Principles of Decision Analysis

NCOSE Symposium

10 August 1994

Dr. C. Robert Kenley

O/M1-30 B/158

Lockheed Missiles & Space Company, Inc.

P.O. Box 3504

Sunnyvale, CA 94088-3504

Phone : (408) 756-6329

Fax: (408) 743-1795

E-mail: Bob_Kenley@qm.is.lmsc.lockheed.com

ABSTRACT

The discipline of decision analysis has been used successfully to assist strategic planners determine the best way to invest major resources in the private and public sector enterprises. It also can be used to determine the maximum amount to be paid to reduce uncertainty through research, experimentation, and testing before making a final decision to develop and deploy a system concept. Many in the field of systems engineering are unfamiliar with the discipline, and this paper presents an overview of the principles of decision analysis.

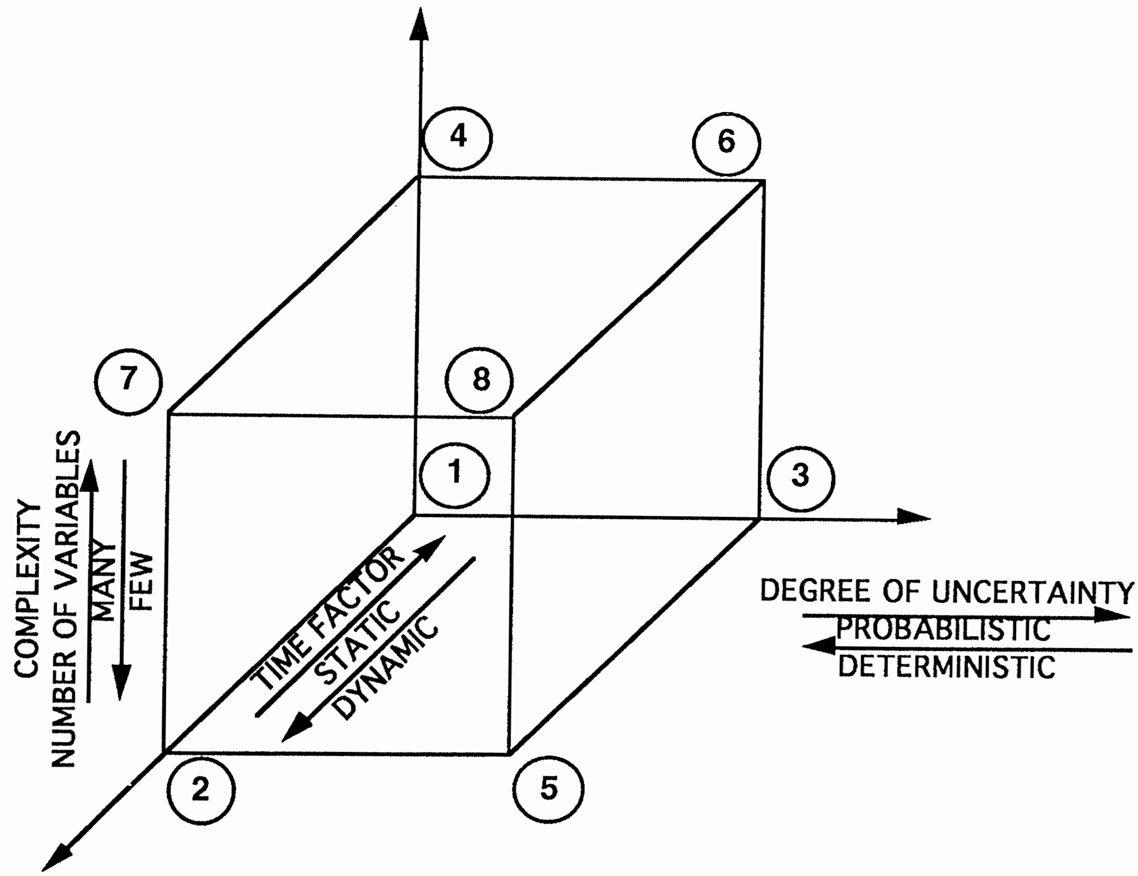

Figure 1. The Problem Space for Decision Making

Newton's calculus is represented by the corner of the cube labeled 1 and models the behavior of a single variable, position, with respect to a single parameter, time. Extension of this to multivariate problems allowed for calculation of geometric areas and volumes. Eventually, complex dynamical problems were codified as differential equations, shown in the corner labeled 2. While Newton, Leibniz, and others were settling the issues of time and motion, others such as Laplace, Gauss, and Bayes were of developing models for uncertainty and simple decision making under uncertainty needed to wisely invest in games of chance, where usually the wisest action is to not invest at all. This is the corner labeled 3. By 1951, a graduate student at Berkeley named Dantzig figured out how to solve a linear program by the simplex method, and linear and non-linear programming in the field of operations research exploded to solve decision problems with many variables, as indicated by the corner labeled 4. Extension of uncertainty to many