



INFORMS TutORials in Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Decision Analysis = Decision Engineering

James E. Matheson

To cite this entry: James E. Matheson. Decision Analysis = Decision Engineering. *In* INFORMS TutORials in Operations Research. Published online: 14 Oct 2014; 195-212.
<https://doi.org/10.1287/educ.1053.0015>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2005, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Decision Analysis = Decision Engineering

James E. Matheson

Department of Management Science and Engineering, Stanford University,
Stanford, California 94305-4026, jmatheson@smartorg.com

Abstract Decision analysis is an engineering approach to helping decision makers reach rational decisions. It incorporates systems engineering and decision theory, enhanced by the ability of modern computation to build value-based models and carry out computations needed to deal with complexity in the number of factors, uncertainties, and dynamics. It also includes the processes for reaching good decisions with real people (and real organizations) and gaining their commitment to carry them out. Decision analysis is not a single approach, but a discipline with underlying principles and procedures that are adapted to diverse situations. Decision analysis uses tools like decision theory, influence diagrams, system dynamics, game theory, etc. to “engineer” good decisions in new product development, business strategy, space system safety, etc. The paper ends with the “10 commandments of decision analysis.” You can guess the first.

Keywords decision analysis; decision engineering; tutorial

Decision Engineering

Suppose we were to ask an experienced and well-qualified engineer to specify the design of a bridge over a river crossing. Would we expect the bridge expert to contemplate for a while and then tell us the specification for a sound bridge? Of course not! The expert would examine the conditions where the bridge is to be built, the loads it must carry, the funds available for construction, etc. He would probably do a preliminary design of several feasible bridge types, calculate stresses and strains on the structural elements of the bridge, and determine which design would allow the construction of the best bridge for this purpose. He would examine the effects of climate, such as flooding and winds. He would use many different skills in assessing the environment of the bridge and the needs of its users in developing possible bridges and in coming to a choice of the best one. He would be likely to consult experts in many specialties to refine his understanding of the needs and desires of the users as he goes through this engineering process of coming to a recommendation. He will need to explain to laymen why he made or recommends certain choices, and he may need to obtain the sponsors’ trade-offs on matters such as cost versus aesthetics.

In contrast, executives are often asked to make critical choices based on gut feelings. Even an experienced engineer would not attempt to build a sound bridge this way, and the executive, even one with a great education and lots of experience, should not be expected to arrive at sound decisions this way either. A new top executive of an organization told his subordinates not to waste time on decision making, but rather to bring their hard decisions to him because he was a decision maker, and he would call the shots for them. Wall Street often plays into this concept by expecting changes in top management to instantly turn around poorly performing organizations. Like the bridge design, a good decision needs to be properly engineered; it is naive to guess the answers directly.

The term “decision analysis” was defined by Ron Howard in a 1965 presentation of a paper, “Decision Analysis: Applied Decision Theory” (Howard [7]). Here he defined decision analysis as “a formal procedure for the analysis of decisions.” By procedure he did not mean

decision theory. Ron Howard defined his original procedure as three phases: a deterministic one, where alternatives are defined and an initial formal model is developed; a probabilistic phase, where probabilities are assessed on the crucial uncertainties, risk attitude is encoded (if necessary), and tentative recommendations are made; and a post mortem phase (later called an informational or appraisal phase) to determine what iterative steps should be taken before acting on a recommendation. This decision analysis procedure is the outline of an engineering process for designing a good decision, analogous to the bridge designer's process for specifying a good bridge. In the paper, he did apply elements of decision theory, particularly the philosophy that probability captures an individual's state of mind about an uncertainty, but his use of decision theory is to fit it into a framework of an engineering approach to decision making.

As a historical note, in the mid-1960s Ron Howard and I were walking with Howard Raiffa on the Harvard campus debating what would be the best name for this profession. Howard Raiffa seemed opposed to the term decision analysis, as it seemed too analytical and academic. Ron Howard actually had similar reservations, but he discarded the term "decision engineering" because engineering has a second meaning of manipulation. A while after this debate, Ron and I were both astonished to discover that Howard Raiffa had changed the name of his new book to *Decision Analysis* (Raiffa [18]). Between the two of them, they cast this term in concrete. However, in the good sense of the word, the field is really about "decision engineering."

The Discipline of Decision Analysis

Why the word "discipline"? I am using the term discipline in the sense of a branch of knowledge or area of teaching. Major universities will organize their engineering activities around disciplines such as mechanical engineering, electrical engineering, chemical engineering, etc. These disciplines gather together all of the specialized knowledge relevant to their discipline and show how knowledge from other specialties fits in as well. Their knowledge base increases over time, but the charter of the discipline remains relatively constant.

Returning to the bridge example, it would be silly to suggest that we design a bridge without using mechanical engineering, to ask, "what are alternate methods we might use to design bridges?" with the idea of testing which method works best. The problem is that mechanical engineering is not a technique, but the field concerned with all of the techniques and knowledge about bridge building. The question does not compare apples with oranges, it compare apples with the bowl holding the apples and oranges.

