MANAGING RISK WITHIN A DECISION ANALYSIS FRAMEWORK

Homayoon Dezfuli[(1)], Robert Youngblood[(2)], Joshua Reinert[(2)]

[(1)] NASA Headquarters, OSMA, 300 E street, SW, Washington DC 20546, USA
Email: hdezfuli@nasa.gov
[(2)] Information Systems Laboratories Inc., 213 Hallock Road, Suite 3b, Stony Brook, New York 11790, USA
Email: ryoungblood@islinc.com; jreinert@islinc.com

(Presented at the Second IAASS Conference, Chicago (USA,) 14-16 May 2007)

ABSTRACT

This paper presents a risk management (RM) approach to managing the risk associated with space systems, and missions pursued using them, throughout their lifecycles. The approach, which is based on an analytic-deliberative decision-making methodology [1, 2, 3], embeds NASA's current Continuous Risk Management (CRM) process in a broader decision analysis framework [4, 5]. The CRM process promotes proactive identification and control of departures from program objectives. The analytic nature of the proposed enhancement to RM promotes formal analysis of the consequences of decision alternatives in terms of performance measures (PMs) relating to program fundamental objectives and explicit treatment of uncertainties. The deliberative aspect of the approach allows the consideration of elements that have not been captured by the formal analysis, and provides an opportunity to scrutinize the modeling assumptions of the analysis and the relevant uncertainties. The use of risk management in this way supports the attainment of the holistic and risk-informed decision-making environment that NASA intends to achieve.

1. INTRODUCTION

The existing paradigm for managing risk at NASA uses the Continuous Risk Management (CRM) process and relies heavily on the application of risk matrices [5]. A brief description of each tool is provided below.

1.1. The CRM Process

CRM is a well-established tool within NASA. As illustrated in Figure 1, CRM is an iterative and adaptive process that in intended to promote the successful execution of program intent.

The steps used in the CRM process are the following:

- Identify – Identify program risk by identifying scenarios having adverse consequences (deviations from program intent).
- Analyze – Estimate the likelihood and consequence components of the risk through Analysis.
- Plan – Plan the Track and Control actions. Decide what will be tracked, decision thresholds for corrective action, and proposed risk control actions.
- Track – Track program performance compared to its plan.
- Control – Given an emergent risk issue, execute the appropriate control action, and verify its effectiveness.
- Communicate and Document – This is an element of each of the previous steps. Focus on understanding and communicating all risk information throughout each program phase.

Figure 1. The CRM Process

As information is gained during program implementation, more is learned about the program's risk. Improved processes may be identified, calling for updates to the risk management plan and to the baseline risk. This adaptive feature of CRM is suggested in Figure 1. CRM is a process in which each step builds on the previous step, initiated at the beginning of the program life cycle and continuing throughout the program.

1.2. Risk Matrices

Risk matrices are widely used in risk management activities. The simplicity of the risk matrix concept and its compatibility with the CRM process are the primary motivations for its pervasive use at NASA. The matrices provide assistance in categorizing and communicating risk issues. Figure 2 shows a setup of likelihood and consequence definitions for a "5x5" risk matrix (very commonly used in space systems RM

applications). In this setup, the likelihood range is subdivided into five likelihood "bins." Similarly, the consequence range is subdivided into five severity levels.

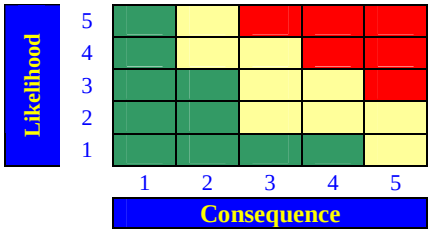

Figure 2. A Typical Risk Matrix

By discretizing likelihood and consequence severity, risk tolerability regimes are then defined. Figure 2 shows three regimes:

- An upper regime (red area) where identified "risks" are considered intolerable, and risk reduction measures are essential.
- A middle regime (yellow area) where costs, disruption of schedule, degradation of performance, or other factors are taken into account as part of risk reduction decision process.
- A lower regime (green area) where identified "risks" are considered negligible, or so small that no risk reduction measures are needed.

