

IDA

INSTITUTE FOR DEFENSE ANALYSES

Evaluation of the Risk Analysis and Cost Management (RACM) Model

Matthew S. Goldberg, Project Leader
Charles A. Weber

August 1998

Approved for public release;
distribution unlimited.

IDA Paper P-3388

Log: H 98-000161

19981006 076

This work was conducted under contract DASW01 94 C 0054, Task T-G7-1353, for the Office of the Director, Acquisition Program Integration. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 1998 Institute for Defense Analyses, 1801 N. Beauregard Street, Alexandria, Virginia 22311-1772 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (10/88).

PREFACE

The Institute for Defense Analyses (IDA) prepared this paper for the Office of the Director, Acquisition Program Integration, under a task entitled "Program Risk Analysis and Management." This publication partially fulfills the task by evaluating the Risk Analysis and Cost Management (RACM) model developed by Lockheed Martin Corporation.

Jerome Bracken, Philip Lurie, and Louis Simpleman of IDA and Stephen Book of Aerospace Corporation reviewed this work.

CONTENTS

I. Introduction and Summary.....	I-1
A. Genesis of RACM.....	I-1
B. IDA's Evaluation of RACM	I-3
C. Logical Basis, Mathematical Rigor, and Numerical Accuracy	I-4
D. Software Implementation	I-5
E. Novelty of RACM Approach	I-7
F. Applicability to Other Defense Contractors.....	I-8
G. Lessons Learned from RACM	I-10
H. Structure of This Report.....	I-11
II. RACM Model	II-1
A. Motivation for RACM	II-1
B. Structure of RACM.....	II-4
C. Incremental Cost Shifters.....	II-7
1. WBS-Element-Specific Cost Shifters	II-7
2. Global Cost Shifters	II-10
D. Correlations	II-13
E. Software Implementation	II-15
1. Platform.....	II-15
2. Interface.....	II-16
F. Possible Enhancements	II-22
1. Platform.....	II-22
2. Other Enhancements	II-23
G. Eliciting the Input.....	II-25
H. Range of Applicability	II-27
1. Program Size	II-27
2. Program Phase.....	II-28
3. Program Maturity	II-28
I. Documentation	II-29

III. Mathematics of RACM.....	III-1
A. Basic Input	III-1
B. WBS-Element-Specific Effects.....	III-1
1. Cost Improvement.....	III-1
2. Rate Increase	III-4
3. WBS-Element-Specific Schedule Adjustment.....	III-4
C. Global Effects	III-10
1. Global Schedule Adjustment	III-10
2. Major Problems.....	III-10
3. Management Effect (Money Allocated Is Money Spent).....	III-11
D. Order Independence	III-22
E. Treatment of Correlations in RACM	III-22
1. Correlations Versus Linear Modeling Relationships	III-22
2. Global Schedule Effect	III-24
3. Management Effect	III-25
IV. Commercial Off-the-Shelf Alternatives to RACM.....	IV-1
A. Crystal Ball and @RISK for Microsoft Excel.....	IV-1
1. General Description	IV-1
2. Correlations.....	IV-4
3. Incremental Overlay Charts	IV-9
4. Summary	IV-11
B. RISK+	IV-11
1. General Description	IV-11
2. Critique	IV-12
3. Distributional Forms	IV-14
C. @RISK for Microsoft Project	IV-15
1. General Description	IV-15
2. Distributional Forms	IV-15
3. Correlations.....	IV-17
4. Incremental Overlay Charts	IV-17
5. Summary	IV-19

V. RACM Field Test at Boeing D&SG	V-1
A. Design of the Field Test	V-1
B. Software Ease-of-Use and Generation of Input Data.....	V-2
C. Numerical Comparison and Validity of RACM Assumptions	V-4
1. Applicability to Proposal Preparation	V-4
2. Applicability to Program Management	V-5
3. Standing Armies.....	V-5
4. Distributional Forms	V-6
D. Summary	V-7
VI. Conclusions.....	VI-1
Appendix A: Lockheed Martin RACM Study.....	A-1
Appendix B: Proof of Optimality of the RACM Budget Allocation.....	B-1
Appendix C: Properties of the Beta Distribution	C-1
Appendix D: Properties of the Triangular Distribution.....	D-1
Appendix E: Using Correlations with Crystal Ball	E-1
Appendix F: Using Correlations with @RISK.....	F-1
Appendix G: Boeing D&SG Review of RACM and its Application to Cost Management, Project Planning, and Proposals	G-1
Appendix H: Book and Young's Budget Allocation Algorithm	H-1
Appendix I: Using Overlay Charts with Crystal Ball	I-1
Appendix J: Using Overlay Charts with @RISK.....	J-1
Appendix K: Overlay Charts in @RISK for Microsoft Project	K-1
Abbreviations	L-1
Bibliography	M-1

FIGURES

II-1. RACM Baseline Input Sheet	II-4
II-2. Number of WBS Elements Specified in MIL-STD-881B	II-6
II-3. RACM Cost Improvement Input.....	II-8
II-4. RACM Rate Increase Input.....	II-9
II-5. RACM Spreadsheet Navigation: Input Ranges.....	II-17
II-6. RACM Spreadsheet Navigation: Output Ranges	II-18
II-7. RACM Cost Tables (left-hand side)	II-19
II-8. RACM Cost Tables (right-hand side)	II-20
II-9. RACM Incremental Overlay Chart	II-21
III-1. Percentage Cost Improvement as a Function of Lot Size and Learning Slope	III-4
III-2. Duration Distributions Along WBS-Element-Specific Parallel Paths	III-5
III-3. Example of Percentage Growth Assumption	III-9
III-4. Density Functions Before and After Transformation.....	III-12
III-5. Cumulative Distribution Functions Before and After Transformation	III-13
III-6. Comparison Between RACM and FRISK Budget Allocation Algorithms.....	III-21
III-7. Correlation Between Two Cost Elements	III-23
III-8. Scatterplot Between Two Transformed Cost Elements	III-27
IV-1. Comparison of Spearman and Pearson Correlations.....	IV-8
IV-2. Crystal Ball Incremental Overlay Chart.....	IV-10
IV-3. @RISK for Microsoft Excel Incremental Overlay Chart.....	IV-10
IV-4. @RISK for Microsoft Project Incremental Overlay Chart.....	IV-18

A-1.	Comparison Between Arithmetic Summing and Statistical Summing	A-8
A-2.	Effects of Statistical Apportionment and Allocation	A-10
A-3.	Accuracy of Monte Carlo Versus Analytical Model	A-15
D-1.	Right-Skewed Triangular Distribution	D-2
D-2.	Iterative Solution for λ Using @RISK	D-10
D-3.	Crystal Ball: Define Assumptions	D-12
D-4.	Crystal Ball: Specify Triangular Distributional Family	D-13
D-5.	Crystal Ball: View Default Parameter Screen	D-13
D-6.	Crystal Ball: Specify Percentile Option	D-14
D-7.	Crystal Ball: View Percentile Parameter Screen	D-14
E-1.	Crystal Ball: Opening Screen	E-3
E-2.	Crystal Ball: Define Assumptions	E-3
E-3.	Crystal Ball: Specify Distributional Family	E-4
E-4.	Crystal Ball: Specify Distributional Parameters	E-4
E-5.	Crystal Ball: Input Correlations	E-5
E-6.	Crystal Ball: Define Forecasts	E-5
E-7.	Crystal Ball: Run Simulation	E-6
E-8.	Crystal Ball: Receive Warning Message, Inconsistent Correlation Matrix	E-6
E-9.	Crystal Ball: Help Screen on Inconsistent Correlation Matrix	E-7
E-10.	Crystal Ball: View Simulation Results	E-7
E-11.	Crystal Ball: Recover the Correlation Matrix	E-8
F-1.	@RISK: Opening Screen	F-3
F-2.	@RISK: Enter Input Variables	F-3
F-3.	@RISK: Enter Output Variables	F-4
F-4.	@RISK: Input Correlations	F-4
F-5.	@RISK: Run Simulation	F-5
F-6.	@RISK: Receive Warning Message, Inconsistent Correlation Matrix	F-5
F-7.	@RISK: View Simulation Results	F-6
F-8.	@RISK: Attempt to Recover the Correlation Matrix	F-6
F-9.	@RISK: Recover the Correlation Matrix	F-7

G-1.	Crystal Ball Graph of D&SG Data.....	G-8
G-2.	@RISK Graph of D&SG Data.....	G-8
I-1.	Crystal Ball: Define Multiple Output Variables	I-2
I-2.	Crystal Ball: Run Simulation	I-2
I-3.	Crystal Ball: Open Overlay Chart	I-3
I-4.	Crystal Ball: Add Forecasts	I-3
I-5.	Crystal Ball: Select the Multiple Output Variables.....	I-4
I-6.	Crystal Ball: View Overlay Chart	I-4
J-1.	@RISK: Define Multiple Output Variables.....	J-2
J-2.	@RISK: Run Simulation	J-2
J-3.	@RISK: Graph the First Output Variable.....	J-3
J-4.	@RISK: Format the Output Graph	J-3
J-5.	@RISK: Change Output Graph Type	J-4
J-6.	@RISK: Change Output Graph Patterns.....	J-4
J-7.	@RISK: Select the Multiple Output Variables	J-5
J-8.	@RISK: View Overlay Chart.....	J-5
K-1.	@RISK for Microsoft Project: Opening Screen	K-3
K-2.	@RISK for Microsoft Project: Specify a Scenario Variable.....	K-3
K-3.	@RISK for Microsoft Project: Specify a Random Variable	K-4
K-4.	@RISK for Microsoft Project: Enter Output Variable.....	K-4
K-5.	@RISK for Microsoft Project: Set the Number of Simulations	K-5
K-6.	@RISK for Microsoft Project: View Simulation Results	K-5
K-7.	@RISK for Microsoft Project: Graph the First Output Variable	K-6
K-8.	@RISK for Microsoft Project: Format the Output Graph.....	K-6
K-9.	@RISK for Microsoft Project: Change Output Graph Type.....	K-7
K-10.	@RISK for Microsoft Project: Change Output Graph Pattern	K-7
K-11.	@RISK for Microsoft Project: Select Multiple Output Variables	K-8
K-12.	@RISK for Microsoft Project: View Overlay Chart	K-8

TABLES

III-1. Shift Factors for WBS-Element-Specific Schedule Adjustment	III-7
III-2. Example of RACM Percentile Calculation	III-17
III-3. Triangular Distributions for Budget Allocation Example	III-19
III-4. Correlation Matrix for Budget Allocation Example	III-20
D-1. Cases for Estimation of Triangular Distribution	D-2
D-2. Parameters of Triangular Distribution	D-5
G-1. Summary of RACM and EELV Estimates	G-2
G-2. Example of RACM Input Data for a WBS Element	G-3
G-3. Decomposition of RACM and EELV Estimates	G-6
G-4. Sequential Distribution Functions of RACM Risk Estimates	G-7
G-5. Risk Ratings from EELV Program Office	G-11

I. INTRODUCTION AND SUMMARY

A. GENESIS OF RACM

Although risk is a ubiquitous feature of defense programs, neither the Department of Defense (DoD) nor its contractors have a standard set of tools to forecast and manage risk. Consequently, proposals submitted by various contractors may not be commensurable because they are implicitly bidding different levels of risk. Indeed, each contractor's proposal may not even be internally consistent if the contractor uses incorrect mathematical algorithms to propagate risk from the individual cost elements up through the entire program.

The improper treatment of risk may affect not only our ability to forecast program performance, but also the actual program performance itself. Suppose, for example, that the winning bidder allocates the entire program budget to the individual cost account managers. If the cost account managers feel compelled to spend their entire individual budgets, then the possibility of a cost underrun is ruled out. Alternatively, a more careful risk calculation may enable the contractor to determine a risk reserve level that is sheltered from the cost account managers unless absolutely needed. In this fashion, it may be possible to underrun the overall program budget, generating savings that can be shared by both the contractor and DoD.

These and related concerns motivated a small team at Lockheed Martin Missiles and Space Company (located in Sunnyvale, California) to develop the Risk Analysis and Cost Management (RACM) concept and associated computer software. Their concerns are expressed first-hand in the following quotation:

The standard methods of developing cost estimates and managing budgets were first questioned in response to the growing concern that costs of many programs had increased significantly over the past 20 years. Most of the increases could not be attributed to inflation or other known factors which normally increase costs. The first explanation was that each element or cost account in the WBS [work breakdown structure] was being "padded." If this were true (and sometimes it is), the "padding" should have eliminated the many overruns that were occurring. They didn't.

The next assumption was that the overruns were caused by those responsible for forecasting costs, i.e., they were not recognizing, in advance, all of the costs which were to be encountered. This may account for the overruns but would not account for the significantly increased costs. The third assumption was an excessive number of changes being incorporated which were beyond the original scope of the contract and outside the original specifications. This establishes a false cost baseline for future programs.

Because these explanations were not adequate in explaining the increase in costs, a study was initiated to develop a model which considered each step in the forecasting and managing of a program, especially the human element. The most important finding was the effect that "hidden" incentives have on the final cost of a program. These incentives are not the usual cost incentives or performance incentives written into a contract or made part of the standard operating procedure. These are the "hidden" incentives that govern the conduct of the personnel performing the tasks.¹

RACM was developed over several years by a small team numbering roughly three to four individuals funded through the Lockheed Martin overhead accounts. They aggressively briefed their evolving concept widely throughout DoD, including the following organizations:

- Office of the Secretary of Defense (OSD), Acquisition and Technology;
- OSD Cost Analysis Improvement Group (CAIG);
- Ballistic Missile Defense Organization (BMDO);
- Defense Advanced Research Projects Agency (DARPA);
- Air Force Materiel Command (AFMC);
- Naval Air Systems Command (NAVAIR); and
- Army, Navy, and Air Force centers for cost analysis.

In particular, RACM came to the attention of the Office of the Under Secretary of Defense for Acquisition and Technology, Directorate of Acquisition Program Integration (API). That office saw considerable potential in the RACM concept, and tasked the Institute for Defense Analyses (IDA) to conduct a thorough evaluation.

¹ Creaghe Gordon, "Risk Analysis and Cost Management (RACM): A Cost/Schedule Management Approach Using Statistical Cost Control (SCC)," Los Gatos, California, June 1997, pp. 10-11. Related information is contained at the following web site: <http://pw2.netcom.com/~chgordon/risk-racm.html> [accessed June 1, 1998].

B. IDA'S EVALUATION OF RACM

IDA's evaluation began in December 1994. IDA's specific tasks were delineated as follows:

- evaluate the logical basis and mathematical rigor of the RACM concept;
- evaluate the numerical accuracy of the RACM concept and associated software implementation;
- conduct a thorough search of the open literature to determine the extent to which RACM differs from and, perhaps, improves upon other tools generally available in the literature or widely used in industry; and
- design and implement an experiment on a particular program in the field to determine the ease with which RACM could be adapted for wider use by DoD program offices or other commercial contractors.

The next four sections of this chapter summarize IDA's findings on each of these tasks.

The tasking also called for IDA to work with Lockheed Martin Missiles and Space Company, to the extent necessary and appropriate, to ensure that IDA had a complete understanding of RACM before issuing our final opinion. The interactions between IDA and Lockheed Martin took the form of a subcontract that called for Lockheed Martin to perform the following specific tasks:

- describe the fundamental assumptions and concerns underlying the RACM model;
- describe the inputs required to operate the RACM model;
- describe the RACM modeling methodology;
- describe how RACM calculates the reserves for each element in the work breakdown structure (WBS) during each phase of the program;
- provide advice on possible application of RACM to current DoD initiatives and policies (e.g., Design-to-Cost, Cost/Schedule Control Systems); and
- maintain and support a working version of the RACM computer program for use on a Microsoft Windows-compatible personal computer.²

To fulfill the subcontract, Lockheed Martin prepared a deliverable that is reproduced as Appendix A of this report. The major sections of Appendix A generally correspond to the six specific tasks listed in the preceding paragraph.

² Microsoft, Windows, Excel, Access, and Project are registered trademarks of Microsoft Corporation.

C. LOGICAL BASIS, MATHEMATICAL RIGOR, AND NUMERICAL ACCURACY

RACM is a spreadsheet-based tool for estimation and management of risk associated with defense contracts. RACM assumes that costs can be arrayed into a hierarchical WBS, so program cost can be summed from the subelements to the elements and from the elements to the entire program.³ RACM requires that the user supply the baseline labor and nonlabor cost distributions for each WBS element, which RACM then combines into a baseline distribution of total program cost. RACM allows for risk factors that shift the cost distribution of each individual WBS element, in turn shifting the distribution of total program cost. RACM also allows for global risk factors that shift the distribution of total program cost directly, rather than operating on the individual WBS elements. The user must supply a range of possible values for the cost implications resulting from each risk factor. Finally, RACM displays the effects of each risk factor sequentially, building up from the baseline program cost to the final program cost that embodies all of the WBS-element specific and global risk factors.

RACM's estimates are based upon some strong assumptions regarding statistical cost distributions as well as contractor behavior:

- the baseline labor and nonlabor costs in each WBS element are normally distributed;
- Money Allocated is Money Spent (MAIMS), so WBS-element managers never underspend any budget allocated to them by central management; and
- costs in a particular activity continue to accrue at a constant rate until *all* parallel activities are completed ("standing armies").

The universality of these assumptions is addressed in Section F of this chapter. Temporarily accepting these assumptions, we determined that most of RACM's mathematical calculations were performed correctly. We note exceptions in the following paragraphs.

RACM contains a formula to compound the effects of multiple risk factors on the standard deviation of total cost. Although an exact formula is known, early versions of

³ The ability to sum costs may appear so obvious as to be universal. To see why this is not the case, suppose that alternative development efforts proceed in parallel, with the first successful arrival being adopted. Total development cost in this case equals the full cost of the successful alternative (whose identity is not known in advance), plus the truncated costs of all the unsuccessful alternatives. Thus, total cost becomes a complex, probabilistically weighted average of truncated random variables.

RACM used an approximate formula instead. At our suggestion, the exact formula was adopted in later versions of RACM.⁴

RACM does not fully account for correlations among cost elements. Positive correlations imply that the various cost elements tend to swing in the same direction: above-average costs in element *A* are associated with above-average costs in element *B*, and below-average costs in element *A* are associated with below-average costs in element *B*. Positive correlations tend to increase the standard deviation of total cost. When the standard deviation is underestimated, so are the high-order percentiles (e.g., the budget level necessary to cover the costs incurred in 90 out of 100 possible replications of the program). RACM only *partially* captures correlations via the global effects, which uniformly scale the cost of each WBS element by a common factor. Early versions of RACM contained a cell range in which the user could explicitly enter the entire correlation matrix. We suggest that the correlation matrix be reinstated.

RACM uses numerical approximation as opposed to Monte Carlo simulation. RACM keeps track of the mean and standard deviation of total cost as each sequential risk factor is applied. However, even given the mean and standard deviation, one must still posit a distributional form in order to compute percentiles for total cost. Except for the MAIMS effect, RACM assumes that total cost is normally distributed at each stage. The MAIMS effect is mathematically more complex, and RACM uses a beta distribution to represent final cost after application of that effect. Although the use of a beta distribution for final cost is not exact, we believe that it is a tolerable approximation.

D. SOFTWARE IMPLEMENTATION

RACM is currently implemented as a Microsoft Excel spreadsheet, programmed to run via formula-linked spreadsheet cells. Although this implementation may have been suitable for a small team of analysts doing their own work, it is not suitable for wider use throughout the defense industry. Migration to a higher-level programming environment, such as Visual Basic or Microsoft Access, would facilitate the following much-needed enhancements:

- user-friendly interface;
- on-line help;

⁴ Our formula is "exact" in the sense that it does not introduce any additional error when compounding the effects of multiple risk factors. Of course, the results of an "exact" formula are no more accurate than the input that the user supplies.

- choice among distributional forms;
- flexible number of WBS elements;
- blank templates for additional, user-defined risk factors;
- greater ease in comprehending, modifying, and customizing code;
- improved database management; and
- ability to export results to project management or earned value software.

RACM documentation is quite poor. In fact, the *only* existing documentation consists of Appendix A of this report, which Lockheed Martin wrote at our behest. Recall that our charter was to evaluate RACM, not to write either a comprehensive user's manual or a tutorial on data elicitation. However, those two elements are essential before RACM (or any other model, for that matter) can be considered for wider adoption throughout the defense industry.

