2010

Monitoring Risk Response Actions for Effective Project Risk Management

Kujawski, Edouard; Angelis, Diana

Wiley Periodicals, Inc.

Systems Engineering, Vol 13, No. 4, 2010
https://hdl.handle.net/10945/44168

Regular Paper

Monitoring Risk Response Actions for Effective Project Risk Management

Edouard Kujawski[1,*] **and Diana Angelis**[2]

[1]*Systems Engineering Department, Naval Postgraduate School, Monterey,* CA 93943
[2]*Defense Resources Management Institute, Naval Postgraduate School, Monterey,* CA 93943

Received 8 October 2008; Revised 5 June 2009; Accepted 17 August 2009, after one or more revisions
Published online 9 November 2009 in Wiley Online Library (wileyonlinelibrary.com)
DOI 10.1002/sys.20154

ABSTRACT

Complex projects typically involve high-consequence, project-specific risks that require detailed analysis and for which risk response actions (RRAs) need to be developed and implemented. The risk picture is dynamic. The sources and consequences of risks evolve and change over the project lifecycle; thus, it is necessary to constantly monitor risk. RRAs that do not keep pace with the changing project situation are a major cause of risk management failures. This paper extends traditional cost risk analysis from a purely macroscopic perspective by evaluating and tracking project-specific risks and RRAs at the microscopic level. The key elements of the method are (i) develop risk scenarios, (ii) model them using generalized decision trees, and (iii) quantify the risks using Monte Carlo simulation. For each risk the probability and cost values are conditional on the specific RRA and the preceding outcomes. The use of fractional factorial design provides a subset of all possible RRA combinations for efficiently determining the preferred total project RRA solution. Risk curves are generated to provide the necessary information to analyze, track, and manage the performance of the selected RRAs over time. Project managers and team leaders can use this information to dynamically manage the RRAs to keep pace with the changing project situation, thereby increasing the probability of project success in a cost-effective manner. The approach is detailed using a realistic but simplified case of a project examined first with one and then expanded to three technical risks. © 2009 Wiley Periodicals, Inc. Syst Eng 13: 353–368, 2010

Key words: risk analysis; technical risk; risk response actions; decision tree; risk curves; risk dependencies; correlation; risk dynamics; microscopic analysis; Monte Carlo simulation; design of experiments

1. INTRODUCTION

Complex engineering projects are susceptible to project-specific risk drivers such as low Technology Readiness Level (TRL), high-design and/or manufacturing complexity, significant requirement changes, sizeable quantity changes, large funding uncertainty, severe acts of nature, serious accidents, and poor management decisions [Sage, 1992; Chapman and Ward, 1997; Conrow, 2003; U.S. Government Accountability Office, 2009]. These risk factors must be identified, assessed, mitigated, and controlled through formal risk management. The latter is an essential and critical discipline that is implemented in many of today's complex engineering projects/programs (referred to simply as "projects" in this paper). Probabilistic Cost Risk Analysis (PCRA) provides a proper framework for handling the many different elements of cost uncertainties, including project-specific, high-consequence outcomes that result from failure to meet technical performance requirements. Figure 1 depicts PCRA as an integral part of the risk management process. To be

* Author to whom all correspondence should be addressed (e-mail: ekujawsk@nps.edu; diangeli@nps.edu).

effective, PCRA should interface with the risk management activities, be regularly updated to reflect change in risk, and be integrated within the project's Earned Value Management System (EVMS) [NDIA-PMSC, 2005].

The emphasis on risk management supports efforts to reduce lifecycle costs of system acquisitions. Analysis of performance, cost, and schedule risks over the lifecycle of a system can yield substantial benefits. Conversely, ignoring important stages of the lifecycle can lead to substantial risk for product development at the beginning of the lifecycle and for product upgrade or replacement at the end [Pennock and Haimes, 2002]. The lack of and/or inadequate dynamic risk management has been identified as a major cause of project failures. The Lockheed Management Student Guide [Waldof, 1998: 53] states:

> Risk management efforts that fail do so because the risk control actions did not keep up with a changing program situation.

Many sources of cost uncertainty in complex engineering projects such as economic/business factors (rates-wages, overhead, vendor/supplier stability, inflation indices, multi-year assumptions, etc.), learning-rate assumptions, and cost-reduction initiatives are well understood within the framework of a macroscopic perspective. These can be effectively modeled with classical Probability Distribution Functions (PDFs) such as the triangular, Beta, lognormal, and Weibull distributions [Garvey, 2000; Vose, 2006]. These classical PDFs, however, do not model some of the negative influences of the acquisition process and human behaviors such as the MAIMS principle, "Money Allocated Is Money Spent" [Gordon, 1997; Kujawski, Alvaro, and Edwards, 2004]. The MAIMS principle has important implications for contingency and risk management [Kujawski, 2007]. Another important effect is dependencies between cost elements. These can be accounted for using correlation coefficients

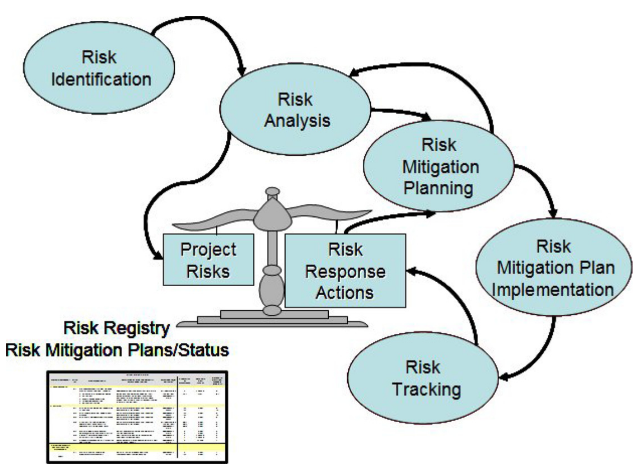

Figure 1. The DoD risk management process modified to explicitly include probabilistic cost risk analysis (adapted from DoD [2006]). [Color figure can be viewed in the online issue, which is available at wilyeonlinelibrary.com.]

[Book, 2000/2001] or explicitly modeled [Garvey, 2000]. However, even accounting for these additional effects, macroscopic factors constitute only a fraction of today's typical project risk drivers and, therefore, cost uncertainty.

Cost, schedule, and technical risks are confounding factors. Problems meeting technical requirements usually result in cost overruns and schedule delays. It is tempting to assume or claim that the classical PDF analysts typically elicit for cost also quantify the risks associated with project-specific, high-consequence events. Analysts often go through the effort of identifying and discussing risk drivers. However, when it comes to quantifying the risks and estimating cost contingencies, they simply apply high/low ranges without thinking about how a particular technical risk driver affects one or more cost elements. We think that this approach is incomplete because it (1) masks valuable information and visibility about important risks and (2) fails to adequately track RRAs.

Best practices implement the identification of risk drivers as the kick-off activity of cost risk analysis [Blanchard, 1998; Smith and Merritt, 2002; INCOSE, 2007]. Hollmann [2007] explicitly accounts for the risk drivers using their Expected Values (EVs) in addition to the macroscopic classical PDFs. He refers to this approach as "Driver-Based Monte-Carlo" (DBM). This represents a significant improvement over today's typical Monte Carlo analysis. However, the use of EV is a major limitation for use in a truly probabilistic framework where Decision-Makers (DMs) may be confronted with low-probability, high-consequence outcomes that influence their psychological biases [Tversky and Kahneman, 1992]. Clemen and Reilly [2001: 133] write:

> Comparing expected values of different risky projects is useful, but in many cases EV inadequately captures the nature of the risks that must be compared. With risk profiles, however, we can make a more comprehensive comparison of the risks.

The objective of mitigating the technical risk drivers is to lower both the probability and consequences of cost overruns and schedule delays while meeting the technical performance. In this paper we look at the cost impact of technical RRAs, which we refer to as Technical Risk Cost (TRC). We characterize TRC in terms of cost risk curves, consistent with the following two definitions of risk:

—"Risk can be quantified by a probability distribution of the potential outcomes, or by the relevant moments of that distribution" [Paté-Cornell, 1996: 96].
—"Risk is a measure of future uncertainties in achieving program performance goals and objectives within defined cost, schedule and performance constraints" [Department of Defense, 2006: 1].

The analysis of specific risk drivers and RRAs requires a microscopic view (referred to in this paper as the "micro view"). It is best carried out with tools such as decision trees (DTs), influence diagrams, or other discrete representations [Marshall and Oliver, 1995; Kujawski, 2002; Kujawski, Alvaro, and Edwards, 2004; Haimes, 2004]. We use the micro

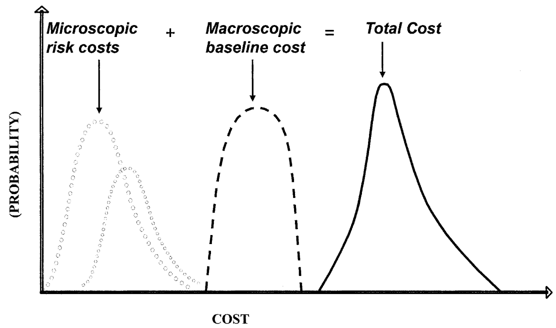

Figure 2. Total project cost as the sum of the macroscopic baseline costs and micro risk costs (adapted from Federal Transit Administration [2004]).

view to explicitly model high-consequence events and RRAs and thereby provide a support tool for better decision-making [Dillon and Paté-Cornell, 2001]. The micro view also assists Subject-Matter Experts (SMEs) to think about credible, high-consequence events and better deal with overconfidence or optimism biases. However, the micro view is too cumbersome to individually analyze every technical risk and source of cost uncertainty. It complements and needs to be integrated within today's typical PCRA. The total project cost is then modeled as the sum of the macroscopic baseline cost (referred to in this paper as the "macro view") and the individual TRCs, as depicted in Figure 2.