variables interacting in time was developed by a Russian named Markov, shown in the corner labeled 5. In the 1950's Wald, Raiffa, and others had developed statistical decision theory to the point where many uncertain variables could be modeled, represented by the corner labeled 6. As engineers attempted to reach the moon, control orbiting spacecraft, and run steel mills with computers, multivariate control theory was refined, as represented by the corner labeled 7. Decision analysts are attacking problems that lie in the realm of problem solving that requires modeling many uncertain variables that evolve over time. They are trained in and borrow from all of the previous techniques described. They make many approximations to help them get to an answer. They find shortcuts to representing complex problems to enable the decision maker to make a good decision in a timely fashion. They also tend to be quite proficient in computer programming. In short, a decision analyst is a systems engineer who works with the business planning staff.

3. Decision Analysis Modeling and Process

The modeling approach for decision analysis is similar to the approach that one would take for any large-scale systems analysis problem. Shown in Figure 2 is the standard block diagram for the decision analysis modeling paradigm.

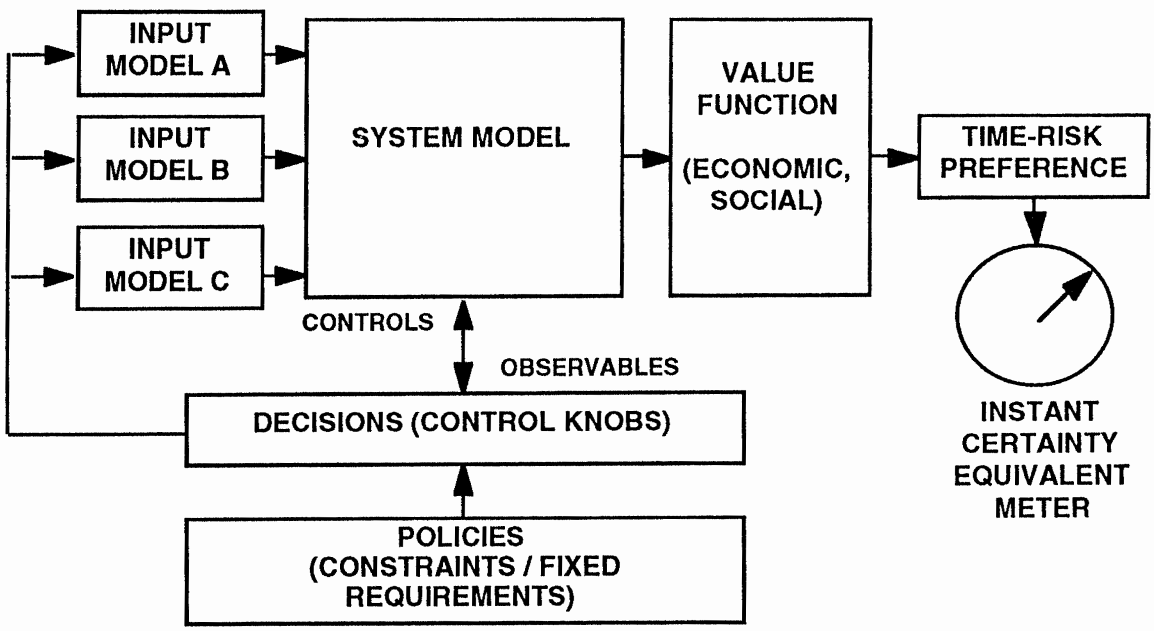

Figure 2. A Decision System Modeling Paradigm

The approach allows for many different input models to be developed and interface with the system model for the decision problem being attacked. Examples of input models are the electrical demand forecast model for the nation of Mexico, the results of a scientific investigation of the response of dental cavities to treatment by various types of preventative therapies, and the force laydown of Soviet intercontinental ballistic missiles. A decision analyst who is too proud to accept the outputs of others is one that is likely to be unemployed.