The risk tolerability regimes defined by risk matrices are based on the concept of iso-risk curves. In Figure 3, the boundaries between the green and yellow areas, and between the yellow and red areas are iso-risk curves. Any point on a given iso-risk curve corresponds to the same expected consequences, or "risk."

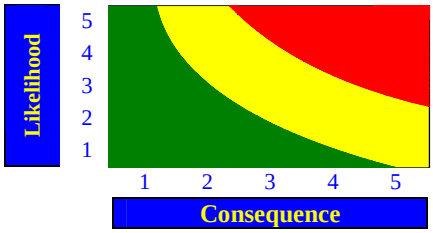

Figure 3. Iso-Risk Curves

Expected consequence *E(C)* or "risk" is calculated as the product of the probability and consequence:

$$Risk \equiv E(C) = Probability \times Consequence \qquad (1)$$

Note that the "expected consequences" definition of "risk" is operationally different from the "triplets" concept of risk, which will be presented later. A given value of *E(C)* could result either from low-probability (P), high-C or high-P, low-C scenarios. The iso-"risk" idea suggests that these are equivalent for decision-making purposes, while the triplets concept allows for the possibility that they are not.

1.3. Observations on the Existing Risk Management Paradigm

The existing RM paradigm places emphasis on identification of individual "risks" and on accountability for action items associated with particular "risks." Given a program decision, individual "risks" to a successful outcome are identified, prioritized according to the ranking scheme of a risk matrix, and assigned to a responsible entity. The existing paradigm is essentially concerned with *management* of individual risks, given a previously formulated program baseline. A more proactive approach would introduce risk management considerations more explicitly into the process of selecting alternatives in the first place.

Additionally, for the following reasons, use of risk matrices in decision-making should be approached with caution.

- If the risk matrix is viewed as the deliverable, rather than the underlying analysis, it is tempting to substitute subjective completion of the form for careful analysis.
- Interaction between risks is not usually considered explicitly. Each "risk" is mapped onto a matrix individually.
- The matrix deals with individual risks, not with aggregate risks (i.e., total risk). This supports assignment of specific "track" and "control" action items to individuals, but does not furnish proper perspective to decision-makers.
- Consequence types are often not discriminated. Inclusive consequence severity levels (e.g., human safety and asset safety together) short-circuit the ability to perform risk-trade studies.
- Uncertainties are not acknowledged and characterized. A risk is assumed to exist within one likelihood range and consequence range, both of which are assumed to be known.
- The desire to balance likelihood against consequence drives RM decisions. A rare but severe risk contributor may warrant a response different from that warranted by a frequent, less severe contributor, even though both have the same expected consequences.

In summary: risk matrices can be used to communicate assessment results for individual risk contributors, but are not an acceptable substitute for analysis, and should not be allowed to overshadow the presentation of the integrated risk profile.

2. A PROACTIVE RISK MANAGEMENT APPROACH

The purpose of risk management is to promote program success in two general ways: by incorporating risk-informed decision-making in the formulation of the program baseline, and by identification and control of departures from the program baseline. This purpose is addressed by embedding risk management in a Decision Analysis framework. This provides risk management a focus on project objectives, an analytical basis for risk management decisions and risk trade studies, a tight coupling between the bases for major decisions and the ensuing management activities, and a framework for dealing with uncertainty [3, 5].

Risk management as presented here is an iterative process to

- Identify the risks associated with implementation of decision alternatives,
- Analyze this risk and provide input to prioritizing work to resolve uncertainty, if warranted,
- Provide input to determining the preferred decision alternative in light of program priorities,
- Plan, track, and control the implementation of the selected alternative,
- Communicate, deliberate, and document the risk, and
- Iterate with previous steps in light of new information.

As discussed here, risk management is not just follow-through on mitigation of perceived threats to success, given a chosen course of action. Rather, it is a proactive, analytic-deliberative, risk-informed approach to enable/ enhance the decision makers' effective selection of key program alternatives, with implementation and follow-through considered prospectively as part of the process of developing decision alternatives.