One could conceive of replacing RACM's numerical approximation approach with Monte Carlo simulation. From a purely mathematical point of view, Monte Carlo simulation is slower but more accurate, because it does not presume that total program cost follows any particular distributional form such as normal or beta. In our opinion, these mathematical considerations do not strongly favor either approach over the other. Although Monte Carlo simulation definitely runs more slowly, it does not run so slowly in this application as to discourage the user from conducting extensive replication or sensitivity analysis. Conversely, the beta distribution for total program cost is sufficiently flexible that we see little gain from the agnosticism of Monte Carlo simulation with respect to distributional forms.

Monte Carlo simulation does have the advantage that commercial off-the-shelf (COTS) software is available to relieve some of the programming burden. Packages such as Crystal Ball and @RISK enhance Microsoft Excel by performing the simulations using built-in mathematical functions that are transparent to the user. They offer an explicit matrix in which the user may enter the correlations among cost elements. They also have the capability to produce incremental overlay charts, displaying the sequential effects of each risk factor on total program cost. In addition, these packages provide the following desirable features:

- user-friendly interface,
- on-line help,

- extensive choice among distributional forms, and
- automatic generation of charts and reports.

These features of the COTS software packages are reviewed in Chapter IV. The COTS software would still operate in a spreadsheet environment. Thus, the user would still have to manually copy and paste many rows of formulas in order to increase the number of WBS elements. Similarly, the user would still have to manually create the templates for additional risk factors. Most critically, the use of COTS software would do nothing to ease the process of eliciting the input from the technical experts (see Section F). Hence, we see little to be gained from recasting RACM as a Monte Carlo simulation.

E. NOVELTY OF RACM APPROACH

The RACM developers drew an important distinction between “arithmetic summing” and “statistical summing.” Using the former approach, one naively adds the percentiles of the various WBS-element distributions to estimate the corresponding (e.g., 90th) percentile of total cost. It is easy to demonstrate the folly in arithmetic summing; we provide a simple counterexample at the start of Chapter II. Instead, RACM uses statistical summing, which estimates the distribution of total cost and then explicitly computes the desired percentile.

Although arithmetic summing may still be prevalent in industry, its mathematical incorrectness has been known in the open literature for some time. The most recent example is a briefing by Stephen A. Book that has been widely circulated throughout the cost-analysis community.⁵ Still, there is some merit in any tool that encourages analysts to switch from an incorrect mathematical approach to a correct one.

RACM is almost unique in containing a module for managing risk reserves. The RACM developers conjectured, and we have rigorously demonstrated (see Appendix B), that an equi-percentile budget allocation is optimal when the MAIMS effect is operative. That is, expected program cost is minimized when central management allocates to each WBS-element manager the same percentile of his or her respective cost distribution. RACM’s reserve-management algorithm computes the common percentile for each WBS element that is consistent with a target percentile for total program cost. However, the

⁵ Stephen A. Book, “Do Not Sum ‘Most-Likely’ Cost Estimates,” The Aerospace Corporation, Los Angeles, California, May 1995.

theoretical justification for this algorithm vanishes if the MAIMS effect is turned off. Because the MAIMS effect may not be universal among all defense contractors (see Section F), the equi-percentile budget allocation is, at best, an appealing heuristic, not necessarily an optimal solution.

An alternative reserve-management algorithm by Stephen A. Book and Philip H. Young has also been widely circulated.⁶ Their algorithm does not presume to optimize any particular objective. However, it too is heuristically appealing. Moreover, we present a test problem in Chapter III for which the RACM and Book/Young algorithms yield essentially equivalent solutions. We would encourage contractors to use either algorithm to help manage risk reserves, although the results are not necessarily optimal and might be overridden by program-specific considerations.

Finally, we note that there is again little advantage to recasting RACM as a Monte Carlo simulation for the purpose of reserve management. Estimating the various percentiles using Monte Carlo simulation would be slightly slower and perhaps slightly more accurate, but we find neither difference particularly compelling. In any case, the user should include the correlations among cost elements in order to avoid understating the standard deviation of total cost and the corresponding percentiles.

F. APPLICABILITY TO OTHER DEFENSE CONTRACTORS

We conducted a field test at Boeing Defense and Space Group (D&SG—located in Seattle, Washington) to determine whether RACM is portable to another defense contractor's site. Boeing had just completed a major proposal for the Evolved Expendable Launch Vehicle (EELV). We briefed them on the RACM concept and software implementation, and we convinced them to retrofit RACM on the EELV proposal. We established a dialogue with the Boeing D&SG Manager of Statistical Analysis and Simulations, Estimating. That dialogue is reproduced as Appendix G of this report and summarized in Chapter V.

In their Request for Proposals (RFP), the Air Force EELV Program Office provided their own taxonomy of risk factors. Although Boeing had never before encountered that particular taxonomy, they managed to structure their proposal around it. By contrast, Boeing found it extremely difficult to map their cost accounts into the

⁶ Stephen A. Book, "Recent Developments in Cost Risk," The Aerospace Corporation, Los Angeles, California, May 1992; Philip H. Young, "FRISK: Formal Risk Assessment of System Cost Estimates," The Aerospace Corporation, Los Angeles, California, February 1992.

categories demanded by RACM. Boeing set two of the RACM risk factors to zero because they were unable to back them out of the baseline cost estimates. For two of the other risk factors, it appears that they did not understand the RACM definitions and simply made up numbers. The lack of documentation on RACM was a common complaint during our dialogue with Boeing.

Boeing D&SG is not the only organization that found it difficult to provide input for RACM. At one point, we contemplated a field test of RACM at a government depot. The RACM developers informed us that the field test could not proceed without them because only they, not the IDA team, could properly elicit the input. This exchange took place after the IDA team had already been involved in the RACM evaluation for a full year. If true, their assertion does not bode well for the possibility of widespread dissemination of RACM.

Boeing D&SG disputed the normality assumption for baseline costs: "All of our cost modeling experience strongly suggests that cost [distributions] are lognormal."⁷ Boeing also disputed two of RACM's most basic assumptions regarding contractor behavior. First, they denied the existence of the MAIMS effect in their plant: "The cost management mechanisms in place at Boeing largely preclude the 'money allocated [is] money spent' phenomenon."⁸ To be fair, we should note that RACM also provides visibility into the distribution of total program cost with the MAIMS effect set to zero. Thus, RACM could still be used to model the effects of the other risk factors (e.g., cost improvement, rate increase, and so on), provided that the input data could be arrayed according to these categories. However, little utility would remain from using RACM to manage risk reserves if the MAIMS effect were turned off. RACM manages reserves using an equi-percentile budget allocation, which is not necessarily optimal unless the MAIMS effect is thought to be operative. And, as we have mentioned, alternative algorithms are available for managing reserves without the RACM superstructure.

In addition, Boeing D&SG denied the "standing army" problem in their plant:

The staffing of a program (contract) is seldom level-loaded, perhaps never, if study contracts are excluded and depending on the definition of "level."...In addition, the target value is just that, a target, it is generally not manned to....This issue could certainly have an effect, if it were not well managed. Much of the cost management effort on a program is

⁷ Appendix G, p. G-23.

⁸ *Ibid.*, p. G-24.

expended to manage out the effect that poor budgeting might have....Budget pressure and the ability to move resources to other parts of a project or to other projects or sites allow the organization to operate without "standing-armies."⁹

The following quotation cogently summarizes Boeing's experience in attempting to use RACM:

RACM is an Excel-based tool with a specific view of the causes and quantified effects of cost, schedule, and technical "risk." As such it, in all likelihood, fits into the cost management and estimating infrastructures from which it was created. RACM does not mesh well with our cost estimating and cost management data sources, estimating procedures, and management requirements nor the management philosophy that I [Manager of Statistical Analysis and Simulations, Estimating] am aware of at the Boeing Company.¹⁰

G. LESSONS LEARNED FROM RACM

As we stated in the very first line of this report, risk is a ubiquitous feature of defense programs. Any tool that encourages a more systematic treatment of risk may lead to more rational defense procurement. However, neither RACM nor any other particular tool can be viewed as a "silver bullet" to remove all risk or prevent all cost overruns on defense programs.

DoD would be well-advised to demand risk estimates, in preference to single point estimates, on all major procurement programs. In addition, RACM's sequential view of the various risk factors is extremely useful to the contractor and the government alike. DoD might consider requiring this sequential view, as was done in one instance by the Air Force EELV Program Office. DoD might impose other general requirements, such as a set of percentile levels at which cost must be reported. However, we discourage DoD from requiring that contractors use RACM or any other particular tool when preparing their proposals. We recommend instead that contractors be allowed to choose the tools that best fit their own accounting systems, estimating methods, and management philosophies. The discussion of risk estimation and management in this report, along with our review of COTS software, should help guide contractors in making these choices.

⁹ The full quotation is found in Appendix G, pp. G-22 to G-24.

¹⁰ *Ibid.*, p. G-9.

H. STRUCTURE OF THIS REPORT

The remainder of this report is organized roughly around IDA's four specific tasks listed previously in Section B. Chapter II is a general description of RACM, covering its motivation and structure, the "look and feel" of the software implementation, the sources of the input, the range of applicability, and the state of the documentation. Later chapters revisit most of these topics in considerably greater detail. For example, Chapter III provides a comprehensive evaluation of RACM's mathematical assumptions, computational algorithms, and numerical accuracy. Chapter IV explores the possibility of replicating RACM's best features in a commercial software environment. Along the way, we comment in Chapter III on the novelty of the RACM concept vis-à-vis other published techniques for risk estimation and management. We also comment in Chapter IV on the novelty of the RACM software implementation vis-à-vis COTS alternatives. Chapter V summarizes the field test that we conducted at Boeing D&SG; a complete transcription of the dialogue between IDA and Boeing D&SG is contained in Appendix G. Finally, Chapter VI presents our conclusions.

II. RACM MODEL

A. MOTIVATION FOR RACM

RACM was developed as an internal tool by a team at Lockheed Martin Missiles and Space Company, Sunnyvale, California. The Lockheed Martin team developed RACM in response to certain problems perceived in DoD hardware procurement. One problem was the method by which safety levels are calculated when preparing program proposals. To consider the simplest possible example, suppose that a program contains two uncorrelated cost elements. Cost of element *A* is normally distributed with mean \$10 and standard deviation \$3. Cost of element *B* is normally distributed with mean \$20 and standard deviation \$4. It is well known that 97.7% of the probability under a normal distribution falls to the left of the mean plus twice the standard deviation, $\mu + 2\sigma$. Thus a 97.7% safety level for the cost of element *A* is \$16, and a corresponding safety level for element *B* is \$28. A naive procedure would simply add the two safety levels, yielding an estimate of \$44 for total program cost.

Given the above assumptions, total program cost is normally distributed with mean equal to the sum of the cost-element means, \$30. The standard deviation is given by the root-sum-of-squares (RSS) formula:

$$\sqrt{3^2 + 4^2} = \sqrt{25} = 5.$$

A 97.7% safety level for total program cost is $\$30 + (2 \times \$5) = \$40$. Thus, simple addition of the two safety levels leads to an overstatement of 10% (i.e., \$44 versus \$40). Put differently, because the estimate \$44 lies 2.8 standard deviations above the mean, simple addition leads to a *true* safety level of 99.7%.

The RACM developers coined the term “arithmetic summing” for simple addition of safety levels, and the term “statistical summing” for calculation of the true safety level of total program cost. In their view, arithmetic summing is the more common practice in industry, resulting in overstated cost estimates. The discrepancy is exacerbated when the number of cost elements increases. For example, suppose that there are 20 uncorrelated cost elements, each funded at a 95% safety level. For any one cost element, a 95% safety

level implies that the budget is adequate to cover realized cost in 19 of 20 cases. Then among the 20 cost elements, we expect on average to see 19 within budget and only 1 overrun. Unless the magnitude of the single overrun is overwhelming, the 19 cost elements within budget should easily suffice to cover the 1 overrun, with money to spare. Thus, the true safety level for the entire program is much larger than the 95% safety level for each individual cost element. Conversely, a 95% safety level for the entire program can be achieved with individual safety levels far below 95%.

The Lockheed Martin team perceived a second, related problem in DoD procurement. They asserted that the winning bidder would always spend the entire, overstated budget, because there is never any incentive to underrun a budget. Thus, the overstated budget estimates become self fulfilling. Moreover, this phenomenon contaminates the databases used for cost estimation by both industry and government. The historical data reflect budget levels and overruns above those budget levels, but not the underruns that might have occurred had behavioral incentives allowed them. The overstated historical data feed into even greater overstatement of cost during subsequent rounds of contracting:

Any historically based model should be used with caution....Two of the most consequential factors [that incorrectly influence the data] are:

- Improper allocation of resources (money allocated is money spent),
- Acceptance by the contractor of changes beyond the scope of the contract without adjustments.

CERs [cost-estimating relationships] are based on historical data and the relationships that were apparent on those programs. These relationships are usually adjusted for each program that they represent. While historical data [are] necessary in trying to predict the future, any use of the data must be dissected, examined, and only then can it be reassembled for use in a model....The impact of accepting changes beyond the scope of a contract creates a false cost baseline. Any future analysis for use in cost forecasting would develop CERs which would assume a product as originally specified without the benefit of knowing what additional tasks are represented by the costs.¹

RACM is a tool intended to overcome these problems. First, it contains a percentile calculator to determine safety levels for the individual cost elements that are

¹ Creaghe Gordon, "Risk Analysis and Cost Management (RACM): A Cost/Schedule Management Approach Using Statistical Cost Control (SCC)," p. 8.

consistent with a target safety level for the entire program. RACM's percentile calculator applies statistical summing rather than arithmetic summing, so the problem of overstated individual safety levels is effectively eliminated. RACM also offers a sequential view of the risk factors that shift the distribution of total program cost.

Second, RACM contains a module intended to assist contractors in allocating budgets to WBS-element managers. Recognizing that budgets are seldom underrun, the idea is to allocate only a portion of the total program budget, and hold the remainder in a management reserve fund. The reserves could be applied to cover overruns in particular WBS elements, but would not be automatically depleted by WBS-element managers in an attempt to spend their individual budgets.

One could also conceive of RACM being used in government program offices during the source selection process. There are several possible scenarios along these lines. For example, a contractor might prepare a bid using RACM, and submit both the RACM input and output to the government. The government evaluators would then have access to the input data underlying the bid. The evaluators could verify that the output was consistent with the input. They could also perform sensitivity analysis on the bid price by perturbing the input. Note, too, that because RACM is an "open" model (i.e., all of the formulas are completely visible to the user), the contractor could easily manipulate the input to achieve a desired set of outputs.

Alternatively, one could reengineer RACM so that the formulas were hidden from the user (i.e., convert RACM into a "black box"). The contractor would submit the RACM input to the government; the government, in turn, would run RACM to determine the corresponding output. However, if the contractor had access to the model, even with the formulas hidden, it could still (with some difficulty) reverse engineer the model to determine a set of input consistent with the desired output. At least one contractor has informally confessed to us having reverse engineered the government-supplied Parametric Review of Information for Costing and Evaluation of Hardware (PRICE-H) and Parametric Review of Information for Costing and Evaluation of Hardware-Life Cycle (PRICE-HL) models. Indeed, they whimsically refer to the reverse-engineered models as ECIRP (i.e., PRICE spelled backwards).²

² PRICE-H and PRICE-HL are described at the Air Force Financial Management web site: <http://www.saffm.hq.af.mil/SAFFM/afcaa/cross/cross.html#PRICEH>. Additional information is available from the vendor: PRICE Systems, 700 East Gate Drive, Suite 200, Mt. Laurel, New Jersey, 08054. Their web site is: <http://www.buyfs.com/index.htm>.

B. STRUCTURE OF RACM

RACM requires that the user provide an extensive set of input. The most basic input is a WBS-element structure. This structure must be both exclusive and exhaustive; every dollar of cost must belong to exactly one WBS element. A WBS-element structure is a standard component of cost accounting on government contracts, so this requirement is not new. A WBS-element structure is hierarchical, so a particular cost element may be composed of lower-level subelements that sum to the higher-level element. For example, Figure II-1 is one of the RACM input sheets. There are five subelements under the element "Missile": "Propulsion," "Payload," "Reentry," "G&C" [Guidance and Control], and "IA&T" [Integration, Assembly and Testing]. Because the subelements are listed explicitly, their costs would be modeled but the higher-level cost ("Missile") would be inferred as their sum. Conversely, the cost element "S/W Eng'g" [Software Engineering] is listed without any of its subelements. Thus, Software Engineering is modeled, but its subelements are not modeled and remain below the resolution of the model. The decision about which elements or subelements to model is at the discretion of the user. However, we note that two or more distinct levels of indenture may be modeled, as long as the user obeys the cardinal rule: *model every dollar of cost exactly once*.

MAJOR PROGRAM WBS ELEMENTS	LABOR							NON-LABOR			
	EQUIV. PEOPLE (E. P.) (AVG MAN LVL)		AVERAGE LABOR RATE	COST ESTIMATE K\$ (CALCULATED)		PROBABILITY OF SUCCESS (%: 0 < P(s) < 100)		COST ESTIMATES K\$		PROBABILITY OF SUCCESS (%: 0 < P(s) < 100)	
	LOW	HIGH	\$/HR	LOW	HIGH	LOW	HIGH	LOW	HIGH	LOW	HIGH
Missile											
Propulsion											
Payload											
Reentry											
G&C											
IA&T											
S/W Eng'g											
Program Mngmnt											
Systems Eng											
ST&E											
Training											
Data											
Support Equip											
Initial Spares											

Figure II-1. RACM Baseline Input Sheet

The RACM developers offer the following guidance on choosing the level of indenture, with which we generally concur:

Any program requiring the use of "level two" or "level three" elements, as defined in MIL-STD-881, would be considered sufficiently large as to have the potential for large errors if arithmetic summing is used instead of statistical summing. The most easily understood analysis technique would be one that uses level two elements (usually 11) [*sic*] with selected level three elements. The total number of elements used for analysis should be minimized and, as a general rule, should not exceed number of level three elements. However, this decision should be dollar-driven. The combination of elements in the WBS should maintain a balance of costs among elements. For example, no single element in a 40-element WBS should represent 20% of the total cost. The ideal would be for each element in a 40-element analysis to represent approximately 2% to 3% of the total cost.³

Notwithstanding the reference to 40 WBS elements, the current implementation of RACM is limited to modeling at most 30 WBS elements at whatever level of indenture. This limitation is solely a consequence of the current software environment, not an intrinsic feature of the RACM concept. A Military Standard, MIL-STD-881B, governs the number of WBS elements on various types of military hardware systems.⁴ As we indicate in Figure II-2, the 30-element limit is more than adequate at WBS level 2. However, nearly 60 elements may be required at WBS level 3.

The column headings in Figure II-1 reveal the level of detail required to model any single WBS element. The costs in any WBS element are first partitioned into labor and nonlabor components. Regarding labor costs, the user must provide two percentile points of the cost distribution. Specifically, the user must provide low and high estimates of full-time equivalent (FTE) people per month, as well as the percentile level ("probability of success") associated with each. Note that the user is free to select the percentile levels, which need not be symmetric (e.g., if the user supplies the 10th percentile as the low estimate, there is no requirement that the high estimate correspond to the 90th percentile).

³ See Lockheed Martin's documentation on RACM, reproduced as Appendix A of this report. Specifically, the quotation is excerpted from p. A-3.

⁴ "Military Standard: Work Breakdown Structures for Defense Materiel Items," U.S. Department of Defense, MIL-STD-881B, 25 March 1993.




Figure II-2. Number of WBS Elements Specified in MIL-STD-881B

The input in Figure II-1 reflects the *baseline costs*, but RACM also models *cost shifters* that incrementally perturb either the individual WBS-element cost distributions or the distribution of total program cost. As we see shortly, one of the cost shifters is uncertainty in the labor rate. For the purposes of Figure II-1, however, the “average labor rate” is interpreted as a fixed, baseline value.⁵ RACM’s strategy is to embed uncertainty in FTEs within the baseline distribution, but to account for uncertainty in the labor rate among the cost shifters. Although this is clearly not the only possible strategy, it appears to be innocuous.

Nonlabor costs include materials and subcontracts, for which a partitioning into FTEs and labor rates is not meaningful. Thus, the user directly provides low and high estimates of nonlabor costs for each WBS element. The nonlabor costs, too, are perturbed by the incremental cost shifters elsewhere in the model. Thus, the entries in Figure II-1 are again interpreted as baseline values. Note that both the average labor rate and the nonlabor costs are fully burdened with overhead.

Although the WBS elements and subelements modeled by RACM are mutually exclusive and exhaustive, one could still entertain the possibility that they are correlated.

⁵ RACM reconciles an hourly labor rate with monthly FTEs using a notional 151-hour work month. This factor is hard-wired into the model, but could easily be changed.