In this paper we develop a practical and realistic integrated micro-macro PRCA approach that project managers and team leaders can use to dynamically determine RRAs and effectively keep pace with the changing project situation. We focus on technical risks and quantify the associated TRCs at timely decision points in terms of risk curves. In Section 2 we discuss the dynamic picture of project risks and the need for an integrated micro-macro approach. In Section 3 we present the use of Generalized Decision Trees (GDTs) and illustrate their application for the analysis of a single risk with two RRAs as the project evolves. In Section 4 we extend the approach to multiple risks and illustrate the method for the realistic but simplified case of a project with three technical risks. We also use a fractional factorial design to determine the preferred TPRRA combination. In Section 5 we discuss and demonstrate the value of monitoring and dynamically managing technical risks. We close the paper with some recommendations for further development.

2. FROM RISK ASSESSMENT TO TOTAL RISK MANAGEMENT

2.1. The Risk Assessment vs. Risk Management Paradigms

"Risk management" is an overloaded term. It is used to denote both (1) the activities following risk assessment and (2) the entire process of risk assessment and its management. Following Haimes [1991], we use the term "total risk management" for the latter. The standard quantitative risk assessment paradigm focuses on the following triplet of questions articulated by Kaplan and Garrick [1981]:

1. What can go wrong?
2. What are the associated likelihoods?
3. What are the consequences?

Once these critical questions have been answered, the greater challenge of risk management is to address and control the following three issues articulated by Haimes [1991]:

1. What can be done and what options are available?
2. What are the tradeoffs in terms of costs, benefits, and risks?
3. What are the impacts of current decisions on future options?

Successful risk management requires that the above three critical issues be properly and continuously addressed to reflect the dynamic character of project risks depicted in Figure 3 [Murphy et al., 1996; Graham, 2004]. The sources and consequences of risk continue to evolve and change over the project lifecycle. The performance of the RRAs needs to be monitored and controlled to ensure they are adequately mitigating risk. Depending on the performance, some RRAs might need to be changed. In general, at any point in time there will be a mix of acceptable and unacceptable results. Consequently, management reserves need to be reviewed on a periodic basis and dynamically allocated where needed to ensure project success [Kujawski, 2007].

2.2. A Dynamic Risk Management Process

Today's typical PCRA includes the following steps [Garvey, 2000]:

1. The cost and/or risk analysts (simply referred to as "analyst" below) and the SMEs jointly identify the individual risk drivers.
2. The analyst and the SMEs jointly screen the identified risks for further analysis and risk mitigation.
3. The analyst and the SMEs account for the technical risks as cost uncertainties within the framework of the classical cost PDFs.

We now modify and extend the approach as follows:

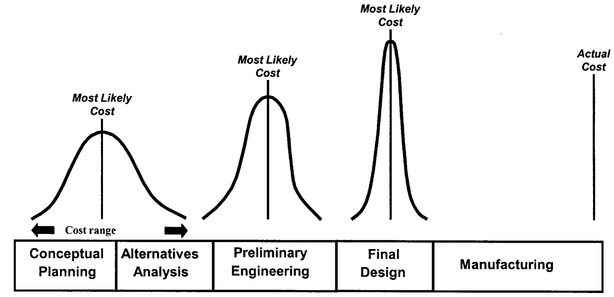

Figure 3. The dynamic picture of project cost risk (adapted from Federal Transit Administration [2004]).

4. The analyst models the TRC of each technical risk and selected RRAs over time using a GDT (see Sec. 3).
5. The analyst works with the SMEs to quantify the value of the decisions and outcomes for each GDT using discrete and continuous PDFs. We favor the Direct Fractile Assessment (DFA) method for data elicitation and fitting the associated cost elements with a three-parameter Weibull distribution [Brown, 2008].
6. The analyst quantifies the GDTs using Monte Carlo simulation. The potential outcomes are modeled in terms of risk curves. This is in contrast to the use of Classical DT (CDT) analysis where decisions are based on EV.
7. Risk curves are used to examine the range of possible outcomes for specific RRAs and provide the data for informed cost-benefit tradeoffs.
8. The analysis is updated based on actual results to effectively monitor and control the performance and selection of RRAs as old risks are retired and new risks arise.

The proposed dynamic risk assessment and management approach provides a mathematically valid as well as practical framework for dealing with the dynamic character of project risks. It builds on the sequential decision analysis model of Stonebraker and Kirkwood [1997] by assessing the changing project situation in terms of risk profiles. It thereby provides DMs with the necessary information to dynamically select or modify RRAs in accordance with their assessment of future outcomes and risk tolerance rather than relying solely on EVs.

3. GENERALIZED DECISION TREES FOR MODELING AND ANALYZING RISK RESPONSE ACTIONS—SINGLE RISK CASE

We model and analyze each screened risk driver and the associated candidate RRAs using GDTs. GDTs follow the standard DT representation. The decision nodes and chance nodes are depicted as squares and circles, respectively. The ordering of the decision nodes corresponds to different temporal deterministic events. The branches that originate with decision nodes represent the available RRAs. The branches that originate with chance nodes represent the possible probabilistic outcomes. There are two major differences between the GDT and the classical DT: (1) PDFs rather than discrete branches are associated with the chance nodes; and (2) the outcomes are specified as PDFs rather than point estimates or EVs. GDTs, therefore, provide a powerful technique for dealing with the complex situations typical of many of today's projects. They avoid bushy trees, generate risk curves, and remove the reliance of decision-making based on EVs [Kujawski, 2002]. The comparison between the GDT and the classical DT is further discussed in Section 3.2.

3.1. Illustrative Example #1—A Single Technical Risk

To illustrate the approach, we consider the single technical risk related to the fabrication of a complex first-of-a-kind module or system. The risk is associated with the probability that the module may fail and have to be redesigned. This, of course, has cost and schedule consequences. We focus on the cost consequences of the following two RRAs: (i) Directly fabricate the module, which is denoted by the branch labeled Fab_A2; or (ii) build a prototype and then fabricate the module, which is denoted by the branch labeled Prototype. The associated GDT is depicted in Figure 4. A descriptive label, a probability, and a cost distribution are associated with each branch. The ordering of the decision nodes corresponds to the different temporal deterministic events in the development and fabrication cycle of the module. The probability and cost values are conditional on the specific RRA and the preceding outcomes. The outcomes are collectively exhaustive and mutually exclusive [Papazouglou, 1998]. They are therefore statistically independent and it would be a mistake to model them as correlated. To be explicit, the outcome "Success_A1" depends on the preceding events (Prototype, Success_A1, and Fab_A1). In contrast, "Success_A1," "Fail-

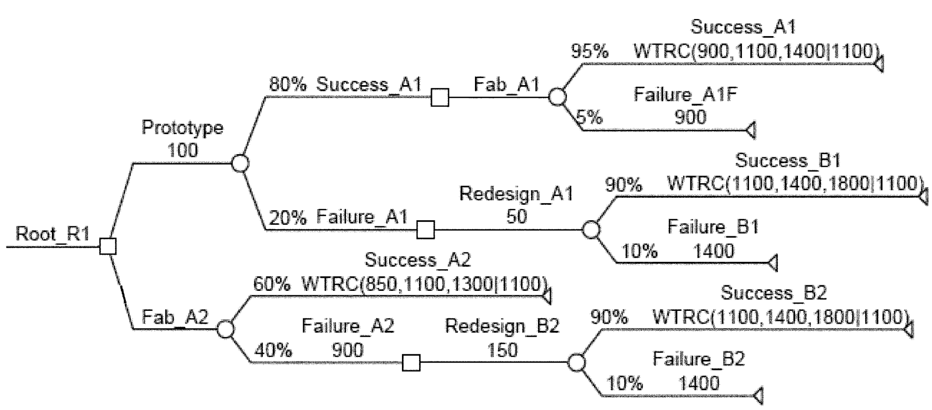

Figure 4. Illustrative example #1. Generalized decision tree for the Risk #1 analysis. It depicts the fabrication of a module. To mitigate the associated technical risk, the project considers two initial RRAs: (i) Directly fabricate the module (branch labeled Fab_A2), or (ii) build a prototype and then fabricate the module (branch labeled Prototype). WTRC(x, y, z | c) denotes the technical risk cost associated with a MAIMS-modified Weibull PDF with the specified 10th, 50th, and 90th percentiles and allocated or baseline cost c. The cost values are in $K.

ure_A1F," "Success_B1," and "Failure_B1" are mutually exclusive events and independent events.