Similarly, we are often asked to compare "decision analysis" with some other technique. From a decision-engineering perspective, this is the same kind of silly question. One problem is that many people think decision analysis is the technique of decision trees. They might really be asking if can we compare decision trees with some other method such as influence diagrams, or even if we can compare normative recommendations with descriptive behavior. These are appropriate questions within the field of decision analysis. Furthermore, our own profession gets caught up in similar questions, such as comparing decision analysis with game theory, system dynamics, or multiattribute utility. From a decision-engineering perspective, all of these items are tools within the discipline, and it is absurd to compare them with the whole discipline.

When I use the term decision analysis in the remainder of this paper, I mean it in this sense of the field of decision engineering.

When to Use Decision Analysis

Decision analysis is for when you do not "know" what to do. You may know what to do from long experience with a repetitive class of decisions. In many of these repetitive situations the human ability to recognize patterns, and in emergency situations to act on them quickly, is the heart of the decision problem. Gary Klein and others have developed a whole field

of “naturalistic decision making” by studying these sorts of situations (Klein [12]). The currently popular book *Blink* looks at similar situations where human pattern recognition dominates, but in the latter part of the book Gladwell shows many instances where these judgments go awry (Gladwell [4]). In an example I studied, the operators dealing with the Three Mile Island nuclear plant incident focused early on an incorrect accident scenario, then rejected all evidence to the contrary: It took a new operator who walked into the room hours later, without this prejudgment, to recognize the real situation and stop the accident. In high-stakes situations, a rapid decision analysis look at the situation might uncover mistakes. Bayes’ rule might have led to the correct diagnosis. Another situation of “knowing” involves decisions close to basic value judgments, such as a decision of whom to marry, or whether to have a child. While I have seen decision analysis used in both of these situations, many people are guided by inner values and “know” what they want. Here introspection often trumps analysis.

Decision analysis is most powerful when focused on high-value decisions dominated by uncertainty, dynamics, and complexity. It is usually more important to uncover better alternatives, cut through the complexity, and assess the biggest uncertainties, than it is to finely optimize decision variables. While there is no clear boundary, decision analysis situations and optimization situations are dominated by different issues and emphasize the use of different tools. In organizational settings, there are many reasons to use decision analysis, even if you think you know what to do. The first is to double-check your intuition and possibly hone it further. Ask if this is a situation surrounded by “fog.” Even the best-trained and experienced pilots cannot fly level in fog without instruments: They are likely to turn their plane upside down and crash. Second, you may need to explain the basis of your decision to others: superiors to justify your actions or recommendations, and subordinates to convey the understanding necessary to execute the appropriate actions and adapt to new information. Third, you may want to bring many participants into the decision process. Often there are experts available in many aspects of the situation. They provide information needed to model and assess various factors. When responsibility for the decision is diffuse, a group may use the explicit structure of decision analysis to work together to a single recommendation. Decision analysis is often used to get a wide range of participants to agree on the wisdom of the decision and to gain the knowledge and motivation to implement it well.

Philosophy

Decisions are a way to gain a degree of control over a reality too complex to understand fully. The decision analysis approach explicitly treats only a bounded portion of this reality. However, it includes tools for understanding or testing if the bounds on the problem are inappropriately set. Limiting reality to the natural bounds (scope or frame) of a particular decision (or set of decisions) helps people to temporarily focus on issues relevant to this decision and motivates them to treat this decision carefully and to reach closure. This focus helps prevent mental and organizational attention from wandering from one thing to another without resolving anything.

Declaring Decisions

For some problems, just formulating a clear decision is a big step. In educating youth, the Decision Education Foundation (DEF), a nonprofit focused primarily on teenagers, has found that getting individuals to see that they have a choice, other than simply reacting to situations, opens up a whole new way of controlling their future (Abbas et al. [1]). They have a hallway poster that says, “stop, think, and decide.” Organizations too often do (or do not do) the same thing, carrying on reactively with the same old habits without recognizing that they have important and valuable decisions to make. Unless something breaks down, organizations are not faced with the need for decisions, yet decisions, such as getting the jump

on the competition, need to happen a long time before competitive breakdown. Declaring times when decisions need to be made (and focusing the appropriate effort on making them) is a critical responsibility of organizational leadership. Many decisions that will ultimately determine the fate of an organization fall in Covey's important, but not critical, quadrant (Covey [2]). Leadership needs to continually set and monitor the decision agenda for their organization to assure that key decisions are "declared" and deliberately made with sound processes.

Bounding the Problem

The world is too complex to comprehend in one frame or model. Therefore, we bound our perspective on the world in different ways for different situations. An electrical engineer would not use quantum mechanics or Maxwell's equations to design a circuit, but he might use them to design a component or a communications system. We can only grasp a small portion of reality at a time, but if we get it right we are able to explain our universe better and better and to develop a whole host of new products and ways of organizing our businesses and our lives. Once we declare a decision we need to place appropriate bounds on the piece of reality we need to capture, and to determine what degree of verisimilitude we need in our logic and models to reach a good or satisfactory decision. We would not use the same degree of precision to order a dinner as we would to decide on a critical medical treatment.