The objective of the decision analyst is to make use of what is available and interface the inputs with the system models in which he or she has specific expertise. The favorite system model for decision analyst up until the 1980's was decision trees, which model uncertainty by use of probability trees, and model preference and values by attaching a function at end of the branches of the tree. Since the 1980's, much of the modeling has switched over to influence diagrams, which help to bridge the communication gap that the decision makers and input model experts have with a decision analyst. Influence diagrams have the same appearance as interaction diagrams described in systems engineering circles, but they have very well-defined rules that allow manipulation of the diagrams to solve problems. For discrete random variables and Gaussian continuous random variables, data structures and algorithms have been defined to solve decision problems using influence diagrams instead of trees. Prior to influence diagrams, the decision analyst would tend to discuss the problem in the form that the computer was solving it, rather than in a natural cause and effect relationship more familiar to scientific experts and decision makers.

The value function most popular is the net present value of corporate profits; however, in political environments, a multiattribute model is used that is similar to those used by systems engineers for trade studies. For a cavities prevention R&D investment problem with the National Institutes of Health, a ratio of net present R&D expenditures to net present cavities reduced was developed to narrow down the alternatives. If a dollar spent saved someone from having a cavity, it was a good deal. If the ratio showed that it cost $2000 to save a cavity, then it was reasonable to question the value of further research on such a treatment.

The time preference function most often used is net present value. Some more sophisticated models have been studied to assess time preference when estate planning and environmental protection decisions are involved. The most popular technique for measuring risk attitude involves assessing a single parameter from a decision maker. The basic question is "For what value of x are you willing to risk a 50-50 chance of losing x or gaining 2x in return?" The answer to this question drives a simple exponential utility model for risk preference. For increasingly larger values of x, the exponential model becomes linear, such that maximizing expected return represents the decision maker's attitude toward taking risks. Most aerospace companies have a very low value of x relative to their commercial counterparts. Within a corporation, the value of x increases as executives are questioned up the management chain to the CEO level. Quite often, big risks are not taken, because corporate officers do not define the corporate-level risk policy (the baseline value of x) and educate the lower level executives responsible for R&D and marketing investments.

Borrowing from systems analysis, the decision analyst provides parametric results as a function of varying the decisions or control knobs, and the input parameters to the various model blocks. This entire process of decision making that includes sensitivity analyses of various types is shown in Figure 3.