We briefly return to our simple example of a program containing two cost elements. Total program cost is, by definition, the sum of the costs of element *A* and element *B*. While there is no overlap between their costs, they could be correlated if they are driven by a common set of factors. For example, fabrication costs for element *A* and element *B* might both be sensitive to the cost of electricity, variations in which would cause both cost elements to swing in the same direction. Early versions of RACM included a cell range in which the user could enter the correlations, if desired. Correlations do not affect the mean of total program cost, but do affect the standard deviation. The same set of correlations would be used when calculating the standard deviation during the baseline analysis, and again after application of each incremental cost shifter. We discuss correlations in more detail later in this chapter.

C. INCREMENTAL COST SHIFTERS

RACM allows for three factors that shift the cost distribution element by element, possibly by different amounts for each element:

- Cost improvement,
- Rate increase, and
- WBS-element-specific schedule adjustment.

RACM also allows for three factors that directly shift the distribution of total program cost, rather than operating on the individual WBS elements:

- Global schedule adjustment,
- Major problems, and
- Management effect.

We briefly discuss the various cost shifters in this section, but defer a detailed mathematical analysis to Chapter III.

1. WBS-Element-Specific Cost Shifters

a. Cost Improvement

The first cost shifter captures improvements to the baseline estimates due to learning effects. In this context, “learning” includes not only improvements due to repetitive labor tasks, but also improvements due to advances in technology or manufacturing process. Although the baseline estimates may already embody some degree of learning, the cost shifter reflects variations in the learning rate above or below

the baseline assumptions. Separate estimates are provided for the labor and nonlabor components of each individual WBS element, as shown in Figure II-3. The input to this adjustment is expressed as a mean percentage plus or minus a symmetrical error (e.g., 10% \pm 6%). The symmetrical interval could contain negative values if there is a possibility that baseline learning rates will not be achieved. However, we argue in Chapter III that the use of symmetrical intervals is unnecessarily restrictive.

PROGRAM WBS ELEMENTS (FROM ABOVE)	ELEMENT COST IMPROVEMENT			
	LABOR		NON-LABOR	
	COST IMPROVEMENT (PERCENT)		COST IMPROVEMENT (PERCENT)	
	CHANGE - %	UNCERTNTY	CHANGE - %	UNCERTNTY
Missile				
Propulsion	20	30	5	5
Payload	10	20	5	5
Reentry	10	20	5	5
G&C	30	20	5	5
IA&T	10	20	5	5
S/W Eng'g	10	20	5	5
Program Mngmnt	20	30	5	5
Systems Eng	20	30	5	5
ST&E	20	30	5	5
Training	25	30	5	5
Data	20	10	5	5
Support Equip	20	10	5	5
Initial Spares	20	10	5	5

Figure II-3. RACM Cost Improvement Input

b. Rate Increase

The second cost shifter captures cost growth due to increases in labor rates. Separate estimates are provided for the labor and nonlabor components of each individual WBS element, as shown in Figure II-4. Labor rates could affect the nonlabor component of WBS cost if, for example, a subcontractor's labor-rate increases are passed along to the prime contractor via an escalation clause. The input to this adjustment is again expressed as a mean percentage plus or minus a symmetrical error.

The final element-specific cost shifter is the schedule adjustment. For each WBS element, and separately for labor and nonlabor, the user enters the number of "critical/parallel paths," an integer between 1 and 15. This value indicates the number of activities that must all finish in order for the labor (or, respectively, nonlabor) portion of spending in the WBS element to terminate. The concept is that when there are more

critical/parallel paths, there are more ways in which the program schedule can slip, resulting in larger costs.

		ELEMENT RATE INCREASE			
		LABOR		NON-LABOR	
PROGRAM WBS ELEMENTS (FROM ABOVE)	RATE INCREASE (PERCENT)		RATE INCREASE (PERCENT)		
	CHANGE - %	UNCERTNTY	CHANGE - %	UNCERTNTY	
Missile					
Propulsion	5	5	0	0	
Payload	5	5	0	0	
Reentry	5	5	0	0	
G&C	5	5	0	0	
IA&T	5	5	0	0	
S/W Eng'g	5	5	0	0	
Program Mngmnt	5	5	0	0	
Systems Eng	0	5	0	0	
ST&E	5	5	0	0	
Training	5	5	0	0	
Data	5	5	0	0	
Support Equip	5	5	0	0	
Initial Spares	5	5	0	0	

Figure II-4. RACM Rate Increase Input

c. WBS-Element-Specific Schedule Adjustment

The precise Lockheed Martin definition of a critical/parallel path is as follows:

It is assumed there is some concept of a schedule flowchart, "PERT" [Program Evaluation and Review Technique] chart, or generic time-phased performance/milestone chart associated with each WBS element included in the RACM description as well as a similar, higher level chart for the entire ensemble. The absence of such information would indicate no potential slippage influence factor to be considered and imply an input of zero maximum slippage across a single path through the respective element/task. A "critical" path is any one that might cause the potential slippage on its own/independently of any other. The "number" of critical paths is then just the simple counting of the number of different paths so identified as critical.⁶

⁶ See Appendix A, p. A-31.

The Lockheed Martin definition of critical/parallel paths is unconventional. In a standard textbook on operations research, Hillier and Lieberman use the following sequence to define a critical path in a PERT network:⁷

- The *earliest time* for an event is the (estimated) time at which the event will occur if the preceding activities are started as early as possible;
- The *latest time* for an event is the (estimated) time at which the event can occur without delaying the completion of the project beyond its earliest time;
- The *slack* for an event is the difference between its latest and its earliest time;
- A *critical path* for a project is a path through the network such that the activities on this path have zero slack. (All activities and events having zero slack must lie on a critical path, but no others can.)

To avoid confusion, we refer to the RACM concept as “parallel paths,” not “critical paths.”

Finally, the user also provides a most-likely duration and a “3-sigma” pessimistic duration for both the labor and nonlabor components of each WBS element. The mathematical treatment of the schedule adjustment is deferred until Chapter III.

2. Global Cost Shifters

The next three factors operate directly on the distribution of total program cost.

a. Global Schedule Adjustment

The user enters the number of global parallel paths, an integer between 1 and 15. This value indicates the number of WBS elements that must all finish in order for the entire program to finish. The precise Lockheed Martin definition of a global parallel path is as follows:

At the “global” or system level, WBS elements play the role of [parallel] paths and the total count of possible elements capable of producing such slippage is the number [of parallel paths] expected in the RACM input.⁸

The WBS elements identified as parallel paths are those that could conceivably delay the entire program, while the remaining WBS elements are presumably of such

⁷ Frederick S. Hillier and Gerald J. Lieberman, *Introduction to Operations Research*, fifth edition, New York: McGraw Hill, 1990, pp. 372–373.

⁸ See Appendix A, p. A-31.

short duration that they could not possibly delay the entire program. Note that the current implementation of RACM restricts the total number of WBS elements to 30 at most, but the number of parallel paths (necessarily a smaller number) to 15 at most.

The user also provides a most-likely duration and a "3-sigma" pessimistic duration for the entire program. Again, the mathematical treatment of the schedule adjustment is deferred until Chapter III.

b. Major Problems

The user enters a list of up to four "major problems," along with their associated probabilities of occurrence and cost impacts. The names of the major problems are provided by the user. Thus, the major problems serve as placeholders for sources of risk that, in the opinion of the user, are not adequately addressed anywhere in the model. Examples might be a labor strike or failure of a test flight. RACM assumes that the probabilities of occurrence of the major problems are statistically independent, and that the cost impacts are strictly additive.

c. Management Effect

The RACM developers gave the following description of the behavior underlying the management effect:

There is also a much more subtle effect; the effect of the "self-fulfilling prophecy." This effect results from hidden incentives to spend whatever budget is allocated in order to achieve the most reliable product. Few, if any, incentives exist on most programs which motivate the design engineer or product development team to finish the project ahead of schedule or to reduce costs and risk not meeting all specifications. This results in excess time spans and excess personnel which translates into excess program costs. In fact, the Budgeted Cost of Work Scheduled (BCWS) in a program using CSCSC [Cost/Schedule Control Systems Criteria] is often interpreted by people working on the program as a goal which must be met. Underruns are often considered as being as serious a [problem] as overruns are. Neither may be a serious problem or could be a significant problem. An instant interpretation by program personnel results in a misuse of the intended purpose of the process. However, a significant number of people interpret this as the "budget line" not to be varied from. There are many more reasons why program costs are higher than they need to be, such as, unwillingness or inability to remove personnel who are no

longer needed, inability of the program management to “see” when a program is completed, etc.⁹

A succinct summary was provided in a follow-on paper by one of the RACM developers:

The incentive structure which exists on most programs [contains] subtle incentives which are incorporated into present day management styles. They are unwritten incentives. That is:

- The inability and reluctance to off-load competent personnel from the program,
- The need to spend the budget provided or “it will not be available next year,”
- The need by the program personnel such as design engineers to provide the most reliable product for the money provided.¹⁰

We should emphasize that the behavior just described represents an observation made by the individuals who developed RACM. It is certainly not an official position of the Lockheed Martin Corporation. Nor, as we see in Chapter V, is this observation shared by the cost-estimating department at another major defense contractor, Boeing Defense and Space Group (D&SG). Therefore, the general validity of the management effect within the defense industry remains unproved.

We henceforth use the term “money allocated is money spent” (MAIMS) interchangeably with the management effect. The intricate mathematical analysis of MAIMS is presented in Chapter III. The essence of that analysis is as follows. At the start of program execution, upper management allocates to each WBS-element manager a fraction of his or her total budget. Under the MAIMS principle, the WBS-element manager will spend *at least* this initial allocation. Thus, the distribution of possible cost outcomes for each WBS element is transformed. In particular, the mean cost outcome is increased and the standard deviation is decreased. Although the cost distributions for the individual WBS elements shift discontinuously, the distribution of total program cost will generally shift smoothly. The latter, too, experiences an increase in the mean and a decrease in the standard deviation.

⁹ *Ibid.*, p. A-6.

¹⁰ Creaghe Gordon, “Risk Analysis and Cost Management (RACM): A Cost/Schedule Management Approach Using Statistical Cost Control (SCC),” p. 6. The first observation (the reluctance to off-load personnel) has been demonstrated empirically in a sample of defense contractors: Matthew S. Goldberg and Thomas P. Frazier, “Employment and Utilization of Engineers Among Defense Contractors,” *Naval Research Logistics*, Vol. 41, No. 7, December 1994, pp. 853–874.

RACM provides the contractor with two distinct capabilities. First, it enables the contractor to account for the shift in cost during proposal development. That is, the contractor anticipates the shift in cost and bids the mean (or some appropriate percentile) of the *transformed* cost distribution rather than the original cost distribution. Thus, bid levels are more realistic because they embody contractor behavior.

The second capability is actually designed to attenuate the MAIMS effect. For example, suppose that the contractor wins the competition with a 95th-percentile total bid. The contractor might decide to allocate only a portion of this total at the start of the program and hold the remainder in a management reserve fund. Although RACM does not calculate the “optimal” reserve level (a difficult concept to define), it does contain an algorithm to spread the budget allocation across WBS elements in a reasonable and defensible manner. In Chapter III, we investigate the properties of RACM’s reserve-management algorithm, and compare it to an alternative that has been proposed in the literature.

D. CORRELATIONS

Early versions of RACM included a cell range in which the user could enter the correlations, if desired. However, the RACM developers later removed this capability from the model. Their rationale for inclusion and subsequent removal is contained in the following two paragraphs:

As RACM views the world, correlation would be included/modeled at the fundamental “basic” WBS-element level and would represent a relationship between two cost elements. If the original WBS breakdown separated two portions of the same basic task into different items in the RACM model inputs, these elements would likely need a correlation statement describing their combined behavior. It is assumed this will be avoided and that the fundamental WBS element structure chosen will reflect independent efforts. Another potential source of correlation would be if the same person or “expert” made the predictions for two fundamentally distinct elements but he himself is biased in his estimation processes. Although provision was included in the original RACM implementation to handle such an instance, it has been decided that determining the proper correlation coefficients is extremely difficult at this time and the chance for error overshadows the potential gain. It is possible that the “placeholder” for correlation should be left in the proposed future

RACM model, but at this time, it has been removed from the current demonstration version.¹¹

They continue:

[Correlations are] not currently an influence factor described/utilized by the RACM model. This is because it was always intended only to represent the possible correlation between the fundamental basic cost estimates and has not turned out to be significant (or even definable) for any cases seen so far. It was a place-holder in the original formulation and has since been removed. Every instance [of correlation] brought to our attention so far has actually been a case of some outside influence which is modeled elsewhere in the simulation and wouldn't be applicable (or worse, double-counted) if assumed as the correlation asked for in the original RACM input list.¹²

We examine these assertions more closely in Chapter III. To some extent, correlations are interchangeable with modeling relationships. At the extreme, suppose that two cost elements always moved in strict proportion, $Y = bX$. Then in the spreadsheet, the two cost elements could be represented in any of the following three equivalent ways:

- treat X and Y as distinct cost elements with maximal correlation of $\rho = 1$;
- link the X and Y cells in the spreadsheet by the formula $Y = bX$, and zero-out their correlation (to avoid double counting); or
- treat Y as a proportional add-on to X , and combine them into a single spreadsheet cell: $Z = X + Y = X + bX = (1 + b)X$.

As we show in Chapter III, the equivalence breaks down when the two cost elements are correlated only imperfectly, $0 < \rho < 1$. In that situation, a modification of the first strategy listed above is obviously correct: simply treat X and Y as distinct but correlated cost elements. By contrast, we show that the second and third strategies lead to an overstatement of $Var(X + Y)$.

RACM actually attempts to capture the correlations using *nonlinear* modeling relationships. One possible mechanism is the global schedule effect that operates multiplicatively on all of the WBS elements. Another possible mechanism is the MAIMS effect: the left-hand tail of each WBS-element distribution is removed, so “high” values are drawn from each distribution, apparently inducing a correlation. However, we show in

¹¹ See Appendix A, p. A-22.

¹² *Ibid.*, p. A-33.

Chapter III that these mechanisms only *partially* capture the effects of correlations on the variance in total cost. Although these mechanisms are automatically in place when using RACM, sole reliance on them without introducing explicit correlations leads to an *understatement* of the variance in total cost. Correspondingly, the correct percentile levels (e.g., the budget level necessary to cover the costs incurred in 90 out of 100 possible replications of the program) will be understated as well.

E. SOFTWARE IMPLEMENTATION

1. Platform

RACM was developed on a Microsoft Excel platform. RACM does not have formal version numbers. The first version of RACM we received was written in Microsoft Excel version 5.0. That version and all subsequent versions of RACM are upwardly compatible to the current Excel versions 7.0 and 8.0. The version of RACM we evaluated (created 20 August 1996) occupies 500 kilobytes of disk storage. Earlier versions fluctuated in size between roughly 500 kilobytes (i.e., 0.5 megabyte) and 1 megabyte.

RACM consists of seven worksheets within a single Microsoft Excel workbook. RACM is *not* programmed in either Visual Basic, Microsoft Access, or even the older Microsoft Excel 4.0 macro language. Instead, it is programmed to run via formula-linked spreadsheet cells. A change to any one of the input cells will feed changes in many intermediate and final output cells. Pressing the "Recalculate" key (F9 on the keyboard) refreshes all linked cells to reflect any changes to the input. Given the current limitation to 30 WBS elements, recalculation is essentially instantaneous.

Although RACM recalculates quickly, it contains no provision for automatically cataloging the sensitivity results when input is changed. However, sensitivity results could be cataloged using the "Scenario Manager" feature of Microsoft Excel. In particular, it would be good practice to save a baseline scenario (under a different name) for downstream comparison to actual costs incurred. RACM can be rerun midway through a program to model total program cost in light of actual costs incurred to date. This procedure amounts to a fresh look at the entire program using updated input.

Previous forecasts and actual costs to date are combined in an *ad hoc* fashion via the updated user input, rather than through a formal statistical model.¹³

It is extremely difficult for an outsider to comprehend the calculations encoded as formula-linked spreadsheet cells. Indeed, our attempts to track down particular calculations often resulted in a “wild goose chase” across several of the RACM worksheets. Our task would have been considerably simpler had RACM been programmed in Visual Basic, especially if structured programming concepts were employed to their fullest extent. It would be equally difficult for an outsider to modify or customize the RACM spreadsheet.

There is no user’s manual for RACM. Nor is there any online help, other than the standard Microsoft Excel help. We discuss documentation in more detail later in this chapter.

2. Interface

Navigation within RACM is achieved using Microsoft Excel named ranges. Pressing the “Go To” key (F5 on the keyboard) brings up the dialog box in Figure II-5. From this dialog box, the user can jump to any of nine input ranges:

- Project title,
- Basic estimates,
- WBS-element-specific schedule effects,
- Cost improvements,
- Rate increases,
- Global schedule effect,
- Major problems,
- Management policy, or
- Management policy variation.

The first seven input ranges are self-explanatory. The eighth input range, “Management policy,” contains cells in which the user enters the bid level and allocation

¹³ A more formal Bayesian approach using Kalman filtering is found in Mark A. Gallagher and David A. Lee, “Final Cost Estimates for Research and Development Programs Conditioned on Realized Costs,” *Military Operations Research*, Vol. 2, No. 2, 1996, pp. 51–65. A software implementation is found in David A. Lee and John A. Dukovich, “Using the Rayleigh Analyzer,” AT701C1, Logistics Management Institute, March 1998.

level, respectively, for the entire program. Thus far, we have only briefly alluded to the bid and allocation levels in our discussion of the management effect. These concepts are explored in greater detail in Chapter III. We also see in Chapter III that RACM assumes an *equi-percentile* budget allocation. That is, each WBS-element manager is budgeted the same percentile of his or her respective cost distribution. The ninth and final input range, "Management policy variation," enables the user to override the equi-percentile budget allocation if desired.

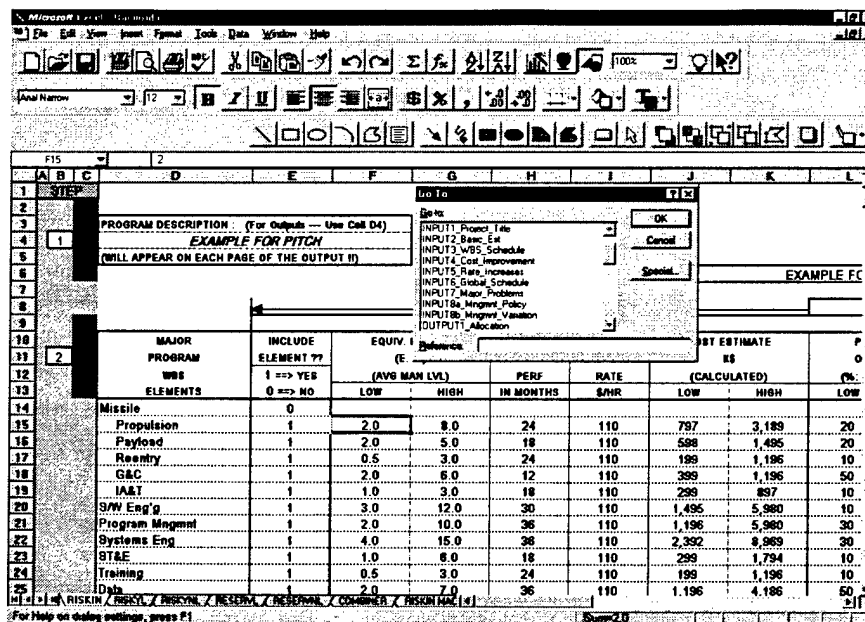


Figure II-5. RACM Spreadsheet Navigation: Input Ranges

By scrolling down the dialog box, as illustrated in Figure II-6, the user can jump to any of nine output ranges:

- Allocation,
- Breakdown,
- Cost tables,
- Cost plot,
- Plot data table,
- Beta distribution,
- Sensitivity plot,
- Profit distribution, or
- Reserves beta.

