



Decision Analysis: Practice and Promise

Ronald A. Howard

Management Science, Vol. 34, No. 6. (Jun., 1988), pp. 679-695.

Stable URL:

<http://links.jstor.org/sici?sici=0025-1909%28198806%2934%3A6%3C679%3ADAPAP%3E2.0.CO%3B2-M>

Management Science is currently published by INFORMS.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/informs.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

DECISION ANALYSIS: PRACTICE AND PROMISE*†

RONALD A. HOWARD

*Department of Engineering-Economic Systems, Stanford University,
Stanford, California 94305*

Decision analysis stands on a foundation of hundreds of years of philosophical and practical thought about uncertainty and decision-making. The accomplishments and promise of the field are impressive, yet it has not become commonplace even in very important decisions. While human nature may pose an ultimate limitation, maintaining clarity of concept and exploiting progress in the realms of scope, skill, and efficiency should lead to more widespread use.

A central conceptual distinction is that between normative and descriptive views of decision-making. We discuss the importance of maintaining this distinction in the face of attempts to compromise.

The procedures for formulating, eliciting, evaluating, and appraising the decision problem are all experiencing major improvements. The strategy-generation table helps in finding creative alternatives. Decision quality concepts permit us to assure both effectiveness and efficiency in analyzing decision problems. The influence diagram provides new clarity to the conversation between decision-maker and analyst, allowing representations that are both easily understandable and mathematically consistent. The clarity test makes sure we know what we are talking about regardless of what we are saying about it. Direct and indirect values illuminate preferences. Generic risk attitude considerations indicate how to relate corporate risk tolerance to the financial measures of the corporation. Spreadsheet, decision tree, and influence diagram programs speed evaluation. Intelligent decision systems realized in computers offer promise of providing the benefits of decision analysis on a broader scale than ever before. Decision analysis is now poised for a breakthrough in its usefulness to human beings.

Introduction

Let us begin by assessing the status of systematic reasoning about human action. Consider the following:

By this theory, we learn to appreciate precisely what a sound mind feels through a kind of intuition often without realizing it. The theory leaves nothing arbitrary in choosing opinions or in making decisions, and we can always select, with the help of this theory, the most advantageous choice on our own. It is a refreshing supplement to the ignorance and feebleness of the human mind.

If we consider the analytic methods brought out by this theory, the truth of its basic principles, the fine and delicate logic called for in solving problems, the establishments of public utility that rest on this theory, and its extension in the past and future by its application to the most important problems of natural philosophy and moral science, and if we observe that even when dealing with things that cannot be subjected to this calculus, the theory gives the surest insight that can guide us in our judgment and teaches us to keep ourselves from the illusions that often mislead us, we will then realize that there is no other science that is more worthy of our meditation.

This was written by Laplace in 1812. With stylistic changes, it could describe an optimistic view of decision analysis today.

The possibility of effective, systematic reasoning about human action has been appreciated for over two hundred years. Laplace's predecessor, Bayes, showed in 1763 that probability had epistemological power that transcended its aleatory uses. In the early 1700s, Bernoulli captured attitudes toward risk-taking in mathematical form.

* Accepted by Robert L. Winkler; received July 30, 1986. This paper has been with the author 10½ months for 1 revision.

† *Departmental Editor's Note.* At the Los Angeles TIMS/ORSA meeting in April 1986, Ronald A. Howard was awarded the Frank P. Ramsey Medal for distinguished contributions to the field of Decision Analysis. This paper is based on his talk, "An Appreciation of Decision Analysis," given at the time the Ramsey Medal was awarded.

The resurgence of the field in modern times began with statistical decision theory and a new appreciation of the Bayesian viewpoint. We see concern with prior-to-posterior operations, conjugate distributions, and utility theory in the modern form. Another seldom-mentioned source is the collective contribution of Jeffreys and Jaynes, who worked in the field of physics rather than in mathematics and statistics and who provided an all-encompassing view of probability, not as an artifact, but as a basic way of reasoning about life, just as had Laplace. Jeffreys and Jaynes developed very clear ways of relating probabilities to what you know about the world around you, ways that provide dramatic insights when applied to the molecular processes that interest many physicists. However, Jaynes also showed that these ideas pay off handsomely when applied to inference problems in our macroscopic world.

These contributions led, in turn, to decision analysis, which we might call “decision engineering.” The development of computers over the last few decades has given us a new capability for handling complexity. While Laplace, working in the days of Napoleon, could have used virtually all the concepts we have now, he would have been very limited in computing problems of practical size. From the systems engineering legacy of the Second World War, we learned procedures for modeling and performing sensitivity analysis, which were added to the ideas of statistical decision theory to give us a powerful way of analyzing decisions.

Today we have a discipline of decision analysis: a systematic procedure for transforming opaque decision problems into transparent decision problems by a sequence of transparent steps. Opaque means “hard to understand, solve, or explain; not simple, clear, or lucid.” Transparent means “readily understood, clear, obvious.” In other words, decision analysis offers the possibility to a decision-maker of replacing confusion by clear insight into a desired course of action.

