Decision Analysis Meets High-Stake Planning
Commentary on Paper by Sven Koenig

Roger Arno

NASA Ames Research Center
MS 244-14
Moffett Field, CA 94035-1000
rarno@mail.arc.nasa.gov

Abstract

Thoughtfully structured, analytical processes should be part of NASA's (or any goal-oriented organization's) decision-making. Decision theory, incorporating schemes for characterizing outcomes, uncertainties, and risk preferences, can be very useful in structuring and formalizing logical decisions. This commentary on Koenig's *High-Stake Planning* paper gives one interpretation of the proper and possible use of structured decision techniques in NASA -- at least in those situations where unquantifiable political factors do not overwhelm the process.

Perhaps the most interesting issue raised by Koenig's paper is whether it is useful to characterize or model NASA's (or any decision-maker's) risk preference in terms of "optimism" and "pessimism." Decision-makers can feel lucky or unlucky, but can such feelings be modeled in any useful or meaningful way? And more importantly, should unfounded hunches or irrational mood swings be included in business and technical decisions? The answer to both questions is very likely "no."

Expected Value Decisions

Decision analysts generally agree that choices can be made on the basis of the likelihood of an event (or "outcome") and the impact of that event. The foundation of most decision theory is that decision options can be compared and prioritized by calculating their "expected values," which are the products of their respective impact magnitudes (benefit or disbenefit of the outcome) and the probabilities of their occurrence. The reason for using "expected value" is that it works. It maximizes gain and minimizes loss over the long term. It is proven every day in casino operations, as well as in businesses that routinely make decisions with profit and loss potential.

In real-life situations, "expected values" frequently must be adjusted to reflect the perceived (subjective) worth of an outcome or the subjective perception of the probability of its occurrence. This kind of subjectivity is widely evident, as most decision makers shy away from extreme losses, or they disproportionately favor relatively sure, albeit smaller, gains. For example, a typical decision-maker might select an assured $200K over a 50/50 chance at $1M. Clearly, this approach would not produce the greatest profit over many opportunities, but this risk-averse behavior is real. It simply says that for this decision-maker, a sure $200K has a utility that is equal to or greater than an uncertain $1M. It does not say that any assured dollar amount is worth more than a 50/50 chance at a million. On the contrary, many people are willing to give up a dollar in hand for a five-million-to-one chance at $1M in the lottery.

It is important to notice that risk aversion is not dependent so much on the probability of winning or losing, as it is on the dollar amounts involved. NASA might be willing to spend thousands or millions of dollars on a very speculative research activity with uncertain payoffs, but just as a casino owner would not risk his entire business on a fair bet, NASA wants to avoid ventures that might jeopardize its credibility and future. This risk-averse nature (which could equally well be called loss-averse nature) can be reasonably well characterized for different decision-makers or organizations. It is frequently displayed as a plot of the subjective utility of dollar amounts against actual dollar values, as shown below.

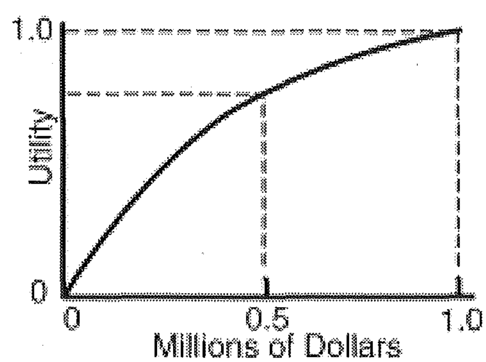

Here in this typical risk-averse utility plot, it can be seen that a million dollars does not have twice the utility of half-a-million dollars. The utility curve for a casino owner might be straight. For a person in desperate need of a million dollars, the curve would probably be bent the other way.

NASA and Risk

At the height of the space race, in the 1960s and 1970s, a national inferiority complex produced money and a willingness to take chances. Early US attempts to launch satellites resulted in frequent failure. The prevailing attitude was that the value of success far outweighed the risk and embarrassment of failure. In decision analysis parlance, the expected value of a successful launch was larger than the expected value of losing the nation's technical preeminence to a foreign power.