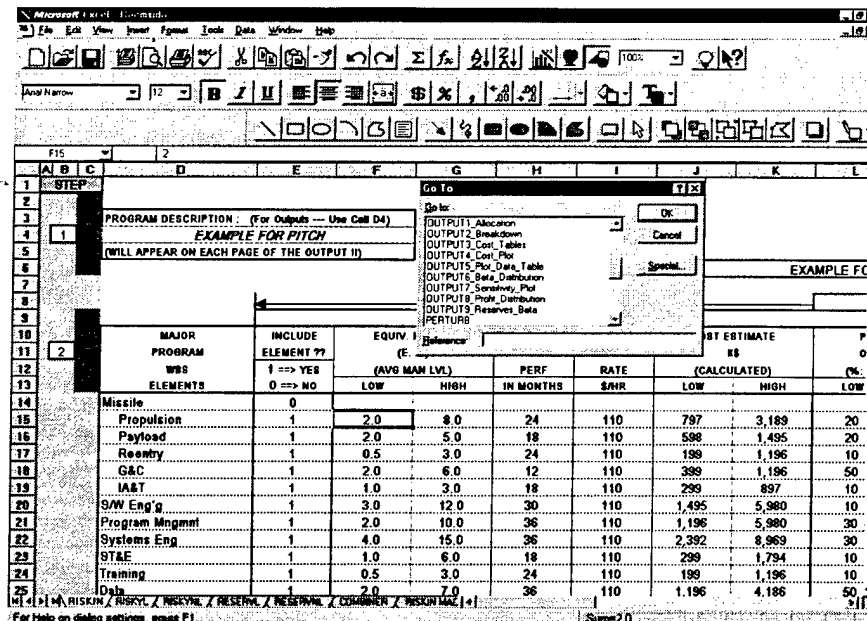


Figure II-6. RACM Spreadsheet Navigation: Output Ranges

Some of the output ranges are difficult to interpret and, again, there is no documentation either in the form of a user's manual or online help. However, the most important output ranges appear to be as follows. First, the range named "Cost tables" contains seven incremental estimates of the cumulative distribution function (CDF) of total program cost. The cost tables, which actually appear side-by-side in the spreadsheet, are displayed here as Figures II-7 and II-8. These tables correspond to the CDFs upon incrementally introducing the following risk factors:

- Basic estimates,
- WBS-element-specific schedule effects,
- Cost improvements,
- Rate increases,
- Global schedule effect,
- Major problems, and
- Management impacts.

The risk factors are introduced in the order indicated by these bullets. We argue in Chapter III that this order is indeed appropriate.




Figure II-7. RACM Cost Tables (left-hand side)

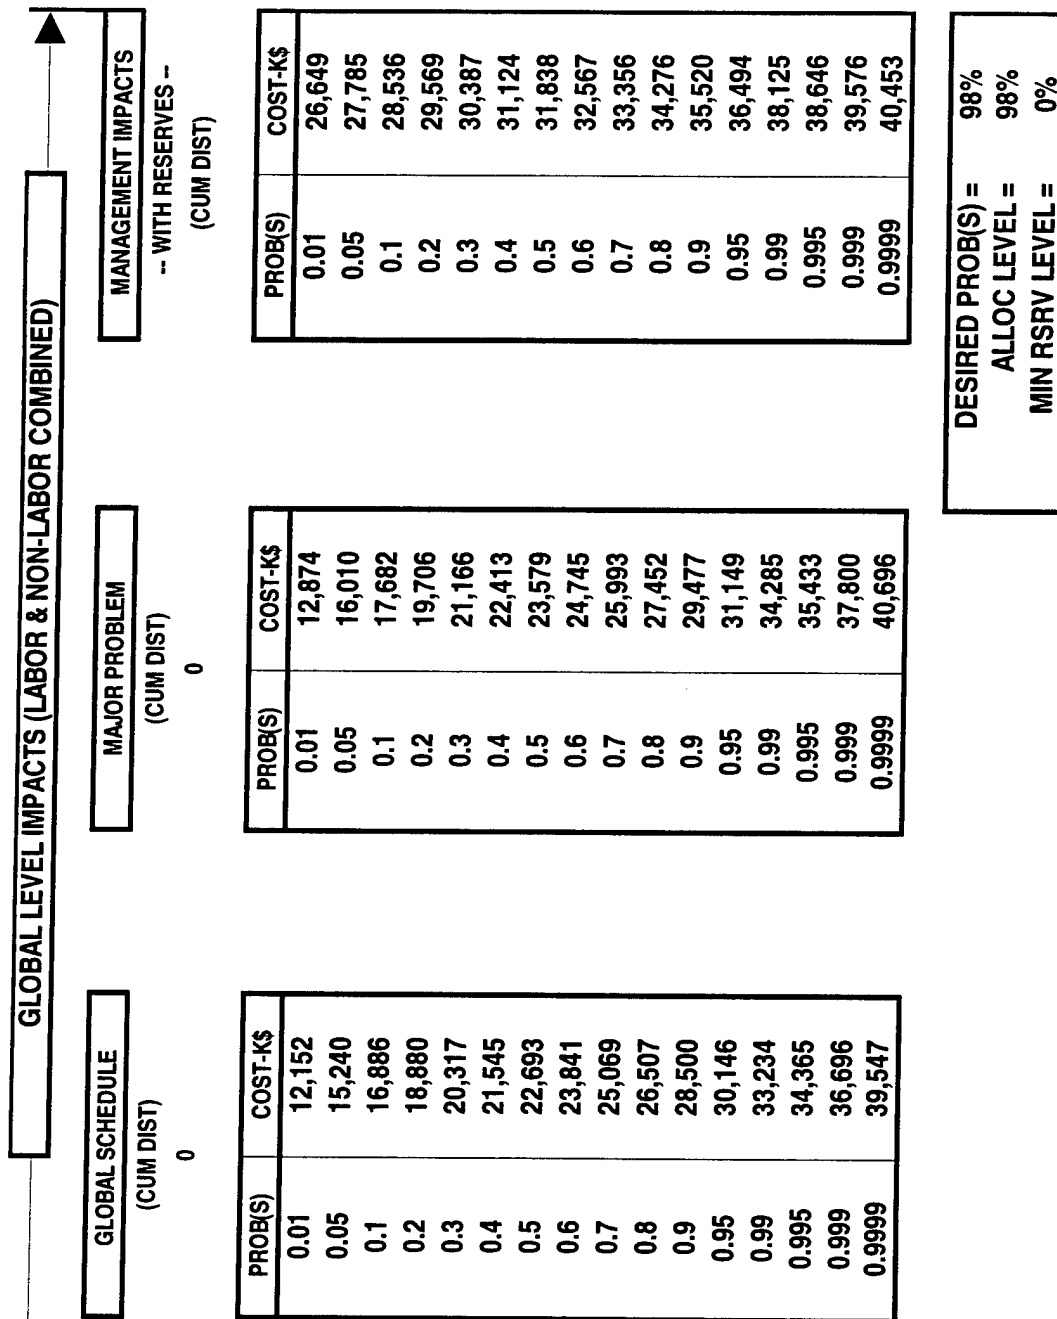


Figure II-8. RACM Cost Tables (right-hand side)

To generate the cost tables, RACM keeps track of the mean and standard deviation of each WBS element as each risk factor is applied. In Chapter III, we present the formulas for updating the means and standard deviations to reflect the risk factors. RACM fits a normal distribution to baseline cost in each WBS element, as well as to cost after application of each risk factor. The only exception is that RACM fits a beta distribution to total program cost after the transformation associated with the MAIMS effect ("Management impacts"). The rationale for fitting a beta distribution at this stage is presented in Chapter III.

The output range named "Cost plot" is an Excel overlay chart of the data contained in Figures II-7 and II-8. The cost plots, displayed here as Figure II-9, are obtained by analytical calculation from the fitted distributions as described in Chapter IV.

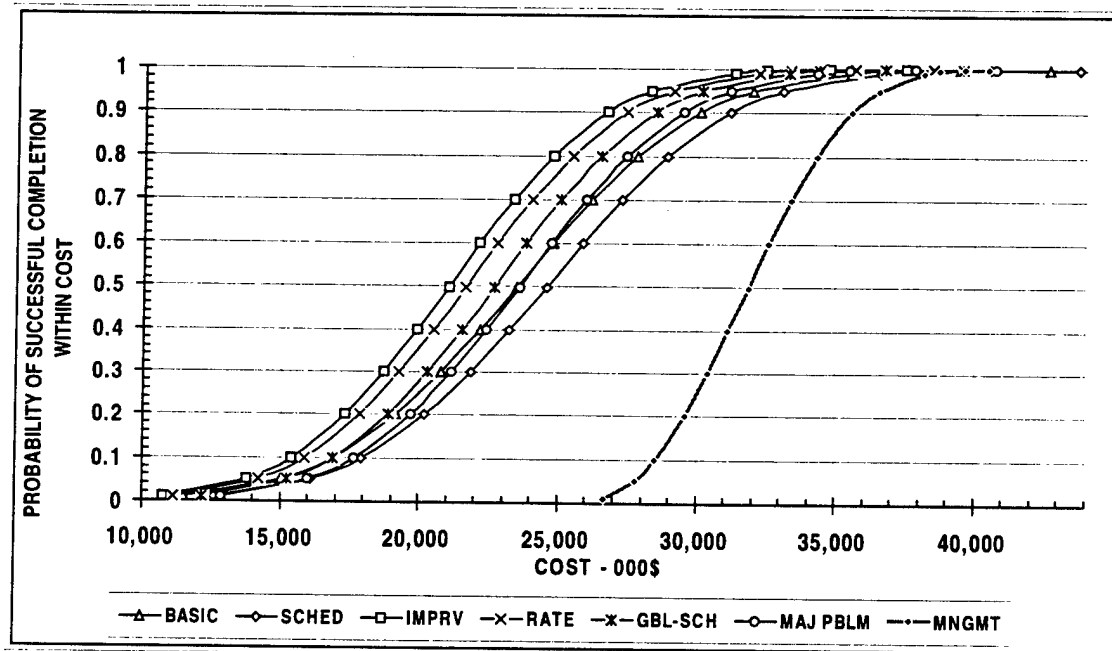


Figure II-9. RACM Incremental Overlay Chart

The sequential view of the various risk factors, illustrated here in Figures II-7 through II-9, is probably the greatest strength of the RACM implementation. As we see in Chapter IV, this capability is difficult to duplicate using COTS software. However, as we discuss later in this chapter, the sequential view poses some problems when eliciting the input.

The output range named "Beta distribution" computes the parameters of the beta distribution for total program cost. RACM's calculations are consistent with the formulas that we report in equations (C-5) and (C-6) of Appendix C. Although this information is used internally, it will not be of any interest to the majority of model users. Nor, in our opinion, are any of the other output ranges that we have not explicitly discussed.

F. POSSIBLE ENHANCEMENTS

1. Platform

To summarize much of the preceding discussion, the RACM implementation in Microsoft Excel suffers from the following shortcomings:

- no version numbers to reflect software enhancements,
- no user's manual,
- no online help,
- limited to 30 WBS elements,
- difficult to comprehend, modify, or customize code (in contrast to a structured programming language), and
- primitive navigation system.

Although it is easy to criticize the RACM implementation, one must recall the history of its development. RACM was originally designed by a small group of Lockheed Martin analysts for their own use. They did not anticipate its use elsewhere within their own company, and they certainly did not anticipate its possible use more widely throughout the defense industry. The limited scope of their effort was surely responsible for their decision to develop RACM on a Microsoft Excel platform.

Virtually all of these shortcomings could be alleviated by rewriting RACM in a structured programming language such as Visual Basic. Under current versions of Microsoft Excel (e.g., version 7.0 or 8.0), the data could still reside in a spreadsheet, but the navigation and much of the calculation would be controlled by programming statements rather than by formula-linked spreadsheet cells. A graphical user interface (GUI) would replace the primitive navigation system. In particular, the user could input the number of WBS elements into a dialog box at the beginning of each model run. The model would then automatically reset the dimensions of all the data arrays. The model would also be simpler to debug or modify because the calculations would be distributed among well-defined subroutines.

A Visual Basic environment would also facilitate some natural enhancements to RACM. For example, a dialog box could offer the user a choice among several probability distributions in addition to the normal, such as the uniform, triangular, beta, or lognormal. A good example of this feature in COTS software is the "Distribution Gallery" in Crystal Ball for Windows.¹⁴

2. Other Enhancements

Another enhancement would be to add a time dimension to the model. The current "Rate increase" factor reflects a one-time perturbation to the labor rates, applicable to the entire duration of the program. Thus, on the first day of the program, the labor rates are determined for the entire duration, possibly at levels different from those embodied in the baseline estimates. A more interesting situation for multiperiod programs is when the labor rates start at known levels, but are perturbed to unknown levels part of the way into the program (e.g., workers receive an annual raise within the program horizon, but the percentage raise is not known at the outset). This situation could be modeled indirectly within RACM by treating the starting labor rates plus the expected annual raise as the baseline value, and treating variations from the expected raise (properly delayed and discounted) as the "Rate increase." This calculation would have to be performed off-line, and would not be auditable within RACM alone. Alternatively, we are suggesting another worksheet that would perform this calculation directly within RACM. This approach would make explicit the user's assumptions about the timing of pay raises and the discount factors. The new worksheet would feed its results to the current "Rate increase" input sheet.

Chapter III shows that, excepting the highly nonlinear management effect, the cost shifters operate either additively or multiplicatively. Yet another enhancement would be to include additional templates for additive or multiplicative effects, beyond those reflected in the current list of cost shifters. The "major problems" cost shifter already allows for up to four additive effects, but these effects apply globally rather than at the WBS-element level. We are suggesting templates for either additive or multiplicative effects, applicable at either the global or WBS-element level. The names of these effects would be provided by the user, and they would serve as placeholders for sources of risk that are not adequately addressed elsewhere in the model.

¹⁴ We discuss the capabilities of Crystal Ball for Windows in Chapter IV. The "Distribution Gallery" feature is illustrated in Appendix E, Figure E-3.

We have already indicated that the cost-improvement effect is overly restrictive because it allows only a symmetrical range of cost improvement around the baseline value. An alternative would be to allow the user to input a symmetrical range for the learning slope relative to an exponential learning curve. As we demonstrate in Chapter III, a symmetrical range for the learning slope implies an *asymmetrical* range of cost improvement. We develop the mathematical basis for this statement in Chapter III, and we suggest that our formulas be incorporated into RACM.

Additional enhancements could improve the database management within RACM. The software currently expects the user to navigate the input ranges by using the "Go To" key and typing the input data into the appropriate blank spreadsheet cells. An alternative would be to allow the user to develop the input in a different software application and import it into RACM in database format (.dbf) or Microsoft Access format (.mdb). We have already indicated input that operates behind the scenes at a deeper level than the input that appears in the RACM input ranges: labor rates by time period, discount factors, ranges of learning slopes, and additive or multiplicative effects at the global or WBS-element level. For large programs, spanning multiple time periods and more than 30 WBS elements, accuracy and convenience would both improve if the input could be developed externally to RACM.

Finally, the utility of RACM would improve if its results could be exported, again in either database or Microsoft Access format, to other software for program management or earned-value management.¹⁵ For example, Performance Analyzer (PA) is a government-owned program-management tool available to DoD offices and their contractors free of charge.¹⁶ Performance Analyzer performs the following function:

PA is a PC-based system to analyze and report cost/schedule data. Earned value, planned cost, and actual cost data are reported by contractors using Cost Performance Reports (CPR) or Cost/Schedule Status Reports

¹⁵ Earned-value management is described at the OSD Acquisition and Technology web site: <http://www.acq.osd.mil/pm/>. Earned-value management contains 32 criteria that replace the 35 criteria comprising the older Cost/Schedule Control Systems Criteria (C/SCSC). A comparison of the two sets of criteria is found at: <http://www.acq.osd.mil/pm/currentpolicy/critcomp.htm> [accessed June 1, 1998]. See also U.S. General Accounting Office, "Major Acquisitions: Significant Changes Underway in DoD's Earned Value Management Process," NSIAD-97-108, Washington, D.C., May 5, 1997.

¹⁶ Performance Analyzer can be downloaded from the Air Force Financial Management web site: <http://www.hq.af.mil/SAFFM/>. The following branches must be taken: "Hot links," "Earned value management," "Tools," "Cost Management Systems," "Performance Analyzer for Windows" [accessed June 1, 1998]. A manual is available from the vendor: Cost Management Systems, Inc., "Performance Analyzer Version 1.3 for Windows," 301 West Maple Avenue, Vienna, Virginia 22180.

(C/SSR). These reports can be loaded electronically into PA using either standard [electronic data interchange] EDI transfer protocols...or PA proprietary transfer files. PA enables the analyst to identify and analyze problem areas/trends (e.g., cost overruns or schedule slippages); support independent cost analysis exercises; produce summary charts and reports for the analyst, management, and executive; and calculate forecasts and indices.¹⁷

With some effort, RACM could be modified to export results either to Performance Analyzer or to commercial program-management software.

G. ELICITING THE INPUT

The RACM developers give the following description of the process for eliciting the input:

The inputs should come from the interview process between the cost analyst who will be running RACM and the individual program managers/cost account managers involved in actually performing the tasks. In practice, we have experienced many variants on this procedure depending upon the experience of the cost analyst, the statistical knowledge of the program personnel, the accumulated knowledge of the interviewers about the specific engineers and program personnel providing the answers, and seemingly unique situations that arise in every application. Almost surely, the input/interview process will be iterative in nature. This, however, can be very valuable in terms of good understanding of the final product and the confidence that there will not be a lot of missing information in the final estimates. To be sure, the final output is only as good as the quality of the input. But this is true no matter what the cost estimating procedure chosen—certainly not a unique problem with RACM! *In fact, because of the possibility for sensitivity analyses that the analytic model provides, it should be considerably easier to identify and resolve potential input errors than with other approaches.* [Emphasis in original.]

Although historical data [have] quite often been thought of as the best source of information for such analyses, caution should be observed. There are many sources for error buried in most historical data bases. New ways of doing business, the “self-fulfilling prophecy” phenomenon [i.e., the management effect or MAIMS], missing documentation, unknown sources of apparent correlation, and management philosophy followed (what probability of success was associated with that bid and

¹⁷ This description is taken from the Defense Technical Information Center (DTIC) web site: <http://www.dtic.mil/dodim/costool.html> [accessed June 1, 1998].

what reserves/fencing of budget were in place, etc.) are only some examples. Historical data may well be the best source available, should always be considered when [they] exist, but should always be used carefully.¹⁸

We have several comments on this process. First, the italicized passage alludes to a greater ease of performing sensitivity analysis using RACM's analytical approach in contrast to a Monte Carlo approach. Although we agree that the analytical approach requires less computational time, we argue in Chapter IV that the time difference is not a major consideration in choosing between the two approaches. Second, we have already acknowledged that historical data may be contaminated by the management effect, because the data do not reflect the budget underruns that might have occurred had behavioral incentives allowed them. However, we believe that the RACM developers are too quick to dismiss historical data in favor of expert opinion. The difficulties in eliciting expert opinion are well-documented and legion.¹⁹

Two major events occurred during the course of our evaluation, leading us to conclude that eliciting expert opinion is the most serious impediment to implementing RACM. At one point we contemplated a field test of RACM at a government depot. The RACM developers informed us that the field test could not proceed without them because only they, not the IDA team, could properly elicit the input. This exchange took place after the IDA team had already been involved in the RACM evaluation for a full year. If their assertion is true, it does not bode well for the possibility of widespread dissemination of RACM.

Although the field test was never actually performed, we did convince the Boeing D&SG to retrofit RACM on a proposal that they had just completed. Their experience in attempting to do so, discussed in greater detail in Chapter V, reveals that even knowledgeable analysts had difficulty interpreting several of the cost categories and cost shifters within RACM. In particular, the Boeing team set the cost improvement and rate increase cost shifters to zero because these effects were already included in their baseline cost estimates. They were unable to back these effects out of the baseline to produce the "bare-bones" baseline expected by RACM. Nor were they able to directly estimate the

¹⁸ See Appendix A, Section 3.2, "Sources of the Inputs," p. A-27. The discussion continues into Section 4.0, "Process for Eliciting the Inputs," pp. A-27 through A-30.

¹⁹ M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge: Cambridge University Press, 1990, especially chapters 6 and 7.

global schedule or major problem effects. It appears that they just made up numbers to test the latter two features of RACM: "Due to the intended depth of comparison, an attempt to develop inputs for these [global schedule and major problem] RACM risk contributors...was not made. However, to demonstrate RACM's sensitivity to these parameters, values were input."²⁰

The inability of Boeing D&SG analysts to provide the necessary input, and the RACM team's lack of confidence in the IDA team's ability to do so, both militate against the use of RACM outside of Lockheed Martin's Sunnyvale plant.

H. RANGE OF APPLICABILITY

The RACM developers identified three factors that determine the applicability of RACM:

- program size,
- program phase, and
- program maturity.

By "program phase," they are distinguishing among development, procurement, and operations and maintenance (O&M). In addition, the applicability of RACM may vary between the domains of proposal preparation and management of existing programs.

1. Program Size

The RACM developers make the following statement about program size, with which we generally concur:

The size and type of the program should be the initial factor[s] used to determine the need for a statistical approach. Size is a function of the number and cost, in dollars, of the WBS elements required to define a program. The greater the number of elements used to define the program and the larger the dollar amount associated with each element, the more applicable is the RACM process. For example, a study contract which is a cost plus contract requires only a simple analysis. However, a contract which requires a great number of WBS elements and large funding can induce considerable errors if a single point estimate approach is used to estimate the cost of the program.²¹

²⁰ See Appendix G, pp. G-5 to G-6.