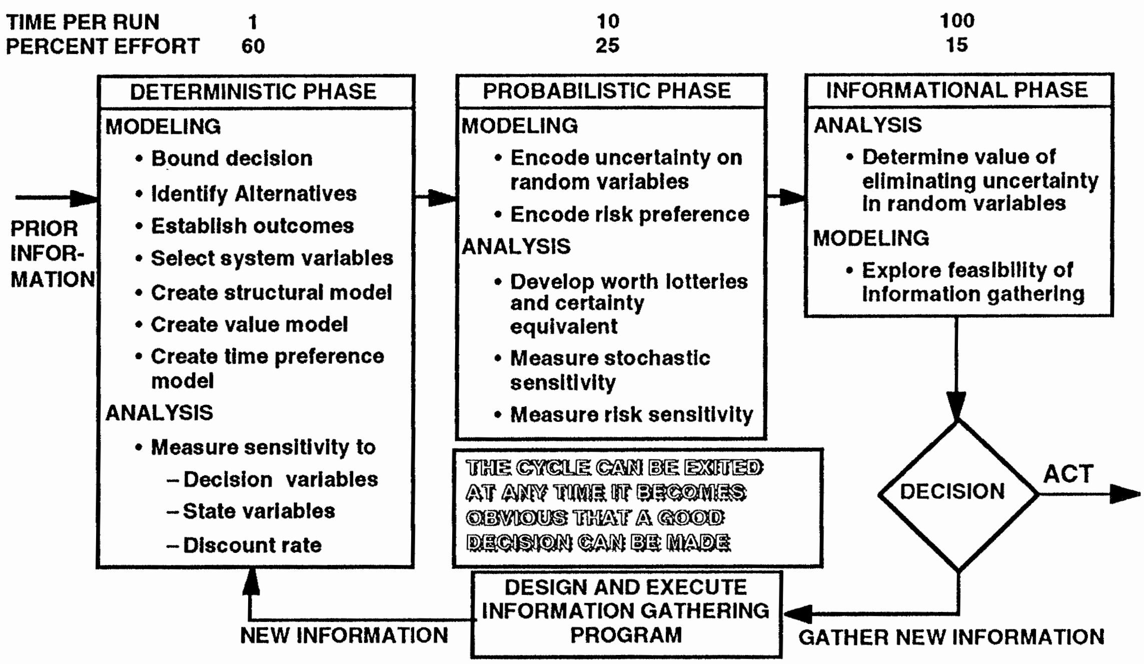

Figure 3. The Decision Analysis Cycle

Because the amount of processing required increases by an order of magnitude when uncertainty is added to a problem, decision analysts spend a significant amount of time analyzing the problem without probability models. This is a significant phase where the "trade space" and relationships among the variables are established. Sensitivity analysis is performed for the purpose of determining if there is a "killer trade" where one decision alternative dominates the others, and the analysis can be stopped and action taken. Sensitivity analysis is performed on the variables that are uncertain to see if they impact the optimal decision or vary the value of the outcome significantly. Those which have a significant impact either on the optimal decision or the total expected value of the decision are modeled probabilistically in the next phase of the process. Those which have little impact are set at their mean value for the next phase of analysis.

The probabilistic phase repeats the previous deterministic phase, except with probability added into the modeling. For many variables, there is insufficient marketing or scientific data to determine an empirical probability distribution of the outcomes. Decision analysts have developed many techniques to help experts encode the uncertainty that they have with respect to outcomes without having to perform exhaustive data analysis. Part of the training of a decision analyst is in interview techniques to help elicit probabilities from people who are uncomfortable putting down numbers. They are similar to the previously described (-x, +2x) technique for measuring attitudes toward taking risks.

In addition to extensive probability modeling and decision modeling, decision analysis has added a unique contribution to making decisions. This is the evaluation of the economic value of eliminating or reducing uncertainty. The concept was developed under statistical decision theory, but decision analysis has built much of its reputation on applying the concept to real world problems. The concept is illustrated in Figure 4.