In spite of these developments, within this century John Maynard Keynes (1964) wrote:

Most . . . decisions to do something positive . . . can only be taken as a result of animal spirits—of a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities.

I doubt that the situation has changed very much up to the present time. Even today, very few decision-makers benefit from the full power of decision analysis. To understand why this is and what might be done about it, let us begin with a brief description of decision analysis.

The Decision Analysis Process

A decision analysis is performed using the process shown in Figure 1. We start with some real decision problem facing a decision-maker, an opaque one if the analysis is to be truly useful. Our intention is to apply a sequence of transparent steps to provide such clarity of insight into the problem that the decision-maker will undertake the recom-




FIGURE 1. The Decision Analysis Process.

mended action. The first step of formulation fits a formal model to the decision-maker's opaque real situation. We call this formal representation of the problem a "decision basis" and shall discuss it further in a moment. The decision basis must be evaluated by a primarily computational process to produce the alternative that is logically consistent with the basis and therefore recommended. Then we perform an appraisal of the analysis to gain insight into why the recommended alternative is not only logically correct, but so clearly persuasive that the person will act accordingly. The appraisal may well reveal some shortcomings of the analysis, requiring a refinement of the formulation to assure that it is truly appropriate to the problem. At some point, the appraisal step will show that the recommended alternative is so right for the decision-maker that there is no point in continuing analysis any further.

The core of this process is the elicitation or synthesis of the decision basis. Figure 2 shows the role of the basis in more detail. The basis has three parts: the choices or alternatives the decision-maker faces, the information that is relevant, and the preferences of the decision-maker. Because detailed discussion of these basis features appears elsewhere (Howard 1983a), let us treat them only briefly to establish a common ground of understanding. The alternatives may be readily apparent or may be generated as a major activity of the formulation using tools we shall discuss later. By information, we mean any models, relationships or probability assignments that may be important in characterizing the connection between decisions and outcomes. The models could be complex and dynamic, or, in some cases, very simple. The uncertainty that remains would be characterized by probability assignments. The preferences of the decision-maker would be represented in at least three dimensions. The decision-maker would have values on one outcome as opposed to another, and time preference considerations on outcomes now versus outcomes later. Finally, the decision-maker would have a risk preference governing outcomes with different degrees of certainty.

With the basis complete, evaluation could begin. The best decision is selected by




FIGURE 2. Elicitation and Evaluation of the Decision Basis.

logical operations. The appraisal provides sensitivities to choice, information, and preferences.

Given that the benefits of decision analysis have been foreseen for centuries, why is it not yet routinely applied to important decisions? I believe the answer can be found in considerations of human nature, concept, scope, skill, and efficiency. We shall examine each consideration and discuss our ability to deal with it.

Considerations of Human Nature

Accepting decision analysis requires a belief in the value of systematic, logical thought as a basis for decision-making. This cognitive style will not be natural to people who prefer to be guided primarily by feelings rather than thought. Research based on the Myers-Briggs Type Indicator shows how people differ in the way they like to perceive and the way they like to judge (Myers 1980). Decision analysts should realize that not everyone sees the world as they do. They should appreciate the special insights they provide that can eliminate the "blind-spots" of those who rely mainly on feelings.

Some who are comfortable with the cognitive style of decision analysis may avoid it in organizational settings because it requires a willingness to be open and explicit in decision-making. Those who fear criticism or who wish to control by concealment will not find decision analysis to their taste.

Conceptual Considerations

Using decision analysis to make decisions usually requires changing the way we think about the world. We must not only use new information, but change the context within which we process this information. This means creating new distinctions that are apparently not natural, but learned.

The most important distinction needed for decision analysis is that between decision and outcome. I tell my students that if they learn nothing else about decision analysis from their studies, this distinction will have been worth the price of admission. A good outcome is a future state of the world that we prize relative to other possibilities. A good decision is an action we take that is logically consistent with the alternatives we perceive, the information we have, and the preferences we feel. In an uncertain world, good decisions can lead to bad outcomes, and vice versa. If you listen carefully to ordinary speech, you will see that this distinction is usually not observed. If a bad outcome follows an action, people say that they made a bad decision. Making the distinction allows us to separate action from consequence and hence improve the quality of action.

Another distinction central to decision analysis is that between descriptive and normative models. A descriptive model characterizes what happens in the world without judgment of it. A normative model states what ought to be the case. It is the judgment of desirability of a human being about proper action in the world. Thus, physical science is descriptive, the rules of law are normative. It is descriptively true that I make mistakes in arithmetic calculation, but I accept the norms of arithmetic. In fact, it is the acceptance of these norms that allows me to recognize a mistake.

The whole idea of a normative model arises when we are not satisfied with our functioning. Because we know we are subject to optical illusions, we believe our measuring instruments rather than our perceptions. Because we know that we become disoriented in flying without visual cues, in bad weather we rely on flight instruments rather than on our body's sense of orientation. We need normative models when what we are tempted to think may not be so. In view of the many easily demonstrated lapses in human decision-making that we can observe, who would want to rely on unaided judgment for a complex and important decision problem?

