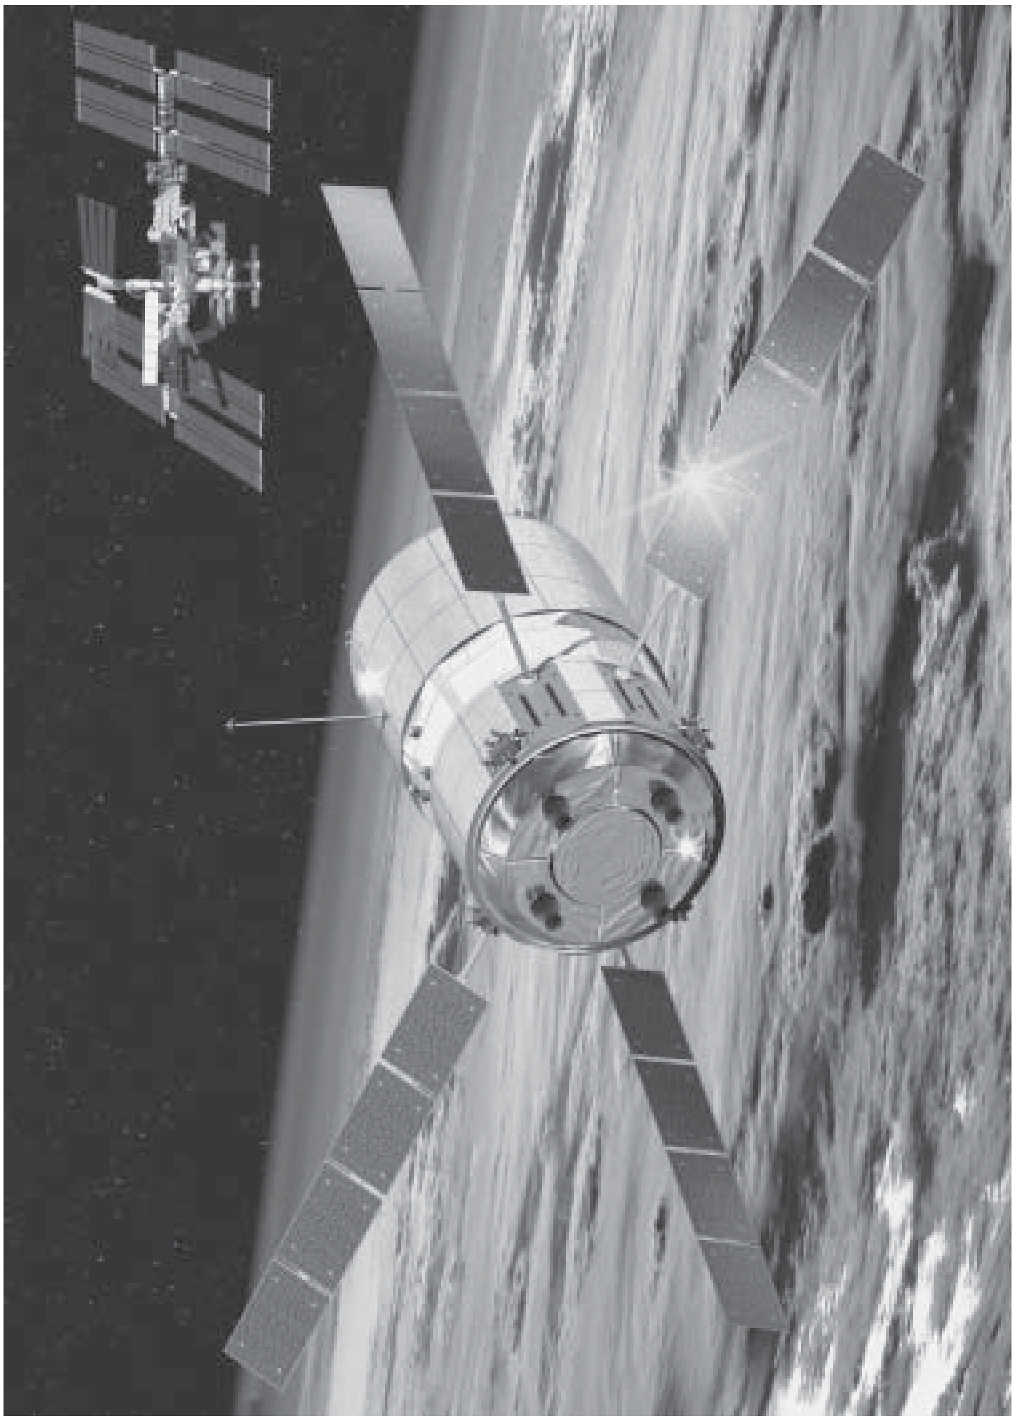

RISK-BASED DECISION SUPPORT TECHNIQUES FOR PROGRAMS AND PROJECTS

Barney Roberts, Clayton Smith, and David Frost

This article is designed for the project management professional who intends to make risk-based decision making a fundamental, integrating principle of the project's operating processes. It is about making decisions using information that relates possible future outcomes to the risk inherent to decisions made. A project manager needs to make two types of decisions: those that relate to the business aspects of a *project* and those that relate to the performance aspects of the *product.* Part 1 details the project-focused tools and techniques and Part 2 details the product-focused tools and techniques. Advanced integrated quantitative techniques and tools that have been proven to have high utility to decision makers are presented.

A project manager needs to make two types of decisions: those that relate to the business aspects of a *project* and those that relate to the performance aspects of the *product*. Decisions that are related to the business aspects are focused on how much things cost or might cost, how long it takes, or may take, to do something. Business aspects of the project are fraught with risks in cost, schedule, fabrication, testing, and production of the product. Decisions that are related to the performance aspects of the product are focused on things like reliability, maintainability, safety, and operations. Performance aspects of the product are fraught with risks in system failures, operational failures, environmental impacts, and overt and covert external threats.

This paper is divided into two parts: Part 1 for the project-focused tools and techniques and Part 2 for the product-focused tools and techniques. The following is a listing of those decision-support products that we have found to be of greatest utility and value to the projects.

1. Project-focused tools and techniques
   - Cumulative Distribution Functions (CDFs) for project completion date
   - CDFs for cost estimate at completion
   - Double Pareto boxes
   - Stochastic Critical Path Analysis

2. Product-focused tools and techniques
   - Bandaid Charts
   - Fussell-Vessely Charts

In each description of a tool or technique, the following format is used:

1. What is it? A description of the specific product,

2. How does it work? A brief overview of the analytical technique, and

3. What is its utility? A few examples of applications in decision making and the value derived.

PART 1: PROJECT-FOCUSED TOOLS AND TECHNIQUES

CUMULATIVE DISTRIBUTION FUNCTIONS FOR PROJECT COMPLETION

What is it? A Cumulative Distribution Function helps us understand the uncertainty or the confidence associated with stochastic variables. The CDF is the mathematical integral of the probability density function (PDF). The PDF represents the probability that different outcomes of a random variable will occur. The sample plot in Figure 1 is a PDF and a CDF for the outcome of a pair of dice thrown a large number of times. The plot is read as