Today, with relatively tight budgets, ambitious projects, and decreased public anxiety about national pride, low cost and low risk are priorities. NASA's current risk-averse, safety first approach is largely derived from the high degree of success it enjoyed decades ago. NASA, through skill, luck, and ample funding, demonstrated impressive space exploration capabilities. The public and congress came to expect nothing less than excellence from the world's greatest space agency. In that context, failure became not only a contradiction to past performance, it was a tremendous embarrassment to NASA and the nation. The loss of the first teacher in space was devastating. The failure of space systems to do what previous space systems did seems inexcusable. The resulting risk averseness is partly an acknowledgement that failure has a cost that exceeds the missed science or money spent.

As the danger of world domination by other nations has, at least temporarily, subsided, decision criteria are more complex. At the same time NASA has proclaimed its willingness and ability to "do more for less," it is advocating the chancier "faster, better, cheaper" philosophy. In addition, lower budgets and ambitious programs have forced some risky international partnerships (e.g., the International Space Station). With no little green men on Mars and no jungles under the clouds of Venus, science goals are more esoteric. These crosscutting philosophical overtones make NASA decisions more political, opaque, obscure, and difficult to model. While there is some doubt that NASA decision processes can be modeled at all, the bigger question may be whether they should be modeled, or whether, perhaps, NASA decisions should be more methodical and structured, with decision criteria openly discussed.

The Rover Operations Example

The application of decision analysis to Koenig's rover is difficult due to a number of unspecified parameters (as of this writing). First, there is no specific value or utility metric for the rover completing its mission in a given time, or for partially completing its mission. Second, there is no limit indicated on the availability of the resource (presumably electrical energy, although it could be computer memory, or something else). Third, there are no penalties identified for getting lost or stuck in the mud. Nonetheless, we can apply decision theory by assuming that minimizing the resource consumption is the ultimate goal.

In this example, resource consumption is proportional to the distance the rover moves. Each location-sensing event provides useful navigation information, but it also consumes some of the resource. Less frequent sensing saves resource, but can cause the rover to stray off the path, resulting in more resource consumption. For each sensing frequency there is a statistical spread of possible resource consumption levels. The following figure shows typical consumption curves for two sensing frequencies.

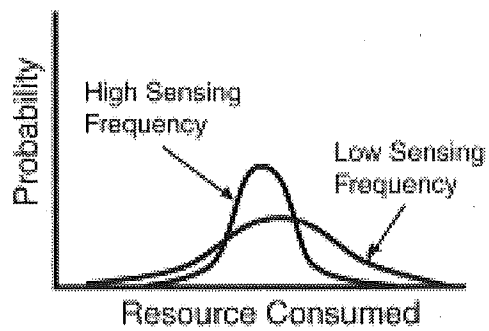

It is easy enough to calculate the expected value (or "expected resource consumption") for each sensing frequency. On the average, high sensing frequency wins, but in a small percentage of trial runs, low sensing frequency produces even lower consumption levels.

Koenig's thesis is that real (e.g., NASA) decision makers are swayed by the variance in the resource consumption curve. He postulates that an "optimistic" decision-maker may well select the option with the highest average and highest possible resource consumption, because it also offers a small chance of a very low consumption. Conversely, he suggests a "pessimistic" decision-maker will select the sensing frequency with the lowest variance and lowest maximum consumption, in order to avoid the option with highest possible consumption. Realistically, without some inordinate advantage for low resource consumption, there is no rationale for selecting anything but the sensing frequency with the lowest expected resource consumption.

In a situation where the resource is limited, and the chance of failure has important negative consequences, the choice becomes clearer. The figure below arbitrarily sets a hypothetical upper limit on resource consumption. Clearly, the decision-maker who selects the lower sensing frequency is not only foolishly "optimistic," but accepting a significant probability of failure.

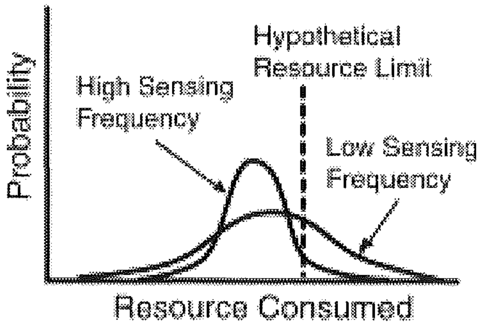

Indeed, some decision-makers feel lucky on some days, and unlucky on others. However, attempting to formally characterize and apply this behavior can be counterproductive. First of all, overlaying intuition and hunches on carefully constructed expected value calculations will ultimately result in reduced overall productivity. Second, attempting to overlay best estimates of risk and uncertainty with additional subjective feelings about the risk is really an admission that the original probability estimates are not believed and need to be adjusted. Placing subjective feelings on the best estimates of probabilities is a bit like a craps player saying, "I know the odds of throwing a seven are truly six out of 36, but today I feel like they are 18 out of 36." Baring fortune telling and supernatural intervention, decision making should be based on rational, honest, thorough assessments of probability and utility. The only time intuition is useful is when it stimulates further assessments of the variables used to calculate the expected values.