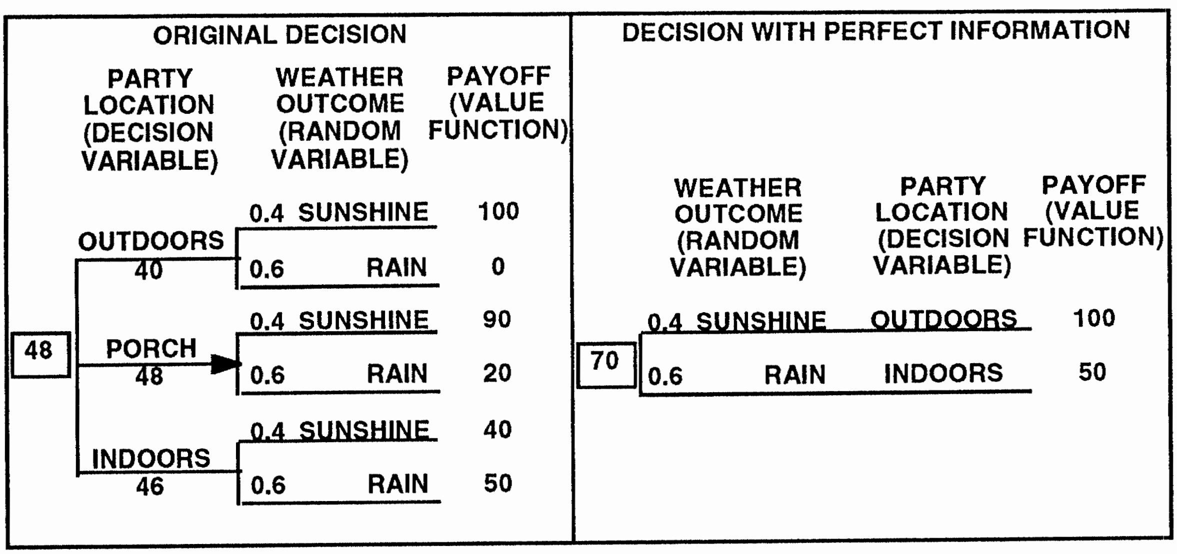

Figure 4. Value of Eliminating Uncertainty Example

The left-hand frame of Figure 4 depicts a simple decision that has to be made regarding the location of a birthday party. There's a 60% chance of rain, and the caterer must be instructed where to set things up. The optimal place to be is in the yard on a beautiful day, so the value of this situation is set at 100. The worst thing is to have the party outdoors in the rain, so it's valued at 0. Having the party on the porch while it's raining is not really good, since the porch is south-facing, and in California the wind comes from the south when it's raining. Other scenarios and the value of their various outcomes are assessed. Using maximum expected value as the decision criterion, the optimal location for the party is on the porch. The right-hand frame shows what would happen if we were guaranteed a 100% correct forecast from an perfect weather forecaster, or if we had the flexibility with the caterer and the partyers to adapt in real-time to the weather. Prior to deciding to purchase the perfect forecast, our information is that there is a 60% chance of rain; therefore, there is a 60% chance that the perfect forecaster will tell us that it's going to rain and a 40% chance of being told that it will be sunny. Once we have the forecast, we would chose to be indoors or outdoors according to this valuable information. In the left-hand frame the expected value of the decision to hold the party on the porch is 48. In the right-hand frame the expected value of our decision has been increased to 70; therefore, the maximum we are willing to pay for the perfect forecast is 22.

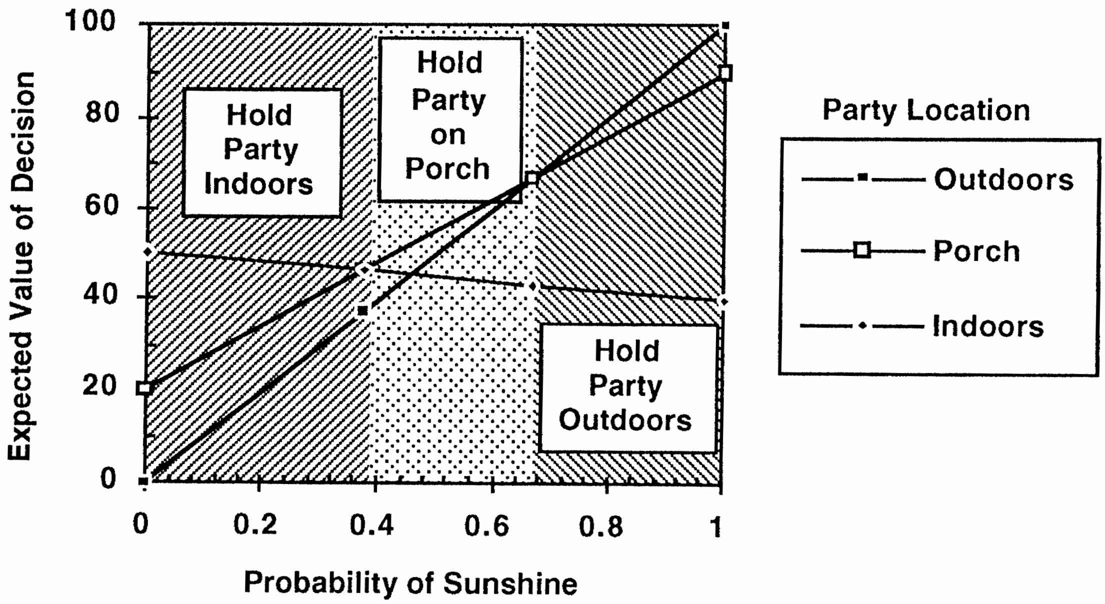

Figure 5. Sensitivity of Value to Probability of Sunshine

The concept of totally eliminating of uncertainty can be extended to evaluating the effect of partial elimination of uncertainty. This is often the case when we collect additional data before proceeding with a project to the next phase. Potential applications in the aerospace field are determining how much to spend for engineering development units, data collection experiments, integrated system tests, independent research and development, and a myriad of items. For the party problem, we will calculate the value of partially reducing uncertainty. First, we need to understand the sensitivity of the original decision to the probability of sunshine, shown in Figure 5.