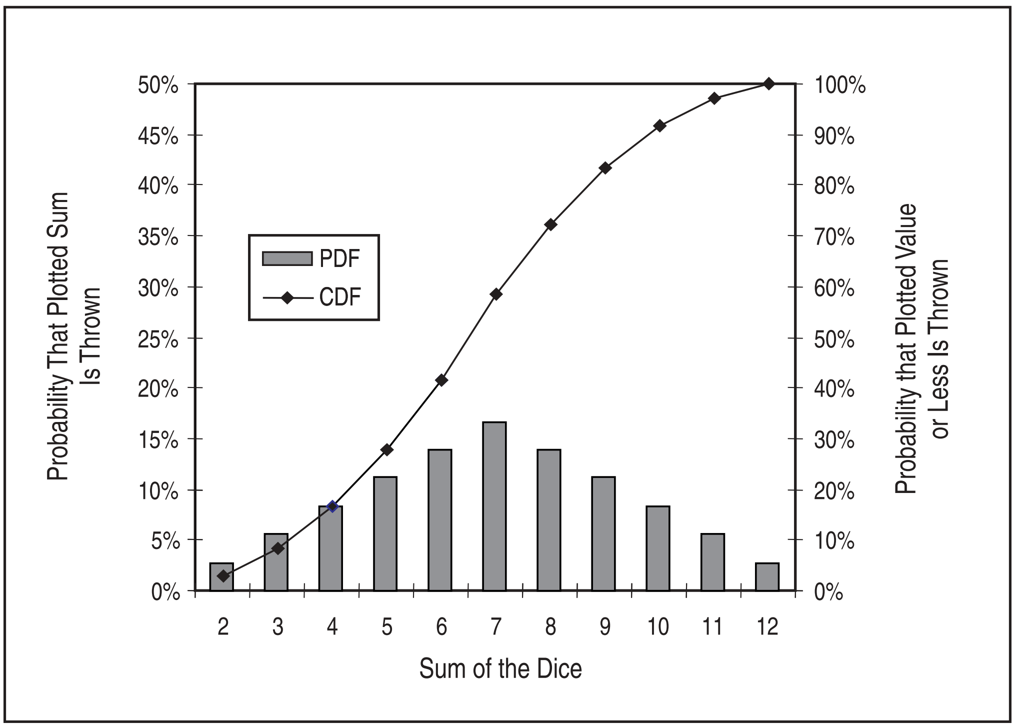

Figure 1. Simple Example of the Probability Density Function (PDF) and Cumulative Distribution Function (CDF) Plots Showing the Outcome for a Pair of Dice Thrown a Large Number of Times

follows: "The probability that an eight is thrown is 14 percent (the left hand vertical axis) and the probability that an eight or less is thrown is 72 percent (the right hand axis). Conversely, the probability that a number higher than eight is thrown is 28 percent.

Utility. The impact of risks on the expected completion date can be plotted. The effects of different mitigation plans can be compared and selected based on the project's propensity for risk. Mitigation actions will, in general, add cost to the baseline project, thus a project manager would want to see actual quantitative information on just how much residual risk would remain as a function of investments to mitigate the risks.

In the sample shown in Figure 2, the project manager can meter the mitigation investments versus the projects propensity for risk. For example, a real "gutsy" project manager may accept the 20 percent probability of completion, especially if the costs for the mitigation options are very high. On the other hand, if the costs are moderate or acceptable, or if the project manager is averse to risk, option A, moderate mitigation, may be chosen. At the other extreme, a very low investment and/or a very risk-averse project manager, Option B, substantial mitigation, may be a better choice.

How are CDFs Created? Here is where we get to the value of this paper and these tools and techniques to different levels of maturity in the project management staff. This analysis cannot be successfully performed if the project does not have an integrated master schedule (IMS); it doesn't have to be perfect, but must be a reasonable semblance of an executable plan that is based on analogous experience.

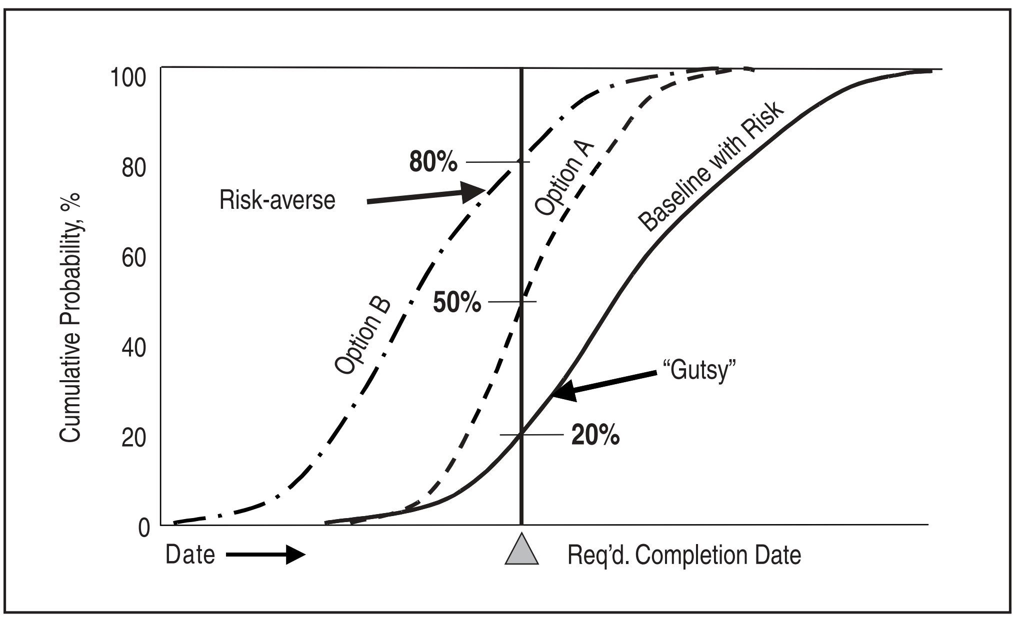

Figure 2. Example Cumulatiave Distribution Function (CDF) and How a Project Can Meter Its Investments to Mitigate Those Risks Versus Its Propensity for Risk

The technique, illustrated in Figure 3, is to collect the project's risks, as defined within the framework of any typical risk management process, understand through analogies or expert opinion the impact to each Work Breakdown Structure (WBS) line item in the IMS, and then perform a Monte Carlo simulation. There exist several commercial software tools that can perform this analysis. The CDF is an output from any of those tools. We have experienced remarkable accuracy from these techniques when no mitigation actions are implemented and the risks are accepted in a set of six predictions that we have done for a major space program. Our predictions, performed 18–24 months in advance, estimated a 5–6 month schedule slip due to risk. No mitigations were taken by the program giving us the opportunity to test the accuracy of the tools. The results had a mean error rate of only 20 days at the 50th percentile.

An Actual Case Example. The examples shown here are from a NASA space exploration mission. In Figure 4, the spacecraft could be launched in either of the two narrow bands within a 6–8 week window occurring about six months apart. A significant amount of flight design, trajectory analysis, and mission planning must be performed to support either launch window. The planned launch was at the beginning of the first band, January 5, 2001, and flight design had begun to support the first launch window. However, the risk analysis showed less than 20 percent chance for project completion in time for the first launch opportunity. Having the risk analysis demonstrate at least an 80