If we picture the situation as a foggy mess, we need to put a circle around the piece of this mess we intend to capture in our decision-making process. Inside these bounds we will need to model the situation to whatever degree of approximation is required to reach the decision (or class of decisions) at hand. We follow Einstein's sage advice to use the simplest model that adequately treats the situation. Otherwise, we get into the game of trying to model the world, which is impossible even within our problem bounds. Simplicity yields more insight and motivation than complexity.

Bounds need to be tested and revised. When solving the problem gets technically difficult, for example by requiring detailed optimization or difficult assessments, I take that as a sign that the bounds need to be revisited. In strategic settings constraints are rarely fixed, even if people initially claim that they are inviolable. Pricing constraints often leads to simpler solutions, and adjusting prices gives insight into the value of letting go of the perceived constraint. The interfaces of the problem bounds with the real world are probability assessments on what might happen "out there," our perception and creation of a range of possible alternatives, and higher-level value assessments. For example, public corporations are presumably run to benefit the shareholder-owners. Therefore, we might be tempted to include shareholder preferences, the way shareholders change through market transactions, price formation in the stock market, and so forth. However, if we are designing a new product, we would typically leave out all of this detail and use maximizing the net present value of cash flow as a surrogate. In this case, we interface the bounded problem with the larger reality through assessment of a discount rate and a risk attitude. In another situation, we might be concerned about the price of energy-intensive raw materials and be tempted to model the energy industry, but it is probably wiser to directly assess the future cost of raw materials, and, if this is not sufficient for our purposes (reaching a specific decision), expand to directly assessing future energy price and how our raw material costs (and other important costs) depend on it. We adjust the problem boundaries to fit the problem structure and peoples' ability to assess value, uncertainties, and alternatives at the interface.

A great example of models going wrong is given by the well-documented saga of a Wall Street hedging firm called Long-Term Capital Management (Lowenstein [13]). The founders included Nobel Prize winners, who set out to make money from "riskless" hedging transactions by finding sophisticated arbitrage opportunities in the financial markets. Over time they came to believe their market models' predictions and to take more and more "naked" risk. In the end the market did not behave as their models predicted, and they came very

close to bringing down the whole U.S. and international financial systems. Economists like to attribute these problems to market imperfection, but they are really due to model imperfection combined with hubris. They came to believe that their model was reality and that departures from their model were defects of the world—not defects of their model. Physics learned and lamented this lesson when the theory of relativity displaced Newton's laws, but if you listen to their talk today—for example, about multiple worlds, they forget this is just an interpretation of a theory, not reality. The lesson learned over and over is that the only perfect model of reality is reality itself, and our model is only useful to describe a piece of reality for a specific purpose—in our case, to reach a specific decision or a small class of decisions. Bounding the problem correctly, and reexamining those bounds carefully, is critical (Matheson [14]).

Admitting and Treating Uncertainty

Some people think that they should guess the future and then make decisions consistent with that guess. Corporate staff members have characterized one of their jobs as telling top management what the future will bring, so that management can make decisions assuming that future. For example, if the staff projects demand growth, management can easily figure out when to add capacity. However, if they were to admit uncertainty, management might have to work out ways of hedging their capacity additions, for example, using contract manufacturing for a period until they are sure demand will be sufficient to fill a new plant. The kinds of alternatives to consider under certainty are more limited than the kinds needed to deal with uncertainty. So, admitting uncertainty is essential to developing a good set of alternatives from which to choose. In addition, with uncertainty comes the pain of risk taking, and we need to consider how to trade off risk versus return.

Once decision makers admit they are facing uncertainty, their choice could have both good and bad outcomes. An excellent decision could have a bad outcome, while a poor decision might have a good one. This means that the decision maker cannot logically be responsible for the actual outcome, but only for making good decisions. Furthermore, because the decision is made in advance of its outcome, judging the quality of the decision is most useful at the time it is being made! Therefore, we need a definition of a good decision at the time it is made, to assure decision makers that they are making good decisions and to allow others to audit the quality of decisions that have been made.


The whole decision analysis process is itself the definition of a good decision! One must ask: Has the process been applied well, were the right participants and experts brought into the process, were the right tools applied correctly, etc.? If not, the decision is suspect. The need to distinguish the quality of the decision from the quality of the outcome is the genesis of decision analysis.

The Nine Principles of the Smart Organization

In 1998, David Matheson and I introduced “nine principles of the smart organization,” illustrated here in Figure 1 (Matheson and Matheson [15]). Each principle represents a philosophy and behaviors that contribute to the support of high-quality decision making. Together they create a culture that is receptive to the decision analysis approach and that motivates its use. They are organized in triads around:

- achieving purpose—continual learning about what is valuable and finding alternatives to create and capture more value;
- understanding the environment—building an outside-in strategic perspective and capturing all of our information and insights with uncertain and dynamic thinking and formal models;
- mobilizing resources—developing disciplined decision-making processes that are mobilized through open information flow and that align and empower all participants in the process.

FIGURE 1. Nine principles of the smart organization.



Space does not allow for an explanation of these powerful principles and their ramifications, so please see the book for further detail.