The data in Figure 4 are as follows:

1. The cost elements are modelled using three-parameter Weibull PDFs fitted to the 10th, 50th, and 90th percentiles determined in accordance with the DFA method. In this paper, these PDFs are denoted by Weib(10^{th}, 50^{th}, 90^{th}).

2. The above PDFs are modified in accordance with the MAIMS principle to account for the fact that account managers rarely underrun their Originally Allocated Budget (OAB). The modified PDFs are proper PDFs with a deltalike function at OAB that account for all values less than or equal to OAB [Kujawski et al., 2002]. In this paper, these MAIMS-modified Weibull PDFs are denoted by Weib(10^{th}, 50^{th}, 90^{th}|OAB).

3. The TRC of each RRA is measured relative to the OAB. It is interesting to note that this is equivalent to the use the semivariance or value-at-risk concept [Markowitz, 1997]. The TRC is then a random variable given by

$$\mathrm{WTRC}(10^{th}, 50^{th}, 90^{th}|\mathrm{OAB}) \equiv \mathrm{Weib}(10^{th}, 50^{th}, 90^{th}|\mathrm{OAB}) - \mathrm{OAB}.$$

4. To be explicit, consider the PDF WTRC(900, 1100, 1400|1100), associated with the top branch labeled Success_A1 in Figure 4. The initial PDF Weib(900, 1100, 1400) corresponds to a three-parameter Weibull distribution with the following parameters: Location: 797.2, Scale: 373.6, and Shape: 1.74. The associated mean is 1130. In contrast, the mean of Weib(900, 1100, 1400|1100) is 1255. This increase is expected because the minimum value in any trial is now 1100.00. The expected TRC is 155. The cost values should be thought of as $K.

The selection of a RRA is a deterministic event and only the outcomes associated with it can be realized. It would, therefore, be incorrect to weigh or combine the outcomes of the two RRAs since they represent two mutually exclusive decisions.

3.1.1. Analysis of the Initial Risk and RRAs

We analyzed each RRA and the potential outcomes using a commercial Monte Carlo simulation Excel add-in (Crystal Ball®, to be specific). The PDFs and risk profiles for each individual RRA at the start of the project are depicted in Figures 5a, 5b, and 5c. The PDFs (Figs. 5a and 5b) are multimodal and cannot be represented using any of the PDFs commonly used in risk analysis [Vose, 2006]. These PDFs provide significant information. The PDF for the direct fabrication RRA (Fig. 5a) has two high-value modes that correspond to the sequence of events in which (i) only the first fabrication fails (labeled "Success_B2") and (ii) the subsequent fabrication following redesign also fails (labeled "Failure_B2"). The PDF for developing a prototype RRA (Fig. 5b) exhibits a high-value mode that corresponds to the outcome in which the fabrication of the module fails (labeled "Failure_A1F"). The other high-value peak is hardly noticeable.

The complementary cumulative distribution functions (CCDF) or "risk curves" shown in Figure 5c provide a more global picture. The exceedance probability is the probability of exceeding a given cost. For example, the Risk_Cost(Fab_A2) curve shows that there is ~30% probability that the TRC will exceed $1100K. For any given value on the "Risk cost" axis, the risk curve that corresponds to the lower exceedance probability represents the lower risk. It is seen from Figure 5c that the "prototype" risk curve is significantly lower than the "fabrication" risk curve for all values greater than ~$200K. The investment of $100K for building a prototype provides a significant reduction in TRC. For the manager trying to decide if it is worthwhile to invest in the prototype option, the answer is to invest as long as the anticipated benefits from the prototype (whether it be cost savings, time savings, information, etc.) exceed $200K and/or if the low-probability/high-consequence costs of direct fabrication are unacceptable.

3.1.2. Analysis of the Sequential Risks and RRAs

We use the GDTs to analyze the performance of the candidate RRAs at key decision points in the project life-cycle. This information is essential for tracking the residual risk exposure versus the cost expended on RRAs and modifying the RRAs as needed to ensure mission success. Figure 6 depicts the risk curves at the "Prototype" RRA start and after the successful demonstration of the prototype. These two risk curves represent the TRC of the risk exposure at two different points in time and thereby provide a metric for the risk exposure as the project evolves. The residual risk following a successful prototype is less than the original risk. As expected, the risk curve moves to the left of the original risk curve and is steeper, which reflects a reduced risk. In contrast, a risk curve that moves to the right of the original risk curve means that the risk exposure is increasing and the selection of RRAs needs to be reconsidered.

3.2. Today's De Facto Decision Tree Analysis

There are significant differences between the approach followed in this paper and today's de facto DT analysis, which we refer to as CDT analysis. While both approaches model decisions under uncertainty using DTs, much of the valuable information available in the risk curves is lost in the CDT "folding back the tree" procedure that calculates the EV of each alternative. The CDT analysis selects alternatives based on the highest or lowest EVs depending on the situation [Clemen and Reilly, 2001: 115]. Many technical managers however do not favor the use of EV for decision-making. They consider risk to depend more on the magnitude than the probability of the undesirable outcomes and therefore not adequately characterized by its EV [Shapira, 1995]. In addition, the perception of risk and therefore the preferred RRAs are specific to the value of the project, the available resources, and the risk attitude of the DMs.

For completeness and clarity, we performed the CDT analysis for illustrative example #1. We used the commercial software, DecisionPro™. This tool also includes Monte Carlo simulation and a limited set of PDFs for use as input values. (Note: DecisionPro™ does not provide the Weibull distribution.) There are other commercial DT analysis tools that are equally applicable and solve decision trees using Monte Carlo

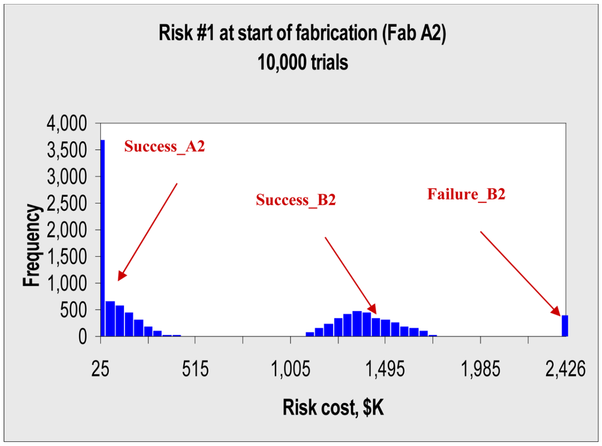

Figure 5a. Illustrative example #1. Probability distribution for Risk #1 outcomes given direct fabrication of the module (sequence starting with branch labeled Fab_A2 in Figure 4). The different possible outcomes are identified. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

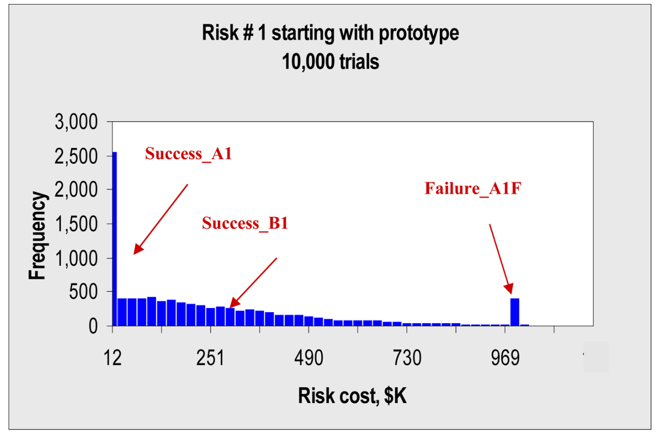

Figure 5b. Illustrative example #1. Probability distribution for Risk #1 outcomes given development of a prototype (sequence starting with branch labeled Prototype in Fig. 4). The different possible outcomes are identified. The other high-value peak is hardly noticeable. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

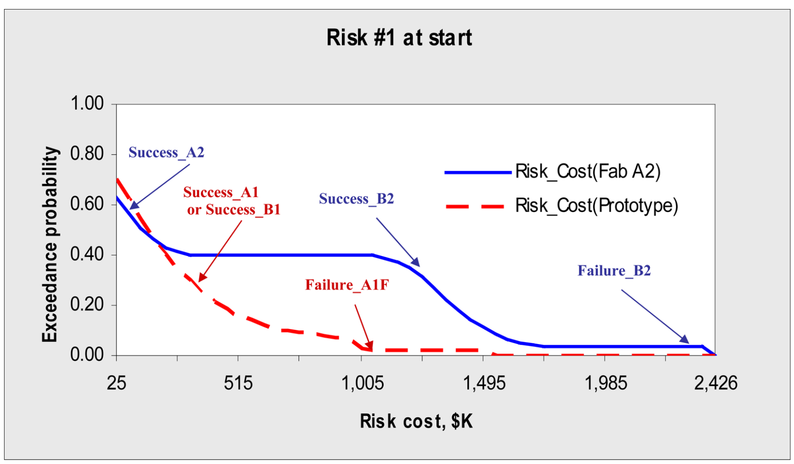

Figure 5c. Illustrative example #1. Risk curves corresponding to the PDFs in Figures 5a and 5b. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