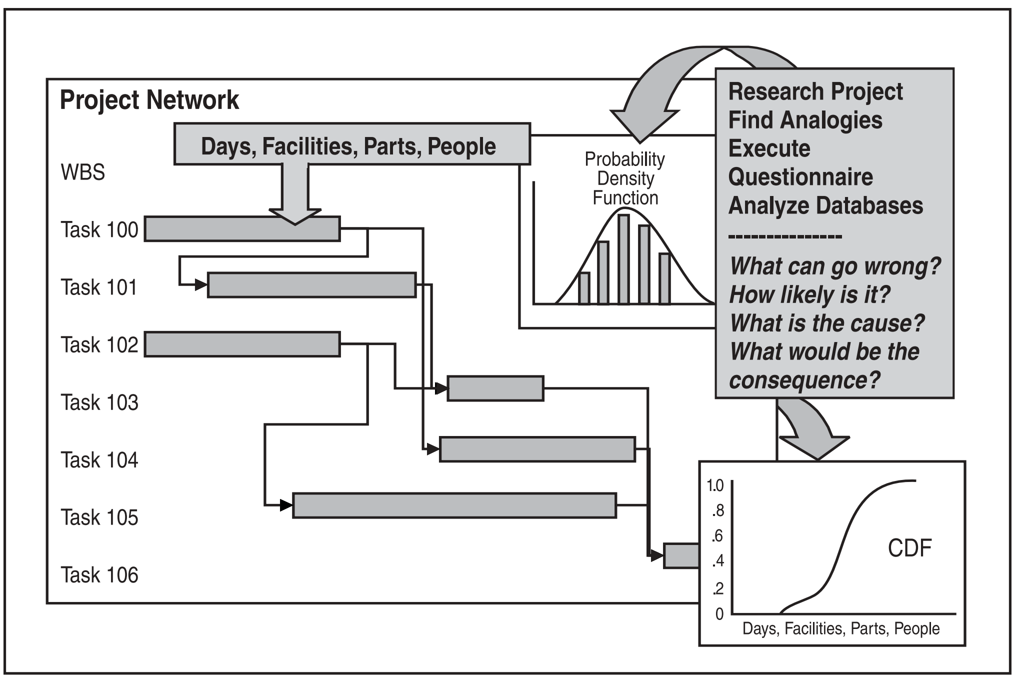

Figure 3. The Cumulative Distribution Functions (CDF) Are Generated from a Monte Carlo Simulation of the Project's Integrated Master Schedule (IMS)

percent chance for completion in the second band, one month later, gave the project the confidence needed to switch the flight design to be consistent with the second launch window saving significant extra cost because of the high probability that the flight design will need to be repeated.

The CDF for completion can also be plotted over time to serve as a tracking metric for the total integrated risk environment of the project. The object is to have a visual display that illustrates a decreasing trend (or not) of the risk environment. One type of plot that accomplishes this is shown in Figure 5. This is from the same project that is referenced above in Figure 4, but the analysis was repeated over time and is plotted in Figure 5 with the 80th and 20th percentiles defining the ranges plotted, and the diamond symbols indication the 50th percentile. Figure 4 correlates to the date line of 3/00 in Figure 5.

Read the chart as follows: The upper horizontal band on the plot is "Ready Early." "Ready On-Time" is the middle band that also spans the launch window. "Ready Late" is the lower band, which means a 6-month slip to the next launch window and all associated costs that go with that slip. The upper line plotted is the deterministic completion date (i.e., no risk) and the lower line plotted with the 20th and 80th percentile confidence bands on the risk-adjusted completion date. The project's objective is to continue to invest in risk mitigation actions until the band and the area of highest likelihood is no longer in the "Missed Launch Period" area of the chart. Note the improving trend over

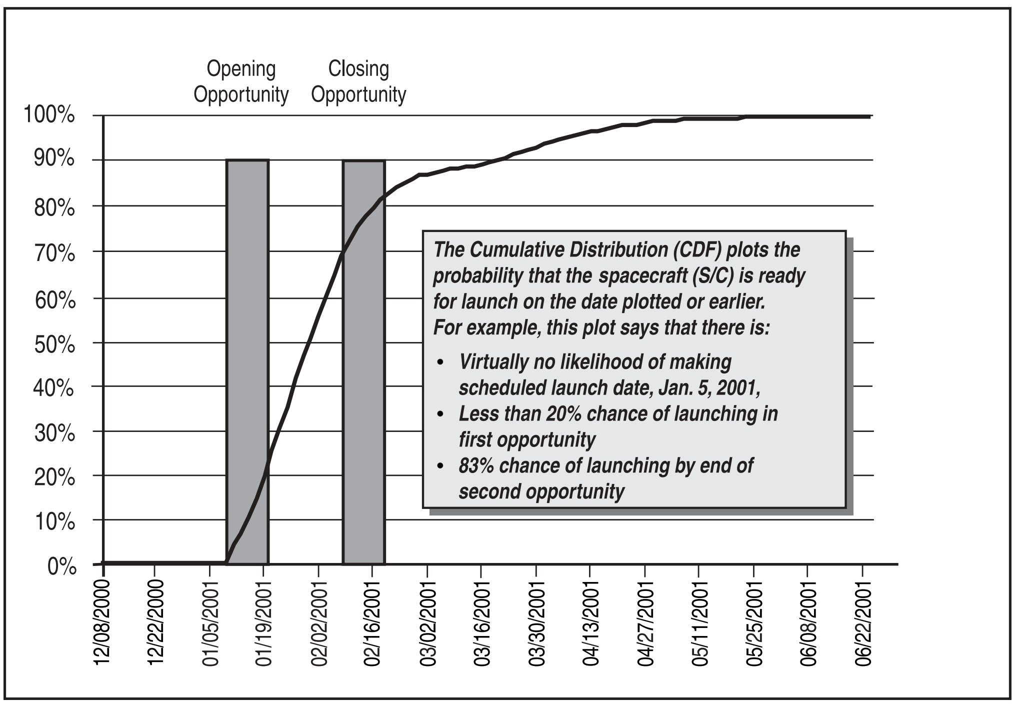

Figure 4. The Risk Analysis Clearly Demonstrated that the Mission Design Should Be Moved to the Second Opportunity

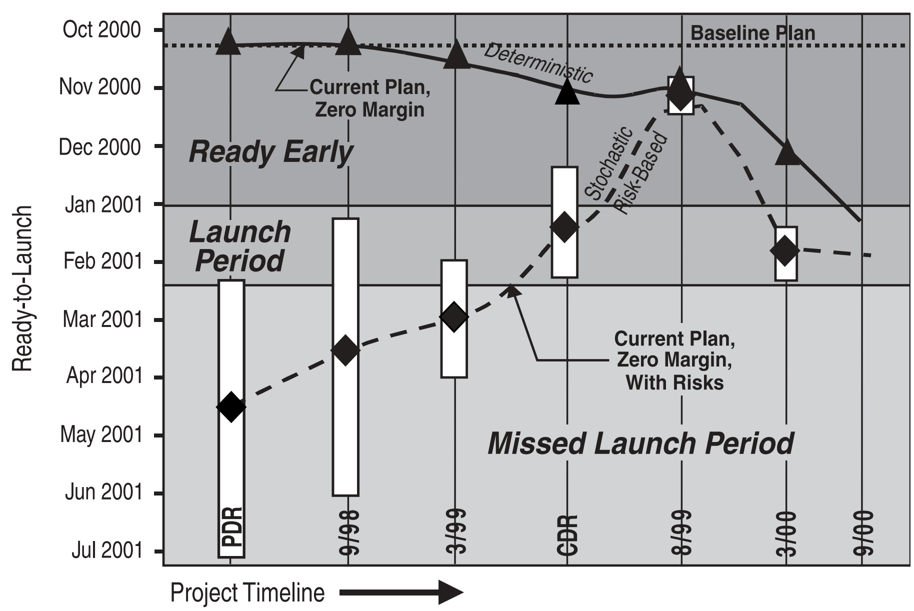

Figure 5. Tracking the Cumulative Distribution Function (CDF) Over Time Throughout the Project Life Cycle, Gives the Project Manager a Trend of the Risk Environment

time indicating the success of the risk mitigation actions as well as some "Accepted" risks passing their exposure window without becoming problems. One should ask, "What happened that caused the downward trend at the end?" A costly mitigation plan had been put in place to deal with a risky component. Seeing the substantial risk-based margin gave the project manager the confidence needed to abandon the mitigation plan, save the money, and still meet the completion date.