3. REVIEW OF SEVERAL KEY TERMS AND CONCEPTS

3.1. Risk

The term "risk", when used without further qualifications, is very general and applies to a large variety of user contexts. This is true even when use of the term is limited to the space program and system domain, because risk concepts are routinely applied to more than one context within the domain, e.g., in relation to safety, program / project management and decision-making, or mission assurance. Risk, in a general sense, is the expression of potential shortfalls relative to explicitly established and stated program objectives. The potential shortfalls may be in any one or more of the three basic program execution domains:

- system technical performance, which is understood to include safety,
- program cost,
- program schedule.

NASA uses the term "safety" broadly, to include human safety (public and workforce), environmental safety, and asset safety [4]. When specifically considering safety-related risk, the shortfalls take the more specific form of adverse outcomes in any of the potential impact areas that have safety relevance:

- human life and/or injury,
- health effects,
- property and/or equipment damage
- environmental damage.

Regardless of the type of risk that may be of interest for specific circumstances, technical assessment and consideration requires the definition and characterization of three basic components of risk:

- A definition of the *scenarios* that may happen. Scenario definition is especially useful when organized in a logical fashion to identify the cause-consequence relationship of events that constitute scenarios.
- A characterization of the *probabilities* of the risk scenarios that have been identified. This characterization can be expressed quantitatively in the form of a probability over some reference period of time or set of activities, or as a "frequency", i.e., a probability per unit of time. The characterization should include uncertainty in the probability.
- A characterization of the severity of the *consequences* associated with the scenarios that have been identified.

The "triplet" concept of risk as defined above is operationally useful because it makes clear that in order to define, assess and manage risk it is necessary to produce three components of risk: undesired scenarios, their probabilities, and their consequences. Mathematically, risk can be written as a set of triplets:

$$Risk \equiv \{< Scenario, Probability, Consequence >\} \quad (2)$$

Analysis of risk in these terms not only supports quantification of performance (discussed below), but also furnishes insights into the importance of different contributors to risk, helping to steer the formulation of improved alternatives.

3.2. Performance Measures (PMs)

A Performance Measure (PM) is a quantifiable metric used to characterize performance of the decision alternatives with respect to a particular fundamental objective. Capability PMs relate to fundamental mission objectives. Examples of PMs for an architecture decision for a lunar surface mission might include mass delivered to lunar surface, mass returned from lunar surface, surface accessibility, usable surface crew-hours, and system availability.

Safety PMs are metrics that provide measures of the safety performance of a system. Safety PMs can be defined in terms of the probability of a consequence type of a specific magnitude (e.g., probability of any general public deaths or injuries) or the expected magnitude of a consequence type (e.g., the number of public deaths or injuries). Probability of Loss of Mission P(LOM) and Loss of Crew P(LOC) are two particularly important PMs for manned apace missions. Because an actuarial basis does not suffice for prediction of these probabilities, modelling will be needed to quantify them.

3.3. Hazard vs. Risk

Hazard is distinguished from risk. Hazard is a state or a set of conditions of a system that has the potential to cause harm. The harm is realized when the hazard, together with the occurrence of certain events in the environment of the system, produces an accident or mishap with consequences adverse to safety. Risk addresses not only the potential for harm, but also the scenarios leading to adverse outcomes and the probability of these outcomes (refer to Equation 2).

3.4. Probabilistic Risk Assessment (PRA)

PRA is a scenario-based risk assessment technique that quantifies the likelihoods of various possible undesired scenarios and their consequences, as well as the uncertainties in the likelihoods and consequences. PRA can be applied to quantify Performance Measures that are very closely related to fundamental objectives (e.g., probability of loss of crew). PRA focuses on development of a comprehensive scenario set, which has immediate application to identify key and candidate contributors to risk. In all but the simplest systems, this requires the use of models to capture the important scenarios, to assess consequences, and to systematically quantify scenario likelihoods. These models include reliability models, system safety models, simulation models, performance models, logic models, etc.