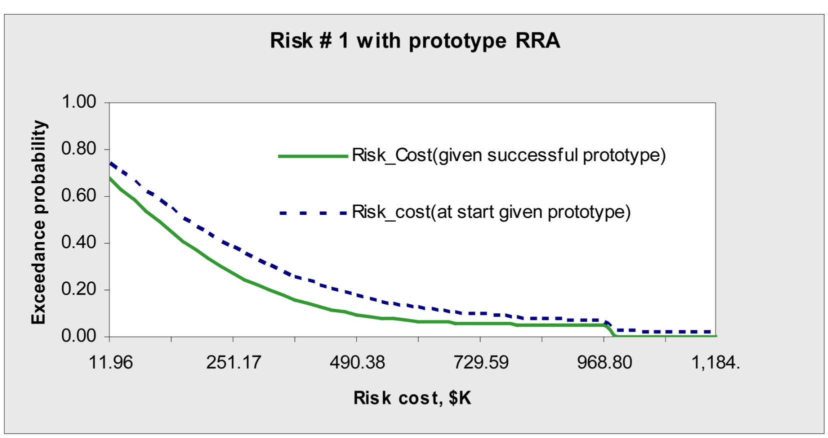

Figure 6. Illustrative example #1. Risk exposure characteristics for Risk #1 with the development of a prototype at the start of risk mitigation and after successful demonstration. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

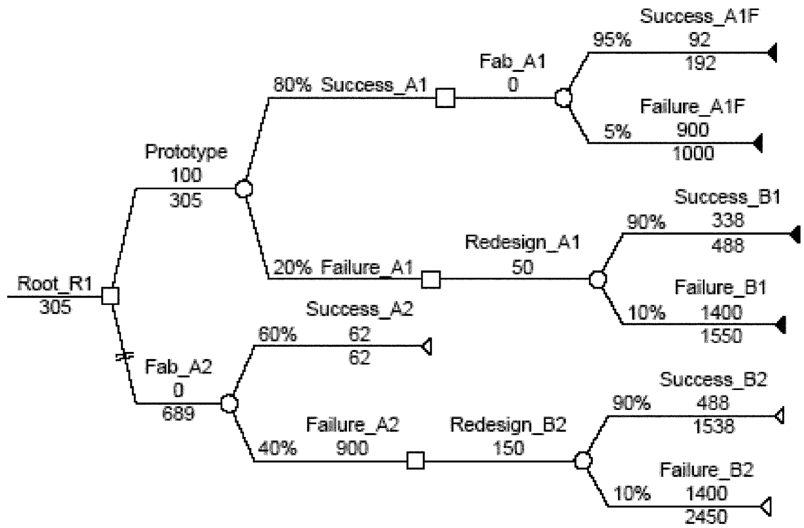

Figure 7. Illustrative example #1. Classical decision tree analysis for the Risk #1 analysis. The structure is identical with the generalized decision tree in Figure 4. Each PDF has been replaced by its expected value, and the "folding back the tree" procedure has been applied to obtain the EVs for the two options.

simulations [Maxwell, 2008]. However, few offer the flexibility of commercial Excel add-ins such as Crystal Ball® and @Risk to implement the GDT analysis presented in this paper. A detailed comparison of DT tools is beyond the scope of this paper.

The resulting CDT is depicted in Figure 7. Its structure is identical to the GDT in Figure 4. The number immediately below the branch name corresponds to the point estimate of the incremental cost associated with the decision or probabilistic outcome. The number below an end branch represents the point estimate of the TRC if the associated outcome is realized. The major differences between the CDT and the GDT are that (1) the PDFs are replaced by their expected values and (2) the "folding back the tree" procedure is applied to provide the EVs rather than the cumulative risk profiles [Clemen and Reilly, 2001: 132]. The CDT analysis thereby discards the alternatives with the inferior expected values regardless of the range of the possible outcomes. In this example, the CDT analysis yields EVs of \$305K and \$689K for the Prototype and Fab_A2 options. We note that these values are the EVs provided by the GDT analysis; this is as it should be. However, the GDT analysis also provides the cumulative risk curves, which is valuable information for the decision-maker.

4. THE QUANTIFICATION OF MULTIPLE PROJECT RISKS

Consider a project with n credible technical risks $\{R_i, i = 1, \ldots, n\}$. Each risk, R_i, is characterized by a probability of occurrence p_i and a spectrum of possible outcomes with a PDF $L_i(X_i = x_i)$, where X_i is a random variable that quantifies either the possible technical performance, cost, or schedule outcomes. This paper focuses on cost including technical TRC. One may then think of this set of risks as a risk portfolio [Kujawski and Miller, 2007] with a generalized discrete PDF,

$$R_S(X_S = X_1 + X_2 + \cdots + X_n = x) = \left\{ \langle p_1, L_1(x_1)\rangle, \langle p_2, L_2(x_2)\rangle, \ldots, \langle p_n, L_n(x_n)\rangle, \langle 1 - \sum_{i=1}^{n} p_i, 0\rangle \right\}. \quad (1)$$

The technical project costs in Eq. (1) may be probabilistically dependent [Dillon-Merill et al., 2008]. The dependencies are project-unique and need to be reassessed as the project situation changes and the risk picture evolves. For example, the technical performance and the cost or time to complete one module may depend on the performance and availability of another module. In this situation, the probability and cost values for a given risk are conditional on one or more other risks. The dependencies between the different risks can be either explicitly modeled or empirically accounted for using correlation coefficients. Explicit modeling of the dependencies eliminates the need for correlation coefficients. Whenever possible, the modeling of the dependencies is preferred to the use of empirical correlation coefficients [Garvey, 2000: 283]. It is tempting to use rank-order correlation because they can be easily input in commercial simulation tools (Crystal Ball®, @Risk, DecisionPro™, etc.); however, this could lead to erroneous results especially when dependencies have already accounted for [Garvey, 2000: 333; Salmon, 2009].

As depicted in Figure 2, the total project cost is a random variable that consists of the sum of the m base cost elements and the explicitly identified TRCs. Depending on the situation, the m base cost elements $\{BC_i, i = 1, \ldots, m\}$ may be modeled as point estimates or random variables $\{Y_i, i = 1, \ldots, m\}$ with continuous PDFs, which we denote by $BC_i(y_i)$. In general, n is not equal to m since some cost elements may have multiple risks while some have none. The total project cost is then a random variable with a PDF that is the probabilistic sum of the m base cost elements and n risk-driver costs:

$$TC(Z = \sum_{i=1}^{m} Y_i + \sum_{i=1}^{n} X_i = z) =$$

$$\sum_{i=1}^{m} BC_i(y_i) + R_S(x_1, x_2, \ldots, x_n), \tag{2}$$

where $BC_i(y_i)$ is the PDF for the base cost element BC_i and $R_S(x_1, x_2, \ldots, x_n)$ is given by Eq. (1).

Equation (2) can be computed using commercial Monte Carlo simulation tools such as Crystal Ball® and @Risk. These tools also provide tornado charts that conveniently quantify the importance of the various risk drivers and their link to the overall cost risk. Projects can use this information to rank the risks and prioritize RRAs. This is in sharp contrast with: (1) the use of point estimates that are at best ambiguous because overly confident staff provide low cost estimates, while others may inflate their cost estimates to make it easier to achieve success; (2) the use of S-curves that provide a limited "black box" view of project risks and cost uncertainty; and (3) decision-making based solely on qualitative assessments.

4.1. Developing Total Project Risk Response Actions

Developing RRAs is a creative process that may require generating new and innovative ideas [Hall, 1998: 141]. There are typically many RRAs for a particular source of risk. The morphological box approach [Sage and Armstrong, 2000] can be used to obtain useful insight into the solution space and identify a set of preferred RRAs. At a minimum, a project should consider the following three categories of RRAs for each risk [Kujawski, 2002]: (1) Accept the risk as is; (2) immediately implement risk reduction actions such as selecting alternative designs, switching or using multiple vendors, modifying the requirements, and/or pursing parallel paths; or (3) obtain additional information by investing in additional analysis, testing, and/or prototyping. These three RRAs adequately characterize the basic RRA options with little or no loss of generality and, if deemed necessary, they can be combined to provide more comprehensive or higher-level RRAs.

If there are many risks each with multiple RRAs, the set of all possible Total Project RRAs (TPRRAs) may be quite large and the analysis can be automated. However, computers are no substitute for critical thinking [Dixit and Nalebuff, 2008: 121]. It is therefore important to examine techniques that help reduce the set of TPPRAs and provide deeper insight into the solution. Design of Experiments (DOE) provides a method that can be used to define a fractional factorial design [Creveling, Slutsky, and Antis, 2003] for use in determining the preferred set of RRAs. This has the potential to greatly reduce the analysis effort and simplify the task of interpreting the results.

A project with multiple risks should at a minimum consider the following three strategies based on the three RRA categories discussed above:

—Strategy 1, accept all project risks as is. The project recognizes the existence of the risks; however, it considers it acceptable to simply monitor them using their standard approach.
—Strategy 2, immediately implement RRAs for every project risk. This is rarely realistic or feasible given schedule and budget considerations. However, it provides information on what it might cost to essentially eliminate technical risk from the project.
—Strategy 3, select RRAs that provide a good Return on Investment (ROI) [Hall, 1998] based on the traditional expected value and/or the low-probability/high-consequence outcomes.