Cumulative Distribution Functions for Cost

What are they? They are simply no more than the same function as the schedule CDF but with cost as the domain. These functions can be generated in a way that is consistent with schedule risks if the IMS is resource loaded. Many low maturity projects carry the cost data in a separate database from schedule information making it very difficult to get a good coordinated cost-schedule risk analysis. As long as the project has at least an acceptable IMS, one way to circumvent this low maturity management approach is to use the cost data and create cost simulators that can be loaded into almost any of the tools used for the IMS (such as MS Project). When that is done, the results from the analysis will be both Cost and Schedule CDFs, *and* they will be consistent, which is a very important consideration in analysing risk.

Utility. One obvious utility is expected cost at completion as a function of risk, and the plots will be very similar to those shown in Figure 4, so that feature will not be discussed. However, they may be used to determine project reserves (see Figure 6). The project manager would do this by first establishing some level of acceptable risk, or propensity for risk, if the customer has not specified it. If not specified by the customer, this is usually done by a brainstorm session wherein the project management staff express their opinions. Sometimes it is set by organizational policy. Suppose that the project manager was risk averse and hence wanted to be 80 percent certain that the project's risk-based cost at completion would not exceed the project's budget. The project manager must iterate the design and/or de-scope the risk profile until the 80th percentile aligns with the planned budget. Then the project manager would want to establish a reserve that is equal to the difference between the 50th percentile and the 80th percentile and hold that amount as a reserve against risk.

The most important use of the cost CDF is the analysis of the effectiveness of mitigation investments. One may create the cost CDFs for several options then compare the investment to the return. Two things can happen: the curve can move to the left, reducing cost; and the slope can increase or decrease indicating a change in uncertainty. Sometimes an investment can be made that does not reduce the expected cost but may be a desirable investment for the project due to a reduction in uncertainty.

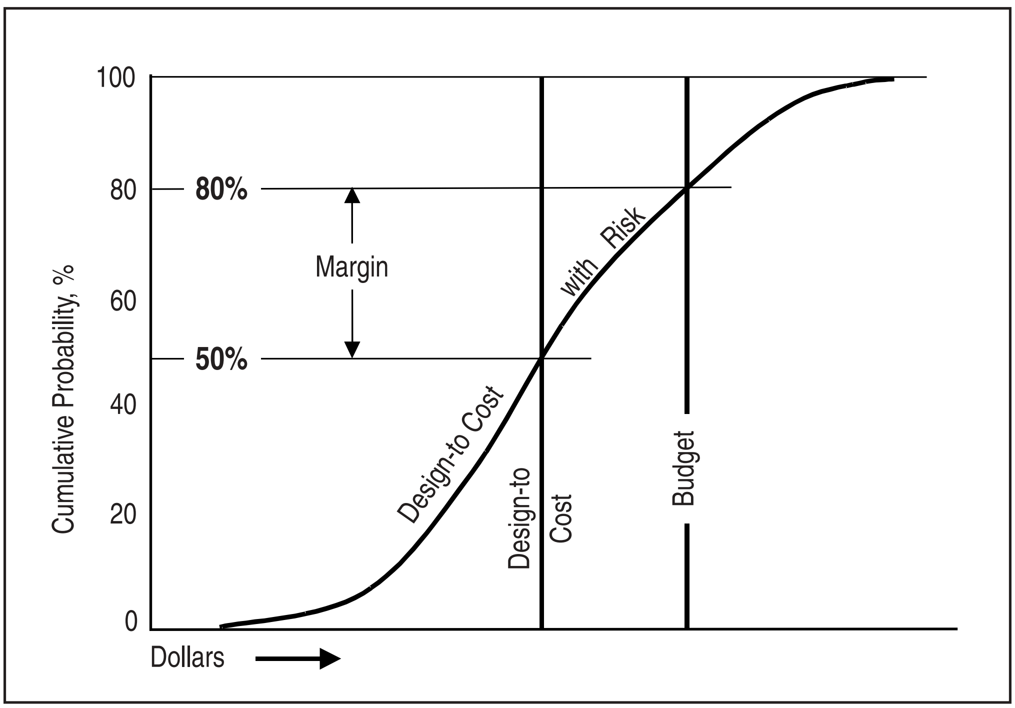

Figure 6. An Illustration In Using Risk and the Cumulatiave Distribution Function (CDF) for Estimating Reserve Funding

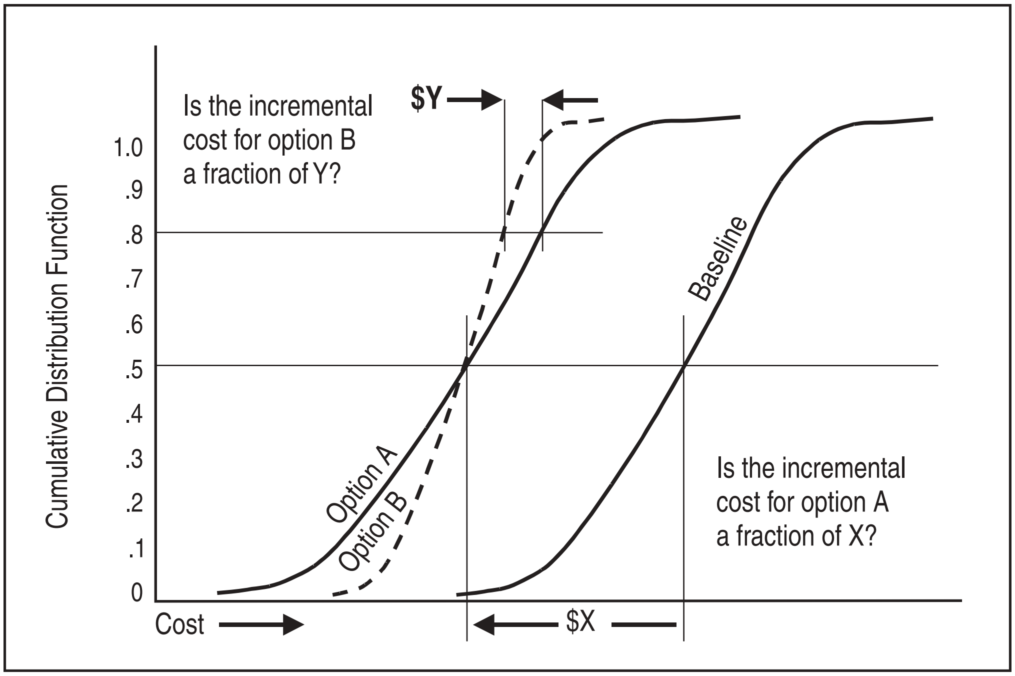

Figure 7. Techniques for Decision Making on Risk Mitigation Investments Illustrate the Value of the Cost Cumulative Distribution Function (CDF)

Figure 7 illustrates these uses. Mitigation option A reduces the expected risk exposure by $X; if the investment to achieve this improvement is some acceptable fraction of $X, then the project manager should accept this option. Suppose that the project manager could invest another $Y and the result is option B. The expected value of the final cost is unchanged, but the reduction of uncertainty has value to the project. One possible judgment would be to compare the at-risk cost reduction at the 80th percentile, and if this value is some multiple greater than the investment to achieve option B, then the project manager ought to make the additional investment.