4. EMBEDDING THE CRM PROCESS IN A DECISION ANALYSIS FRAMEWORK

Figure 4 shows the interface between Decision Analysis and the elements of CRM. As illustrated here, Risk Management is an iterative activity that includes formulation and analysis of key program alternatives, as well as tracking implementation of the selected alternatives within CRM. This activity is performed throughout the product life cycle.

The first step in Figure 4 ("Formulate Objectives Hierarchy / PMs", top middle of figure) culminates in selection of a set of program-specific Performance Measures (PMs) that both scope and prioritize much of the risk management activity. Program alternatives are formulated in light of these PMs (next box on Figure 4), and a preferred alternative is ultimately selected by the decision-maker based in part on quantification of expected PM performance using risk analysis. Deliberation is used as appropriate in selecting the final alternative.

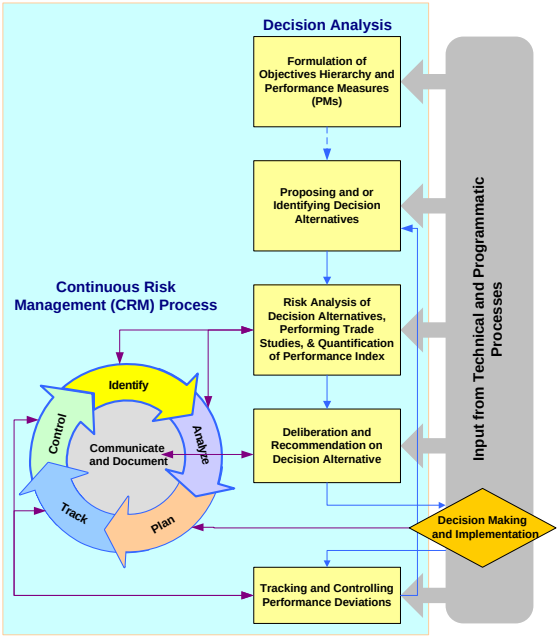

Figure 4. Embedding the CRM Process in a Decision Analysis Framework

The risk analysis of decision alternatives not only guides selection of a preferred alternative, it also carries out the "Identify" and "Analyze" steps of CRM. Selection of a preferred alternative is based in part on an understanding of the risks associated with that alternative. Alternative selection is followed immediately by a planning activity in which key implementation aspects are addressed, namely, risk

tracking and control, including risk mitigation if necessary. Also shown conceptually on Figure 4 is the interface between risk management and other technical and programmatic processes.

In discussing each process step below, reference will be made to the coupling between the decision analysis steps and the corresponding elements of CRM.

4.1. Formulation of Objectives Hierarchy / PMs

In this step, a program Objectives Hierarchy is developed, and PMs are determined [2, 4, 5]. An example of an objectives hierarchy is shown on Figure 5. Details will vary from program to program, but a construct like Figure 5 is behind the program-specific objectives hierarchy.