4.1.1 Applying ROI to the Selection of RRAs

Hall [1998: 143] defines the risk management ROI as "the savings for all managed risks divided by the total cost of the risk management activities." Since we are interested in both the expected values and the low-probability/high-consequence outcomes, we define the ROI at different exceedance probabilities [Gogolkiewicz, 2007] as the ratio of the conditional expected value [Haimes, 2004: 308] of the TRC reduction divided by the cost of implementing the TPRRAs:

$$\mathrm{ROI}[\mathrm{TRC} \geq c] = \left[\frac{\int_c^\infty xp(x)dx}{\int_c^\infty p(x)dx}\right] \div (\mathrm{Cost\_of\_TPRRA}), \tag{3}$$

where $p(x)$ is the PDF for the TRC. When $c = 0$, Eq. (3) reduces to the traditional ROI. The bracketed expression in the right-hand side of Eq. (3) represents the conditional loss or TRC associated with outcomes greater than or equal to c.

There is some danger in applying the ROI for ranking RRAs since it does not explicitly identify the RRA costs. The ROI may mislead the unwary to a RRA that is beyond a reasonable investment [Melese, 2007]. Similarly, comparing RRAs at a single confidence level or exceedance probability is not a robust method because it does not consider all of the complete range of the outcomes. Another approach is to directly compare the risk PFDs and/or CCDFs. It is rare that a RRA exhibits first-order stochastic dominance. However, these distributions provide DMs full visibility into all of the potential outcomes. DMs can rank or select the RRAs depending on their attitudes toward risk and. budget constraints.

4.2. Illustrative Example #2—A Project with Three Technical Risks

Consider the hypothetical project with the three technical risks scenarios depicted in Figures 4, 8a, and 8b. This example is both rich and simple enough to illustrate: (1) several diverse RRAs and their analysis, (2) the dynamic nature of the risk picture, (3) the importance of continuous risk monitoring, and (4) the dynamic allocation of resources as a critical aspect of continuous risk management.

The Risk #1 scenario is identical to illustrative example #1. The associated GDT is depicted in Figure 4. The Risk #2 scenario with its associated GDT depicted in Figure 8a may be thought of as a prime contractor who subcontracts the engineering and fabrication of a complex module. The technical risk is associated with the probability that the subcontractor fails to deliver a module that performs to the technical requirements and the cost of the rework. We examine the TRCs of the following two RRAs: (i) Subcontract to a single contractor A, denoted by the branch PDR_A; and (ii) subcontract to two contractors and select the best one for fabrication of the module at a contractual design review, denoted by the branch PDR_AB. By initiating two different contractors with different offerings, the prime significantly reduces the probability of failure. The RW branches represent the different costs associated with rework. We do not invoke the MAIMS principle when dealing with rework. The cost of each of rework branch is modeled with a three-parameter Weibull distribution specified in terms of the 90th, 50th, and 10th percentiles, which are provided by SMEs or based on relevant historical data.

The GDT for Risk #3 is depicted Figure 8b. The Risk #3 scenario may be thought of a prime contractor who considers two different Verification and Validation (V&V) strategies as a means of risk reduction. The technical risk is associated with the probability that the design problems are not identified prior to fabrication and the cost of the rework. The branches labeled VVS_1_Start and VVS_2_Start represent implementing the standard approach with planned expenditures of $300K and a very thorough V&V effort with planned expenditure of $1000K, respectively. The rework is assumed to be inversely related to the V&V effort and it is modeled with a three-parameter Weibull distribution. We note that V&V deserves a thorough analysis rather than simply considering two RRAs as assumed in this illustrative example. Engel and Barad [2003] have proposed a methodology for developing V&V strategies and selecting an optimal one. Their approach can be integrated within the framework presented in this paper.

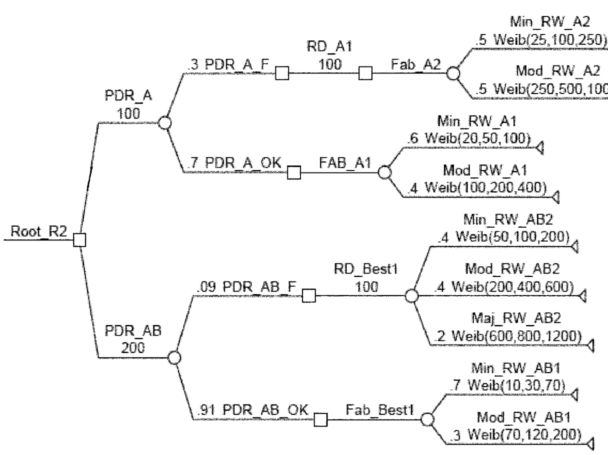

Figure 8a. Illustrative example #2. Generalized decision tree for the Risk #2 analysis. It depicts a prime contractor who subcontracts the engineering and fabrication of a complex module. To mitigate the associated technical risk, the project considers two initial RRAs: (i) Subcontract to a single contractor A (branch labeled PDR_A), or (ii) subcontract to two contractors and selecting the best one for fabrication at a contractual design review (branch labeled PDR_AB). Weib(x, y, z) is the three-parameter Weibull distribution specified in terms of the 10th, 50th, and 90th percentiles for rework costs.

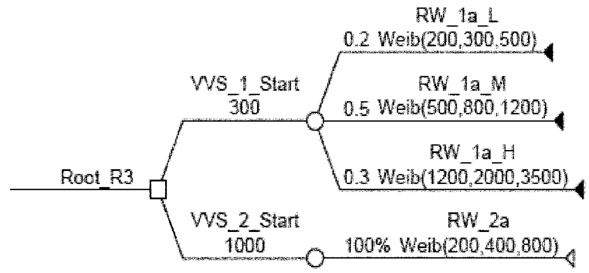

Figure 8b. Illustrative example #2. Generalized decision tree for the Risk #3 analysis. It depicts a prime contractor who considers two different V&V strategies as means for risk reduction: (i) use of the standard V&V approach (branch labeled VVS_1_Start), or (ii) use of a very thorough V&V approach (branch labeled VVS_2_Start). Weib(x, y, z) is defined in Figure 8a.

We assume that the three risks in this illustrative example are associated with distinct technical risk factors and nonconflicting resources. It is therefore proper to model them as independent. In the case of probabilistically dependent risks, the dependencies should be explicitly modeled or, if necessary, accounted for using correlation coefficients as discussed in Section 4.1.

4.2.1. Analysis of the Initial Set of Risks and RRAs

Given three technical risks each with two potential RRAs, there are eight possible initial total projects RRAs (TPRRA) that the project may opt to implement. Figure 9 depicts the morphological chart for the three risks and the identified RRAs. We analyze each of the individual three RRAs as described in Sections 3.1 and 4.1. The Risk #1, Risk #2, and Risk #3 profiles and ROI values for the mean and the 50th and 80th exceedance probabilities are presented in Figures 10a, 10b, and 10c, respectively. Based on these results, we consider the three TPRRAs in the alternatives matrix depicted in Figure 11:

Risk ID	RRA#1	RRA #2
Risk #1	Fab_A2	Prototype
Risk #2	PDR_AB	PDR_A
Risk #3	VVS_2_Start	VVS_1_Start

Figure 9. Illustrative example #2. Morphological chart for the three hypothesized technical risks.

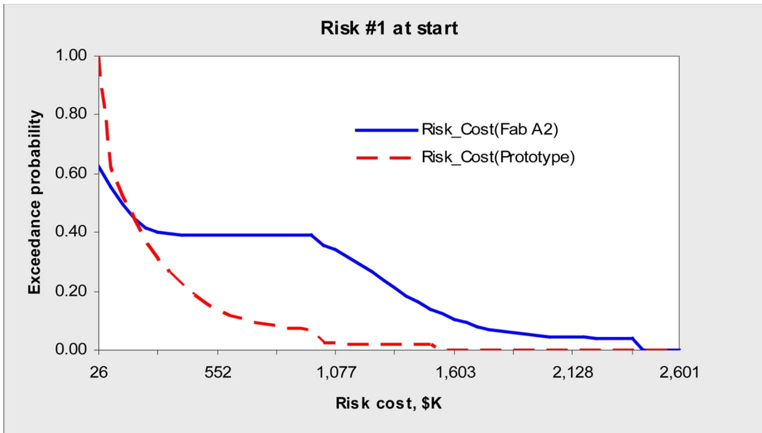

	TRC, $K		
	Fab_A2	Proto	ROI
80th percentile	1,392.04	451.72	9.65
Mean	626.99	304.19	3.23
50th percentile	151.52	170.38	-0.19

Figure 10a. Illustrative example #2. Risk curves and ROI table for Risk #1 at start. The risk curves are identical to Figure 5c. The risk curves are replicated for convenience. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Strategy #1. Use of the lowest cost option for each risk, which is equivalent to proceeding as normal; i.e., do not implement specific RRAs for any of the three risks. This corresponds to: (1) Risk #1, directly fabricate the module; (2) Risk #2, proceed with a single contractor; and (3) Risk #3, implement the standard V&V effort.

Strategy #2. Use the most effective RRA for each risk. It is the approach that a risk-averse project manager would favor if he had sufficient funding. This corresponds to: (1) Risk #1, develop a prototype; (2) Risk #2, proceed with two contractors for; and (3) Risk #3, implement the very thorough V&V effort.