²¹ See Appendix A, p. A-3.

2. Program Phase

The RACM developers argue that their model may be applicable in all three program phases: development, procurement, and O&M.

Development programs by their very nature have great uncertainty in the costs associated with meeting performance requirements, the quality of personnel involved, and the schedule time-to-complete. For example, achieving performance requirements may not be very uncertain if the requirement is a variation on known results (e.g., a minor extension of range). However, if the performance requirement involves new technology (e.g., a major extension of range involving new propellant not yet developed), the uncertainties can be large and important.

In the manufacturing phase of a program, the degree of uncertainty is dependent upon the maturity of the program, the complexity of the product, the work force, etc. If the objective is the production of a stable, high-volume item such as pills, the uncertainties associated with the cost projections would be minimal. However, the uncertainties associated with most weapons procurements are great because weapons are typically made in low volume and are a new product with an unknown history. This is the type of situation in which a statistical approach would be most useful and would introduce a greater degree of reliability in the estimating process.

The operations and maintenance phase should be part of the original analysis as well as any analysis continued during the development of any new system for two reasons: the O&M phase is usually the most costly; and the unknowns accompanying an undeveloped system resulting from factors involving such things as training, storage, and quality problems which may be recognized only after deployment, cause each WBS cost element to have large uncertainties.²²

The applicability of RACM may also vary between the domains of proposal preparation and management of existing programs. We return to this point in Chapter V, where we discuss the reactions of Boeing D&SG upon attempting to use RACM.

3. Program Maturity

The RACM developers make the following statement about program maturity:

During the development of the cost model for the program...each estimate has distributions associated with [it]. During the conduct of the program, information can be obtained from the CSCSC system that will be able to

²² *Ibid.*, pp. A-4 to A-5.

reduce the [variance] of the distributions and adjust the nominal. As the program continues to mature, the distributions should continue to be adjusted. For example, a subcontractor may have been on a cost plus fixed fee contract during the initial development phase, however, the subcontractor may be put on a firm fixed price contract during the final development phase. Another reduction might occur from an analysis of CSCSC. If, after 50% of the program is completed, the Earned Value shows no significant variance from that projected in the model, the [variance of the] distribution can be significantly reduced.²³

I. DOCUMENTATION

In our opinion, the state of RACM documentation is quite poor. As we have already discussed, the limited scope of the RACM team was responsible for their decision to develop RACM on a Microsoft Excel platform. Similarly, in view of their limited scope, they did not foresee the need to write much documentation. In fact, their only documentation was created at the behest of IDA to facilitate our evaluation, and is reproduced as Appendix A of this report.

Although we encourage our readers to make their own determination, our opinion is that Appendix A is insufficient to prepare even experienced analysts to correctly use RACM. We believe that our report, taken in its totality, improves the state of documentation somewhat. However, it was not within our charter to write either a comprehensive users' manual for RACM or a tutorial on data elicitation. Those two elements are essential before a model such as RACM can be considered adequately documented.

²³ *Ibid.*, p. A-5.

III. MATHEMATICS OF RACM

This chapter gives a detailed analysis and critique of the mathematics of RACM. We begin with the basic input and proceed to the WBS-element-specific effects and then the global effects. We also discuss the order in which the various effects are calculated and the treatment of correlations in RACM.

A. BASIC INPUT

For each WBS element, and separately for labor and nonlabor, RACM begins with two percentile points of the cost distribution. That is, x_1 is a percentile point such that the area to the left is p_1 , and x_2 is another percentile point such that the area to the left is p_2 (e.g., the 10th and 90th percentile points). RACM also assumes that the baseline cost for each WBS element is normally distributed. Letting Φ denote the standard normal CDF, we have two equations to determine the mean (μ) and the standard deviation (σ):

$$p_1 = \Phi\left(\frac{x_1 - \mu}{\sigma}\right) \quad \text{and} \quad p_2 = \Phi\left(\frac{x_2 - \mu}{\sigma}\right). \quad (\text{III-1})$$

These equations may be solved for μ and σ .

$$\sigma = \frac{x_2 - x_1}{\Phi^{-1}(p_2) - \Phi^{-1}(p_1)} \quad \text{and} \quad \mu = x_2 - \sigma \Phi^{-1}(p_2). \quad (\text{III-2})$$

The solution involves the inverse of the normal CDF, which is preprogrammed into Microsoft Excel and calculates essentially instantaneously.

B. WBS-ELEMENT-SPECIFIC EFFECTS

1. Cost Improvement

The input to this adjustment is expressed as a mean percentage plus or minus a symmetrical error (e.g., $10\% \pm 6\%$). This adjustment factor is treated as having a mean as stated, and a standard deviation equal to one-third of the (one-sided) deviation; in the above example, $3\sigma = 6\%$ or $\sigma = 2\%$. The rationale is that, for a normal distribution,

99.7% of the probability lies within ± 3 standard deviations of the mean. The assumption of normality is not used in any other way, except for the budget allocation algorithm described later in this chapter.

Denote cost prior to application of this factor as C and cost after application as $C \times (1 + f)$. Assume that C and f are statistically independent. Under these conditions, it may be shown *exactly* that (without normality or any other particular distributional assumption):¹

$$\begin{aligned} E[(1 + f) \times C] &= [1 + E(f)] \times E(C), \\ \text{Var}[(1 + f) \times C] &= [E^2(1 + f) \times \text{Var}(C)] + [E^2(C) \times \text{Var}(f)] + [\text{Var}(f) \times \text{Var}(C)]. \end{aligned} \quad (\text{III-3})$$

However, early versions of RACM used the following approximation for the variance, based on Taylor series:²

$$\text{Var}[(1 + f) \times C] \approx [E^2(1 + f) \times \text{Var}(C)] + [E^2(C) \times \text{Var}(f)]. \quad (\text{III-4})$$

By omitting the final (positive) term in equation (III-3), RACM systematically understated the variance induced by compounding the adjustment factors. We note that the RACM developers have subsequently replaced equation (III-4) with the exact equation (III-3).

Cost improvement, particularly for labor cost, may arise from an exponential learning curve. Lot cost has a particular value under baseline or expected learning assumptions. Lot cost has a different value under alternative learning assumptions. The difference between these two assumptions gives a range of cost improvement, bracketing the value zero though not necessarily symmetrically.

To pursue these ideas, assume that unit cost follows a power function:

$$C(X) = a X^{-b}. \quad (\text{III-5})$$

¹ See Leo Goodman, "On the Exact Variance of Products," *Journal of the American Statistical Association*, Vol. 55, 1960, pp. 708–713; or M. Granger Morgan and Max Henrion, *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge: Cambridge University Press, 1990, pp. 189–190.

² See F. Seiler, "Error Propagation for Large Errors," *Risk Analysis*, Vol. 7, 1987, pp. 509–518; or Morgan and Henrion, *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis*, chapter 8.

The "learning slope" is defined as the ratio of unit costs for two units whose positions in the production sequence differ by a factor of two:

$$\rho = C(2X)/C(X) = 2^{-b}.$$

It is generally assumed that $0 \leq b < 1$, so $1/2 < \rho \leq 1$. Given this specification for unit cost, we may compute the cumulative cost of a lot consisting of X units:

$$F(X; b) = \int_0^X a y^{-b} dy = \frac{a}{1-b} X^{1-b}.$$

Cost improvement may be measured by the percentage difference between lot cost with the baseline or expected learning rate (i.e., learning parameter b^*), and lot cost with an alternative learning rate (i.e., learning parameter b):

$$\begin{aligned} G(X; b) &= \frac{F(X; b^*) - F(X; b)}{F(X; b^*)} = 1 - \frac{F(X; b)}{F(X; b^*)} \\ &= 1 - \left[\left(\frac{1-b^*}{1-b} \right) X^{(b^*-b)} \right]. \end{aligned} \tag{III-6}$$

Figure III-1 depicts this relationship for a range of lot sizes and three different learning slopes. With a smaller learning slope, more learning occurs and cost improvement is positive; with a larger learning slope, cost improvement is negative. The range in possible values for the learning slope could arise from studies of the historical costs of similar systems. Often, the range will take the form of a symmetric interval centered around the baseline learning slope. For a given lot size, however, a symmetric interval for the learning slope implies an asymmetric interval for percentage cost improvement. Thus, RACM should be modified to allow an asymmetric range of cost improvement.

Figure III-1 can also be used to estimate the effect of lot size on total lot cost at various learning rates. Although the contractor usually bids on a fixed lot size, the RFP may specify consideration of several alternative lot sizes if funding levels are uncertain.

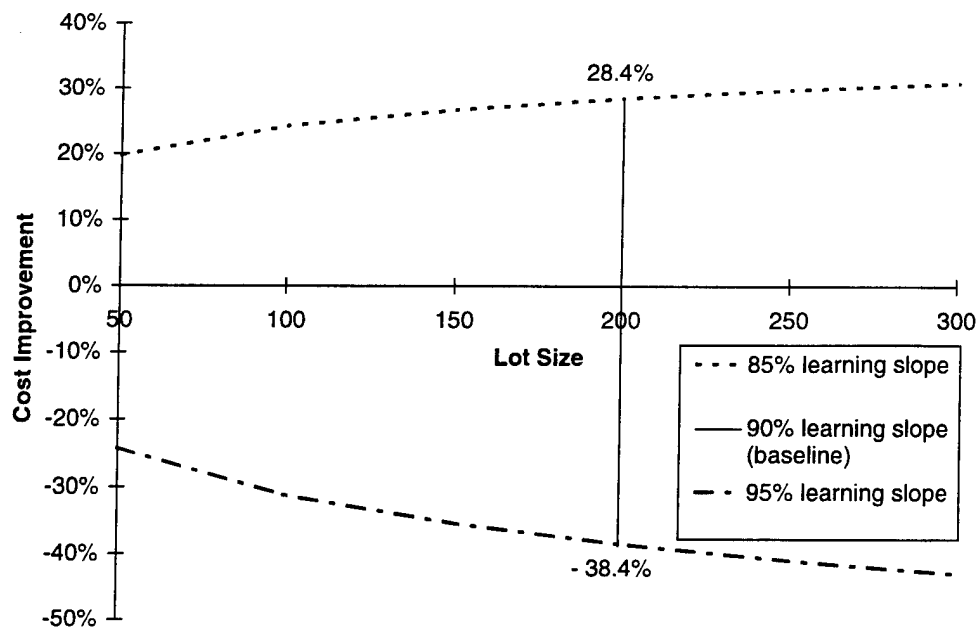


Figure III-1. Percentage Cost Improvement as a Function of Lot Size and Learning Slope

2. Rate Increase

This factor captures cost growth due to increases in labor rates. The mathematical analysis here is similar to that for the cost improvement. In particular, the exact update formula in equation (III-3) should be used rather than the approximation in equation (III-4).

3. WBS-Element-Specific Schedule Adjustment

a. Derivation of Adjustment Factors

For each WBS element, and separately for labor and nonlabor, the user enters the number of parallel paths, an integer between 1 and 15. This value indicates the number of activities that must all finish in order for the labor (or, respectively, nonlabor) component of spending in the WBS element to terminate. The precise Lockheed Martin definition of a parallel path is given in Chapter II. RACM assumes that the durations along the various parallel paths are statistically independent and identically distributed. The user also provides a most-likely duration and a “3-sigma” pessimistic duration for both the typical (identically distributed) activity comprising the labor component of the WBS element and the typical (identically distributed) activity comprising the nonlabor component.

The assumed duration distributions for each identically distributed activity are depicted in Figure III-2. For labor activities, the distributions are triangular with the right-hand endpoint three times as far from the mode as the left-hand endpoint. Thus a schedule overrun is assumed to be three times as likely as a schedule underrun.³ The user provides the mode (i.e., most-likely duration) and the right-hand endpoint (i.e., pessimistic duration). Sigma is determined as one-third the difference and, finally, the left-hand endpoint (i.e., optimistic duration) is computed as the mode minus sigma. Note that sigma is not actually the standard deviation of this triangular distribution; sigma is merely a parameter to allow calculation of the left-hand endpoint that the user does not directly provide.⁴

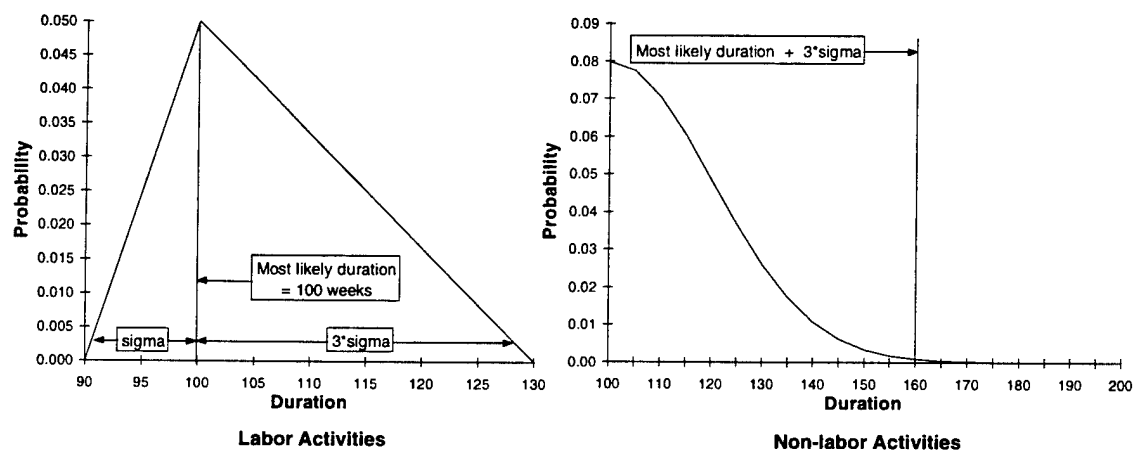


Figure III-2. Duration Distributions Along WBS-Element-Specific Parallel Paths

For nonlabor activities, the distributions for each parallel path are assumed to be half-normal. The rationale for this distribution is that nonlabor activities (i.e., subcontracts) will most likely complete on time, will sometimes complete late, but will *never* complete early. The user again provides the most-likely and pessimistic durations, and sigma is determined as one-third the difference. Note that sigma is not actually the

³ The probability of underrun is depicted in Figure III-2 as the area of the (sub)triangle with endpoints 90 and 100; the probability of overrun is depicted as the area of the (sub)triangle with endpoints 100 and 130. The area of a triangle equals half the product of its base and height. Because both triangles have the same height, the areas (i.e., probabilities) are proportional to the width of the base. Thus an overrun is three times as likely as an underrun.

⁴ This triangular distribution has mode m , low value $m - \text{sigma}$, and high value $m + (3 \times \text{sigma})$. It can be shown that the standard deviation is equal to $.850 \times \text{sigma}$.

standard deviation of this distribution either.⁵ Also, the half-normal is not the only possible distribution for representing completion times with this general shape. One alternative, for example, is a right-triangular distribution (i.e., a triangular distribution with the left-hand endpoint equal to the mode).

The mean and variance of labor and nonlabor cost are individually updated, and then combined to yield the mean and variance of total cost within the WBS element. We now describe the update methodology. Let m denote the most-likely duration, σ (sigma) the delay parameter from the previous two paragraphs, n the number of parallel paths, and f_n and g_n the mean and sigma factors (which depend on n) selected from Table III-1. Also let C_0 denote the mean cost and V_0 the variance in cost prior to the schedule adjustment. RACM uses the following equations for the mean and variance of cost, adjusted for the maximum duration along the n parallel paths:

$$\begin{aligned} E(C) &= C_0 + \left[\left(\frac{C_0}{m} \right) \times f_n \sigma \right], \\ \text{Var}(C) &= V_0 + \left[\left(\frac{C_0}{m} \right) \times g_n \sigma \right]^2. \end{aligned} \tag{III-7}$$

To understand these equations, we begin by writing adjusted cost as the sum of unadjusted cost and an additive adjustment:

$$C = C_0 + \left(\frac{C_0}{m} \right) \times t. \tag{III-8}$$

In equation (III-8), the random variable t represents the schedule delay, and the term (C_0/m) represents the “burn rate” or average cost per unit time. Note the assumption that, during the delay period, costs continue to accumulate at the same average rate estimated

⁵ If sigma is thought of as the standard deviation of the original (two-sided) normal distribution, then the area beyond $3 \times \text{sigma}$ under the “folded” half-normal distribution is 0.0027. Thus the 3-sigma point essentially completes the right-hand tail of the half-normal distribution. The standard deviation of the half-normal distribution is actually equal to $.603 \times \text{sigma}$.

for the most-likely duration; we return to this assumption shortly. Equation (III-8) is consistent with the equations in (III-7), provided that we can show the following:

$$E(t) = f_n \sigma ,$$

$$Var(t) = (g_n \sigma)^2 . \quad (III-9)$$

Table III-1. Shift Factors for WBS-Element-Specific Schedule Adjustment

Number of Parallel Paths	Labor Cost			Nonlabor Cost	
	Approximate Mean Factor	Exact Mean Factor ^a	Approximate Sigma Factor	Approximate Mean Factor	Approximate Sigma Factor
1	0.667	0.667	0.850	0.798	0.603
2	1.151	1.150	0.770	1.128	0.602
3	1.419	1.416	0.705	1.325	0.587
4	1.599	1.592	0.648	1.466	0.571
5	1.719	1.720	0.602	1.567	0.557
6	1.819	1.819	0.562	1.655	0.547
7	1.901	1.898	0.530	1.722	0.537
8	1.962	1.962	0.505	1.784	0.528
9	2.014	2.017	0.482	1.832	0.520
10	2.063	2.064	0.463	1.881	0.513
11	2.102	2.105	0.447	1.923	0.507
12	2.145	2.140	0.431	1.960	0.502
13	2.172	2.172	0.416	1.995	0.496
14	2.201	2.201	0.401	2.025	0.492
15	2.227	2.226	0.387	2.053	0.488

a. Source: handwritten notes provided by Stephen A. Book.

We first consider the simplest case by setting the number of parallel paths $n = 1$ and also setting $\sigma = 1$. In this case, the triangular distribution of the single labor activity in Figure III-2 has a mode of 0 and endpoints -1 and +3; the half-normal distribution of the single nonlabor activity has mode 0 and pessimistic (i.e., 3-sigma) value +3. A triangular distribution with these parameters has mean $2/3$ ($= 0.667$) and standard deviation 0.850, exactly the factors f_1 and g_1 in the first row (corresponding to $n = 1$) of Table III-1. Similarly, a half-normal distribution⁶ with these parameters has mean $\sqrt{2/\pi}$ ($= 0.798$) and standard deviation $\sqrt{1 - (2/\pi)}$ ($= 0.603$), again the factors f_1 and g_1 in the

⁶ The properties of the half-normal distribution are found in Norman L. Johnson and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions-1*, New York: John Wiley, 1970, pp. 81-83.

first row of Table III-1. Thus the equations in (III-9) hold in the special case considered in this paragraph.

The equations in (III-9) continue to hold when we relax the assumption $\sigma = 1$. A value of $\sigma \neq 1$ just rescales the probability distribution, and the factor σ in (III-9) correctly reflects the effect of rescaling on the mean and variance. All that remains is to show that the factors f_n and g_n in Table III-1 correctly adjust for the further relaxation $n \neq 1$. These factors represent the mean and variance of the maximum of n identically distributed random variables, each having the standardized distribution described in the previous paragraph (i.e., triangular with mode 0 and endpoints -1 and +3, or half-normal with mode 0 and pessimistic value +3). The RACM developers arrived at these factors through Monte Carlo simulation. We independently verified the factors through our own Monte Carlo simulation, using the Crystal Ball simulation software with 10,000 replications. Our simulation results agree to at least two decimal places with the factors in Table III-1. Finally, Stephen A. Book has calculated the *exact* factors for the shift in the mean labor cost. These exact factors, too, agree to at least two decimal places with those in Table III-1. We are not aware of any derivation of exact counterparts to the other three columns of factors. However, the accuracy of the approximate factors for mean labor cost suggests that the values in Table III-1 are close enough for virtually all uses.

b. Percentage Growth Assumption

We now turn to the implications of using burn rates to extrapolate cost into the delay period. We first observe that, when using burn rates, the RACM additive adjustment equation in (III-7) may equally well be viewed as a multiplicative adjustment:

$$E(C) = C_0 \times \left[1 + \left(\frac{f_n \sigma}{m} \right) \right]. \quad (\text{III-10})$$

The RACM model apparently works under the assumption that a given percentage delay (i.e., expected delay time, $f_n \sigma$, relative to the most-likely duration, m) implies cost growth *of the same percentage*. If there are, for example, 15 parallel paths, then all 15 must finish before spending on this WBS element terminates. However, the model assumes that spending on all 15 activities will continue until the 15th activity is finished, even if the 14 other activities have already finished. This assumption may be at odds with reality. A more appropriate assumption may be that spending for each activity terminates when that particular activity finishes. Under the latter assumption, an additional delay in the 15th activity would increase the cost of that activity, but not the cost of the 14

completed activities. Thus the percentage cost growth could be considerably smaller than the percentage schedule delay.