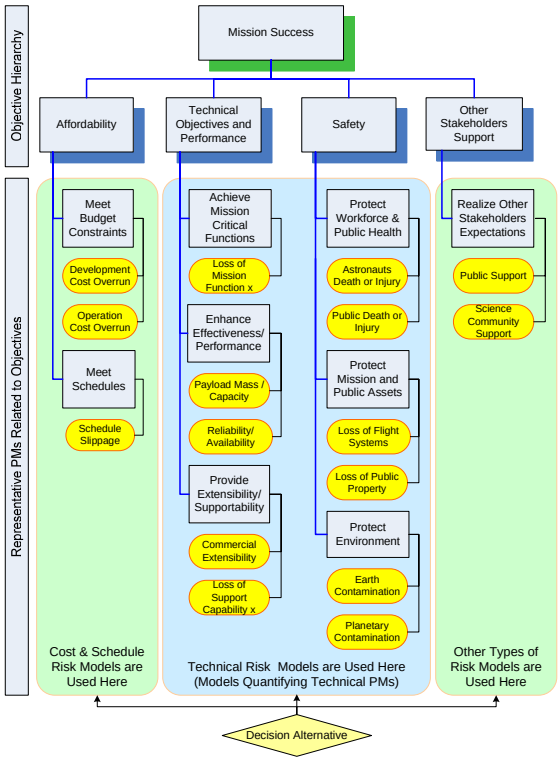

Figure 5. An Example of an Objectives Hierarchy

The top tier of this hierarchy is "Mission Success." The idea is to evaluate the impact on this ultimate objective of each decision alternative listed in the diamond at the bottom of the figure. Since "Mission Success" is very general, a hierarchical approach is employed to develop quantitative metrics that will measure the achievement of this top-level objective. The next tier in the tree lists the general objective categories that constitute mission success, i.e., "affordability," "technical objectives and performance," "safety," and "other stakeholder support." At the next tier, these categories are elaborated upon further by listing a number of objectives. Thus, the category "safety" becomes the three objectives: "protect workforce and public health," "protect mission and public assets," and "protect environment." For each third tier objective a set of PMs are identified. For example, two PMs for the objective "environment" are: "planetary contamination" and "earth contamination." This hierarchical, tree-like structure shows the objectives that the decision maker values in making the decision.

Some PMs cannot be reasonably quantified using analytical models. In these cases, a constructed scale should be developed [6]. For example, "Stakeholders Support" has no natural measurement. Therefore, a discrete scale consisting of degrees of stakeholder support should be developed.

Once the PMs are established, the utility to the decision-makers within the range of each PM must be determined [1, 3, 5, 6]. The purpose of this step is to allow for the performance of decision alternatives to be measured with a single metric. Utility is a number ranging from zero to one. For example, the utility of different amounts of the Schedule Slippage PM must be determined. The maximum reasonably foreseeable amount of slippage is assigned a utility of zero. The minimum reasonably foreseeable amount of slippage is assigned a utility of one. Intermediate amounts of slippage are assigned utilities by the decision-makers within the range of zero to one.

4.2. Proposing and/or Identifying Decision Alternatives

Proposing and/or identifying decision alternatives is the step of creating a set of alternatives that can potentially achieve the goals and objectives of the system and are good enough to warrant the investment of the analytic resources needed to rank them. This step depends on understanding the system's functional requirements and operational concept [5]. No approach is guaranteed to produce a set of alternatives that includes the globally-best choice, but consideration of the objectives and PMs during the formulation of alternatives helps to produce better alternatives. Running an alternative through an operational time line or reference mission is a useful way of determining whether it can plausibly fulfill these requirements. Sometimes it is necessary to create separate behavioral models to determine how the system reacts when a certain stimulus or control is applied, or a certain environment is encountered. This provides insights into whether it can plausibly fulfill time-critical and safety requirements. Defining plausible alternatives also requires an understanding of the technologies available, or potentially available, at the time the system is needed.

4.3. Risk Analysis of Decision Alternatives, Performing Trade Studies and Ranking (CRM Steps: "Identify," "Analyze")

The goal of this step is to carry out the kinds and amounts of analysis needed to characterize the risk for two purposes: ranking decision alternatives, and performing the "Identify" and "Analyze" steps of CRM.

To support ranking, trade studies may be performed. PMs that can affect the decision outcome are quantified for all alternatives, with quantification to include uncertainty as appropriate. As discussed later, these PM results are rolled up into integrated "Performance Indices" (PIs) on the basis of which the decision alternatives can be ranked. These results are then input to the deliberation process indicated on Figure 4.

Refer to Figure 6 . Risk analysis can take many forms, ranging from qualitative risk identification (essentially scenarios and consequences, without performing detailed quantification of likelihood using techniques such as fault trees), to highly quantitative methods such as PRA [4]. The analysis stops when the technical case is made; if simpler, more qualitative methods suffice, then more detailed methods need not be applied.