Decision analysis is the normative practice of decision-making. The acceptance of its norms allows me to recognize my decision mistakes. Using decision analysis I can be

sure that the decision I make follows the rules I would use to reason in very simple, and therefore transparent, situations, as well as the rules I would use to reduce complex situations to transparent ones.

Some decision theorists have questioned the normative concepts. They desire to weaken the norms until the normative behavior agrees with the descriptive behavior of human beings, to construct theories of decision-making that are both normative and descriptive. A moment's reflection shows that if we have a theory that is both normative and descriptive, we do not need a theory at all. If a process is natural, like breathing, why would you even be tempted to have a normative theory?

Weakened norms permit contradictions of many of the valued features of decision analysis. Such features include the sure-thing principle, independence of irrelevant alternatives, no willingness-to-pay to avoid having an alternative, nonnegative value of information, stochastic dominance, transitivity of preference (and hence no possibility of being a money pump), and no dependence of preference on what you might have received, as opposed to did receive, in any deal (no regret). The price is high.

For example, consider the following one-time decision. A die will be rolled and you will receive, at no cost to you, payoffs from each number that might result according to your choice of payoff schemes A and B given in this table:

Die Outcome:	1	2	3	4	5	6
Payoff from Scheme A	\$600	\$700	\$800	\$900	\$1,000	\$500
Payoff from Scheme B	\$500	\$600	\$700	\$800	\$900	\$1,000

The example was presented by an eminent decision theorist at a national meeting. He said that he would prefer A to B in the problem as posed because he would do better in five out of the six cases. However, he said, if the deal were for hundreds of thousands of dollars instead of for hundreds of dollars, then he would prefer B because of the regret he would feel with A should the number 6 result.

I asked him whether this was a hypothetical position or the preference instruction he would give me should I act as his agent in making actual decisions. He confirmed that he would want these preferences to govern my action.

When I returned from the meeting, I put this example to at least a dozen decision analysts at many levels of experience. All thought I was joking or that I had made a mistake in posing the problem. They could see no difference between A and B since they assigned the same probabilities to the same prizes.

The fact that a major decision theorist is advocating principles that are at variance with what I believe to be those of the overwhelming majority of practitioners illustrates the importance of conceptual considerations. I have no hesitation in using for myself and in recommending for others a decision rationale based on the axioms of orderability, continuity, substitutability, monotonicity, and decomposability (Howard 1983b), which in turn implies stochastic dominance, the sure-thing principle, etc. The weakening of the structure may yield interesting mathematics but, in my opinion, has yet produced only inferior guides to action.

How can we help to disseminate the distinctions, concepts, and principles of decision analysis? We must provide general education in the subject, not just in graduate school and in college, but in high school and elementary school. I recently heard an executive with major responsibility for investing billions of dollars ask what an expected value was and how it was calculated—175 years after Laplace and 300 years after Huygens defined expectation. Here was a man ill-served by his formal education.

Considerations of Scope

Is decision analysis too narrow for the richness of human decision? How can we assure that we are working on the right problem? How can we deal with problems

without well-defined alternatives? How can we know that our analysis is of high quality, and what does that mean? We would proceed at peril without satisfactory answers to these questions, but happily we are making progress.

Framing

We can think of framing as a way to avoid errors of the third kind; namely, working on the wrong problem. Framing deals with the phenomenon Freud called the “presenting problem.” He noticed that when he asked new patients what was bothering them, their concerns seldom turned out to be their fundamental problems as revealed after the beginning of analysis. When you set out to help a decision-maker, it is only too easy to think that the opaqueness of his or her problem is correctly stated. You must often probe deep to make sure that the framing of the decision is really appropriate, that the underlying decision problem has really been uncovered. Framing is the most difficult part of the decision analysis process; it seems to require an understanding that is uniquely human. Framing poses the greatest challenge to the automation of decision analysis.

Creating Alternatives—The Strategy-Generation Table

The most important idea in creating alternatives that I have encountered is the strategy generation table. The strategy table is almost self-explanatory. It shows how a total strategy can be specified by selecting among decisions in each of many areas, which we may think of as individual decision variable settings.

A typical strategy table appears as Figure 3. In this illustration, a conglomerate is selecting a strategy for the various parts of its business, such as an electric utility division, exploration and production division, an oil field services division, etc. The strategy table also shows a dividend decision as part of the total strategy. When you discuss possible strategies with executives, you find that not all combinations make sense, that a certain decision in one area implies or at least indicates particular decisions in other areas. After consideration in this case, the executives developed an alternative strategy requiring holding or restricting investment in the utility area, selling or milking the cash of the exploration and production area, etc., as shown by the connecting lines in the diagram. This strategy was readily described by the executive team as a “service

Strategy Theme	Utility	E & P	Oilfield Services	Forest Products	Coal	Acquisition	Dividends	Debt/Equity Ratio
Baseline	Aggressive Supply Buildup	Increase Exploration Budget	Aggressive Expansion and R&D Program	Hold-Improve Earnings	Purchase Additional Reserves	None	70%	1/1
Service Business	Hold/Restricted Investment	\$800M Investment	Modest Expansion	Add Timberlands	Joint Venture Synfuels	Service Business	50%	1.5/1
Resource Acquisition								
P/L Emphasis			Hold		Hold	Resource Business	25%	2/1
Gradual Liquidation	Severe Capital Constraint	Sell/Milk	Milk	Sell	Milk	P/L	0%	

Each combination of strategy elements defines a strategic alternative.