DOUBLE PARETO BOXES

What are they? The Double Pareto boxes are two-dimensional arrays, the rows are the WBS line items that are impacted by risks, and the columns are each individual risk. The cells of the matrix contain some attribute of the project that is important to risk-manage, usually dollars or days. Any spreadsheet software that supports sorting is a suitable tool for this analysis. Once the data are extracted from the Monte Carlo network analysis described above, and loaded into the cells, the sorting functions are used to sort the highest cell values into the upper left corner. Then the matrix is sectioned, or truncated, at the row (the WBS items) where the cumulative summation of the cell values equals 80 percent and at the column (the risks) under the same condition; hence "Double Pareto" box. The result is the sectioning off of those few risks that are causing 80 percent of the problem and those few WBS line items that are "receiving" that 80 percent of the impact.

How are they created? The Monte Carlo IMS analysis tool that was used above is used to perform this analysis but is run for each individual risk, one at a time. The data are used to fill the cells of the matrix. The cells can contain either the risk impact in days or dollars, or, in fact, any resource or metric considered to be of value to the project. We recommend that the project hold a brainstorm session in the early phases to determine "What is to be risk-managed."

Utility. Program and project resources are precious and should not be spent on trivial issues. In the cases where we have used the Double Pareto box, we have found a great reduction in the number of risks that need to be mitigated and tracked and the number of WBS line items that are threatened by risks. For the Space Station, the "worry" risks were reduced by an order of magnitude and the "worried" WBS line items were on the order of a dozen. This also provides the program manager or project manager with a tool to deal with "whiners," being able to quickly weed them out by checking the Double Pareto box to see if they made the cut-box.

Impacted Task Title	Risk Drivers				Cumulative Contribution
	Baseline	Risk 10 Late Software	Risk 27 Star Tracker	Risk 61 Sequence Timer	
GN FSW BUILD 4.0 Delivery to ATLO – Science	19.8	19.8	0.0	0.0	22.3%
ALTO SCHEDULE MARGIN – Denver (Used to model ATLO overrun)	17.7	0.0	3.2	7.4	42.2%
GN FSW BUILD 3.0 Delivery to 4.0 for ACS testing (MST 3)	16.7	16.7	0.0	0.0	61.0%
Star Tracker FLT Design, Purchase, Recieve, and Test	9.1	0.0	9.1	0.0	71.2%
FSW Phase 5.0 Delivery to ATLO – Launch	9.0	9.0	0.0	0.0	81.3%
Total	88.9	53.9	12.3	11.0	
Cumulative Contribution		60.6%	74.4%	86.8%	

The Double Pareto box greatly focuses the project's attention on the few risks that cause 80 percent of the problem and the few WBS lines that are receiving this 80 percent impact. The cells in this graphic contain either delta-dollars or delta-days due to each risk.

Figure 8. The Double Pareto Box

An example of an actual Double Pareto box from a NASA space science mission is shown in Figure 8. In this example, the left-most column is the baseline case with all risks, following that to the right are the single, driver risks. As few as three risks, of the project's approximately 34, constitute the 80 percent and only five tasks, out of 1000, bear that impact. In this case, the project manager's span of attention was greatly reduced.

STOCHASTIC CRITICAL PATH ANALYSIS

What is it? The Stochastic Critical Path has been the most valued product that we have produced for the project manager. It is a specific portrayal of the schedule network of a project wherein an additional piece of information is added. First the deterministic critical path is highlighted and then the various stochastic critical paths are highlighted to produce a visual image for decision making. The deterministic critical path is the critical path that is determined without the consideration of risk; it is the collection of tasks that determine the total completion time of the project. Due to risk, there is a probability that some other tasks not on the deterministic critical path may increase in duration so that they increase the total completion time of the project. The risks put those tasks on a probabilistic critical path with an associated probability.

In the absence of the risk analysis, the deterministic critical path will be the one to which a project manager will place the maximum amount of attention and, hence, management resources. The stochastic critical paths are all of the other critical paths that

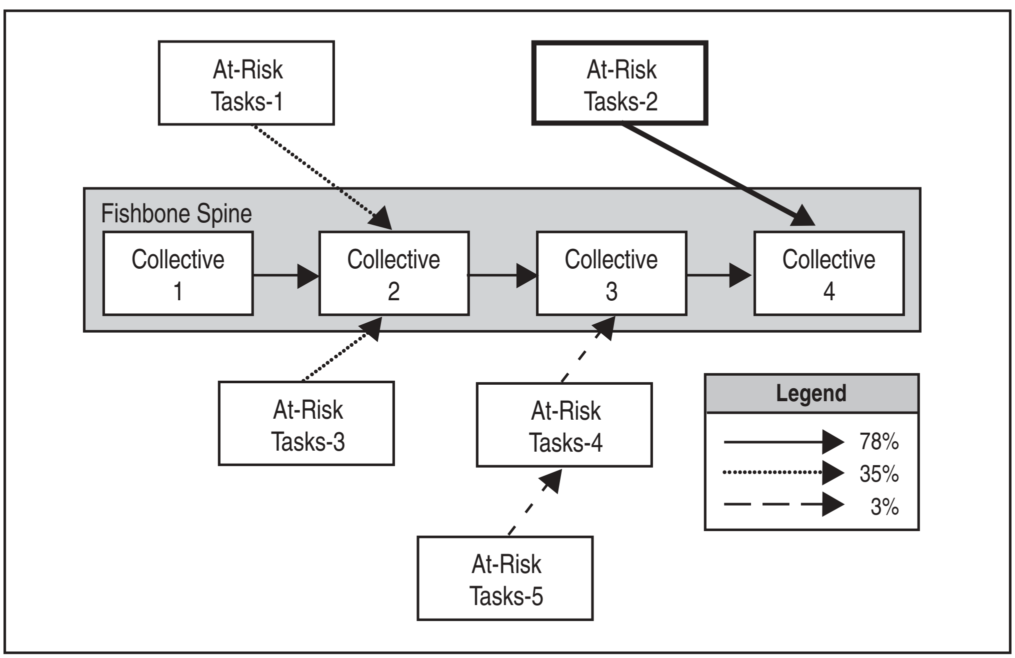

Figure 9. The Stochastic Critical Path Chart

may be critical depending on risk and, hence without the analysis, may never catch the attention of the project manager. Being dependent on the outcome of a risk, they are "critical" with some probability.

How is it created? The stochastic critical path is a result of the Monte Carlo schedule network solution illustrated in Figure 3. As the network is analysed and sampled via the Monte Carlo analyser, the critical path is recorded for each iteration and the software captures the frequency that a task is on the critical path. Figure 9 illustrates the result.

Imagine the diagram to be a fishbone. Let the spine of the fishbone be a representation of the deterministic critical path. Greatly simplify the activities on the spine collecting activities into groups such that they are single but no more than two collectives between the bones that come into the spine. The bones that come into the spine are the alternative stochastic critical paths. Color, or shade the bones, (i.e., stochastic critical path activities) to correspond to a legend that specifies probability of being on the critical path. You need to know that there are other terms that are used by commercial software analysis packages to describe the probability of being on the critical path, those are (1) criticality and (2) diversity.

Figure 9 is a simplified Stochastic Critical Path diagram used to illustrate the fundamental features and Figure 10 is a diagram from an actual project.