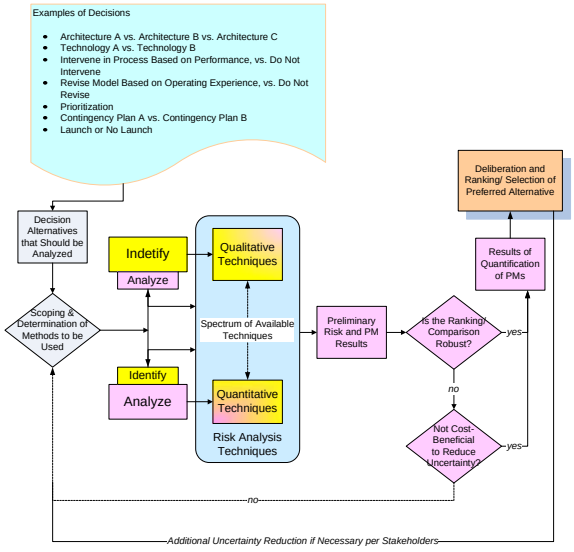

Figure 6. Risk Analysis of Decision Alternatives

Selection and Application of Appropriate Methods

The nature and context of the problem, and the specific PMs, determine the methods to be used. In some projects, qualitative methods are adequate for making decisions; in others, these methods are not precise enough to appropriately characterize the magnitude of the problem, or to allocate scarce risk reduction resources. In the latter case, the improvement in decision-making that would result from a more precise analysis justifies the increase in costs of the improved analysis [3, 5]. The risk analyst needs to decide whether risk identification and judgment-based characterization are adequate, or whether the improved quantification of PMs through more detailed risk analysis is justified. In making that determination, he/she must balance the cost of risk analysis against the value of the additional information to be gained. The concept of "value of information" [3] is central to making the determination of what analysis is appropriate and to what extent uncertainty needs to be quantified. Value of information is discussed briefly in Reference 5.

In addressing safety, qualitative methods are used to scope out the hazards that need to be addressed, and then the hazards are addressed more quantitatively, as necessary and appropriate. Although quantitative accuracy is not an early priority, completeness in the set of hazards identified IS an early priority. Completeness is never assured, but it is very important to identify probabilistically significant influences on risk. Operating experience is valuable in some areas, but in general, data sufficient to determine the likelihoods of the top-level adverse consequences of interest (e.g., loss of crew) are not available. If accidents occur as a result of a design flaw, inadequate testing, or analysis error, such causes are eliminated after the fact, but for high-consequence adverse outcomes, the preferred method is to identify the mechanisms a priori and prevent them. There is no way to be certain of completeness, but a systematic approach to hazard identification and extensive use of applicable operating experience are two important ingredients for promoting completeness.

Quantitative Methods

If the stakes associated with a decision are high, or if a risk-informed decision requires careful consideration of uncertainty and/or careful quantification of risk metrics, then more detailed and correspondingly expensive methods are implied. For example, NPR 8715.3 (NASA General System Safety Requirements) [4] requires that, for manned missions, a PRA be performed to quantify the probability of Loss of Crew. PRA adopts a systematic approach to scenario development, in whatever detail is necessary for quantification. Understanding of the structure of the scenarios supports identification of specific risk issues, and formulation of corresponding risk mitigation strategies. PRA applies state-of-knowledge methods for quantifying probabilities and consequences (including uncertainty). Quantification of the probabilities may call for significant use of operating experience data from NASA or other program areas, human error modeling, and,

depending on the technical scope of the analysis, simulation of phenomenological behavior may also be needed. The resources needed for PRA are justified by the importance of the consequences modeled. In general, the stopping rule for risk analysis is that the cost in time and resources of further analysis is no longer justified by the expected benefits.

Quantification of Performance Index (PI)

PMs are quantified for each alternative and used to quantify an overall "Performance Index" (PI) [2, 5] for each alternative. These results are then used for ranking alternatives. In general, the PMs are not valued equally by the decision maker. This leads to the assignment of weights w_i to the PMs. These weights represent the decision maker's judgments and values. They can be assigned directly or, preferably, using structured methods such as the Analytic Hierarchy Process (AHP). These weights must be normalized to unity, i.e.,

$$\sum_{1}^{N_{PM}} w_i = 1 \tag{3}$$

Where N_{PM} is the number of performance measures. The impact of each decision alternative on each PM is assessed in terms of the corresponding utility. The weights and utilities are used to calculate the performance index for each decision alternative which represents the overall average utility of that alternative.