FIGURE 3. The Strategy-Generation Table.

business” strategy for the company. As such, it had a clear differentiation in their minds from, for example, a “resource acquisition” strategy, which is not indicated on the diagram. The idea of a strategy theme enables people to discuss a few significantly different strategies rather than a combinatorially exhaustive and exhausting list. Typically, five to seven theme strategies might be selected for further analysis.

When I first came across the strategy generation table, it seemed rather simplistic to me, particularly from a technical point of view. I had criticisms, such as, “We are not doing a complete search of the space of alternatives.” Yet I found that there were very few ideas in decision analysis that responded to the multiplicity of possible strategies in a strategy problem. As a result, I came to regard the strategy-generation table not as a quick and dirty approach, but rather as a very useful tool for helping people think their way through problems where there were literally thousands of possible strategies.

To illustrate this point, my son Robert recently drew to my attention a letter in a local paper. The letter commented on the advertisement of a fast food hamburger chain that boasted that you could get 1,024 different hamburgers at their establishments. The reason was that their hamburgers could either have or not have each of ten possible ingredients, like lettuce, pickle, etc. As the letter writer pointed out, one of the burgers is the “nullburger.” In addition, of course there is the burger consisting only of lettuce, the burger consisting only of onion, etc. That’s why the 1,024 possibilities that immediately leap to the mind of a mathematician do not accurately describe the real possibilities faced by human beings ordering hamburgers.

Decision Consulting

What is the area of responsibility of one who would help another make a decision? Is it simply to provide proper framing, elicitation, evaluation, appraisal, etc.? Or is there a responsibility beyond this role to make sure that there is some kind of quality in the whole decision-making process? This means quality not just in elicitation and evaluation, but quality in the content of the alternatives and information, and quality in the clarity of the preferences.

At Strategic Decisions Group, we have been using the term “decision consulting” to describe this expanded role that extends beyond the traditional view of the decision analyst. The decision consultant is responsible not only for proper elicitation and evaluation, but also for making sure there is quality of content in each part of the process, possibly by serving as a broker who locates new information sources for alternatives, models, or assessments.

Measuring Decision Quality

We have been trying out the form shown in Figure 4 as a means of assessing decision quality. The form is intended to evaluate a decision made by any method, not just the formal method we have been discussing. The form would be completed with written statements. For example, opposite “decision framing” on the form we would discuss whether the proper problem was analyzed. We could check to see whether there was informational excellence, whether the most cost-effective information sources were used. As for creativity, we would consider whether there was an adequate search for promising alternatives. This does not mean that in a medical decision problem the decision consultant would sit down and invent new medical treatments, but rather that the consultant would interview several medical experts to see whether the broadest practical range of treatment alternatives had been considered. The decision consultant could tell the patient that the course of treatment recommended was based on the best medical advice, a claim that goes beyond that of the traditional decision analyst.

To be sure whether values were clear would require intimate knowledge of the decision-making setting. Integration and evaluation with logic has been the strong point

Decision Framing "The Right Challenge"		
Decision Basis	Informational Excellence	
	Creativity -- Significantly Different Alternatives	
	Clear Values	
Integration & Evaluation with Logic		
Balance of Basis		
Commitment to Action		

FIGURE 4. The Elements of Decision Quality.

in decision analysis. Here the question would be whether any mistakes were made, whether appropriate sensitivities were developed, and, in general, whether the right conclusions were drawn.

The "balance of basis" heading refers to the allocation of effort within the analysis. If the analyst spent too much time clarifying values when what was needed was better alternatives, there would be a problem in basis balance.

Finally, if you want to judge the whole operation as an example of decision quality, you would want to see whether the right course of action was so clearly indicated and communicated that the decision-maker developed a clear commitment to action in that direction. Of course, the decision consultant cannot compel action. But any inaction should be the result of a failure of will, not the result of unpersuasive recommendations.

Our concern with the definition and achievement of decision quality is just beginning. Yet, it is clear that asking these questions can only lead us in the right direction.

Considerations of Skill

Even when working with the right people on the right problem with the right concepts, the discipline of decision analysis could fail because the processes of communication and elicitation place excessive demands on the decision-maker or the expert. We are experiencing rapid progress in easing these demands.