Figure III-3 is an example of the percentage growth assumption. Suppose that there are 15 labor activities, with the duration of each given by the triangular distribution shown earlier in Figure III-2; this distribution has a mean duration of 107 weeks and a standard deviation of 8.5 weeks. We sampled 15 activity lengths from this distribution, and sorted these activities from shortest to longest. The longest activity in this sample takes 128 weeks to complete. Now suppose that some perturbation delays the longest activity from 128 weeks to 132 weeks, or by 3.1%. RACM assumes a fixed burn rate, so that cost growth for the entire WBS element is equal to this same 3.1%. Under the alternative assumption, cost growth on the longest activity equals 3.1%, but neither duration nor cost increase for the 14 shorter activities. To estimate the magnitudes involved, assume that all activity-weeks are equally costly. The total number of weeks in Figure III-3, before the additional delay, is equal to the sum of the lengths of all 15 bars, or 1,630. Thus percentage cost growth in this WBS element is only 4 weeks over a base of 1,630 weeks, or 0.2%.

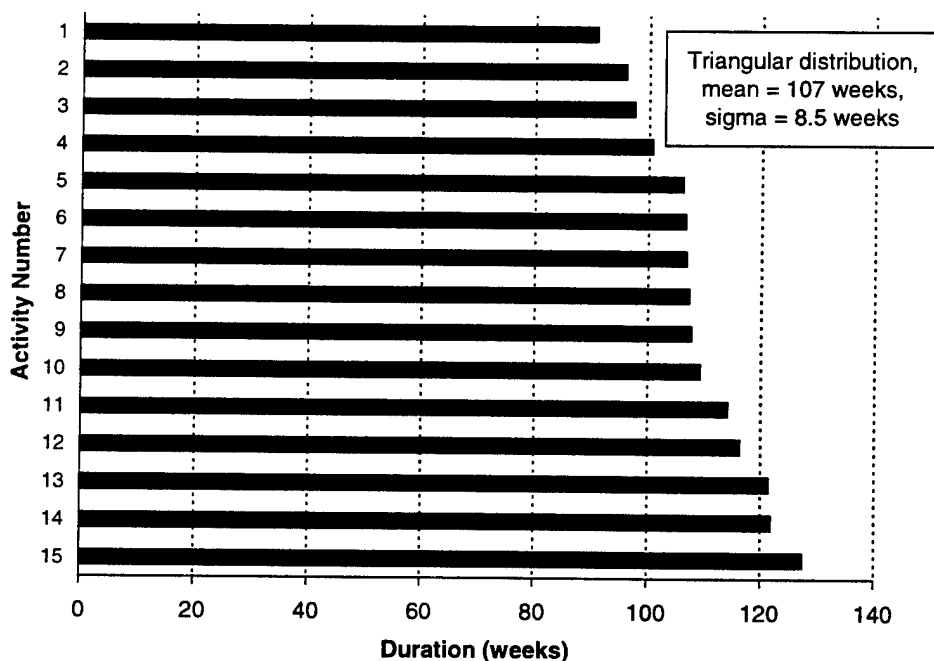


Figure III-3. Example of Percentage Growth Assumption

C. GLOBAL EFFECTS

1. Global Schedule Adjustment

The user enters the number of global parallel paths, an integer between 1 and 15. This value indicates the number of WBS elements that must all finish in order for the entire program to finish. The precise Lockheed Martin definition of a global parallel path is given in Chapter II. The user also provides a most-likely duration and a “3-sigma” pessimistic duration for the entire program. The global schedule adjustment is then applied to only the *labor* component of total program cost. The rationale is that delays on the part of subcontractors may cause cost growth in the corresponding WBS elements (already accounted for in the model), but will not cause additional cost growth for the entire program.

The global schedule adjustment again uses the equations in (III-7) and the factors presented in Table III-1. Thus mean cost is multiplied by the factor $1 + (f_n \sigma / m)$, where now m represents the most-likely duration of the entire program, σ represents one-third the difference between the most-likely and pessimistic durations, and n is the number of global parallel paths.

2. Major Problems

The user enters a list of up to four “major problems,” along with their associated probabilities of occurrence and cost impacts. The model computes the probabilities of every possible combination of these four major problems (none may occur; there are four ways that exactly one will occur; there are six ways that exactly two will occur; there are four ways that exactly three will occur; or all four may occur). The model also computes the cost impact of every possible combination, which is simply the sum of the cost impacts of the events that comprise that particular combination. An implicit assumption is that the costs of the major problems are indeed additive rather than, for example, multiplicative.⁷ Finally, the model computes the expected cost impact as the weighted sum of the cost impacts across all the combinations, with weights equal to the respective probabilities. Note that the adjustment for major problems is applied to only the *labor* component of total program cost.

⁷ Suppose that laborers in one trade declare a strike just after failure of a test flight. The costs of redesign and rework could easily be super-additive, if not multiplicative.

Although RACM's calculations are correct, they could be simplified. RACM explicitly considers the $2^4 = 16$ combinations of the four major problems, and works with the probabilities of each combination; these probabilities necessarily sum to 1.0. Alternatively, one could work with the underlying four major problems rather than the 16 combinations. Let P_i $\{i = 1, \dots, 4\}$ denote the probability of the i th major problem, and note that more than one major problem may arise concurrently so that the $\{P_i\}$ may sum to more than 1.0. Let B_i denote a Bernoulli-distributed random variable that takes the value 1 if the i th problem arises and 0 otherwise. The random variable B_i has mean P_i and variance $P_i \times (1 - P_i)$, and the variables $\{B_i\}$ are assumed statistically independent.

Let C_i denote the incremental cost associated with the i th problem. Finally, let C_0 denote the mean cost and V_0 the variance in cost prior to the adjustment for major problems. Then adjusted system cost may be represented as the following random variable:

$$C = C_0 + \sum_{i=1}^4 B_i C_i .$$

The mean and variance of adjusted system cost are given by:

$$\begin{aligned} E(C) &= C_0 + \sum_{i=1}^4 P_i C_i , \\ \text{Var}(C) &= V_0 + \sum_{i=1}^4 P_i (1 - P_i) C_i^2 . \end{aligned} \tag{III-11}$$

Equations (III-11) require only four terms to perform the adjustment, in contrast to the sixteen terms currently used in RACM.

3. Management Effect (Money Allocated Is Money Spent)

a. Budget Allocation Theorem

The total cost of a program is equal to the sum of the costs of each WBS element. At the start of program execution, upper management allocates to each WBS-element manager a fraction of his or her total budget. Under the MAIMS principle, the WBS-element manager will spend *at least* this initial allocation. Let $f(x)$ denote the density function and $F(x)$ the CDF of cost, in a particular WBS element, prior to budget allocation. Let $f^*(x)$ and $F^*(x)$ denote the corresponding functions after budget allocation.

If we let θ denote the budget level, the two sets of functions are related in the following way:

$$\begin{aligned} f^*(x) &= 0, & F^*(x) &= 0 & \text{if } x < \theta; \\ f^*(x) &= F(\theta), & F^*(x) &= F(\theta) & \text{if } x = \theta; \\ f^*(x) &= f(x), & F^*(x) &= F(x) & \text{if } x > \theta. \end{aligned} \quad (\text{III-12})$$

The effect of this transformation is to fold all of the probability to the left of θ onto a mass point at θ . The density function to the right of θ is unchanged by the transformation. Figure III-4 provides an illustration. Allocating a budget equal to the 69th percentile leads to a mass point at that value (i.e., at \$45), with height equal to 0.69 (the mass point actually extends well beyond the vertical scale of the right-hand chart).




Figure III-4. Density Functions Before and After Transformation

Figure III-5 illustrates the CDFs before and after transformation. The transformed CDF lies along the horizontal axis to the left of the mass point, because the probability of observing cost below the budget allocation is assumed to be zero. The original and transformed CDFs coincide to the right of the mass point. Note that the CDF for this single WBS element does not shift in its entirety. However, the CDF for total program cost will generally shift entirely to the right as the discontinuous shifts among individual WBS elements are smoothed out.

As the CDF for each WBS element shifts, so do other distribution parameters such as the mean and standard deviation. The mean of total cost is the sum of the WBS-element means, so it too will shift. The manner in which WBS-element budget levels are allocated will determine the magnitude of the shift in the mean of total cost. RACM

assumes an *equi-percentile* budget allocation. That is, each WBS-element manager is budgeted the same *percentile* of his or her respective cost distribution. For example, each WBS-element manager may be given his or her respective 69th percentile. Note that this is quite different from giving each WBS-element manager the same *percentage* of the most-likely cost, the pessimistic cost, or any other benchmark measured along the cost (as opposed to probability) axis.




Figure III-5. Cumulative Distribution Functions Before and After Transformation

The RACM developers conjectured that an equi-percentile budget allocation is in some sense “optimal.” The nature of that optimality was derived by the IDA team, and is detailed in Appendix B. We show there that, starting from an equi-percentile budget allocation, any zero-sum budget reallocation (i.e., the total of the budget increases equals the total of the budget decreases) leads to a further rightward shift in the mean of total cost. Thus, if MAIMS is assumed and if the objective is to minimize the mean of total cost, an equi-percentile budget allocation is indeed optimal.⁸

This theorem offers no indication of the optimal, common percentile. For example, it offers no basis to choose between giving each WBS-element manager his or her 65th percentile versus giving each one his or her 70th percentile. The theorem merely states that any given, total budget is better allocated as equal percentiles than as unequal percentiles.

⁸ If the objective is broader, including the standard deviation or other higher-order moments of total cost, then the equi-percentile budget allocation is not necessarily optimal. The optimal budget allocation relative to a broader objective does not appear to have been worked out in the literature.

Determining the optimal percentile goes beyond the scope of either Lockheed Martin's model or IDA's evaluation. A larger percentile increases the minimum threshold as well as the mean of total cost. Thus it might appear that a lower percentile is always better. Pushing this argument to its logical extreme, why not set the percentile, and thus the initial budget allocation, to zero? This extreme solution would not work, because WBS-element managers require seed money to initiate subcontracts, hire workers, purchase materials, and so on. WBS-element managers would immediately appeal to upper management for budget relief, so the initial zero allocations would be extremely transitory. Moreover, the process of adjudicating requests for budget relief is itself costly. Small initial allocations and small budget increments could lead to numerous appeals and excessive management costs.

The optimal budget allocation would balance the costs of upper management intervention against the wasteful spending that occurs when WBS-element managers feel obliged to exhaust their individual budgets. It would be interesting to develop a complete theory of budget allocation, including not only the initial allocation but also the frequency and size of budget increments. However, that theory lies beyond the scope of this evaluation.

b. Percentile Calculation

RACM applies the MAIMS principle, and the resulting budget allocation theorem, in two distinct modes. First, it enables the contractor to account for the shift in cost during proposal development. That is, the contractor anticipates the shift in the cost distribution illustrated in Figures III-4 and III-5, and bids the mean (or some appropriate percentile) of the *transformed* cost distribution rather than the original cost distribution. Second, RACM contains a module designed to help manage program-wide reserve levels during program execution.

RACM requires the user to enter the "probability of success" for the entire program. The probability of success is simply the percentile level of the total cost distribution that the contractor intends to bid. For example, bidding at the 65th percentile would provide a large enough budget to cover the costs incurred in 65 out of 100 possible replications of the program. Thus, the probability of success is exactly 65 percent.

To estimate the shift in each WBS-element cost distribution, RACM must first allocate the 65th percentile of total cost among the various WBS elements. In accordance with the budget allocation theorem, RACM applies an equi-percentile budget allocation.

That is, RACM solves for the common percentile for each WBS element that is consistent with a 65% probability of success for the entire program.

Suppose that each WBS element X_i ($i = 1, \dots, n$) is normally distributed with mean μ_i and standard deviation σ_i . Consider any budget allocation $\{B_1, \dots, B_n\}$ with total budget $B^* = \sum_{i=1}^n B_i$. The individual WBS-element percentile levels are given by:

$$P_i = \Pr(X_i < B_i) = \Pr\left[\frac{X_i - \mu_i}{\sigma_i} < \frac{B_i - \mu_i}{\sigma_i}\right] = \Phi\left[\frac{B_i - \mu_i}{\sigma_i}\right], \quad (\text{III-13})$$

where Φ again represents the standard normal CDF. But with an equi-percentile budget allocation, the percentiles $\{P_1, \dots, P_n\}$ must all be equal. Because Φ is a monotonic function, the expressions $(B_i - \mu_i)/\sigma_i$ must all be equal as well:

$$(B_i - \mu_i)/\sigma_i = \alpha \quad (\text{III-14})$$

for all $i = 1, \dots, n$ and a common value of α . Indeed, we may write simply $P_i = \Phi(\alpha)$.

Next consider total program cost, $X^* = \sum_{i=1}^n X_i$. The mean of total program cost is $\mu^* = \sum_{i=1}^n \mu_i$. The standard deviation depends upon those of the individual WBS elements, as well as the correlations among the WBS elements:

$$\sigma^* = \sqrt{\sum_{i=1}^n \sigma_i^2 + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n \rho_{ij} \sigma_i \sigma_j}, \quad (\text{III-15})$$

where ρ_{ij} is the correlation between the i th and j th WBS elements and $|\rho_{ij}| \leq 1$.

Using equation (III-14), we may write the program-wide budget level as:

$$B^* = \sum_{i=1}^n B_i = \sum_{i=1}^n \mu_i + \alpha \sum_{i=1}^n \sigma_i. \quad (\text{III-16})$$

Finally, the program-wide percentile level is given by:⁹

$$\begin{aligned}
 P^* &= \Pr (X^* < B^*) = \Pr \left[\frac{X^* - \sum \mu_i}{\sigma^*} < \frac{B^* - \sum \mu_i}{\sigma^*} \right] \\
 &= \Pr \left[\frac{X^* - \sum \mu_i}{\sigma^*} < \frac{\alpha \sum \sigma_i}{\sigma^*} \right] \\
 &= \Phi \left[\frac{\alpha \sum \sigma_i}{\sigma^*} \right].
 \end{aligned} \tag{III-17}$$

Given any target value for the program-wide percentile, equation (III-17) may be solved for α . The common WBS-element percentile is then determined as $P_i = \Phi(\alpha)$.

We can also show that the program-wide percentile exceeds the common WBS-element percentile when the latter is greater than one-half, $P^* > P_i > 0.5$; the opposite condition holds when the common WBS-element percentile is less than one-half, $P^* < P_i < 0.5$. To see these points, recall that the standard normal CDF is monotonically increasing with median zero. Thus, $P_i = \Phi(\alpha) > \Phi(0) = 0.5$ implies $\alpha > 0$. We can also show that the term that multiplies α in equation (III-17) exceeds 1.0. It is equivalent, but slightly easier, to show that the square of that term exceeds 1.0:

$$\left[\sum_{i=1}^n \sigma_i \right]^2 = \sum_{i=1}^n \sigma_i^2 + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n \sigma_i \sigma_j \geq \sum_{i=1}^n \sigma_i^2 + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n \rho_{ij} \sigma_i \sigma_j = (\sigma^*)^2. \tag{III-18}$$

It follows that, given $\alpha > 0$, $P^* = \Phi[\alpha \sum \sigma_i / \sigma^*] > \Phi(\alpha) = P_i > 0.5$. A corresponding argument establishes the opposite case, $P^* < P_i < 0.5$.

Table III-2 illustrates these calculations. We first concentrate on the columns under the heading "Proposal development." The target probability of success for the entire program (i.e., the program-wide percentile) is shown in the final row of the table as

⁹ The sum of standard deviations that appears in equation (III-17) should not be confused with the standard deviation of the sum; the correct expression for the latter was given in equation (III-15). One author who fell into this trap is A. D. Kazanowski, "A Quantitative Methodology for Estimating Total System Cost Risk," in *Management of Risk and Uncertainty in Systems Acquisition: Proceedings of the 1983 Defense Risk and Uncertainty Workshop*, Defense Systems Management College, Fort Belvoir, Virginia, July 1983, pp. 135-163. A complete analysis of this situation is found in Matthew S. Goldberg, "Some Fallacies in Cost-Risk Analysis," *PHALANX*, September 1996.

95%. We use the approach described after equation (III-17) to calculate the common probability of success for each WBS element. The value of α turns out to be 0.6292, implying a probability of success equal to $P_i = \Phi(0.6292) = 0.74$ (i.e., 74%). Given the WBS-element means and standard deviations, the 74th-percentile bid levels may then be calculated for each WBS element. The sum of these bid levels, \$22,606, is exactly the 95th percentile of the total cost distribution.

Table III-2. Example of RACM Percentile Calculation

WBS element	Proposal development		Reserve management		WBS-element parameters	
	Bid level	Bid probability	Allocation	Allocation probability	Mean	Sigma
Propulsion	1,816	74%	1,470	56%	1,365	716
Payload	1,021	74%	875	56%	831	301
Reentry	740	74%	602	56%	560	286
Guidance and control	497	74%	358	56%	316	287
Integration, assembly & testing	622	74%	537	56%	512	175
Software engineering	4,975	74%	4,164	56%	3,917	1,683
Program management	3,189	74%	2,288	56%	2,013	1,870
Systems engineering	4,793	74%	3,607	56%	3,245	2,461
Test equipment	1,110	74%	896	56%	831	443
Training	689	74%	555	56%	515	278
Data	1,793	74%	1,203	56%	1,023	1,224
Support equipment	895	74%	736	56%	688	328
Initial spares	466	74%	418	56%	404	98
	Total bid	Program probability	Total allocation	Program probability	Overall mean	Overall sigma
	22,606	95%	17,716	65%	16,220	3,882

Table III-2 may also be used for reserve management during program execution. Even if the contractor wins the competition with a 95th percentile total bid, it might decide to allocate only a portion of this total at the start of the program. Suppose that the contractor decides to allocate at the 65th percentile of the total cost distribution. The same sequence of calculations as above yields a value of α equal to 0.1474, implying a WBS-element percentile of $P_i = \Phi(0.1474) = 0.56$. The 56th-percentile allocation levels may be calculated for each WBS element, and the sum of these allocation levels, \$17,716, is exactly the 65th percentile of the total cost distribution.

c. Estimating the Shift in Mean Cost

RACM requires the user to enter not only the probability of success for the entire program, but also the (presumably lower) percentile of the total cost distribution that it intends to allocate at the start of the program. RACM uses this information to estimate the shift in the mean of total cost due to the MAIMS effect.

The mean of the transformed cost distribution is derived in Appendix B. RACM also computes the variance of the transformed cost distribution. As suggested by Figure III-4, the transformed distribution may look decidedly non-normal. Rather than imposing a normal distribution at this point, RACM instead fits a beta distribution. The latter distribution is extremely flexible in fitting a large variety of distributional shapes for random variables that range over a finite interval.¹⁰ The transformed distribution is, strictly speaking, discontinuous at the left-hand mass point. However, a tolerable approximation is available using a backwards J-shaped beta distribution with an infinite left-hand asymptote (i.e., the density approaches $+\infty$ as the cost variable approaches the mass point from the right).¹¹ In addition, the beta distribution has a finite right-hand tail (i.e., there is a finite maximum cost above which the probability equals zero), in contrast to the infinite tail depicted in Figure III-4. An approximate solution is available by truncating the infinite tail at a cost level equal to the mean plus four standard deviations.

The RACM developers have compared the fitted beta distribution to Monte Carlo simulations of the transformed cost distribution. They report that for any percentile between the 50th and the 95th, the corresponding cost levels derived from the fitted beta and Monte Carlo distributions generally agree within $\pm 1\%$. That is, the Monte Carlo CDF is contained within the band obtained by laterally displacing the fitted beta CDF in the amounts -1% and $+1\%$. Their findings are summarized in Figure A-3 and the surrounding discussion in Appendix A.¹²

¹⁰ Philip M. Lurie, Matthew S. Goldberg, and Mitchell S. Robinson, "A Handbook of Cost-Risk Analysis Methods," Paper P-2734, Institute for Defense Analyses, April 1993.