An Example

A simple example is provided here to show how the PI is calculated. This example is incomplete and stylized for the purpose of this presentation.

The decision to be made is an architecture decision for a science mission. For this example, three PMs are considered, adapted from those in Fig. 5:

- Planetary Contamination
- Payload Mass
- Cost Overrun

Using a structured method or judgment, as described above, the weights (w_i) of these PMs are determined as shown in Table 1 below:

Table 1. Performance Measures and Their Weights

PM	w_i
Planetary Contamination	0.70
Payload Mass	0.15
Cost Overrun	0.15

The utility to the decision-makers within the range of these PMs is determined by the decision-makers. Utilities are dimensionless scales that represent the value to the decision maker of reasonably foreseeable levels of impact, normalized to range from 0 to 1. The following simplified, discrete utility functions are assumed for this example:

Table 2. Utilities for the Performance Measure

PMs and Associated Ranges (Levels)	Utility (u_{ik})
Planetary Contamination	
Widespread Contamination	0.0
Localized Contamination	0.7
No Contamination	1.0
Payload Mass	
1000 kg	0.0
1500 kg	0.6
2000 kg	1.0
Cost Overrun	
120% of Budget	0.0
110% of Budget	0.85
100% of Budget	1.0

Two decision alternatives (Alt) are identified, A and B. Through risk analysis, it is determined that these alternatives perform as follows, where uncertainty in performance is represented by a probability mass function over the range of performance:

Table 3. Results of Risk Analysis

PMs and Associated Levels	Prob. of Alt A Performing at this Level (p_{ik}^j)	Prob. of Alt B Performing at this Level (p_{ik}^j)
Planetary Contamination		
Widespread Contamination	0.0	0.1
Localized Contamination	0.2	0.8
No Contamination	0.8	0.1
Payload Mass		
1000 kg	0.7	0.0
1500 kg	0.3	0.1
2000 kg	0.0	0.9
Cost Overrun		
120% of Budget	0.7	0.0
110% of Budget	0.3	0.1
100% of Budget	0.0	0.9

The PI is a function of the utilities, weights, and performances according to the following equation:

$$PI^{j} = \sum_{i=1}^{N_{PM}} w_i \left(\sum_{k=1}^{Ni} u_{ik}\, p_{ik}^{j}\right) \tag{4}$$

Where,

PI^j is the Performance Index for the jth alternative
N_{PM} is the number of PMs
w_i is the weight of the ith PM
N_i is the number of values that the ith PM can take
u_{ik} is the kth utility bin of the ith PM
p_{ik}^j is the probability that the jth alternative occupies the kth bin of the ith PM.

For Decision Alternative A, the calculations are as follows:

$$\begin{aligned} PI^{A} &= w_1(u_{11}\, p_{11}^{A} + u_{12}\, p_{12}^{A} + u_{13}\, p_{13}^{A}) \\ &+ w_2(u_{21}\, p_{21}^{A} + u_{22}\, p_{22}^{A} + u_{23}\, p_{23}^{A}) \\ &+ w_3(u_{31}\, p_{31}^{A} + u_{32}\, p_{32}^{A} + u_{33}\, p_{33}^{A}) \end{aligned}$$

$$\begin{aligned} &PI^{A} \\ &= 0.7(0.0 * 0.0 + 0.7 * 0.2 + 1.0 * 0.8) \\ &+ 0.15(0.0 * 0.7 + 0.6 * 0.3 + 1.0 * 0.0) \\ &+ 0.15(0.0 * 0.7 + 0.85 * 0.3 + 1.0 * 0.0) \end{aligned}$$

$$PI^{A} = 0.658 + 0.027 + 0.03825 = 0.72325$$

The same equation for Decision Alternative B results in:

$$PI^{B} = 0.75375$$

Based on the current state of knowledge, Alternative B is preferred over Alternative A due to its higher PI. However, it is noteworthy that in this example, the strengths and weaknesses of the two alternatives are very different. “B” performs worse on contamination, but better on payload mass and cost. The performance uncertainties are such that an improved state of knowledge (reduced uncertainty) could actually reverse the ranking (refer to Figure 6).