The Influence Diagram

The greatest advance I have seen in the communication, elicitation, and detailed representation of human knowledge is the influence diagram (Howard and Matheson 1983). Because all of my students and all of the people at the company with which I am associated use influence diagrams so regularly, I have come to believe that they are more widely used and appreciated than is, in fact, the case. It is my experience with influence diagrams that even people whose knowledge of them is far from technically

deep find them to be an extremely important and useful tool for the initial formulation of decision problems. Yet, when discussing the subject with an eminent professor of decision analysis at another university recently, I learned that he had heard about influence diagrams but had never actually used them or taught them to his MBA students. At the end of our discussion, he told me that he would have to look into influence diagrams much more deeply than he had.

I have personally seen that MBAs have little difficulty in using influence diagrams for formulation, that it is not necessary to have a doctorate in mathematics to use this structuring tool effectively. In fact, influence diagrams are the best tool I know of for crossing the bridge from the original opaque situation in the person's mind to a clear and crisp decision basis. The reason is that the influence diagram is very easy for people to understand, regardless of their mathematical facility. There are no arcane symbols, like equal signs, or sideways Ms call sigmas. Such language is not natural for most people who need help with decisions. Yet, if you study influence diagrams, you will find that they are technically as well-defined as decision trees and that they can be evaluated with at least as great ease and efficiency.

A portion of a typical influence diagram is shown in Figure 5. Here the quantities in circles or ovals are considered uncertain. Arrows entering circles mean that the quantities in the circles are probabilistically dependent on whatever is at the other end of the arrow. The rectangles are decisions under the control of the decision-maker. Arrows entering such decision nodes show the information that is available at the time of the decision. This particular diagram is rather small; in practice, it is common for diagrams to have perhaps 30 nodes, and in the early stages of formulation some have more than 100. However, since a large diagram usually reveals insufficient insight as to what is important in making the decision, various forms of sensitivity analysis are used to reduce the diagram to its essential elements.

As the diagram shows, decisions about the research and development strategy affect the chances of technical success, which has to be technically defined. We shall later discuss how to be clear in such definitions. Technical success might be the development of a new jet engine or a new treatment for a viral disease. The level of technical performance will depend on whether you have technical success as well as on the research and development strategy you have used. Technical performance would be a measure like the efficiency of the jet engine or the effectiveness of the disease treatment. When the decision-maker selects a market strategy, both the event of technical success and the level of technical performance given technical success will be known. The extent of market performance will depend on the market strategy, on technical success, on technical performance, and on many other factors not shown in the diagram, such as the products and prices of competitors, the state of the economy, etc. Similarly, net commercial value would depend on many of these same factors, as well as on the product's own market performance.




FIGURE 5. A Typical Influence Diagram.

This influence diagram, although simple, is extremely useful. It can be explained both to a businessperson and to an analyst who is going to perform the ultimate calculations. Few other tools are so general and so useful. We often find that we can draw an influence diagram with a client and then leave it with the client for further consideration and possible change upon reflection. The client can think of new factors that were not included or possibly new interactions that must be taken into account. The diagram is, thus, an extremely helpful representation for extracting information in technically useful form, the core of the elicitation process.

The Clarity Test

When we began the process of probability assignment over 20 years ago, we had overly simple ideas about the nature of the problem. We often skipped over the step of carefully defining the event or variable to which probabilities were to be assigned. We now understand that much of the difficulty in probability assignment arises when a person assigning a probability to an event is unconsciously changing the definition of the event in the course of the encoding. A similar difficulty arises if different people are thinking of different events when they are supposed to be assigning probabilities to the same event. To overcome these problems, we have developed the concept of the “clarity test.”

Consider a clairvoyant who knew the future, who had access to all future newspapers, readings of physical devices, or any other determinable quantity. The clarity test asks whether the clairvoyant would be able to say whether or not the event in question occurred or, in the case of a variable, the value of the variable. In the case of the influence diagram we just discussed, the event “technical success” does not meet this test for whatever product or process we are talking about. To meet such a test, technical success would have to be defined in such terms as “able to operate x hours under conditions y without failure” or “solubility less than z at temperature t .” Once the events and variables in the problem meet the clarity test, there is hope that the probability assignments obtained from individuals will make sense.

In the practical use of the influence diagram, we begin without much concern for the strict definitions of events and variables and make sure that the individuals concerned are comfortable with the structure of the problem. Once the structure is satisfactory, then we formalize all the definitions so the clarity test can be passed before proceeding to probability encoding. This is a typical refinement in concept and in practice that has improved both the efficiency and effectiveness of the discipline.

Important advances have also taken place in probability assignment technique as the result of both insights from psychology and the use of the influence diagram. The author’s paper, “Knowledge Maps,” discusses how to use influence diagrams to elicit and capture the often diffuse knowledge that a person has on any topic. We have become much more effective in accurately representing states of information than we were in the early days of the discipline.