¹¹ Norman L. Johnson and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions-2*, New York: John Wiley, 1970, chapter 24. The distribution that they depict in the upper-left corner of p. 42 provides the best notional fit to our Figure III-4. In terms of the parameterization given in Appendix C, the backwards J-shaped beta distribution is characterized by $0 < \alpha < 1$ and $\beta > 1$.

¹² Figure A-3 actually demonstrates the close agreement between a fitted *normal* distribution and the Monte Carlo simulation; a fitted *beta* distribution should be more flexible and yield even closer agreement.

d. Alternative Budget Allocation Algorithm

An equi-percentile budget allocation is not only optimal in the sense defined above, but may also be perceived as equitable by WBS-element managers, who each receive the same percentile of their respective cost distributions. However, other budget allocation algorithms have been proposed in the literature. For example, Book and Young incorporate a budget allocation algorithm in their FRISK model.¹³ Book and Young define "risk dollars" as the difference between some relatively high percentile of the distribution of total cost (e.g., perhaps the 70th percentile) and the most-likely estimate of total cost. Their algorithm allocates to the WBS elements all of the risk dollars available on the program. Unlike RACM, however, the FRISK algorithm does not appear to optimize any particular objective function. We describe the mathematics of the FRISK algorithm in Appendix H.

In order to assess the practical import of choosing between these two budget-allocation algorithms, we applied both to a numerical example found in Book and Young's series of papers. In their example, a system consists of seven correlated WBS elements, each following a triangular distribution. The parameters of the seven triangular distributions are reproduced here in Table III-3, and the correlation matrix is reproduced in Table III-4. Note that the correlation matrix is necessarily symmetric, so that we show only the upper half.

Table III-3. Triangular Distributions for Budget Allocation Example

WBS Element	Minimum	Mode	Maximum
1	10	20	50
2	20	40	60
3	20	20	80
4	10	40	50
5	20	30	70
6	10	50	90
7	20	40	80

¹³ Stephen A. Book, "Recent Developments in Cost Risk," The Aerospace Corporation, Los Angeles, California, May 1992; Philip H. Young, "FRISK: Formal Risk Assessment of System Cost Estimates," The Aerospace Corporation, Los Angeles, California, February 1992.

Table III-4. Correlation Matrix for Budget Allocation Example

WBS Element		Correlation
#1	#2	
1	2	0.5
1	3	0.7
1	4	0.0
1	5	0.0
1	6	0.0
1	7	0.0
2	3	0.3
2	4	0.0
2	5	0.0
2	6	0.0
2	7	0.0
3	4	0.0
3	5	0.0
3	6	0.0
3	7	0.0
4	5	-0.2
4	6	0.9
4	7	0.0
5	6	-0.4
5	7	0.0
6	7	0.0

The mean and variance of each triangular distribution may be computed using equations (D-1) and (D-2) of Appendix D. The mean and variance of total cost may then be computed using equation (III-15), incorporating the correlations shown in Table III-4. The mean and variance turn out to be 276.7 and $(36.5)^2$, respectively. Book and Young fit a lognormal distribution to the latter two quantities. The 70th percentile of their fitted lognormal distribution is \$293.83. We fit, instead, a beta distribution, using equations (C-5) and (C-6) of Appendix C. The 70th percentile of our fitted beta distribution is \$295.94. Although the difference between the beta and lognormal distributions is small in this example, it can be large in other situations. In particular, the lognormal distribution is not flexible enough to accommodate a distribution with negative skew, which may occasionally arise in practice.

FRISK allocates budgets by first partitioning the \$295.94 into the sum of a most-likely estimate and a total pot of risk dollars. Alternatively, we used RACM to allocate the \$295.94 directly, finding the common percentile for all seven WBS elements that exactly exhausts this sum.

The results of this comparison are shown in Figure III-6. RACM allocates the 61st percentile to each WBS element, whereas the FRISK percentiles range between the 41st and 76th. However, if the MAIMS principle is accepted, the more interesting comparison is between the post-transformation expected cost resulting from the two allocation algorithms. Equation (D-21) of Appendix D gives the mean of a transformed triangular distribution. We evaluated this equation for each individual WBS element and then computed the sum, representing the expected value of total program cost. By allocating \$295.94 using RACM's algorithm and then allowing for the MAIMS effect (i.e., overruns but never underruns of this initial allocation), expected program cost equals \$320.13. Using instead the FRISK algorithm, expected program cost equals \$322.03. The RACM value is necessarily lower, because we proved mathematically in Appendix B that RACM's algorithm minimizes post-transformation expected cost. But the magnitude of the difference (less than one percent) suggests that the optimum is not very sharp, so that the relative advantage using RACM for this purpose may be quite small.

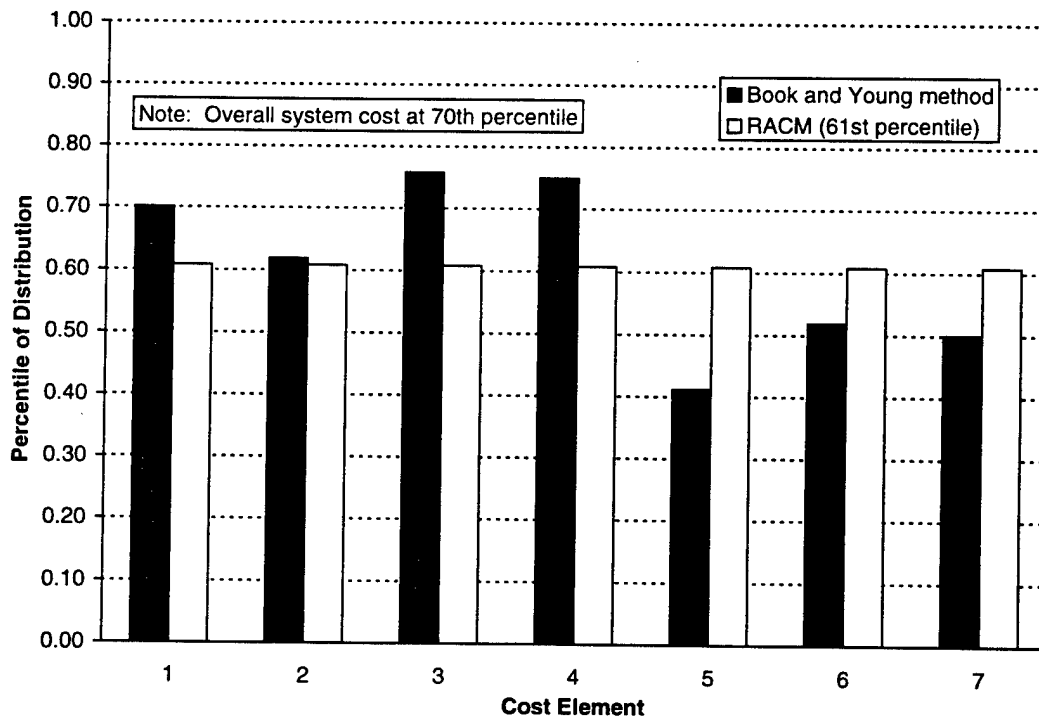


Figure III-6. Comparison Between RACM and FRISK Budget Allocation Algorithms

D. ORDER INDEPENDENCE

The three WBS-element-specific effects operate multiplicatively on cost. We saw in equation (III-3) that the adjustments for cost improvement and rate increase are each of the form $C \times (1 + f)$. Similarly, we saw in equation (III-10) that the WBS-element-specific schedule adjustment is of the form $C \times [1 + (f_n \sigma / m)]$. Because multiplication is a commutative operation, the order of these three adjustments has no effect on the final cost distribution.

The management effect is applied next:

Management policy/reserve consideration impacts are acting only upon the final accumulation of WBS element level factors (right after the [WBS-element-specific schedule adjustment, cost improvement, and] “Rate” term), but are displayed at the end since this is (when properly combined with the global level terms) the “final answer.”¹⁴

This sequence again seems appropriate. The management effect operates on the individual WBS elements through the incentives of the element managers, and should properly precede the global effects.

The global schedule adjustment is again of the form $C \times [1 + (f_n \sigma / m)]$ under the appropriate interpretation of the parameters m and n . Major problems are additive after the global schedule adjustment. Although the management effect is actually computed at the WBS-element level, it is displayed last in the overlay chart:

In the final, global phases of the estimation process, the...global schedule slips and major problem effects are carefully isolated and removed before the “management effect” transformations are applied, and then [the “management effect” is] reapplied afterward to achieve the overall, final cost distributions displayed.¹⁵

E. TREATMENT OF CORRELATIONS IN RACM

1. Correlations Versus Linear Modeling Relationships

As we discussed in Chapter II, correlations are to some extent interchangeable with modeling relationships. To further explore this point, suppose first that two cost

¹⁴ See Appendix A, p. A-34.

¹⁵ *Ibid.*, p. A-21.

elements X and Y satisfy the following *linear* relationship: $Y = bX + u$, with $b > 0$. Under this representation, b is a regression coefficient and u is an error term reflecting the random variation around the regression line. The correlation, ρ , between X and Y is an exact function of the regression coefficient, b , and the standard deviations of X and u :

$$\rho = \left[1 + \left\{ \sigma_u / (b \times \sigma_x) \right\}^2 \right]^{-1/2} > 0. \quad (\text{III-19})$$

Equation (III-19) is plotted in Figure III-7.




Figure III-7. Correlation Between Two Cost Elements

Now consider the following alternative modeling strategies. First, one could treat X and Y as distinct but correlated cost elements. It is necessary to carry along their correlation because, without it, the variance of $X + Y$ (a component of the variance of total cost) would be understated. In fact, we may write explicitly:

$$\sigma_x^2 + \sigma_y^2 \leq \text{Var}(X + Y) = \sigma_x^2 + \sigma_y^2 + 2\rho\sigma_x\sigma_y \leq \sigma_x^2 + \sigma_y^2 + 2\sigma_x\sigma_y. \quad (\text{III-20})$$

Alternatively, one could treat Y as an *exact* linear function of X , and link their cells in the spreadsheet by the formula: $Y = bX$.¹⁶ By making X and Y proportional, this

¹⁶ Equivalently, one could treat Y as a proportional add-on to X , and combine them into a single spreadsheet cell: $Z = X + Y = X + bX = (1 + b)X$.

strategy implicitly assumes a maximal correlation of $\rho = 1$. Then $Var(X + Y)$ is estimated by the right-hand side of inequality (III-20), a clear overstatement. The magnitude of the overstatement is suggested in Figure III-7. The correlation approaches its maximal value of $\rho = 1$, justifying this strategy, only when σ_u is small relative to the normalizing term $b \times \sigma_x$.

We conclude that a linear modeling strategy leads to an overstatement of $Var(X + Y)$, although the overstatement diminishes under the conditions just indicated. Alternatively, a strategy of carrying X and Y as distinct but correlated cost elements leads to an exact estimate of $Var(X + Y)$ (assuming, of course, that the underlying quantities are measured exactly).

2. Global Schedule Effect

As we discussed in Chapter II, the RACM developers have argued against including correlations among the basic cost estimates. Instead, RACM uses a *nonlinear* modeling strategy in its attempt to capture the effects of correlations on the variance of total cost. Specifically, they contend that: "Every instance [of correlation] brought to our attention so far has actually been a case of some outside influence which is modeled elsewhere in the simulation and wouldn't be applicable (or worse, double-counted) if assumed as the correlation asked for in the original RACM input list."¹⁷

We now investigate whether the correlations are captured by the global effects that operate uniformly on all of the WBS elements. For example, the global schedule effect serves to multiply the cost of each WBS element by the factor $1 + (f_n \sigma / m)$, which we will abbreviate henceforth as $1 + \lambda$. We showed in the discussion of equations (III-3) and (III-4) that RACM's approximation systematically understates the variance introduced by the adjustment factors. We will now argue that RACM also ignores the effect on the variance of induced correlations among the adjusted cost elements.

Consider two cost elements, X and Y , both adjusted by a common factor of λ , and assume that X , Y , and λ are all statistically independent. We are interested in the variance of the sum of adjusted costs, $Z = (1 + \lambda)X + (1 + \lambda)Y$. From basic definitions, the variance is equal to:

$$Var(Z) = Var[(1 + \lambda)X] + Var[(1 + \lambda)Y] + 2 \times Cov[(1 + \lambda)X, (1 + \lambda)Y]. \quad (III-21)$$

¹⁷ See Appendix A, p. A-33.

The variances of $(1 + \lambda)X$ and $(1 + \lambda)Y$ may be obtained *exactly* from equation (III-3). Moreover, it is not difficult to derive the covariance:

$$\text{Cov}[(1 + \lambda)X, (1 + \lambda)Y] = E(X) \times E(Y) \times \text{Var}(\lambda). \quad (\text{III-22})$$

The presence of λ in both adjusted costs induces covariance between them, in turn increasing the variance of the sum (because both adjusted costs now tend to move in the same direction). Combining equations (III-21) and (III-22), we obtain the exact variance of the adjusted sum:

$$\begin{aligned} \text{Var}(Z) = & \{ [E^2(1 + \lambda) \times \text{Var}(X)] + [E^2(X) \times \text{Var}(\lambda)] + [\text{Var}(\lambda) \times \text{Var}(X)] \} \\ & + \{ [E^2(1 + \lambda) \times \text{Var}(Y)] + [E^2(Y) \times \text{Var}(\lambda)] + [\text{Var}(\lambda) \times \text{Var}(Y)] \} \\ & + 2 \times E(X) \times E(Y) \times \text{Var}(\lambda). \end{aligned} \quad (\text{III-23})$$

Early versions of RACM underestimated the variances of $(1 + \lambda)X$ and $(1 + \lambda)Y$ by using the approximation in equation (III-4), thereby omitting the third bracketed term in the first and second lines of equation (III-23). Although this problem has been repaired, even the current version of RACM omits the covariance term, i.e., the entire third line of equation (III-23).

We reach two conclusions:

- RACM's understatement of the variance is compounded by the global schedule effect, and
- RACM's treatment of the global effects does not fully capture the correlations.

The latter conclusion is particularly serious because, again, the developers have argued against including correlations among the basic cost estimates. Yet the correlations are not handled correctly in the variance calculation either. It appears that the correlations may not be fully taken into account *anywhere* in the model.

3. Management Effect

The RACM developers have also argued that the correlations may be captured by the management effect. This effect, too, increases the cost of every WBS element, although not by a common factor. Instead, the left-hand tail of each WBS-element distribution is collapsed onto a mass point, so "high" values are drawn from each distribution, apparently inducing a correlation. The RACM developers describe this

phenomenon as: “the simultaneous ‘collapsing’ of the distributions which can occur under the self-fulfilling, ‘allocated implies spent’ premise.”¹⁸

In fact, this argument does not stand up to close scrutiny. The density function of a transformed cost element was given earlier in equation (III-12). If two cost elements, X and Y , are statistically independent prior to transformation, then their joint density function prior to transformation is simply the product of their marginal density functions:

$$g(x, y) = f_x(x)f_y(y).$$

The transformations applied separately to each cost element do nothing to break their independence, so the joint density function after transformation is simply the product of their transformed marginal density functions:

$$g^*(x, y) = f_x^*(x)f_y^*(y),$$

where $f_x^*(x)$ is the transformed density function defined in equation (III-12).

Our assertion may also be verified using simulation. We began with two independent log-normal variables, each having mean 37.5 and standard deviation 20.0. We drew 1,000 values from each distribution using Microsoft Excel’s random number generator. The empirical correlation between the two variables was 0.042; this correlation would presumably converge to zero if we were to increase the number of simulations beyond 1,000.

Next, we transformed each distribution to create a mass point at the 65th percentile, or the value 40.0. This situation was depicted earlier in Figure III-4. Finally, we computed the empirical correlation between the transformed variables. The empirical correlation was equal to 0.045, or essentially the same value observed prior to transformation. The scatterplot between the transformed variables is shown in Figure III-8. Although the mass points for both distributions are quite prominent (containing 65% of the respective probabilities), there remains basically no correlation between the two variables.

¹⁸ *Ibid.*, p. A-34.




Figure III-8. Scatterplot Between Two Transformed Cost Elements

This example illustrates that the management effect does not capture correlations among the cost elements. Once again, it does not appear that the correlations are fully taken into account *anywhere* in the model.

IV. COMMERCIAL OFF-THE-SHELF ALTERNATIVES TO RACM

This chapter evaluates commercial off-the-shelf software that can be used to obtain much of the same functionality found in RACM. As described in Chapter II, RACM fits a normal distribution to the baseline cost in each WBS element, as well as cost after application of each incremental cost shifter. The only exception is that RACM fits a beta distribution to total project cost after the transformation associated with the MAIMS effect. RACM fits the various distributions by keeping track of the means and standard deviations as each cost shifter is applied. The CDF plots are obtained by analytical calculation from the fitted distributions.

Alternatively, one could generate CDF plots through Monte Carlo simulation. Under this approach, a probability distribution is specified for each cost element in a WBS, or for the cost and duration of each project task in a Gantt chart. During any iteration of the simulation, a random draw is taken from each probability distribution. The random draws are combined to determine the total project cost or project duration for that iteration. Many such iterations are performed, and statistics are calculated by comparing across iterations. In particular, the empirical CDF may be plotted for any cost element (including the entire project), or for the cost or duration of any project task. Also, the entire process may be repeated first for the baseline, then with the addition of each incremental cost shifter.

Two of the software products discussed in this chapter, Crystal Ball and @RISK for Microsoft Excel, are add-ons to Microsoft Excel. These products replace fixed WBS-element costs with random variables. The other two software products, RISK+ and @RISK for Microsoft Project, are add-ons to Microsoft Project. These products replace fixed costs or durations of project tasks with random variables.

A. CRYSTAL BALL AND @RISK FOR MICROSOFT EXCEL

1. General Description

Crystal Ball is available from the following vendor:

Decisioneering Inc.

1515 Arapahoe Street, Suite 1311

Denver, Colorado 80202
1-800-289-2550
Fax: 303-534-4818
www.decisioneering.com

As of July 1997, the prices for Version 4.0 were \$495 for new users and \$149 for upgrades from a previous version.

@RISK for Microsoft Excel is available from the following vendor:

Palisade Corporation
31 Decker Road
Newfield, New York 14867
1-800-432-7475
Fax: 607-277-8001
www.palisade.com

As of July 1997, the price for Version 3.5e was \$395.

Neither product, taken by itself, is directly related to either proposal development or project management. They are simply tools to replace spreadsheet cells with random variables. However, Crystal Ball or @RISK for Microsoft Excel could be of considerable value if applied to a spreadsheet that depicts the cost structure of a project. As in RACM, one could design a spreadsheet in which the rows represent the cost elements in a WBS, and the columns represent the various categories of cost (e.g., labor, nonlabor; baseline, shifted; low, most-likely, high; and so on). We have already described RACM's approach of tracking means and standard deviations, then fitting a normal or beta distribution to the cost in each WBS element. We will refer to this approach as "analytical." An alternative is to estimate the CDFs empirically through Monte Carlo simulation. Either Crystal Ball or @RISK for Microsoft Excel would provide the simulation engine to generate the necessary random values.

The analytical approach suffers a slight disadvantage in terms of mathematical accuracy. The exact distribution of a sum of random variables is not generally of either the normal or beta form, although these distributions may provide tolerable approximations. As we report in Chapter III, the fitted beta distribution (including the MAIMS effect) falls within $\pm 1\%$ of a Monte Carlo simulation.

RACM uses normal distributions prior to application of the MAIMS effect. It is well-known that the sum of normal distributions is exactly another normal distribution. The sum of non-normal distributions will approach a normal distribution under conditions stated in the Central Limit Theorem (CLT) of statistics. Unfortunately, the

CLT is often applied rather cavalierly without verifying the requisite conditions. One condition is that, as the number of elements being added increases, no single element contributes more than a negligible fraction to the overall variance of the sum ("non-domination").¹ In addition, most variants of the CLT require statistical independence (implying zero correlation) among the elements being added. If some correlations are present, it may be possible to group the elements into aggregates that absorb the correlations, such that there are no correlations among the aggregates. However, one or more of the aggregates may then become too large, violating the non-domination condition.

The RACM developers recognized many of these points:

The only necessary descriptions of each element along the way from "basic" through "rate-affected" estimates are the mean and variance (the first two moments) of their distributions. When they are combined at each stage, although it is only for visibility and rough impact comparisons, the Central Limit Theorem is invoked to imply that the overall distribution is adequately approximated by a "Normal" distribution...Final adequacy judgments for the overall distributional forms displayed, if deemed necessary, should come from a full blown Monte-Carlo replication of the RACM process (which has been done on example cases as sanity/validity checks).²

On the other hand, the analytical approach is generally faster than Monte Carlo simulation. With the limitation of 30 WBS elements, the RACM spreadsheet recalculates essentially instantaneously. The recalculation includes fresh CDF plots as well as all other summary statistics. Depending, of course, on the sample size drawn, Monte Carlo simulation of a similarly sized spreadsheet may take several minutes. This point was not lost on the RACM developers:

The RACM model is able to provide an analysis in a matter of seconds. The Monte Carlo models will take an hour or more on some of the large programs. This speed is essential for one very good reason. An understanding of any model's forecast requires many "runs" of the model using different inputs. This allows management to understand better what elements in the WBS should be watched and examined further. If each

¹ William Feller, *An Introduction to Probability Theory and Its Applications*, Volume II, New York: John Wiley, 1971, p. 262.