When we wrote the book, we devised an “organizational IQ” test that measures an organization’s conformance to these nine principles. However, we could not prove that these principles created more value for the organizations that embraced them. After publication we administered thousands of organizational IQ tests, which to our pleasant surprise showed that organizations in the top quartile of organizational IQ approach a 50% chance of also being in the top quartile of performance, while organizations in the bottom quartile of organizational IQ have only a 10% chance of being in the top quartile of performance (Matheson and Matheson [16]). *The “high-IQ” organization has about five times the chance of being a top performer than its “low-IQ” counterpart.*

Decision Quality

So, what is a high-quality decision anyway? The ideal time to determine the quality of a decision is when you make it. A retrospective, after the results are in, may shed light on something that should not have been overlooked in the decision-making process, but usually they are used to “punish the innocent” for a bad outcome. Of course, the only time you can correct a poor-quality decision is before you commit the resources. Over the years we have evolved six elements of decision quality and arranged them into a chain as in Figure 2 (Spetzler and Keelin [20]). To determine which link is most important, think of hanging over the edge of a precipice by this chain—the most important link is the weakest one. On the other hand, links that are overbuilt have wasted resources that might have gone into reinforcing the weaker links, but it is too late now.

This leads us to a spider diagram version, as in Figure 3. When reviewing the status of a decision-making process, we should review each element for sufficiency to support a good decision. A common mistake is for decision makers to spend excessive effort gathering information without taking stock of the appropriate frame or considering the range of alternatives they might pursue—information becomes a security blanket to the detriment of a balanced decision analysis.


FIGURE 2. The decision quality chain.



The first link is framing, or bounding the problem, which we have discussed above. More problems fail here than elsewhere. The next three links, called the “decision basis” (alternatives, information and values), are the building blocks of the model of the situation and the assessments that support it. The fifth link is sound reasoning to derive the logical conclusions and insights indicating a correct course of action. Lastly, there should be commitment to action. We have all experienced personal “decisions” to diet, quit smoking, etc., which fail through the lack of will to carry them out.

This brings us to a philosophical question of whether a decision is a commitment to action or the action itself. For example, we buy a house or a car by signing legal papers. Here it

FIGURE 3. The decision quality spider diagram.



seems sensible to call that act the decision. However, in some states you have three days to change your mind on such decisions, so is the real decision not revoking the original commitment over the three-day period? In organizational settings the link between the commitment and the action becomes more tenuous. For example, an executive may decide to build a new factory. This decision results in many more decisions to pick the best site, hire contractors, buy equipment, hire and train employees to run the factory, etc. It is impossible to identify any one irrevocable commitment to build the factory. In fact, the executive decision makers almost never take the irrevocable actions that implement their decisions—they never write the checks themselves.

So, in an organizational setting it is often useful to separate the decision from the actions that implement that decision. A decision maker commits his or her organization to carry out the required actions. This distinction creates the possibility of implementation failure for many possible reasons: lack of will, resources becoming unavailable, new economic conditions, subordinates not following through, etc. A good organizational decision-making process includes mechanisms for increasing commitment to action—increasing the probability that the consequent actions will actually be carried out.

Processes and Tools

A process often used in reaching high-level business decisions is called the dialogue decision process (DDP), which was developed by the Strategic Decisions Group (SDG) and evolved through many client engagements. It is targeted at large-scale organizationally intensive decisions involving both cross-functional teams and several layers of management. It is designed to bring the right people into the process to obtain the necessary knowledge and creativity to understand the decision situation. The process is also designed to motivate and commit these people to carry out the committed decision.

Figure 4 shows a decision process “snaking” between a decision team, the people who bring corporate perspective and the power to execute the decisions, and a strategy team, the people who bring content information and who are likely to participate in implementation—carrying out the “commitment to action.” I have avoided calling the decision team the executive team because executives and managers may serve on both teams, and on different teams for different decisions. Sometimes a few decision team members even change hats and serve as content experts for the strategy team.

FIGURE 4. The dialogue decision process: Deciding.






FIGURE 5. The dialogue decision process: Implementing.



In this design, the decision team members participate in the decision-making process from the beginning, and are not simply judges who review the output of a staff analysis. They are told that if they withhold issues or information during the process, they are out of order bringing them up at the end. On a decision team, they cannot use the common ploy of absenting themselves from the process while sitting in judgment at the conclusion. These two groups carry out an orderly decision process, which goes through four major steps before reaching the decision—commitment to action—at the end of Step 4. To insure the committed action is actually carried out, a second phase with two additional steps is devoted to planning and implementing, shown in Figure 5.

Another process is the classical decision analysis cycle of Figure 6, which is used in this form for problems that are less organizationally intensive. This process emphasizes bounding and framing the problem in Step 1, understanding the important factors and relationships in the deterministic analysis of Step 2, formally dealing with uncertainty and risk attitude to reach recommendations in Step 3, and appraising whether to act now or refine aspects of the analysis further in Step 4. In a major DDP, a complete first pass through the decision analysis cycle might be included in the DDP assessment phase and a second pass included in the next two DDP phases leading up to recommendations. These two processes are different ways of adapting to various organizational settings as well as different ways to look at the decision analyst's job.