Direct and Indirect Values

In constructing the preference area of the decision basis, we have also benefited from refined concepts; for example, the distinction between direct and indirect values. Outcomes are directly valued if you are willing to pay for them in and of themselves. Other outcomes are indirectly valued since you place value on them only because they lead to outcomes that are directly valued. For example, if you are concerned almost exclusively with the present value of cash flows in operating your automobile, you might place an indirect value on a fuel economy measure like miles per gallon because of its ultimate effect on cash flows. However, if you would enjoy having “bragging rights” about how efficient your car is, then you might place a direct value on its miles per gallon figure. By

such distinctions, we can begin to sort out value issues and, hence, bring clarity to the analysis.

Generic Risk Attitude

An advance in the area of preference is the use of generic considerations in establishing the risk preferences of an individual or organization. It should be clear that few people or companies would want a utility curve with kinks and bumps in it even though such a curve could be permissible under the theory. By introducing generic considerations such as having a risk tolerance that is linear in wealth, you can both simplify the assessment process and assure that the resulting utility curve will not display any unforeseen sources of embarrassment. Of particular note here is the use of exponential utility curves to satisfactorily treat a wide range of individual and corporate risk preference.

While the ability to capture risk preference is an important part of our conceptual view of decision-making, I find it is a matter of real practical concern in only 5 percent to 10 percent of business decision analyses. Of course, the situations that require risk preference, such as bidding or portfolio problems, use it very seriously. In working on these problems, we have found both the use of generic forms for risk preference and experience with similar situations to be most helpful.

Corporate Risk Tolerance

To illustrate, let me share some of my experience, which is less than a detailed study, but perhaps still indicative of what might be more generally true.

Several years ago, we worked for three companies who were partners in a joint venture to assist them in making decisions about that venture. Because in this decision risk tolerance was an issue, we began by assuming that their risk preferences were exponential and then assessed the risk tolerances by interviewing top executives in each company. We assessed the exponential risk tolerances by the usual simple procedure: We found the sum of money such that the executives were indifferent as a company investment to a 50–50 chance of winning that sum and losing half of that sum. The result is a close approximation to the exponential risk tolerance. We also reviewed the annual reports to obtain the net sales, net income, and equity of the three companies.

Relating Corporate Risk Tolerance to Financial Measures. All this information is assembled in Table 1, where the companies are called A, B, and C. Here, entries are in millions of dollars, so you can see that we are dealing with some very large companies indeed. Notice that the risk tolerance appears to be growing roughly proportional to all the other numbers. The relationship is even more clear in Table 2, where we have shown the ratio of risk tolerance to each of the other quantities. Note that the ratio of risk tolerance to sales is very close to about 6 percent. Risk tolerance to net income shows more variation, but is about one to one and one-half. Finally, risk tolerance to

TABLE 1
Relating Corporate Risk Tolerance to Financial Measures (\$ millions)

Company	A	B	C
Net Sales	2,300	16,000	31,000
Net Income	120	700	1,900
Equity	1,000	6,500	12,000
Risk Tolerance	150	1,000	2,000

TABLE 2
Corporate Risk Tolerance Ratios

Company	A	B	C	Average
Risk Tolerance/Sales	0.0652	0.0625	0.0645	0.064
Risk Tolerance/Net Income	1.25	1.43	1.05	1.24
Risk Tolerance/Equity	0.150	0.154	0.167	0.157

equity is once more very consistent, with risk tolerance being about one sixth of equity. While all these companies were in the same general oil and chemicals industry, I find the consistency amazing.

You might now ask how general these results are. Would it be possible to advise companies, at least in a particular industry, on approximately what their risk tolerance should be and use this as a starting point for further refinement?

Illustration of the Ratios in Use. I had an opportunity to try this idea a few years later when I got involved with another company, Company D, in a somewhat similar industry. The time pressure for this project was such that I had to make a presentation to the chairman and the president on a recommended course of action without the opportunity to interview them on the question of risk tolerance and with a clear need for the incorporation of risk tolerance in this particular problem. Remembering the earlier results, I proceeded as in Table 3. I multiplied the average ratios from Table 2 by the corresponding measures from Company D's annual report to obtain three suggested risk tolerance numbers: \$212 million, \$188 million, and \$181 million. The average was \$194 million, and I used \$190 million for purposes of the presentation. During that presentation, I had an opportunity to interview the chairman and the president, and they revealed a risk tolerance of \$200 million for the company.

You might say I "lucked out," but I believe that something more fundamental was at work. As I said earlier, I have not pursued this descriptive study of corporate risk preference any further, but the situations I have investigated informally with clients seem to bear out the general ratios we observed in Table 2. This suggests that we might, at least in certain industries, be able to use financial statements to develop guidelines for the establishment of risk preference.

Considerations of Efficiency

The issues we have discussed focus on the effectiveness of decision analysis. Even if it is effective, it can be too costly. We shall discuss current improvements in efficiency which have been most notable in the process of basis evaluation.