Figure 5 shows how there are various break points as the probability of sunshine proceeds from 0 to 1. The original assumption was that the probability of sunshine was 0.4 and the optimal decision was to hold the party on the porch. Previously , we used the payoff values of the best decision for the probability of sunshine equal to 0 and to 1 to determine the expected value of a perfect forecast. When uncertainty is reduced rather than eliminated, it changes our probability of sunshine from 0.4 to numbers other than 0 or 1, depending on the forecast accuracy of the new data and the actual reported data as shown in Figure 6.

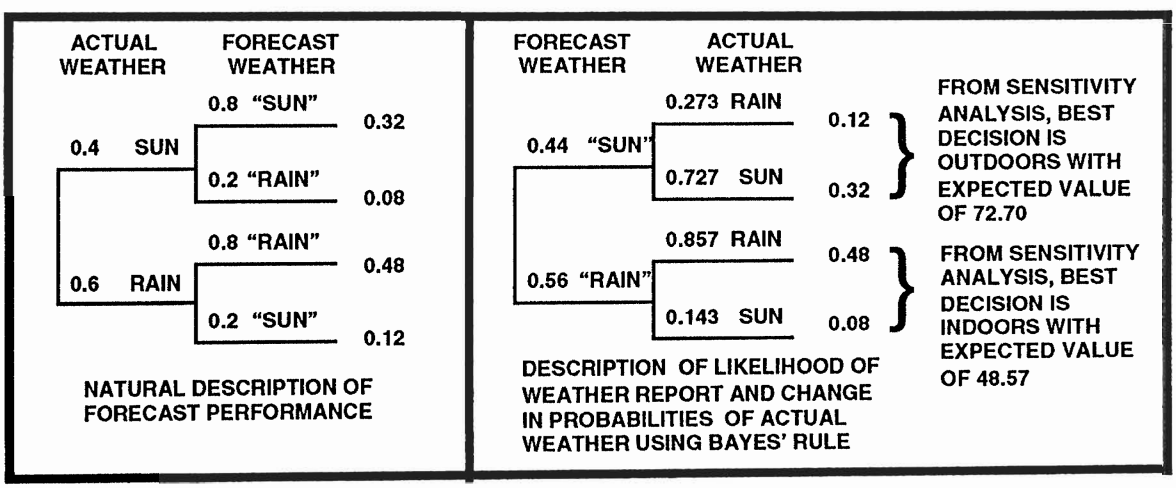

Figure 6. Value of Imperfect Information Example

In Figure 6, the opportunity to buy a weather forecast that is 80% accurate is analyzed. In the left-hand frame, we show that the actual weather according to our prior information has a 40% chance of being sunny. If it turns out to be sunny, there is an 80% chance that the imperfect forecast would have predicted sunny, as shown across the top path across the tree in the left-hand frame. Another possibility is that the weather actually turned out to be rainy, but the imperfect forecast predicted sun, as shown in the bottom path across the left-hand tree. The total probability that the 80% accurate forecast will predict that it is sunny is 0.4 x 0.8 + 0.6 x 0.2 = 0.44, which is shown in the right-hand frame. The left-hand frame is the natural, or cause-and-effect way of describing the problem, and the right-hand frame is the time order in which the events actually happen with the forecast being given before the truth is revealed. The right-hand frame shows that there is a 44% chance that the forecast will be for sun, and that given this forecast, there is a 72.7% chance of it being sunny. Notice that the value is not 80%. This is because in the calculus of probability, we do not totally discard our original opinion that the chance of rain is 40%. This is an important feature of decision analysis, that allows different information sources of various accuracies to be combined. To calculate the value of the imperfect information, the optimal decisions in light of the 80% correct forecaster reports are determined. To calculate the optimal decisions, the information from the sensitivity analysis is required. The expected value of using the 80% accurate forecast is 0.44 x 72.70 + 0.56 x 48.57 = 59.19. Without the forecast, the optimal choice was to hold the party on the porch, which had an expected value of 48. The value of the imperfect information is then 59.19 - 48 = 11.19.

The amount of calculation required to consider the possibility of obtaining more information is significantly more than the calculation required to make the original decision. This calculation for making the original decision is shown in the left-hand frame of Figure 4. The calculation in all of Figures 5 and 6 are required to determine the value of eliminating or reducing uncertainty. That is why the amount of time per run for the informational phase of decision analysis is an order of magnitude greater than for the probabilistic phase, as shown in Figure 3.

