exercise helps the decision maker to ensure that all relevant factors have been considered, to integrate all available information in a consistent and sound way, to reflect on how to trade off different components of value, and to justify the decision to him/herself and others. An important support to this process is sensitivity analysis, in which inputs to the model are systematically varied to observe the impact on the model's recommendations. Sensitivity analysis and the structured modeling process help the decision maker to understand the reasons for the model's recommendations and to adjust the model to make sure it incorporates all important concerns. While the model and the modeling process provide inputs to the decision, the ultimate responsibility for the decision lies with the decision maker.

ADDITIONAL ISSUES IN DECISION ANALYSIS

The prototypical application of decision analysis is a situation in which a trained decision analyst works with a single decision maker to build a model for a major, one-time decision problem. As training in decision analysis becomes more widespread and as software tools become more accessible, decision analysis can be practiced without the intervention of the analyst, and becomes cost-effective for more routine decision problems. An active area of research is the development of technology for specifying reusable template models and model components for commonly recurring problem types. While subjective expected utility theory is a single-actor theory of optimal decision-making, decision analysis is often applied in situations involving more than one actor. Structured elicitation methods have been developed for eliciting a consensus decision model from a group of stakeholders. Ideas and methods from decision analysis have been applied in the field of artificial intelligence to develop inference, prediction, diagnosis, and planning systems.

BIBLIOGRAPHY

Bunn, D. (1984) *Applied Decision Analysis*. New York: McGraw-Hill.

Clemen, R. T. *Making Hard Decisions: An Introduction to Decision Analysis*, 2nd edition.(1997) Pacific Grove, CA: Duxbury Press.

Howard, R. A. (1988) Decision analysis: practice and promise. *Management Science* **34**, 679–685.

Keeney, R. (1992) *Value-Focused Thinking: A Path to Creative Decision Making*. Cambridge, MA: Harvard University Press.

Keeney, R. and Raiffa, H.(1976) *Decisions with Multiple Objectives*. New York: Wiley.

Morgan, G. and Henrion, M.(1990) *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge: Cambridge University Press.

Plous, S. (1993) *The Psychology of Judgment and Decision Making*. New York: McGraw-Hill.

Raiffa, H. (1968) *Decision Analysis*. Reading, MA: Addison-Wesley.

Raiffa, H. and Schlaifer, R. (1961) *Applied Statistical Decision Theory*. Cambridge, MA: Harvard University Press.

von Winterfeldt, D. and Edwards, W. (1986) *Decision Analysis and Behavioral Research*. Cambridge: Cambridge University Press.

Watson, S. R. and Buede, D. M. (1987) *Decision Synthesis: The Principles and Practice of Decision Analysis*. Cambridge: Cambridge University Press.

KATHRYN B. LASKEY
George Mason University,
Fairfax, VA