TABLE 3
Company D Risk Tolerance Using Ratios (\$ millions)

Measure	Ratio (Risk Tolerance/Measure)	Company D Annual Report	Risk Tolerance (Measure × Ratio)
Net Sales	0.064	\$3,307	\$212
Net Income	1.24	152	188
Equity	0.157	1,153	181

Average: \$194

Company D Risk Tolerance Used in Analysis: \$190

Combining Decision Trees and Spreadsheets

One of the most important contributions to efficiency in recent years has been the development of spreadsheet programs for microcomputers. Truly amazing modeling work can now be done by people who are relatively untrained compared to those who wrote programs 10 to 20 years ago. We have also seen the development of very powerful probability and decision tree programs that can efficiently perform all tree operations, including rollback, calculation of value of information (both with and without risk aversion), and sensitivity analysis, and that can also be combined with spreadsheet programs to provide a modeling capability in the same package (such as SUPER-TREE). This means that people without extensive training in decision analysis can represent sophisticated problems involving uncertainty resolution and also rather complex modeling of outcomes.

This new capability slipped up on us almost unnoticed. What we find in practice is that analyses that would have taken us weeks 15 years ago can now be done in days. The improvement in efficiency means either that the analysis can be carried to much greater levels of refinement or that many more decision problems can be analyzed in the same time. For example, in the area of research and development, we now perform procedures called "lab sweeps" in which many projects in the laboratory undergo a brief analysis that attains a rather high level of quality in terms of insights. Such sweeps would not have been practical if each project required one month of analysis. This means we can now perform portfolio analyses of research and development laboratories at a high insight-per-dollar ratio.

Computation with Influence Diagrams

We are now seeing the beginning of influence diagram programs that can compute at least as efficiently as tree programs, and sometimes considerably more efficiently (Olmsted 1983, Shachter 1986). One of the battles I fight with myself is whether it is time to teach influence diagrams instead of trees. Since I did not work with influence diagrams until after I knew trees, it is difficult for me to imagine a "treeless" world. I do not know whether influence diagrams would have the same clarity for me if I were not already familiar with trees. But I can imagine that in the future I will no longer speak of decision trees but will conduct formulations and evaluations solely in the language of influence diagrams. This is a shocking thought because so many people in the world think of decision analysis as the theory and practice of trees. I have asked students of "decision science" what they have learned in their classes, and they told me that they knew how to roll back trees. When I asked them if they had learned how to go into the world of the decision-maker and draw a tree, they often said "no." Well, the news is that what they know how to do is no longer necessary, and what they do not know how to do is almost as difficult as it ever was. But perhaps the influence diagram will help.

Tornado Diagrams and Bubble Charts

The spread of decision analysis ideas can be seen by changes in the language used to describe them. One of the outputs of the decision tree and modeling programs is the type of ordered sensitivity analysis shown in Figure 6. In the past, such diagrams would have been computed with difficulty and laboriously plotted by a draftsman. One of our clients recently gave this diagram a new name: he called it a "tornado diagram," a name that is immediately appropriate when you stare at it for a moment. When you think about it, tornado diagram is a perfect name for this chart in two senses. First, because it is descriptive, and, second, because such a name takes the threat out of using the chart in management presentations. Often when our clients take influence diagrams back to their companies, they call them "bubble charts." And we are very happy with this




FIGURE 6. A Tornado Diagram.

terminology for the same reason. Discussions about bubble charts and tornado diagrams are bound to be both less threatening and more entertaining than those using formal terms. We have thus seen a cultural change both in the efficiency of the computations and in the acceptance of these concepts within many parts of business.

The Promise of the Future

Turning to the future, we might ask how we could extend the advantages of decision analysis to almost everyone interested in using it. One answer might be using computers in new ways, and in particular the ideas of artificial intelligence as represented primarily by expert systems.

Expert Systems—Descriptive Aids

I think of an expert system as shown in Figure 7. A human expert in the domain of a problem provides rules for decision-making to the expert system. The system, in turn, uses these rules to make recommendations in the decision problem. For example, in medicine, one rule might be, "If the patient is able to stand up and is strong enough to call the doctor, give the advance 'take two aspirins and call me in the morning'."

The test of whether such an expert system has been properly designed is whether the recommendations of the system are those that the human expert would make in the




FIGURE 7. An Expert System.

same situation. In other words, a good expert system is one that describes the action of the human expert. For many problems, particularly those without uncertainty and dynamics, the descriptive expert system may work well. However, we must realize that the expert system as we have defined it in this case does not have normative power. Its recommendations can be no better than the recommendations of the expert who trained it.

Intelligent Decision Systems—Normative Aids

We might ask how we could enjoy both the efficiency advantages of the expert system and the normative power of decision analysis. One way would be to construct what we call an intelligent decision system (Holtzman 1985). I view such a system in the manner shown in Figure 8. Here we use the descriptive ideas of expert system theory to construct a decision analysis expert system by applying these ideas at a higher level. The decision analysis expert system would manage the decision analysis process of Figure 1.

We understand that the decision analysis expert is knowledgeable in the process of decision analysis, in what it means to formulate, evaluate, and appraise a decision problem, but not in the content of the area in which a decision is made. The decision analysis expert system would not be knowledgeable about forest fires, or vaccines, or alternatives for a new product introduction, but would know that the most sensitive variables are the most deserving of probabilistic treatment, and that the presence of stochastic dominance can avoid the need to assess risk preference. The rules of this system would be rules for the creation of decision models, rules for their evaluation, and rules for explaining their implications.