Utility. The project manager now has a quantitative representation as to where

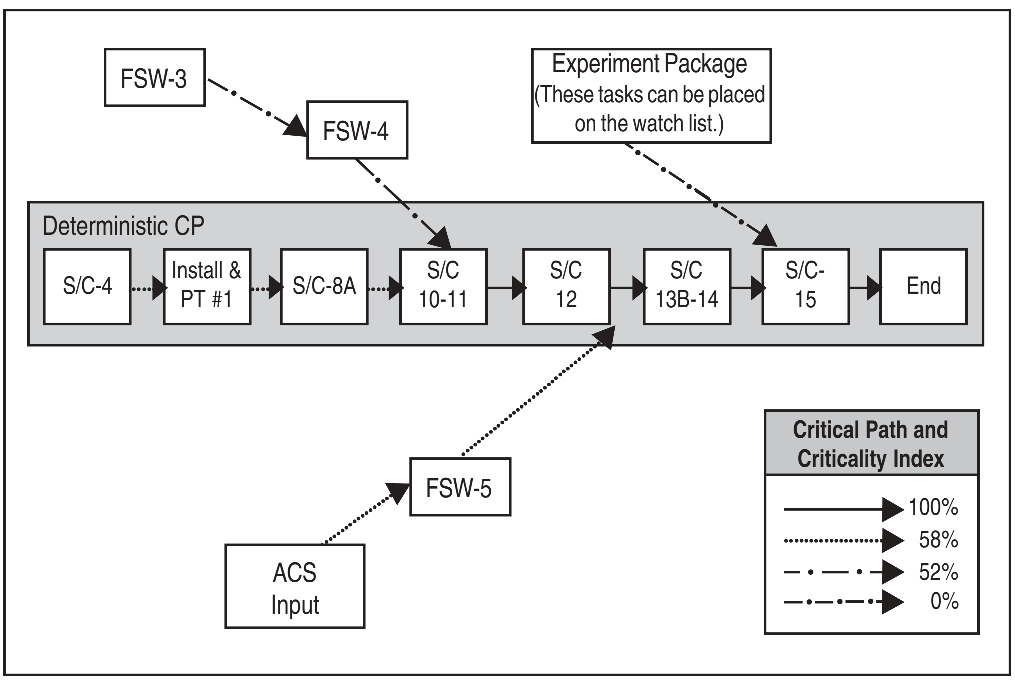

Figure 10. A Simplified Version of the Stochastic Critical Path Developed for an Actual Project

risk-management resources need to be invested. For example in Figure 9, at-risk tasks 4 and 5 should be ignored, even if they are high risk items, they just can't "catch up" with the critical path because it is driven so hard by at-risk tasks 1 and 2. Also note that without the stochastic critical path (CP), the project manager will be managing to the deterministic CP and will be focused on the wrong things. It also tells the program manager that dollars to mitigate risk in at-risk task 2 have twice (or 2.23 times) the value of dollars used to mitigate risk in at-risk task 1.

An Actual Case Example. Figure 10 is the stochastic critical path for a NASA space exploration mission. The actual names of the tasks have been replaced with generic names in some cases and a few other simplifications were made to get the graphic to fit in this paper. First note that the Experiment Package was a high-risk item but never appeared on the critical path because the others drove it so hard. Thus, the Experiment Package could be put on the watch list. There are three other probabilistic critical paths that are competing with each other almost equally; being 50 percent probable. Also, unobservable here because of the simplifications, there are many tasks on the deterministic critical path prior to the task labelled SC-4 that were never on the critical path when risk is considered, thus the project could relax its vigil there as well. Also note the dominance of the Flight Software (FSW) packages on the various stochastic critical paths thus providing a primary focus for risk mitigation. It's also important to note that schedule must be something the program wants to risk-manage if this is to be useful.

PART 2: PRODUCT-FOCUSED TOOLS AND TECHNIQUES

BANDAID CHARTS AND IMPORTANCE VALUE CHARTS

What are they? There are several useful risk-based decision support products that are extractable from a probabilistic risk assessment (PRA) of the product being developed and delivered by the project. Of those products, the ones seeming to have the greatest utility are the Bandaid Charts and the Importance Value Charts. The Bandaid Charts are named for their appearance in that they look very similar to bandages produced by brand-name companies such as Band-Aid. They are a

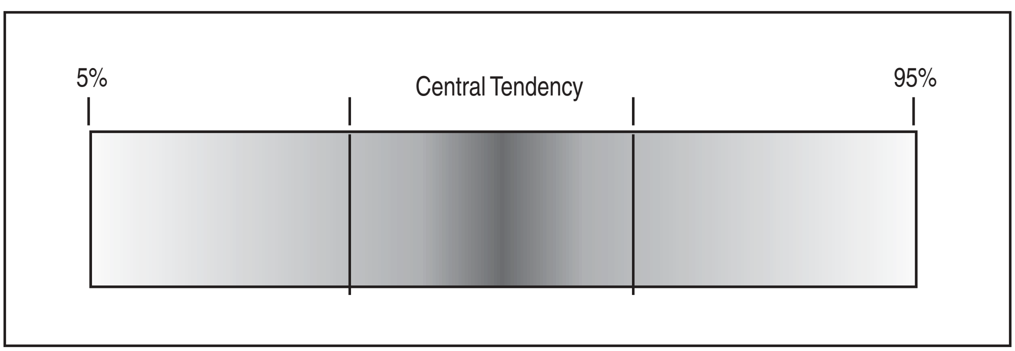

Figure 11. Sample Bandaid Chart

spread of probable outcomes that appear on the chart as a "band" of values with one end being the 5th percentile, the other end the 95th percentile and the center is marked or densified to indicate the central tendency, hence the Band-Aid appearance. One of these is produced for each end state of interest. A sample Bandaid Chart is shown in Figure 11.

The Importance Value Charts are derived from a probabilistic risk assessment technique. They are a result of metrics collected as the PRA is performed that reflect the contribution of selected systems or components to the overall failure rate. They are often normalized to some specific parameter of the decision to avoid the aforementioned problem of everyone focusing on the failure rates rather than the decision information. The Importance Value Chart is a bar graph with each bar representing a specific system, subsystem, or component's contribution to the overall probability of an undesirable event. They are arranged in order with the "tallest" bar on the left and subsequently shorter bars progressing to the right. Figure 12 is a typical example.

To support a project's decisions, one may mark the point where 80 percent of the total loss rate is accumulated by the systems. This quickly draws the decision maker's attention to those few items that need to be subjected to design improvement or additional testing and verification.

How are they created? Both the Bandaid Charts and the Importance Value Charts are outputs from post-analysis of data from a probabilistic risk assessment. PRAs can be created in two ways: bottom-up and top-down. The bottom-up approach can be very costly in that many components are analyzed and modeled to a level of detail that may not affect the end result. Working from the top-down

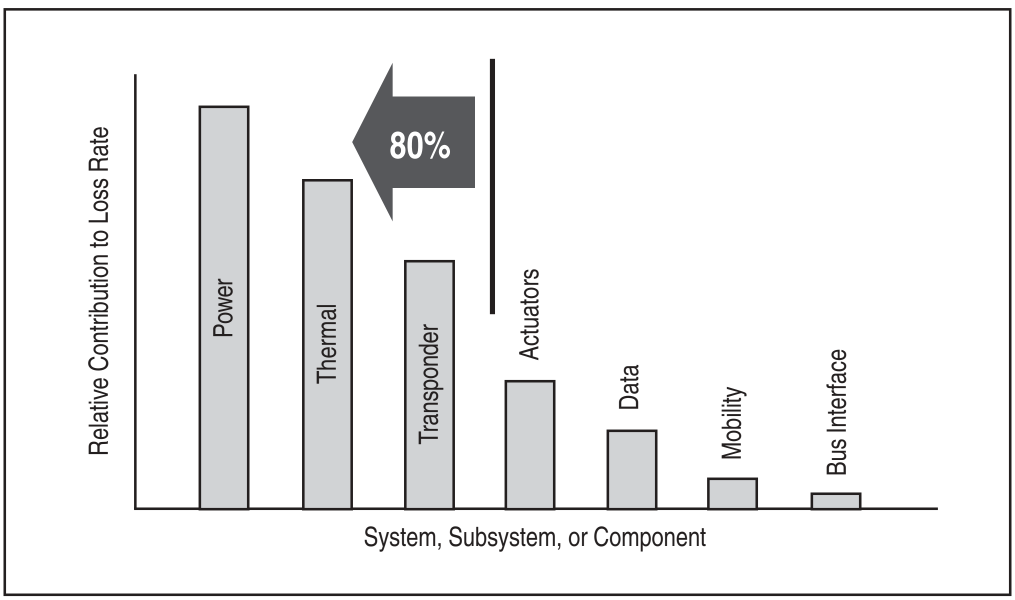

Figure 12. Importance Value Chart

has, so far in our experience, been able to produce the decision-support information at a fraction of the cost of a bottom-up. Of course, if the project can afford the bottom-up analysis, then it would be more thorough and probably the best approach.