FIGURE 6. The four-phase decision analysis cycle.



I will now review a few of the key tools in the framework of the decision analysis cycle. My goal is to show how they fit into the decision-engineering process, rather than to explain the tools in depth. Good overviews and case studies are available in (Matheson and Matheson [15]) and the tools and methods are detailed in *Decision Analysis for the Professional* (McNamee and Celona [17]) as well as in extensive professional literature.

The Decision Hierarchy

A failure mode is to try to solve too many levels of a decision problem simultaneously. For example, many corporate executives are highly skilled at dealing with operation problems, where detailed planning is critical. However, detailed planning sabotages strategy because it severely limits the number of strategies that can be explored. Many so-called corporate strategies are nothing more than multiyear plans, usually with no consideration of uncertainty. Detailed planning is carried out in the second phase of a DDP (look back at Figure 5) after the major strategic direction has been set. However, for setting the strategic decision we need visions of potential landscapes, not detail.

The decision hierarchy is a great tool for setting and communicating the frame of a decision. Figure 7 illustrates the decision hierarchy for a manufacturing plant modernization problem. We take as given the statement called policy decision at the top of the hierarchy, here the policy that we will continue to manufacture—we will not be questioning that decision. However, if all strategies turn out to be poor, we might return to question that policy and revise our frame. In addition, from the bottom of the diagram we see that we will not be concerned with the design of products to be manufactured at this facility, detailed manufacturing operations or making marketing plans, as these are all tactical (or implementation) decisions. All of these items will be optimized after the strategic decisions are reached. However, we may need to put some placeholder tactical decisions in place to evaluate the strategies. We are asserting that the placeholder tactics are sufficient for selecting the best strategies, an assumption we can and should test with sensitivity analysis. The concerns of this decision analysis are the items listed in the center—the strategic decisions—the basic configuration and capabilities of the plant along with high-level strategy for marketing the products it produces.

The decision hierarchy specifies the strategic decision areas that are to be detailed by one or more columns in a strategy table, as illustrated in Figure 8. In Figure 9 we develop four potential strategies by the selection of one option from each column. A real application typically would have 10 to 30 columns, possibly specified in a hierarchy of strategy tables. It would not make logical sense to try to enumerate and model all of the strategies that could

FIGURE 7. A decision hierarchy.




FIGURE 8. The decision hierarchy determines decision columns in a strategy table.




FIGURE 9. Example strategy table.

Strategy Alternatives	Plant Configuration and Location	Technological Stretch	Product Range	Quality and Cost Position	Marketing Strategy
Aggressive Modernization	Current	State of art	Full line	Quality and cost leadership	Sell quality and influence market growth
Moderate Modernization	Close #1	Proven	One basic line and specialties	Improved quality; deferred cost reduction	Sell quality
Consolidation	Close #1; build domestic greenfield	Current	Value-added specialties only	Minimal quality improvements	Current
Run Out	Close #1; build foreign greenfield				


be specified in such a table. Rather, we ask the participants to develop about three to six significantly different alternatives that make up coherent sets of selections to evaluate and test—we can add others based on the insights provided by the initial set.

It is critical that each option in the strategy table represents an actionable, although high-level, specification of direction that the participants could plan in detail if it were selected—so we often ask participants to write an essay of a few paragraphs describing each one. The decision model will contain cost and other features of each possible choice.

The Influence Diagram

The influence diagram furnishes a transition from more qualitative framing to quantitative modeling. Figure 10 is an example influence diagram. On the surface it looks like just another block diagram, but it actually can be used to precisely specify an entire decision model and all of its assessments (Howard and Matheson [9]). Each rectangle represents a decision, which is informed by arrows from any nodes preceding it. You can think of the alternatives as inside the rectangle, to be viewed by opening the box (double-clicking). Each

FIGURE 10. Influence diagram.



oval represents an uncertainty, containing a set of probability distributions conditioned by the graphical nodes immediately preceding it. A special eight-sided node represents a value model. Sometimes deterministic nodes are added, usually indicated by a double oval as a degenerate form of an uncertainty, which allows easy specification of model relationships and even time series. Influence diagrams are a breakthrough representation for capturing uncertain decision situations. On the surface they are a simple picture of the elements of the situation, but they enable automatic model generation, specification of decision trees, and new algorithms for directly solving influence diagrams themselves. They have become a ubiquitous tool for decision analysis, artificial intelligence, and other fields as well.


The Decision Model

Usually after relationships are specified and understood in an influence diagram, an economic model is either automatically generated or manually constructed to calculate first deterministic, then probabilistic results. An example strategic business model is illustrated in Figure 11, which we sometimes refer to as a “boiler diagram.” The center is a deterministic model of the business, markets, competitors, dynamics, etc.—“the boiler”—often implemented as an Excel spreadsheet. Feeding into it from the bottom are all of the decisions we might make, here shown as paths through a strategy table. Feeding into it from the left are all of the uncertain variables, such as those specified by the ovals in the influence diagram. Out of the boiler come results, in business usually in the form of uncertain cash flows, that get discounted to the present, adjusted for risk, and measured by the shareholder value meter. In principle, we just select the decision that maximizes the reading on this meter. However, in practice, we proceed in steps, learning and refining the model and assessments as we go along.