Lest it sound like real as opposed to artificial intelligence, we should realize that we are likely to be able to build such systems only within a particular domain of application. The human decision analyst is too general-purpose to be captured everywhere.

Normative Power

The intelligent decision system would have normative power. The decision analysis expert system would not be providing a direct recommendation, but rather building and evaluating a decision model that is custom-fitted to the decision-maker. The decision-maker would follow the recommendation of the intelligent decision system not because it mimicked some expert in the domain, but because it provided the best recommendation for the decision-maker given the decision basis that had been constructed by the decision analysis expert system. Thus, we are using expert systems “one level up” from their usual function and adding to their computational efficiency the normative power of decision analysis.




FIGURE 8. An Intelligent Decision System.

Interaction Management

We can present a more detailed view of intelligent decision system operation using the diagram in Figure 9. Here we see the intelligent decision system managing the interaction with the decision-maker and carrying out the formulation, evaluation, and explanation functions. The decision-maker provides alternatives, preferences, and information and receives, in turn, recommendations and insights.

Domain Characterization

Each domain in which an intelligent decision system operates would have its own characterization. First, terminology of the domain would be specified. Reaction rates, cost per thousand, whatever were the distinctions of use in this domain would be available. The knowledge base would contain the rules for the domain, such as describing allowable ways of depreciating non-ACRS property. The data base would provide facts, such as the current tax rate in different states. There would be information about structure, such as the models that would be available for capturing relationships in the domain. Information would also be provided in the form of the probability assessments of experts on uncertainties likely to arise in the decision, such as the chance of dying under surgery. All information could be provided with a pedigree to show "who says so." The domain could also have advice, for example, in the form of suggested preferences. Our earlier discussion of corporate risk preference might provide the starting point for suggesting a risk preference function to the decision-maker. The system could also show the reasons why this preference is being suggested, and if the suggestion is rejected, the system would be able to use a risk preference module to obtain a sufficient assessment for purposes of this decision problem.

We also envision having on-call human experts. The experts could provide assessments needed for the bases that were not originally present in the domain characterization. In some instances, these assessments could be provided by the decision-maker, but in others, for example medical situations, external human experts would be necessary.

The Promise

The power of decision analysis transforms opaque decision situations into transparent ones. The intelligent decision system offers the promise of providing this power to individuals, and to lower-level corporate decision-makers who could not afford it at




FIGURE 9. Intelligent Decision System Interactions.

present and in situations requiring very rapid decision-making. This is an exciting promise. Think of all the people facing operations who could benefit from the clarity offered by a professional decision analysis and who might achieve close to this level of help at a cost of hundreds rather than thousands or tens of thousands of dollars. Think also of all the decisions in corporations presently made with inconsistent information and preferences that might be made uniformly more appropriate and profitable by the use of intelligent decision systems.

Conclusion

The discipline of decision analysis has already demonstrated that it can provide persuasive clarity to those few who have the resources and the will to use it. The new challenge to the discipline is to provide this help at an affordable price to all who could benefit from it.¹

¹ I wish to thank my colleagues at Strategic Decisions Group, and especially Sam Holtzman, Jim Matheson, and Carl Spetzler, for the many stimulating and creative conversations we have had on the topics in this paper.

References

- HOLTZMAN, S., "Intelligent Decision Systems," Ph.D. dissertation, Engineering-Economic Systems Dept., Stanford University, Stanford, CA, March 1985.
- HOWARD, R. A., "The Evolution of Decision Analysis," *Readings on the Principles and Applications of Decision Analysis. Vol. I*, Strategic Decisions Group, Menlo Park, CA, 1983a.
- , "Risk Preference," *Readings on the Principles and Applications of Decision Analysis. Vol. II*, Strategic Decisions Group, Menlo Park, CA, 1983b.
- , "Knowledge Maps," submitted to *Management Sci.*
- and J. E. MATHESON, "Influence Diagrams," *Readings on the Principles and Applications of Decision Analysis. Vol. II*, Strategic Decisions Group, Menlo Park, CA, 1983.
- KEYNES, J. M., *The General Theory of Employment, Interest, and Money*, Harcourt, Brace & World, Inc, New York, 1964, 161.
- LAPLACE, P. S., *Philosophical Essay on Probabilities*, Paris, 1812, Chapter 18.
- MYERS, I. B., "Introduction to Type," Consulting Psychologists Press, Palo Alto, CA, 1980.
- OLMSTED, S. M., "On Representing and Solving Decision Problems," Ph.D. Dissertation Engineering-Economic Systems Dept., Stanford University, Stanford, CA, December 1983.
- SHACHTER, R. D., "Evaluating Influence Diagrams," *Oper. Res.*, 34, 6 (November/December 1986), 871-882.
- SUPERTREE, Strategic Decisions Group, 3000 Sand Hill Road, Menlo Park, CA.