Strategy #3. Select the RRAs based on the ROI as described in Section 4.1. This corresponds to: (1) Risk #1, develop a prototype; (2) Risk #2, proceed with two contractors; and (3) Risk #3, implement the standard V&V effort.

The three strategies are compared in Figure 12. Strategy #1 is practically dominated by Strategy #3. This is consistent with the Risk #3 risk curves and low ROI values depicted in

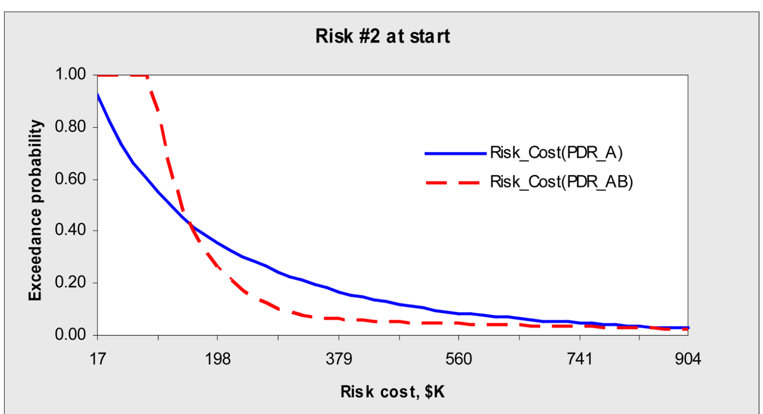

	TRC, $K		
	PDR_A	PDR_AB	ROI
80th percentile	345.34	230.09	1.98
Mean	221.65	208.57	0.13
50th percentile	133.82	151.11	-0.17

Figure 10b. Illustrative example #2. Risk curves and ROI table for Risk #2 at start. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

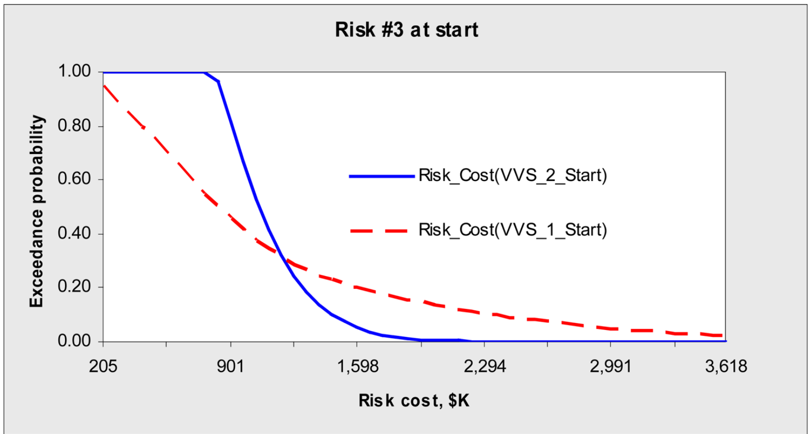

	TRC, $K		
	VVS_1_S	VVS_2_S	ROI
80th percentile	1,645.56	1,333.05	0.45
Mean	1,133.72	1,154.25	-0.03
50th percentile	869.90	1,093.27	-0.32

Figure 10c. Illustrative example #2. Risk curves and ROI table for Risk #3 at start. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Risk ID	**Strategy #1**	**Strategy #2**	**Strategy #3**
Risk #1	Fab_A2	Prototype	Prototype
Risk #2	PDR_A	PDR_AB	PDR_AB
Risk #3	VVS_1_Start	VVS_2_Start	VVS_1_Start

Figure 11. Alternatives matrix of the three TPRRAs selected for further analysis.

Figure 10c. The project could therefore eliminate Strategy #1 from further consideration.

4.2.2. Analysis of Risk Dynamics

We now focus on the analysis of the dynamics of the RRAs and how this information may influence the selection and management of the RRAs. Without loss of generality we illustrate the process using Strategies #1 and #2. The identical analysis applies to Strategy #3; however, Strategy #1 emphasizes the shortcomings of implementing the very thorough V&V effort.

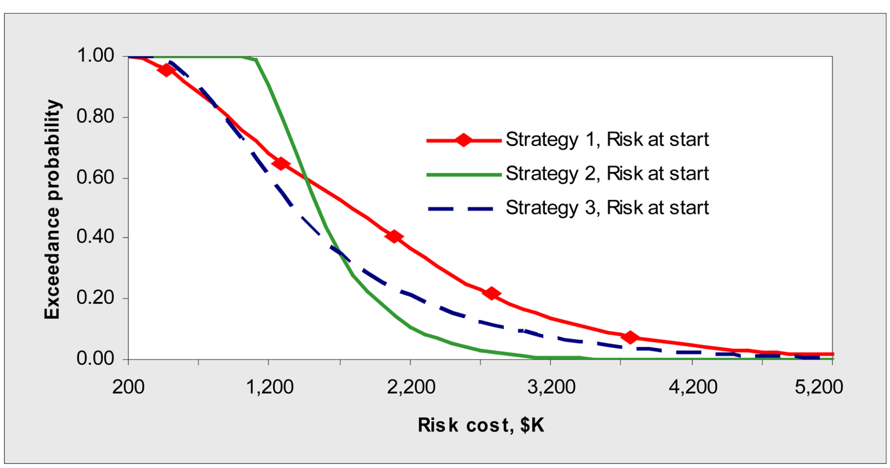

Figure 12. Illustrative example #2. Comparison of three strategies for the initial selection of RRAs. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

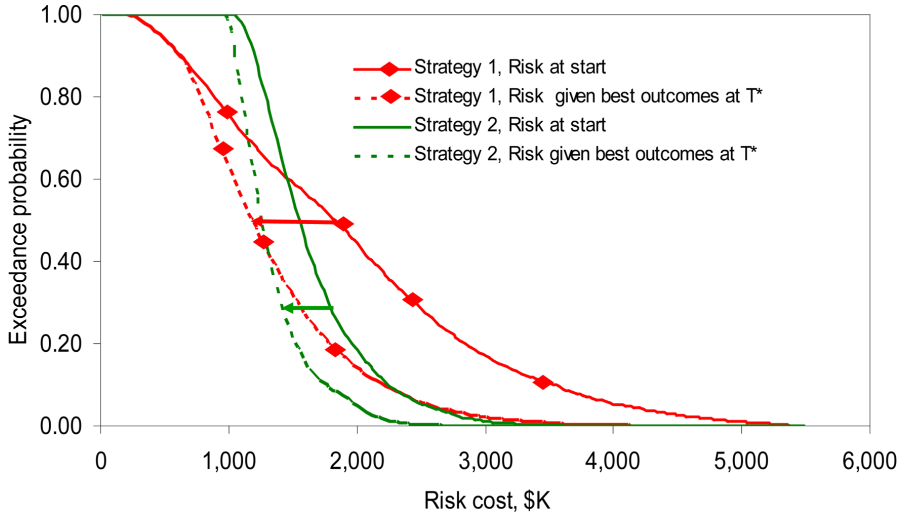

Figure 13. Illustrative example #2. Risk exposure characteristics for Strategies #1 and #2, assuming the best possible outcomes at the 2nd decision points; i.e., good luck prevails on the project for the initial set of risk response actions. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Optimistic view. For each strategy, we assume the best possible outcomes for the first probabilistic nodes, which for convenience we identify by the time of occurrence, T^*.

—The Risk #1 prototype and the Risk #2 review PDR_AB succeed.
—Risk #3 is continuously monitored and the risk reduction is directly accounted in the magnitude of the rework.

Pessimistic view. For each strategy, we assume the worst outcomes for the first probabilistic nodes, which for convenience we identify by the time of occurrence, T^*.

—The Risk #1 prototype and the Risk #2 review PDR_AB fail.
—Risk #3 is continuously monitored and the risk reduction is directly accounted in the magnitude of the rework.

Figures 13 and 14 depict the initial and residual risks under Strategies #1 and #2, assuming the best and worse outcomes, respectively. These data provide bounds for the risk range that may threaten the project following implementation of the initial set of RRAs. Figure 15 depicts these bounds. This is useful information that can be used to avoid analyzing the full factorial set of RRAs and to focus on fractional factorial sets of RRAs including orthogonal arrays [Huynh et al., 2009].