The Tornado Diagram

In the deterministic phase, we assess the continuous uncertainties by estimating the range of probability, most often expressed as the 10-50-90 percentiles. Early in decision analysis practice we would quickly assess these ranges and return to more carefully assess them in the probabilistic phase. Today, we recommend that particular care be taken to remove biases from these range assessments by the same procedures that have been established

FIGURE 11. Key elements of a strategic business model.




for complete probability assessment. There are two reasons for this advice: First, we have learned that biases can be so strong that even deterministic sensitivity can be untrustworthy; and second, in rapid analysis these ranges are often used to create the initial probabilistic assessment—for example, by assigning them discrete probabilities of 0.25, .05, and 0.25, so it is important that they be trustworthy assessments. For more information on biases, see the original *Science* article of Kahneman and Tversky [5], the probability assessment procedure developed with their help (Spetzler and Staël van Holstein [21]), and the more popular book *Decision Traps* (Russo and Shoemaker [19]).

Once the uncertainties have been assessed in consistent 10-50-90 ranges, the deterministic model is used to vary each variable over its range while holding the other variables at their medians, and record the value measure, typically the NPV of cash flow. If these sensitivities are monotonic, the high and low values occur at the end of each range. For each variable we record the difference between the high and low values (the swing) and order the variables in descending order of these differences. If we plot these results as a bar graph, we end up with a characteristic tornado shape, as shown in the example tornado diagram of Figure 12. This figure and the next are taken from Chapter 9 of (Matheson and Matheson [15]), which provides additional discussion of this case.

Looking at the tornado diagram, we see that the first few variables contribute most of the uncertainty. In a simple additive independent model the width of the overall uncertainty bar would be the square root of the sum of the squares of these individual bars, amplifying the dominating contribution of the first few bars. Usually we find the decision-making team has been neglecting one of the first few variables and wasting a lot of time on the last few, in this case the second-to-last bar. While we intend to proceed to a probabilistic phase, we can make important observations by studying the tornado. For example, if we can easily get information about the first few bars, and if that information would change our decisions, we should probably go get it—it has high value of information. If we have a way of favorably controlling an uncertainty, say by purchasing a technology or even a small company, then

FIGURE 12. Sensitivity analysis: Tornado diagram.




we should seriously consider doing it—it has high value of control. If there is no way to get information about one of these top variables, then quit worrying about it—it is existential uncertainty you have to live with, just like your own time of death.

Proceeding to the probabilistic phase, we carry out a complete probabilistic analysis, which will automatically include more combinations of these variables than in the individual sensitivities of the tornado diagram. If we have been quick and dirty in making the range assessments or if we see ways to get better information on the top variables, for example, by interviewing an easily available expert, we should carry out that reassessment. However, if we have been careful assessing and debiasing the 10-50-90 ranges, we may wish to convert these ranges into probability distributions for a first cut at probabilistic analysis.

A set of final probability distributions for the HDTV example is shown in Figure 13. The original three strategies were aggressive analog, evolutionary improvement, and infor-

FIGURE 13. Evaluation of several strategies



mation age, which presented a risk-return dilemma. However, with the insights of the first analysis a new hybrid strategy, sooner and later, was invented, which captured the low end of evolutionary improvement and the high end of information age. This alternative almost stochastically dominates the other alternatives and was accepted as the new corporate strategy, without the need to deal further with risk attitude. It is surprising how often this is the case in business problems with big, discrete alternatives, but of course in problems where you are deciding how much to invest in a venture, or how to bid, risk attitude usually is a critical input.

Value-Based Modeling

Coming from the business decision-making context, we have developed an approach called value-based decision making. The first principle is to make decisions based on monetary value today. If you are spending money today, then you need to explicitly or implicitly justify this expenditure as being worth this value. Implicit judgments, or vague assessments like point-scoring systems, easily lead you astray and into the realm of social/political pressures and fuzzy thinking. Measuring monetary value provides a depersonalized standard of objectivity. A Dilbert cartoon character asked, “Why don’t we just figure out what creates the most value and do that?” This is what we advise!

How do we get to this measure? If we first think deterministically, we often model scenarios that produce a cash flow over time (before or after taxes, free cash flow, etc.). We ask the decision maker, “If you had certain title to this cash flow, what would you trade for in cash today?” We call his answer the present equivalent of the cash flow stream. The answer does not have to be a linear function of the cash flows or stable over time, and I have seen cases in private companies where it is not. However, in most public corporations you will get the answer, “Just discount the cash flow by the corporate cost of capital (risk free) and call that the present equivalent.” This is how we end up using net present value (NPV) as the present equivalent. Also, in the presence of a linear bank it can be shown that this procedure creates the most wealth—it creates the biggest pie.