Performing the top-down analysis, one must first determine, then model the most undesirable outcomes, such as Loss of Crew, Loss of Vehicle, if one were analyzing a space program with crewpersons present; or it could be Loss of Science or Reduced Science Quality if one were analyzing a robotic space exploration mission; or it could be Lethality or Survivability if one were analyzing a Department of Defense (DoD) weapon system or system of systems.

Whatever those undesirable outcomes, there will exist logical scenarios triggered by initiating events. Those scenarios will contain response actions based on the relevant initiating event. The Master Logic Diagram (MLD) is used to identify the initiating events and to put them in context with each other. The MLD is a top-down logical representation of the system (see Figure 13).

Then for each element in the MLD, one would build event sequences that could cause the scenario to be executed. The event sequence diagrams start with an initiator event and end with many end states. The events ask about redundancy, repair, operational workarounds, other consequences, and responses of the system. Note that these scenarios are developed "given that the initiator has occurred." For each relevant element in the event sequence, Fault Trees are constructed that describe the failure events in the systems, subsystems, or components in the product.

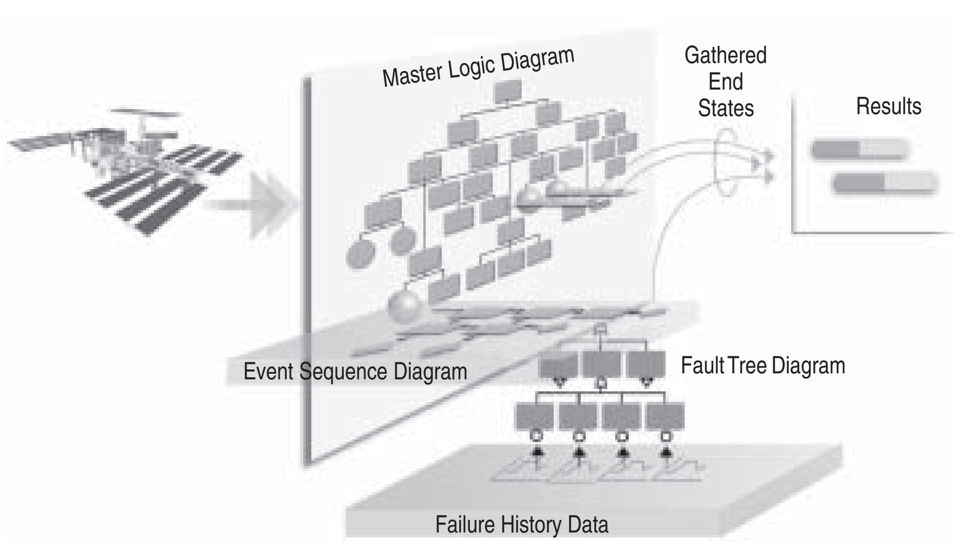

The PRA approach is most effective and efficient when scenario driven, modeled from the top-down, and decomposed no lower than the level that affects the decision information.

Figure 13. The Master Logic Diagram (MLD)

Only those fault trees that are related to the undesirable outcome are modeled. In addition, the model is scenario driven to account for all system and operator intermediate actions that result from the initiating event.

At the very bottom of the model is the database that feeds the fault trees. Another cost-saving exercise is employed here in that one could stop at the system or subsystem level and not need to punch all the way down to a component level should the system or subsystem not prove to be a significant driver. This can be done because of the approach of modeling the elements of the model with probability density functions. For example, a power system of a specific type can be modeled by all analogous power systems available from previous projects. This will, of course, result in a broad probability density function because of the variables that are present at the next level of decomposition. However, if it makes no appreciable difference in the decision-support information, why bother with further decomposition?

There exists a good suite of PRA tools such as QRAS, SAPHIRE, or Monte Carlo simulators that are add-ons to spreadsheets: for example, Palisade Software's @Risk for Excel or PrecisionTree.

Utility of the Bandaid Chart. It is important to note here that the output of these types of analyses are probabilities of failure, a specific number that often attracts so much attention that the utility of the analysis is lost. Of course, the number can be unacceptably low requiring a significant redesign. But once we get beyond that point, it is better to try to understand what the model is telling us about weaknesses in our system, rather than focus on the number. So many times have we all been engaged in heated arguments and

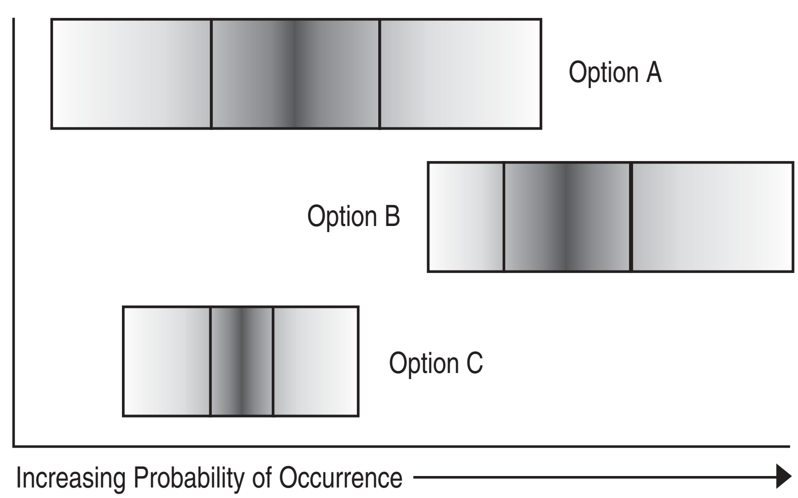

The Bandaid Charts are useful in decision making for selection of options when the input information is fraught with uncertainty.

Figure 14. The Bandaid Charts

prolonged discussions about the specific value of the probability of failure and, probably in some cases, to cover up the number, discredit it, or change the inputs to produce a more favorable outcome to present to the stakeholders. Hence, one should attempt to normalize the results to some specific baseline to avoid these discussions and assist the project manager in making good decisions.

The real utility in the Bandaid charts is decision making about options. A sample Bandaid Chart is shown in Figure 14. In any PRA analysis the actual values for the failure rates of the systems and subsystems in the fault tree are seldom known as precise point-values but more often are probability distributions that represent the uncertainty of the information available for a system or subsystem. Hence, the result of the analysis will also be an uncertain number.

In Figure 14, Option C is clearly better than Option B, having a lower probability of occurrence over its entire uncertainty region (Statistically Dominant). Option A and Option C overlap, thus Option C is not Statistically Dominant over Option A. There exist a fair number of possible outcomes of Option A that could be better than Option C. One might still want to select Option C because of its much lower range of uncertainty that will provide a more stable planning environment.