Problems with Decision Analysis in NASA

In most high-level, high-stake NASA decisions, political considerations (congressional support, fear of failure, threatened credibility, degraded image, budget threats, etc.) are dominant. Attempts to formalize and systematize the process would require that decision-makers identify and quantify personal and organization agendas that they prefer to keep hidden. For decisions that fall below political visibility, decision analysis can be fruitfully applied. However, there are still barriers that seem to escape the insight of the naïve decision analyst.

One reason that traditional decision analysis and utility theory cannot be applied directly to NASA decisions is purely mechanical. The curves plotted to show the nonlinear utility of money as a function of its numerical value do not readily translate to assessments of the worth of mission successes and failures. Imagine, for example, NASA mission option A, which collects 4 GB of data and is judged to have twice the value of mission option B, which collects 3 GB of data. There is an obvious nonlinearity of utility to data quantity, since the last gigabyte adds a disproportionately large amount of utility. The important underlying message is that the assessed mission option values already include the nonlinear subjectivity. If option valuation is done correctly and credibly, further adjustments are neither appropriate nor technically meaningful.

Extending the above example, assume that the probability of success for option A is 50%, and that for option B is 100%. The two options thus have identical expected values (i.e., 50% x 2.0 = 1.0 = 100% x 1.0). Adhering to decision analysis theory and practice, NASA should see the options as equivalent. Suppose, however, that a decision-maker refused to pick option A unless it had at least a 75% probability of success. The "expected value" decision method has either failed or needs some adjustment.

A likely explanation for the decision-maker's seemingly irrational behavior is that there are other factors not formally unaccounted for. These might include the negative impacts of mission failure. Formal decision analysis is of little value if important considerations are missing. If there are significant potential financial, science, personal or political losses, they must be incorporated, along with their respective probabilities of occurrence. The same is true for political or other payoffs that might come with option selection (e.g., money to a contractor in an important congressional district) or mission success (e.g., the favorable publicity of an elder astronaut hero in space).

Another possible reason for the decision-maker's inexplicable attitude is a subtle variation of the previously discussed problem. It may be that he/she has misrepresented the relative utility of the two options. Perhaps option A actually has a subjective value (or utility) of 1.25 or 1.33, or 1.5 (not 2.0) times that of option B. If that is the case, it would be silly to attempt plotting some new set of perceived values of the mission options against the original perceived subjective values. Piling subjectivity upon subjectivity is just an admission that the original subjective values were incorrect and not the ones that the decision-maker wants to use. The solution is to correct the 2.0 valuation or relative utility.

Another misused concept was addressed earlier. It is the notion that somehow risk preference is about the subjective reassessment of probabilities of outcomes. This line of thinking is to the effect that choices are based on the subjective assessment of the probabilities of success and failure. A decision-maker, it may be postulated, will avoid situations where the probability of winning is low or the probability of losing is high. However, lotteries and high-value business and personal decisions clearly show that risk averseness is determined more by the magnitude of gain and loss than by their associated likelihoods. Plotting utility curves of the subjective interpretation of probability against rational estimates of probability will not produce useful predictions of risk preference behavior. A different curve would have to be plotted for each and every relevant value of gain and loss.

Conclusion

Koenig's rover operations example is basically a straightforward exercise in engineering performance comparison. Since there is no apparent limit to the resource, and no penalty for its use, there is no reason to select any option other than the one that provides the lowest expected total consumption. The systematic construction and incorporation of intuition or feelings of luck would provide no advantage, and would most likely result in less than maximum productivity.

The characterization of risk preference and risk-aversion as "optimism" and "pessimism" may be useful in predicting the irrational behavior of decision makers, but there is no useful purpose in incorporating irrational predilections where the goal is to maximize gain and minimize loss. It is understandable that a decision-maker would be risk averse, but that trait can and should be reflected through the full and honest assessment of option probabilities and possible outcomes (both positive and negative) and the application of decision analysis methods.