Now what about uncertainty? First we reduce time streams to present equivalents (or NPVs), so our decision maker is now facing lotteries on present money, an easier situation to contemplate. We then ask the decision maker, “If you faced a particular lottery, how much would you trade it for in a certain payment?” We call this amount the certain equivalent. (A historical note: The reason we avoid the East Coast “certainty equivalent” is that presently equivalent does not make grammatical sense, so we chose parallel terms.) Like the linear discounting case, there are many good reasons to capture risk attitude by an exponential utility function (constant risk attitude), at least approximately. In a dynamic decision situation and using these approximations, we would roll back a decision tree by discounting over time to present equivalents, reducing lotteries to certain equivalents, discounting again, and repeating this process until we work back to the present. For consistency with utility theory, one must increase future risk tolerances by the compounding rate used in the NPV.

What about nonmonetary considerations? First, usually nonmonetary results must be paid for, whether it is keeping employment in the local town, or maintaining a safe environment. Usually these items can be treated by willingness to pay. In a corporate setting, where money is being invested to create value, willingness to pay added to the monetary cash flow seems to work better than multiattribute utility theory. It gives a clearer line of sight to value creation.

Instant Decision Analysis

In the modern world of the Internet and hyperactivity, everyone wants and expects instant gratification. There is little time to stop and think. Can decision analysis deliver results in this new paradigm?

Yes and no! Deep complex problems, especially ones with high organizational (social-political) involvement take time to work through. On the other hand, computer power and the Internet have put analysis and information retrieval on a new footing. In repetitive classes of problems, where participants can be trained to understand their roles and expectations in the decision-making process, there is great potential for rapid high-quality decision making. Repetitive decision classes allow formulation of reusable decision models, quick links to appropriate information sources, and standing interrelationships among people who know their roles.

Importantly, the customer has changed over the years. When I began doing decision analysis in 1964, few managers or executives had even heard of decision or probability theory, and computers were still huge mysterious machines in large air-conditioned rooms. These executives needed training and convincing every step of the way. Today, most executives have been exposed to decision and probability theory, use their own computers every day, use the Internet fluently, and are predisposed to accept decision analysis methods and results if delivered on their terms. Tom Friedman has just published a provocative book on the revolutionary changes taking place, called *The World is Flat* (Friedman [3]).

However, the new paradigm is very different from the old decision-consulting paradigm. In brief, some of the key shifts are:

- from decision doctor to capability builder,
- from client to customer,
- from consultant knows best to customer knows best,
- from managing large projects to empowering teams,
- from “custom tailored” to “off the rack”,
- from time out for a DDP to decision quality in the work flow, and
- from empowered consultants to empowered customers.

For the last five years, I have been part of a software and training startup, SmartOrg, Inc., which has been learning how to deliver into this new paradigm. We have been successful in helping many organizations with dozens of innovation decisions (e.g., R&D, product development) and with balancing portfolios of decisions. The new paradigm can work, but delivering on this promise is very demanding, and the customers are just learning to differentiate between data-rich operational systems and more judgmentally oriented strategic systems. However, the early adopters will attain a new source of competitive advantage: rapid high-quality decision making.

The Professional Decision Analyst

Whom Do You Work For?

In theory, a decision analyst serves a decision maker. On occasion, you might be engaged by an individual allocating his or her own resources to ends you consider ethical. This single individual has the choice of making all of the assessments or of delegating some of them to others. The axioms and principles of decision analysis demand that he ultimately accepts inputs from others as his own, such as probability and value assessments. For example, if he engages an expert to assess a probability distribution on sales of a new product given a certain pricing policy, he should review the expert's thinking and assessments and then assert from everything he knows, including his confidence in the expert and the decision analyst who made the assessment, that he believes these assessments represent his own judgments as he makes this decision. On occasion, decision makers do reject expert assessments.

However, when you are engaged by an organization, who is the decision maker? If you are lucky, there will be a single fiduciary responsible for the decision at hand. Here the situation is similar to the personal one, except that both you and the decision maker are ethically obligated to serve the organization who has engaged you, not decision maker self-

interests. Because of bounded rationality, you may not have a clear line of sight to the total organizational situation, but you and the decision maker should be aligned in this purpose. If the decision maker is furthering his or her own agenda at the expense of the organization, I advise you not to take the assignment.

My philosophy is to always work toward the best interest of the organization that engages me. When the single decision maker is not so clear, I make sure the team is working together in the interests of their organization, and that they are willing to make their arguments through a logical decision analysis rather than organizational politics. The arguments should be about the structure and assessments, not directly about picking the alternative they may initially advocate. Sensitivity analysis often shows that what seem to be major differences lead to identical decisions. It directs them to explore the most important difference, often leading to gathering additional information.

The ultimate limitation is the desire of people to be rational. If that desire is not sufficiently present, this analysis is likely to be professionally and ethically unsatisfactory.

Scope of Responsibility

Another concern is the responsibility of the decision analysis for the content of the analysis. Should a decision analyst provide some of the content, or vouch for the content? When a decision analyst recommends a decision, most decision makers do not clearly distinguish a logical recommendation,—“this is the best you can do with your information”—from an absolute recommendation,—“I recommend that this action is best for you.” I have had many intense discussions with colleagues on this issue.