The use of scenarios as a component of risk is required to support the risk management strategy discussed in this paper. Given a scenario-based model, the significance of classes of contributors can be assessed for potential improvements. In this example, the decision might be improved by better information (reduced uncertainty) regarding the expected performance, (value of information). It might also be possible to improve the designs themselves to achieve better performance (value of control). Sensitivity studies can be performed to bound the change in the PI that could result from a change in the PM values of a particular alternative. Depending on the level of detail at which the scenarios are developed, importance measures could be calculated for finer details in the model, such as for component failure rates, or for the rate of a failure mechanism in a group of components.

In this example, absent design improvements or uncertainty reductions, selection of Alternative B for implementation would be accompanied by high-priority risk management activities to prevent widespread planetary contamination. An integrated model of scenarios leading to planetary contamination would provide a basis for focusing on particular areas to manage this particular risk. Some attention would also be warranted in the areas of cost and payload mass.

4.4. Deliberation

As shown on Fig. 4, deliberation is the final step before decision-making. Some elements of the decision may not have been fully considered in the analysis. For example, there may be a metric that should be considered in the decision, but for which a decision-maker will be hesitant to explicitly produce a utility function. These elements may be considered less formally during deliberation.

There may also be modeling assumptions for which there was no consensus. Deliberation provides a forum for these assumptions to be scrutinized outside of an analytical process.

The final block of Fig. 4, “Tracking and Controlling Performance Deviations,” is not discussed in this paper as it is acceptably handled in the existing process.

5. SUMMARY

This paper advocates an approach to risk management that embeds the CRM process, on which current practice is based, in a broader decision analysis framework. The expected benefit of doing this is to improve the analytical basis for program decision-making of many kinds, including architecture decisions, prioritization of modeling and analysis, prioritization of research to reduce uncertainty, risk management decisions (setting PM thresholds, evaluation of PM trends), and program management decisions.

This is not to take issue with CRM as far as it goes. CRM is seen to have important strengths in its institutional approach to accountability for individual risks. However, CRM does not contain within itself an

objective analytical framework with which to establish the adequacy of a given set of risk management measures in a given program. CRM's emphasis on accountability is not matched by the emphasis on technical strength. Moreover, the use of "risk matrices" for communicating the significance of individual "risks" can appear to substitute for the analysis process; assignment of a risk to a bin may substitute for deeper analysis. Use of risk matrices is not a substitute for the integrated risk profile.

In principle, decision analysis can help in these areas. Also, decision analysis contains a "stopping rule:" a basis for deciding when enough analysis has been done to support the decisions being made. Accordingly, the approach discussed here essentially applies CRM steps in implementation of a standard Decision Analysis approach.

6. ACKNOWLEDGMENTS

This paper has benefited substantially from work performed for the recent update of the Systems Engineering Handbook, especially group discussions among contributors to its sections on Decision Analysis and Technical Risk Management.

7. REFERENCES

[1] National Research Council (1996) *Understanding Risk: Informing Decisions in a Democratic Society*, Commission on Behavioral and Social Sciences and Education, Washington, D.C., National Academy Press

[2] Stamatelatos, M., H. Dezfuli, and G. Apostolakis, "A Proposed Risk-Informed Decision-Making Framework for NASA," Proceedings of the 8th International Conference on Probabilistic Safety Assessment and Management, May 14-18, 2006, New Orleans, LA.

[3] Clemen, R. (1995) *Making Hard Decisions – An Introduction to Decision Analysis*, Belmont, California: Duxbury Press

[4] NASA Procedural Requirement 8715.3, "NASA General Safety Program Requirements," 2006.

[5] NASA Draft Handbook SP-6105, "Systems Engineering Handbook," 2007.

[6] Keeney, R. (1992) *Value-Focused Thinking: A Path to Creative Decisionmaking*, Harvard University Press, Cambridge, Massachusetts.