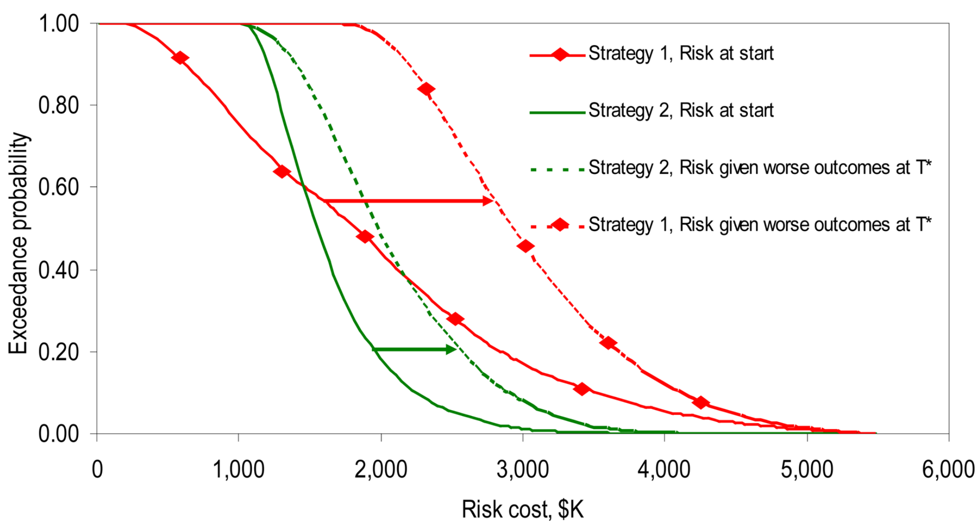

Figure 14. Illustrative example #2. Risk exposure characteristics for Strategies #1 and #2 assuming that assuming the worse possible outcomes at the 2nd decision points; i.e., Murphy's Law prevails on the project for the initial set of risk response actions. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

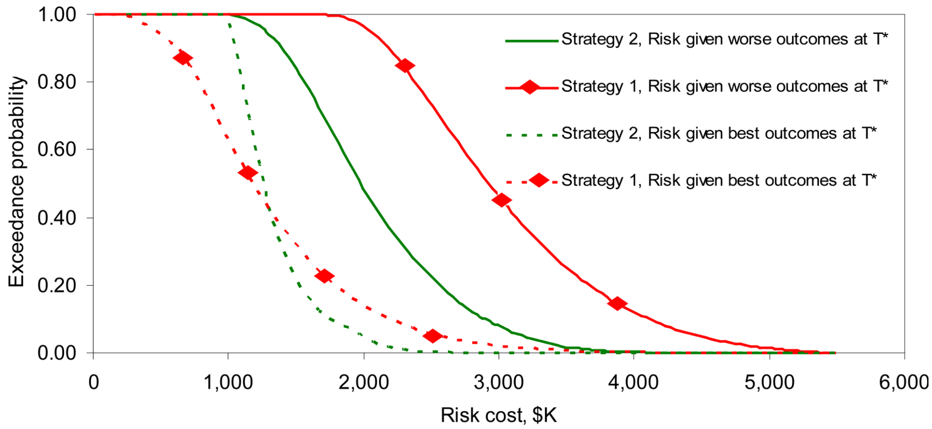

Figure 15. Illustrative example #2. Bounds on the project TRC following implementation of Strategies #1 and #2. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

The dynamic picture of risk and the value of monitoring risk are further discussed in the next section.

5. THE VALUE OF MONITORING RISK

Figures 13 and 14 illustrate the value of monitoring the performance of Strategies #1 and #2 in terms of the TRC. When the best outcomes are realized, both strategies reduce the TRC. As depicted in Figure 13, both T^* curves move to the left and become narrower than the start risk curves. Likewise, if the worst outcome prevails as assumed in Figure 14, then both strategies actually increase the cost risk exposure of the project. Graphing risk curves over time thus provides a metric to measure the success of risk mitigation efforts.

Figures 13, 14, and 15 provide additional quantitative information that can help the project rationally choose the preferred strategy. Under the best-case scenario (Fig. 13), at the start of the project Strategy #1 offers a lower risk exposure below $1500K while Strategy #2 offers a lower risk exposure above that value. Both strategies are equal in terms of exceedance probability (60%) at the "breakeven" point of $1500K. What does the project manager gain by extending the analysis to time T^*? He or she gains the information that the "breakeven" point is lower ($1200K) and the risk at that point is also lower (40%). So, which one is the best choice? The optimistic project manager most likely assumes that the best outcome would be realized and thereby makes a choice based on his/her anticipated benefits. But, of course, there is no such assurance, so let's examine the worst-case scenario.

Figure 14 shows the results of implementing each strategy over time assuming the worst outcome (Murphy's Law). As expected, the "accept the risk as is" Strategy #1 significantly increases the project cost risk exposure when things go bad; but the more conservative Strategy #2 is much less sensitive to bad outcomes. In fact, at T^* Strategy #2 dominates Strategy #1; i.e., it has a lower risk for any value. For pessimists, the choice is simple: Strategy #2 is especially effective in providing insurance against the worst outcomes. Which strategy is chosen depends on the DMs risk aversion and available resources. Is he/she an optimist or a pessimist? We believe examining risk information in this way provides useful insight and helps project managers make better choices.

6. CONCLUSIONS

6.1. Lessons Learned

This paper presents a method for evaluating and tracking project-specific risks at the micro level as well as cost overrun at the macro level. This type of analysis, as opposed to the purely macro-level risk analysis, is essential for risk management. While the macro level provides some information about total cost risk, the micro level allows the project manager to plan and control risk response actions that influence total cost risk. A simplified project with three technical risks was used to illustrate (1) the portfolio aspect of multiple RRAs and their analysis, (2) the dynamic nature of the risk picture, and (3) the monitoring of individual risks and allocation of management reserves.

This paper demonstrates the use of GDTs to model the evolution of the potential RRAs, and risk curves to evaluate the cost risk. The authors believe risk curves are better than the EV results given by CDT analysis because they contain all the risk information both in terms of probabilities and the value at risk. Risk curves derived from Monte Carlo simulation of GDTs are particularly useful for comparing different risk-mitigation strategies. The "breakeven" points help the risk manager understand the conditions under which each strategy is most appropriate. Combined with scenario analysis, it offers an opportunity to make cost-benefit tradeoffs among strategies. This thorough approach allows management to consider what they mean by "acceptable" risk and explicitly models the tradeoff between risk and benefit for a given RRA. Tracking the performance of RRAs over time is key to understanding the dynamic nature of risk management and defining essential changes in strategy. It enables a project to shift from allocating resources to each risk and the initially selected set of RRAs to making the decisions based on the observed development. This dynamic allocation of resources reduces the total risk project by eliminating the need to assume best cases, worst cases, or expected values for the outcomes for the technical risks. The proposed approach provides the detailed information that project managers need and want when they face hard decisions on projects. There is a cost; but it is small considering the potential benefits.

6.2. Future Research and Implementation

The proposed approach is both practical and mathematically valid. The use of fractional factorial designs based orthogonal arrays [Huynh et al., 2009] offers a method to greatly reduce the analysis effort and simplify the task of determining the preferred set of RRAs. To the best of our knowledge, no current commercial tool fully automates the generation and analysis of the proposed generalized decision trees and the time dependence. For additional details, we refer the interested reader to the recent decision analysis software survey by Maxwell [2008]. One of the referees suggested that the proposed approach can be readily implemented using VBA in Excel. We concur that this is a very viable approach. However, we consider software development beyond the scope of this paper.

For a given risk, the ordering of the decision nodes in a DT corresponds to different temporal deterministic events in the development and fabrication cycle of the module. The probability and cost values are conditional on the specific RRA and the preceding outcomes and account for these dependencies. The use of correlation coefficients would double count correlations. There may also be dependencies between different risks. Whenever possible, the modeling of the dependencies is preferred to the use of empirical correlation coefficients. Correlation coefficients should only be used to account for identified dependencies that cannot be explicitly modeled. The modeling of the dependencies among different risks deserves close scrutiny and further investigation**.**

We note that dynamic risk assessment and management has been proposed to enhance the safety of complex systems [Paté-Cornell and Regan, 1998] and enhance situational awareness [Kobylski et al., 2008]. We think that there are likely to be complementarities between these risk areas and the management of technical project risks that are worth investigating.

The next phase is to further develop the methodology and an integrated tool to aid project managers select a portfolio of risk reduction activities that maximizes the cost and schedule benefits of such activities in accordance with the available resources and his/her risk attitude. The dynamic nature of the approach allows the program manager to monitor and manage the performance of risk reduction activities over time and optimize the allocation of risk management resources over the life-cycle of the project. The challenge is to start implementing these more refined cost models and risk management practices.

7. ACRONYMS

CCDF Complementary Cumulative Distribution Function
CDT Classical Decision Tree
DBM Driver-Based Monte-Carlo
DFA Direct Fractile Assessment
DM Decision-Maker
DOE Design Of Experiments
DT Decision Tree
EV Expected Value
GDT Generalized Decision Tree
MAIMS Money Allocated Is Money Spent
OAB Originally Allocated Budget
PCRA Probabilistic Cost Risk Analysis
PDF Probability Distribution Function
ROI Return on Investment
RRA Risk Response Action
SME Subject Matter Expert
TPRRA Total Project RRA
TRC Technical Risk Cost
TRL Technology Readiness Level
V&V Verification and Validation

ACKNOWLEDGMENTS

The research presented in this paper was supported in part by the Acquisition Chair of the Graduate School of Business & Public Policy at the Naval Postgraduate School. The authors also thank the anonymous referees whose constructive comments and suggestions helped to substantially improve the paper.

REFERENCES

B.S. Blanchard, System engineering management, Wiley, New York, 1998.

S.A. Book, Estimating probable system cost, Crosslink 2(1) (Winter 2000/ 2001), 12–21.

C.L. Brown, Improved methodology for developing cost uncertainty models for naval vessels, MSSE thesis, Naval Postgraduate School, Monterey, CA, 2008.