Personally, I come down with mixed advice. Providing his own detailed assessments leads the decision analyst away from objectivity. Once you provide inputs, you have a credibility need to defend them, and you might spend too much time on your own area of responsibility rather than conducting a balanced analysis. However, being generally knowledgeable about the field in which you are working helps with communication and rapid analysis, as long as you can and do still ask the “dumb” rhetorical questions that lead to new thoughts. An expert usually cannot do that.

However, when I make a recommendation, I have moved from the position of logical recommendation to a holistic one. A decision analyst is not just an observer, but also an integral part of bounding the problem, developing a sound model, evoking assessments, etc. I am willing to say that having conducted a decision analysis, as far as I know we have appropriately framed and modeled the problem, considered an adequate set of alternatives, captured reliable information from trustworthy and knowledgeable sources, come to well-reasoned recommendations, and provided a setting to achieve commitment to action. The ultimate test in a business situation is that given what I now know (and the decision analyst often ends up with the most complete perspective on the situation), would I be willing to bet my own money this way? If I have truly engineered a good decision, I should be eager to make this bet.

The 10 Commandments of Decision Analysis

1. Decision analysis is the one master discipline of decision engineering—do not get lost following fads.
2. A single decision maker is the fiduciary responsible for declaring the decision and accepting delegated analyses and assessments—work for the decision maker.
3. Strive to construct a single value measure, in monetary terms—nobody ever put a score in the bank.
4. Have no regret—produce the biggest pie for the organization for whom you are working—do not covet the pie you did not get or the other guy’s pie.
5. Beware of internal rates of return—nobody ever put a rate of return in the bank either.

6. Beware of difference lotteries—they never happen.
7. Beware of triage—you can probably add more value to the big “no-brainers” than to the borderline cases where the choice has little impact on value.
8. If the problem is hard, change your frame—avoid picking constraints as problem bounds.
9. Start simple and iterate—use the simplest model that gives insight and decision clarity.
10. Change with the times and keep up with the new paradigm—the world is flat, after all.

References

- [1] Ali Abbas, Dave Reiter, Carl Spetzler, and Steve Tani. Teaching teens how to make good decisions. *OR/MS Today* 31(4), 2004.
- [2] Stephen R. Covey. *The Seven Habits of Highly Effective People*. Covey Leadership Center, Fireside, NY, 1989.
- [3] Thomas L. Friedman. *The World Is Flat: A Brief History of the Twenty-First Century*. Farrar, Straus and Giroux, New York, 2005.
- [4] Malcolm Gladwell. *Blink: The Power of Thinking Without Thinking*. Little, Brown & Co., New York, 2005.
- [5] D. Kahneman and A. Tversky. Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–31, 1974.
- [6] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–291, 1979.
- [7] Ronald A. Howard. Decision analysis: Applied decision theory. D. B. Hertz and J. Melese, eds. *Proceedings of the Fourth International Conference on Operational Research*, Wiley-Interscience, New York, 55–71, 1966.
- [8] Ronald A. Howard. On making life and death decisions. R. C. Schwing and W. A. Albers, eds. *Societal Risk Assessment, How Safe Is Save Enough?* General Motors Research Laboratories, Plenum Press, New York, 89–113, 1980.
- [9] Ronald A. Howard and James E. Matheson. Influence diagrams. Department of Engineering-Economic Systems, Stanford University, Stanford, CA, 1979.
- [10] Ronald A. Howard and James E. Matheson, eds. *Readings on the Principles and Applications of Decision Analysis*. Strategic Decision Group, Menlo Park, CA, 1983.
- [11] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives*. John Wiley & Sons, New York, 1976.
- [12] Gary Klein. *Sources of Power*. MIT Press, Cambridge, MA, 1998.
- [13] Roger Lowenstein. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House, New York, 2000.
- [14] David Matheson. When should you reexamine your frame? Ph.D. dissertation, Department of Engineering-Economic Systems, Stanford University, Stanford, CA, 1990.
- [15] David Matheson and Jim Matheson. *The Smart Organization: Creating Value Through Strategic R&D*. Harvard Business School Press, Boston, MA, 1998.
- [16] David Matheson and James E. Matheson. Smart organizations perform better. *Research Technology Management* 44(4):49–54, 2001.
- [17] Peter McNamee and John Celona. *Decision Analysis for the Professional*, 3rd ed. SmartOrg, Inc., Menlo Park, CA, 2001.
- [18] Howard Raiffa. *Decision Analysis*. Addison Wesley, 1968.
- [19] J. Edward Russo and Paul J. H. Schoemaker. *Decision Traps: The Ten Barriers to Brilliant Decision-Making and How to Overcome Them*. Fireside, New York, 1989.
- [20] Carl S. Spetzler and Tom Keelin. Decision quality: Opportunity for leadership in total quality management. Strategic Decision Group, 1992.
- [21] Carl S. Spetzler and Carl-Axel S. Staël van Holstein. Probability encoding in decision analysis. *Management Science* 22:340–358, 1975.
- [22] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.