4. Good Engineering Practice

Some of the lessons learned by decision analysts during the practice of their discipline have analogous applications to systems engineers.

The first rule they have is to identify and isolate components of the decision. This follows along the lines of defining system functions and hierarchies, understanding team skills and background, and other good systems engineering practices. By isolating the components and developing structural models, they have found that it locates and resolves sources of disagreement, similar to the results of the bulk of activity in systems engineering. Just as systems engineers allocate responsibility for various parts to the problem to different personnel depending on their background, the decision analyst has a similar process. Defining alternatives to be traded usually is the responsibility of senior personnel close to the design and execution of the project. Information and probability assessment usually is done with the aid of technologists, financial analysts, and marketing personnel. The modeling of preferences is best left to the highest level of decision maker practical. For defense applications, this may mean going to the customer's program manager or even the bosses of the government program manager to understand the priorities of the real decision maker when using decision analysis for a major trade study, such as determining the amount of system testing required. When the objective is to select decisions to maximize corporate return on investment, the usefulness of the decision analysis exercise is directly proportional to the level of the corporate officer involved in determining risk and time preferences.

The second rule is to restrict language to clearly and commonly understood terms. This is equivalent to writing clear, concise requirements in systems engineering, and developing glossaries of terms to assist in requirements definition.

The third rule is to establish the economic scope. Analysts advise their clients to spend less than 1% of the economic scope (program cost or potential profit) on decision analysis. Decision analysis is analogous to the completion of key trade studies during the early phase of a program. The traditional systems engineering standing army averages about 10% on aerospace programs. The 1% figure for decision analysis might match the amount spent on key system trade studies during the life of the program. This says that 90% of the time systems engineers are just pushing paper around to keep the system from falling apart! Actually the other 90% of the time spent by systems engineers is the vital activity of system synthesis. Just recently, Henry Mintzberg pointed out the equivalent concept to strategic planners (another name for decision analysts) who try to increase their role beyond its value to the organization. He stated that ,"Strategic planning isn't strategic thinking. One is analysis, and the other is synthesis." An alternative to the third rule might be that it is vital to balance "paralysis by analysis" against "extinction by instinct".

The fourth rule is to stress detachment. By analogy, "a lawyer who represents himself has a fool for a client". Even in an aerospace systems engineering environment, this makes sense. First, the job of the systems engineer is to make sure the customer is getting what they ask for in terms of performance, cost, and utility of the system. Second, because most systems engineers are on the front end of a project where the loss of a contract is highly likely, detaching ourselves to some degree will help when the program is canceled or the program is lost to the competitor.

The fifth rule is to avoid retrospective studies - there is no value in 20/20 hindsight. At first glance, this seems to not apply to systems engineering. Certainly, no systems engineer has ever been asked to perform a trade study that will justify a design that cannot be changed!

5. Summary

Decision analysis is a discipline that closely parallels the analysis activities used in aerospace systems engineering. It would be well worth the time for current systems engineers to obtain training in the discipline to expand their knowledge base for two reasons. It will make them better systems engineers for the classical technical problems that they are working, and it will give them skills to proved significant help to corporate executives as they attempt to sell aerospace technology in the commercial arena. One immediate application area to current systems engineering is determination of the maximum amount to be paid to reduce uncertainty through research, experimentation, and testing before making a final decision to develop and deploy a system concept. Training to apply the discipline usually takes one year of graduate level training. At Stanford, this includes one quarter of probabilistic analysis, one quarter of beginning decision analysis, and one quarter in a project class where an actual decision analysis is performed (including the search for a client). An experienced systems engineer already has picked up the many of good engineering practices that enable the technology to be transferred to the environment of the customers they are supporting.

6. References

1. IEEE Transactions on System Science and Cybernetics, Vol. SSC-4, No. 3, September 1968

2. Readings in Decision Analysis, 2nd edition, SRI Decision Analysis Group, 1977 (Out of print - available at Stanford and Berkeley libraries)

3. EES 231 Decisions Analysis Notes, Ronald A. Howard, Stanford University

4. Principles and Application of Decision Analysis, Strategic Decisions Group, Menlo Park, CA, 1984

5. "Gaussian Influence Diagrams", Ross D. Shachter and C. Robert Kenley in Management Science, Vol. 33, No. 5, May 1989, pp. 527-550.