C. Chapman and S. Ward, Project risk management: processes, techniques and insights, Wiley, Chichester, UK, 1997.

R.T. Clemen and T. Reilly, Making hard decisions with DecisionTools®, Duxbury, Pacific Grove, CA, 2001.

E.H. Conrow, Effective risk management: Some keys to success, 2nd edition, American Institute of Aeronautics and Astronautics, Reston, VA, 2003.

C.M. Creveling, J.L. Slutsky, and D. Antis, Jr., Design for Six Sigma in technology and product development, Pearson Hall PTR, Upper Saddle River, NJ, 2003.

Department of Defense (DoD), Risk management guide for DoD acquisition, 6th edition (Ver. 1.0), 2006, https://acc.dau.mil/CommunityBrowser.aspx?id=108201, accessed April 12, 2008.

R.L. Dillon-Merrill, G.S. Parnell, D.L. Buckshaw, W.R. Hensley Jr., and D.J. Caswell, Avoiding common pitfalls in decision support frameworks for Department of Defense analyses, Mil Oper Res 13(2) (2008), 19–31.

R.L. Dillon and M.E. Paté-Cornell, PRAM: An advanced programmatic risk analysis method, Int J Technol Policy Management 1(1) (2001), 47–65.

A.K. Dixit and B.J. Nalebuff, The art of strategy: A game theorist's guide to success in business & life, Norton, New York, 2008.

A. Engel and M. Barad, A methodology for modeling VVT risks and costs, Syst Eng 6(3) (2003), 135–151.

Federal Transit Administration, FTA risk assessment, Construction Roundtable, Newark, NJ, 2004, http://www.fta.dot.gov/documents/FTA_-_New_Start_Risk_Assessment_Process.ppt, accessed September 10, 2008.

P.R. Garvey, Probability methods for cost uncertainty analysis: A systems engineering perspective, Marcel Dekker, New York, 2000.

D.J. Gogolkiewicz, Leveraging the power of information to achieve portfolio risk management, INCOSE Virginia Chapter, 2007, http://www.incose.org/hra/riskconf2007/Speaker%20Files/David%20Gogolkiewicz/Strategic%20Thought%20Portfolio%20Risk%20Mgt.pdf, accessed March 20, 2009.

C. Gordon, Risk analysis and cost and cost management (RACM): A cost/schedule management approach using statistical cost control (SCC), 1997, http://www.geocities.com/mtarrani/RiskAnalysis_and_CostManagementOverview.pdf, accessed April 10, 2007.

D.R. Graham, Application of continuous cost-risk management at NASA, AIAA 2004-6069, Space 2004 Conference and Exhibit 28–30 September 2004, San Diego, California, http://pdf.aiaa.org/preview/CDReadyMSPACE2004_1014/PV2004_6069.pdf, accessed October 1, 2009.

Y.Y. Haimes, Total risk management, Risk Anal 11(2) (1991), 169–171.

Y.Y. Haimes, Risk modeling, assessment, and management, 2nd edition, Wiley, New York, 2004.

E.M. Hall, Managing risk: Methods for software systems development, Addison-Wesley, Reading, MA, 1998.

J.K. Hollmann, The Monte-Carlo challenge: A better approach, Proc 2007 AACE Int Trans, Risk.03.1, 2007, http://www.c4ce.com/AACE_Risk_Hollmann_paper.pdf, accessed October 1, 2009.

T. Huynh, B. Connett, J. Chiu-Rourman, J. Davis, A. Kessler, J. Oravec, M. Schewfelt, and S. Wark, Architecting a system of systems responding to maritime domain terrorism by using orthogonal array experiment, Nav Eng J 1 (2009), 79–100.

INCOSE, INCOSE systems engineering handbook, version 3.1, International Council on Systems Engineering, Seattle, WA, 2007.

S. Kaplan and B.J. Garrick, On the quantitative definition of risk, Risk Anal 1(1) (1981), 11–27.

G.C. Kobylski, D.M. Buede, J.V. Farr, and D. Peters, The use of dynamic decision networks to increase situational awareness in a networked battle command, Mil Oper Res 13(2) (2008), 47–63.

E. Kujawski, Selection of technical risk responses for efficient contingencies, Syst Eng 5(3) (2002), 194–212.

E. Kujawski, Dynamic cost-contingency management: A method for reducing project costs while increasing the probability of success, Proc Fourth Annu Acquisition Res Symp, Naval Postgraduate School, Monterey, CA, 2007, pp. 546–554.

E. Kujawski and G.A. Miller, Quantitative risk-based analysis for military counterterrorism systems, Syst Eng 10(4) (2007), 273–289.

E. Kujawski, M. Alvaro, and W. Edwards, Incorporating psychological influences in probabilistic cost analysis, Syst Eng 7(3) (2004), 195–216.

H.M. Markowitz, Portfolio selection: Efficient diversification of investments, Blackwell, Cambridge, 1997.

K.T. Marshall and R.M. Oliver, Decision making and forecasting with emphasis on model building and policy analysis, McGraw-Hill, New York, 1995.

D.T. Maxwell, Decision analysis: Find a tool that fits 2008, survey of D.A., OR/MS Today 35 (October 2008), http://www.lionhrtpub.com/orms/orms-10-08/frsurvey.html, accessed December 14, 2008.

F. Melese, Outsourcing for optimal results: Six ways to structure an evaluation of alternatives, Fourth Annu Acquisition Res Symp, NPS-AM-07-004, Monterey, CA, 2007, p. 395.

R.L. Murphy, C.J. Alberts, R.C. Williams, R. P. Higuera, A. J. Dorofee, and J.A. Walker, Continuous risk management guidebook, Software Engineering Institute, Carnegie Mellon University, Pittsburgh, 1996.

National Defense Industrial Association—Program Management Systems Committee (NDIA-PMSC), Integrating risk management with earned value management, Washington, DC, August 2005, http://www.ndia.org/Content/ContentGroups/Divisions1/Procurement/Integrating_RM_with_EVM.pdf, accessed September 10, 2008.

M.E. Paté-Cornell, Uncertainties in risk analysis: six levels of treatment, Reliab Eng Syst Safety 54 (1996), 95–111.

M.E. Paté-Cornell and P.J. Regan, Dynamic risk management systems: Hybrid architecture and offshore platform illustration, Risk Anal 18(4) (1998), 485–496.

I.A. Papazoglou, Mathematical foundations of event trees, Reliab Eng Syst Safety 61 (1998), 169–183.

M.J. Pennock and Y.Y. Haimes, Principles and guidelines for project risk management, Syst Eng 5(2) (2002), 89–108.

A.P. Sage, Systems engineering, Wiley, New York, 1992.

A.P. Sage and J.E. Armstrong Jr., Introduction to systems engineering, Wiley, New York, 2000.

F. Salmon, A formula for disaster, 17 Wired (March 2009), 74–79.

Z. Shapira, Risk taking: A managerial perspective, Russell Sage Foundation, New York, 1995.

P.J. Smith and G.M. Merritt, Proactive risk management: Controlling uncertainty in product development, Productivity Press, New York, 2002.

J.S. Stonebaker and C.W. Kirkwood, Formulating and solving sequential decision analysis models with continuous variables, IEEE Trans Eng Management 44(1) (1997), 43–53.

A. Tversky and D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty, J Risk Uncertainty 5 (1992), 297–323.

U.S. Government Accountability Office, Decisions needed to shape Army's combat systems for the future, GAO-09-288, Washington, DC, March 2009, http:// www.gao.gov/htext/d09288.html, accessed March 20, 2009.

D. Vose, Risk analysis: A quantitative guide, Wiley, Chichester, UK, 2006.

M. Waldof, Risk management: A course for program teams addressing risk management, Student Guide, Lockheed Martin Corporation, Egan, MN, 1998.

Edouard Kujawski is an associate professor in the Systems Engineering Department at the Naval Postgraduate School. His research and teaching interests include the design and analysis of high reliability/availability systems, risk analysis, and decision theory. He received a Ph.D. in Physics from MIT, following which he spent several years in research and teaching physics. He has held lead positions at General Electric, Lockheed-Martin and the Lawrence Berkeley National Laboratory. He has contributed to the design of particle accelerators and detectors, space observatories, commercial communication systems, the Space Station, and nuclear power plants. He was a participant and contributor to the Lockheed Martin LM21 Risk Management Best Practices and the original INCOSE Systems Engineering Handbook. He is a member of the San Francisco Bay Area Chapter of INCOSE, where he serves on the board of directors.

Diana I. Angelis is an associate professor at the Naval Postgraduate School in Monterey, CA, assigned to the Defense Resources Management Institute with a joint appointment to the Department of Systems Engineering. Her current research interests include the implementation of activity-based costing in government organizations, cost estimating and cost risk analysis, the valuation of R&D through options theory, the effect of transaction costs on acquisition estimates, and business reforms in defense management. She has consulted with a number of government organizations on Activity Based Management, including the Defense Contract Management Agency, the Air Force Materiel Command, the Air Force Security Assistance Center, and the Air Force Flight Test Center. She received her Ph.D. in Industrial and Systems Engineering from the University of Florida in 1996. She is a Certified Public Accountant and a Lieutenant Colonel in the US Air Force Reserve.