An Actual Case Example. In this actual case (Figure 15), a new upgrade system was proposed for a NASA launch vehicle. The upgrade system actually could demonstrate that when the median

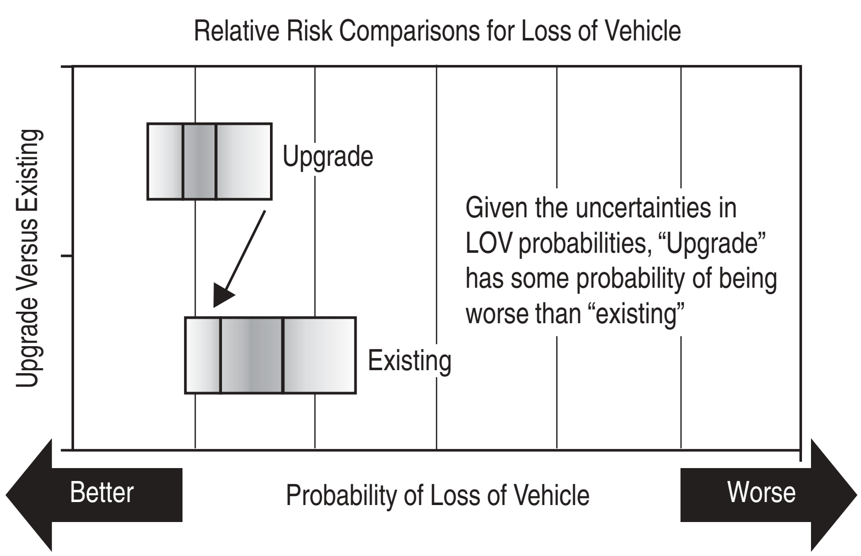

The Bandaid chart was produced demonstrating that the uncertainty of the information indicated that the upgrade has some probability of being worse than the existing system. The overlapping probability densities (not shown here) equated to a 15 percent chance of being worse than the existing system.

Figure 15. Actual Bandaid Chart Example

value for the probability for loss of vehicle was compared with the existing design, there was a notable improvement on the order of 65 percent reduction in loss of vehicle. However, when the analysis was done considering the uncertainty in the information, there was significant overlapping of the two probability density functions. Calculations showed that there was a 15 percent chance that the upgrade would actually be worse in contribution to loss of vehicle than the existing design. Considering the investment costs versus the risk that the upgrade may perform worse than the existing design led to the decision to retain the existing design.

Importance Values

Utility. The best way to illustrate the utility is to imagine the case that a PRA has been completed and you can present to the decision maker either (1) "The project has a 73 percent chance of success" or (2) "These three systems contribute 80 percent of the threat of loss." Both answers have use to the decision maker, but answer (2) provides much more opportunity to make effective and efficient decisions for improvement.

Project decision makers use the Importance Value Charts to refocus early design activities as well as midcourse corrections as the design matures. It also permits the planning for test and verification to focus on systems that are threatening to success. It should also be noted that these threats are strongly dependent on how the product is operated. This assists in the operational planning or the design of support systems in that the operational scenarios can be designed to focus on ways to desensitize

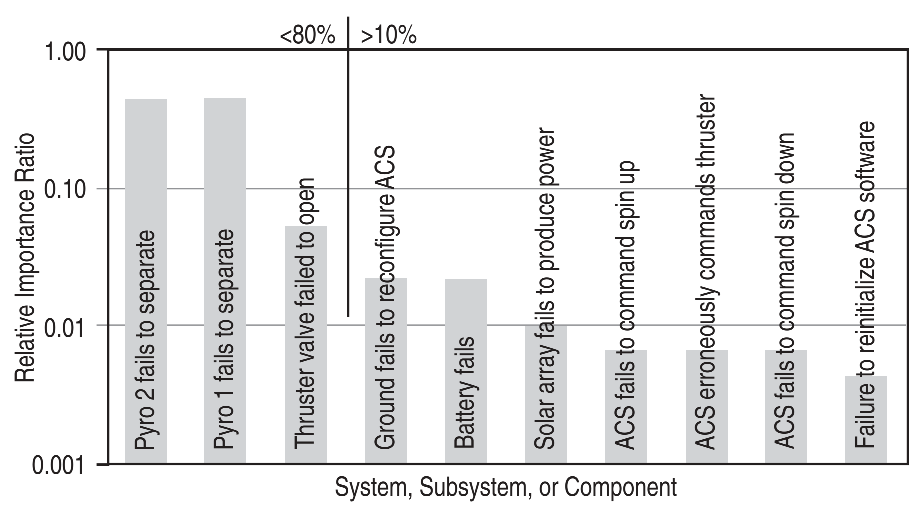

The Importance Value Chart helps the project to focus design, test, and operations planning resources on those few systems, subsystems, or components that are contributing the most to the failure to achieve an objective.

Figure 16. The Importance Value Chart

the impact of failures in these systems or operate them differently to reduce stress and hence failure rate.

Actual Case Example. The project chosen for this case example was a robotic space exploration mission. The PRA was performed to support decisions to be made at the project's Critical Design Review. One Importance Value Chart is presented in Figure 16. In this chart, the pyrotechnics used to separate the solar arrays and permit them to be deployed after launch were the primary drivers along with a thruster failure being a minor contributor.

In this specific project, the decision makers were surprised that the pyrotechnic devices were the major driver because of the exceptional high reliability of the devices. Indeed, they are high reliability, but they were one of the very few critical items that were single-point failures. It was not that they were "low" reliability, but they were less reliable than all the other subsystems most of which were functionally redundant. Knowing that there exists a distribution of reliability data in the estimates, they took actions to assure that they were getting the best of the lots; actions such as increasing the quality assurance measures on the pyrotechnics, independent inspections, additional testing, etc. Similar charts were used to define the operational test and simulation procedures, to inject failures into the simulations that represented the "tall-pole" failures identified by the Importance Values, and to develop contingency operational procedures should they occur.

CLOSING COMMENTS

We recommend that projects consider the value and utility of these products and implement them where appropriate. Should the project manager worry about affordability of this type of analysis, it is best to remember that actions taken based on these analyses avoided expenditures that exceeded the cost by a factor of 20. The simple act of directing the mission design to be done to support the second launch opportunity, as shown in Figure 4, rather than according to the original plan, saved the project more money than the cost of the analysis for the entire project life cycle.

ACKNOWLEDGMENT

The authors wish to acknowledge the contribution of the Double Pareto Box concept by Richard Kitterman.

Barney Roberts is the director of the Futron Risk Management Center of Excellence. Prior to coming to Futron, he was on the engineering staff of Johnson Space Center for more than 30 years. He was involved in project management and system engineering for Apollo, Space Shuttle and advanced programs. Roberts has earned many awards including the NASA Exceptional Service Medal and a patent for an aerobraking orbital transfer vehicle.

(E-mail address: broberts@futron.com)

Clayton Smith is managing a number of high technology assessment projects for Futron Corporation. He is currently managing the development and use of the International Space Station's Probabilistic Risk Assessment in support of the program's decision-making process. Smith has nearly 20 years experience working to analyze large complex systems from a reliability/availability and safety point of view. His clients have consistently recognized him as an outstanding performer.

(E-mail address: csmith@futron.com)

David Frost is a senior consultant and brings 14 years of space program experience to Futron's Risk Management discipline. He has been instrumental in the unique risk based decision analysis activity for the Space Shuttle upgrades program and the risk analysis for the Jet Propulsion Laboratory (JPL) Mars Exploration Rovers project. Frost was awarded the "Silver Snoopy" for his outstanding contributions to the Space Shuttle Program.

(E-mail address: dfrost@futron.com)