² See Appendix A, p. A-20.

analysis takes an hour or more, there will not be enough time to make all of the runs necessary for a full understanding of the modeling results.³

Although Monte Carlo simulation is definitely slower, we believe that the RACM developers have overstated their case. Simulation of a WBS at the RACM level of detail should take at most a few minutes, not "an hour or more" as they claim.

Substitution of Monte Carlo simulation for RACM's analytical approach would not obviate the need to develop a spreadsheet representing the WBS and the various categories of cost. Nor would it ease the process of eliciting the input from technical experts. However, the Monte Carlo software products offer two distinct advantages over the RACM's approach. The first advantage is a more thorough treatment of correlations among cost elements. The second advantage is a more automatic facility to produce incremental overlay charts. Although correlations and overlay charts have been available in RACM, their implementation is not as attractive as that offered by the Monte Carlo software.

2. Correlations

In Chapters II and III, we discuss the philosophical issues surrounding correlations among cost elements. Early versions of RACM included a cell range in which the user could enter correlations, if desired. The RACM developers later deleted this cell range on the grounds that explicit correlations were inappropriate to cost modeling. It would be a simple matter to reinstate this cell range to accommodate those who feel comfortable using explicit correlations.

One potential pitfall in using correlations is that the correlation matrix must be logically consistent; in mathematical terms, it must be *positive definite*.⁴ For example, the following matrix is *not* positive definite:

$$\begin{vmatrix} 1.0 & .9 & .9 \\ .9 & 1.0 & .5 \\ .9 & .5 & 1.0 \end{vmatrix}.$$

³ *Ibid.*, p. A-15.

⁴ The requirement of positive definiteness is discussed, for example, in Matthew S. Goldberg, "Some Fallacies in Cost-Risk Analysis," *PHALANX: The Bulletin of Military Operations Research*, September 1996; or in Philip M. Lurie and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions," *Management Science*, Vol. 44, No. 2, February 1998.

Costs elements #1 and #2 have a substantial correlation of .9, as do cost elements #1 and #3. Thus one would expect cost elements #2 and #3 to have a high correlation as well. The displayed correlation of only .5 is logically inconsistent.⁵ It turns out that the smallest possible correlation between cost elements #2 and #3 in this situation is .619.

The early versions of RACM allowed correlations but did not test for logical consistency. The test may be performed either using eigenvalues or, what is computationally much simpler for low-dimensional problems, using the principal minors of the correlation matrix.⁶ One merely checks that the 2×2 and 3×3 subdeterminants (in general, through the $k \times k$ determinant) anchored in the upper-left corner of the matrix are both positive. The matrix shown above fails the test:

$$\begin{vmatrix} 1.0 & .9 \\ .9 & 1.0 \end{vmatrix} = .190 > 0, \quad \begin{vmatrix} 1.0 & .9 & .9 \\ .9 & 1.0 & .5 \\ .9 & .5 & 1.0 \end{vmatrix} = -.060 < 0.$$

A much more difficult problem is how to adjust the correlation matrix if the original matrix turns out to be logically inconsistent. Both Crystal Ball and @RISK for Microsoft Excel test the correlation matrix for logical consistency and, if it fails the test, apply the Davenport and Iman algorithm.⁷ Appendices E and F show the instructions and associated screen captures for the two respective software products.

The Davenport and Iman algorithm potentially adjusts *all* of the elements in the correlation matrix (except, of course, the unit diagonals) in an effort to achieve logical consistency. In the example above, Crystal Ball returns the following logically consistent correlation matrix:

$$\begin{vmatrix} 1.0 & .871 & .871 \\ .871 & 1.0 & .517 \\ .871 & .517 & 1.0 \end{vmatrix}.$$

⁵ As evidence of logical inconsistency, if cost elements X_1 through X_3 each have unit variance, then the linear combination $.772 X_1 - .449 X_2 - .449 X_3$ has a negative "variance" of $-.047$.

⁶ See Matthew S. Goldberg, "Some Fallacies in Cost-Risk Analysis." A primary source on principal minors is George Hadley, *Linear Algebra*, Reading, Massachusetts: Addison-Wesley, 1961, pp. 259–263.

⁷ James M. Davenport and Ronald L. Iman, "An Iterative Algorithm to Produce a Positive Definite Correlation Matrix from an Approximate Correlation Matrix," Technical Report SAND-81-1376, Sandia National Laboratories, 1981.

Logical consistency may be verified using the principal minors:

$$\begin{vmatrix} 1.0 & .871 \\ .871 & 1.0 \end{vmatrix} = .242 > 0, \quad \begin{vmatrix} 1.0 & .871 & .871 \\ .871 & 1.0 & .517 \\ .871 & .517 & 1.0 \end{vmatrix} = .00026 > 0.$$

@RISK for Microsoft Excel, with apparently a slightly different implementation of the Davenport and Iman algorithm, returns instead the following (also logically consistent) correlation matrix:

$$\begin{vmatrix} 1.0 & .860 & .860 \\ .860 & 1.0 & .478 \\ .860 & .478 & 1.0 \end{vmatrix}.$$

Yet another problem arises if the user has different degrees of confidence in the original correlations. For example, suppose that the user is relatively certain of the correlation between cost elements #1 and #2, as well as the correlation between cost elements #1 and #3. However, suppose that the user is relatively uncertain of the estimate (.5) for the correlation between cost elements #2 and #3. In this situation, the user might prefer an algorithm that leaves the first two correlations unchanged, but changes only the third correlation (i.e., the estimate .5) in order to achieve logical consistency.

Although the Davenport and Iman algorithm does not meet this objective, the more recent Lurie and Goldberg algorithm does.⁸ The latter algorithm offers a weighting scheme by which the user expresses the relative confidence in each of the original correlations. Applying the Lurie and Goldberg algorithm with all of the weights equal yields an adjusted correlation matrix that agrees (to three decimal places) with Crystal Ball's implementation of the Davenport and Iman algorithm:

$$\begin{vmatrix} 1.0 & .871 & .871 \\ .871 & 1.0 & .517 \\ .871 & .517 & 1.0 \end{vmatrix}.$$

⁸ Philip M. Lurie and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions."

However, by applying one thousand times as much weight to the two relatively certain correlations, the following adjusted matrix is obtained instead:

$$\begin{vmatrix} 1.0 & .900 & .900 \\ .900 & 1.0 & .619 \\ .900 & .619 & 1.0 \end{vmatrix}.$$

This matrix leaves the two relatively certain correlations unchanged (to three decimal places), and achieves the objective of logical consistency by changing only the one relatively uncertain correlation. In fact, the uncertain correlation of .500 is replaced with the value .619 which, as we pointed out earlier, is the smallest possible correlation in this situation.

The Lurie and Goldberg algorithm is not available in either Crystal Ball or @RISK for Microsoft Excel, although the developers of the algorithm have offered it to both software vendors. The algorithm is not difficult to implement, and the one complex calculation (a nonlinear minimization step) can be performed using the “Solver” feature that is already bundled into Microsoft Excel. The Lurie and Goldberg algorithm should be included in any serious risk-analysis package.

Finally, it should be mentioned that both Crystal Ball and @RISK for Microsoft Excel generate random values with specified Spearman or “rank” correlation, not the more familiar Pearson or “product-moment” correlation. Figure IV-1 illustrates the difference between these two types of correlation. The figure actually depicts four different data sets. All four data sets share the following two points in common: (1.0, 1.0) and (2.0, 2.0). However, the four data sets differ in the final point that they contain. One of the data sets (the one identified as “rho = 1.000”) contains a final point (3.0, 3.0) that lies along the line determined by the first two points. Thus the Pearson correlation equals 1.000 for this data set. However, the Spearman correlation equals 1.000 for *all four* data sets, even though the final points do not lie along the line and the Pearson correlations are as low as 0.904. The Spearman correlation merely indicates that higher X-values are associated with higher Y-values, but it does not require that the relationship be linear.




Figure IV-1. Comparison of Spearman and Pearson Correlations

A difficulty arises in eliciting the Spearman correlation from technical experts. Many experts would look at the three data sets depicted in Figure IV-1 (other than the one identified as " $\rho = 1.000$ ") and conclude, based on the lack of linearity, that the "correlation" is less than 1.000. However, both Crystal Ball and @RISK for Microsoft Excel, which employ the Spearman correlation, expect the user to input precisely the value 1.000 in this situation.

The software vendors have been using the Spearman correlation because, until recently, there was no established algorithm to generate random values with specified Pearson correlation. Instead, they used the Iman and Conover algorithm to generate random values with specified Spearman correlation.⁹ Recently, Lurie and Goldberg have developed an algorithm that works with the more familiar Pearson correlation.¹⁰ They have also offered their algorithm to both software vendors. Adoption of Lurie and Goldberg's algorithm would circumvent the difficulties in eliciting Spearman correlations. Unlike their algorithm for adjusting the correlation matrix, Lurie and

⁹ Ronald L. Iman and W. J. Conover, "A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables," *Communications in Statistics—Simulation and Computation*, Vol. B11, No. 3, 1982, pp. 311–334.

¹⁰ Philip M. Lurie and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions."

Goldberg's simulation algorithm requires intricate programming in a scientific language such as FORTRAN. Thus, unless it is adopted by the software vendors, it would be difficult for individual spreadsheet modelers to implement this algorithm.

3. Incremental Overlay Charts

In Chapter II, we discuss RACM's ability to produce incremental overlay charts. RACM keeps track of the mean and standard deviation of every cost element as each incremental cost shifter is applied. It then performs analytical calculations to plot the CDF of baseline project cost, baseline plus the first cost shifter, baseline plus the first two cost shifters, baseline plus the first three cost shifters, etc. RACM stores the incremental information in contiguous, side-by-side regions of the spreadsheet. The plot range for Microsoft Excel is the union of columns selected from within these contiguous regions. An example of a RACM overlay chart is shown in Figure II-9.

We have been discussing the feasibility of substituting Monte Carlo simulation for RACM's analytical approach. In so doing, it would be desirable to maintain the ability to produce incremental overlay charts. One possibility would be to preserve essentially the entire RACM spreadsheet structure, but simply replace the analytical calculations with function calls to either the Crystal Ball or @RISK for Microsoft Excel simulation engine. Alternatively, the overlay features built into these two software products could be used to automate the process so that the spreadsheet design could be streamlined somewhat.

The latter approach is demonstrated in Appendices I and J, which show the instructions and associated screen captures for the two respective software products. The final screens are reproduced here as Figure IV-2 and Figure IV-3. The only requirement for overlay charting is creation of a summary region in the spreadsheet. This region must contain the distributional parameters associated with the CDF of total cost after application of each incremental cost shifter. If RACM's distributional assumptions were maintained, total cost at each stage would be approximated by a normal distribution, for which the necessary parameters are simply the mean and standard deviation. The only exception would be the final stage, representing the MAIMS effect. Total cost at that stage would be better approximated by a beta distribution, requiring low and high endpoints in addition to the mean and standard deviation. All of these quantities are already computed within RACM, so the extra effort to array them in a summary region is negligible.

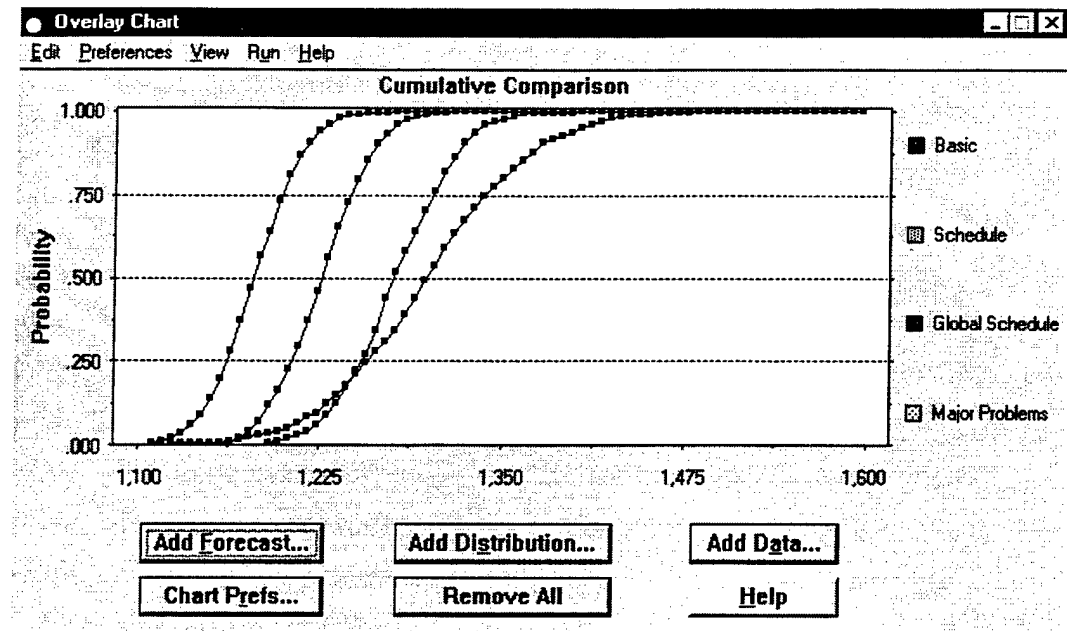


Figure IV-2. Crystal Ball Incremental Overlay Chart

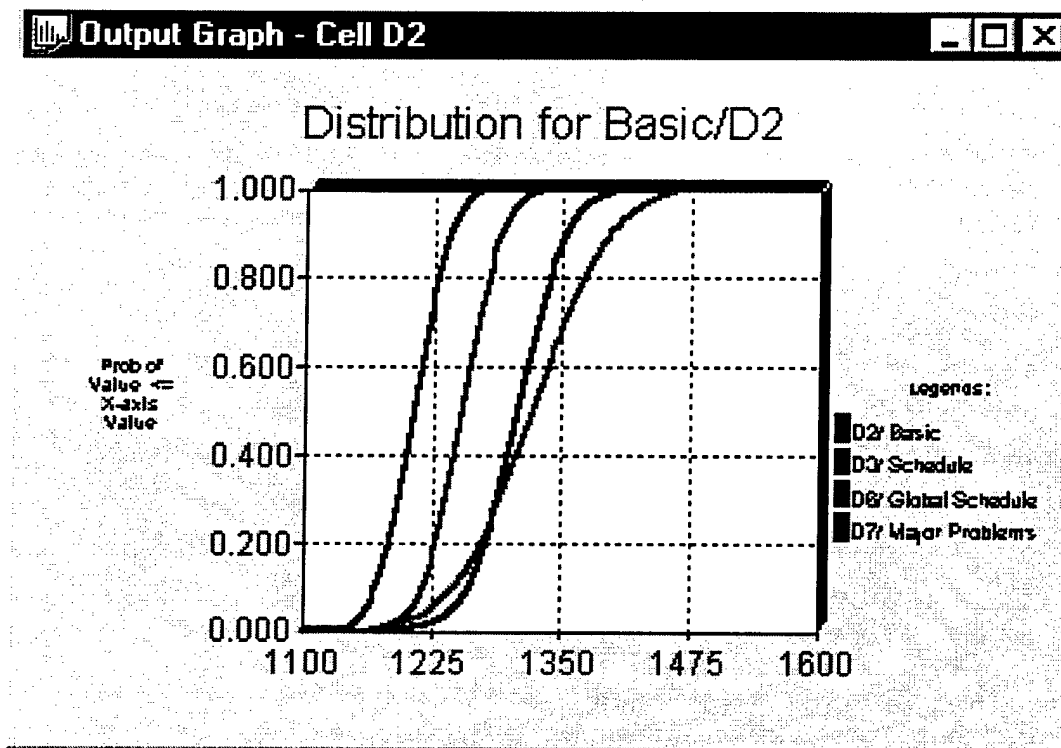


Figure IV-3. @RISK for Microsoft Excel Incremental Overlay Chart

4. Summary

Substitution of Monte Carlo simulation for RACM's analytical approach would yield few benefits. It would be a simple matter to reinstate correlations within RACM. The test for logical consistency of the correlation matrix could easily be programmed into the spreadsheet. Lurie and Goldberg's algorithm for adjusting an inconsistent correlation matrix could also be programmed into the spreadsheet, though it involves more than mere recalculation of a cell formula. Instead, the user would have to explicitly execute the "Solver" feature of Microsoft Excel. However, it may be possible to design an Excel macro function to perform the latter operation.

The two software products under consideration contain the automatic facility to produce incremental overlay charts. However, the spreadsheet must still be designed to keep track of the distributional parameters after application of each incremental cost shifter. Moreover, customization of the overlay charts (e.g., rescaling the axes, or modifying the axis labels) may be easier (or, at least, more familiar) when the charts are drawn directly within Microsoft Excel, rather than within the add-on software.

Add-ons to Microsoft Project, as opposed to Microsoft Excel, bring a scheduling dimension into play that is absent from both RACM and the simulation-driven variations thereof. We consider RISK+ in the next section, and @RISK for Microsoft Project in the concluding section of this chapter.

B. RISK+

1. General Description

RISK+ is an add-on to Microsoft Project. Specifically, we evaluated the version of RISK+ that is designed to run with Microsoft Project version 4.1.¹¹ Many of the observations we make here have already been noted by Book, Blackshire, and Young in their evaluation of an earlier version of this product.¹²

¹¹ Program Management Solutions, Inc. (PMSI), *RISK+: Risk Analysis for Microsoft Project 4.1*, Redondo Beach, California, September 1996.

¹² Stephen A. Book, O. F. Blackshire, and Philip H. Young, "Validation Report on RISK+ Risk Modeling Software for Microsoft Project 4.0 (U)," The Aerospace Corporation, Los Angeles, California, October 1995, Contractor-Proprietary.

RISK+ is available from the following vendor:

Program Management Solutions, Inc. (PMSI)
553 N. Pacific Coast Highway
Suite B-177
Redondo Beach, California 90278
(310) 374-0455
Fax: (310) 374-2090
www.cs-solutions.com

As of June 1998, the price of RISK+ was \$349.

Microsoft Project is itself a well-known product for modeling project schedules and resource consumption. RISK+ is a seamless add-on that replaces the costs or durations of project tasks with random variables. RISK+ uses Monte Carlo simulation to propagate these random variables throughout the entire project. It then constructs statistical summaries (means, standard deviations, percentiles, etc.) of total project cost and project duration.

RISK+ produces three major graphical outputs:

- Estimated Cost Histogram;
- Estimated Completion Date Histogram; and
- Risk Gantt Chart, a custom view added to Microsoft Project. In particular, a field is displayed to the left of each Gantt bar indicating the percentage of time that a task spent on the critical path during the risk-analysis simulation.

The Cost Histogram is produced for only a single "preview task," most logically chosen as the entire project. The Completion Date Histogram is produced for the preview task plus as many additional "reporting tasks" as desired, conceivably including every task in the project.

2. Critique

Book, Blackshire, and Young reported roundoff problems when project durations were specified in units other than minutes (the smallest unit accepted by Microsoft Project):¹³

The Microsoft Project scheduling engine, which does all scheduling calculations, accepts numerical inputs representing durations of activities

¹³ *Ibid.*, p. 2.

comprising a given schedule in units of minutes, hours, days, or weeks. However, all these units are first translated into minutes, in which units all calculations are then processed. After processing, results of the calculations are then translated back to the original input units, *rounded to the nearest input unit*, and then the results are displayed. [Emphasis in original.]

Although this procedure seems correct, apparently the early implementation was deficient. According to PMSI, these problems have since been fixed.

RISK+ does not allow correlations among the costs of various project tasks, nor among their durations.

The network structure of the project (precedence relationships, leads, lags, etc.) has no effect on the final cost outcome; total project cost is simply the sum of the (simulated) costs across all of the project tasks. Nor is there any correlation between cost and duration for a particular task. One extreme assumption is that tasks consume resources at a roughly constant daily "burn rate," so that we would expect an above-average duration to be strongly correlated with an above-average cost. We have questioned the validity of this assumption in Chapter III, and we return to this point again in Chapter V. However, even if the user accepts the assumption of constant burn rates, RISK+ contains no provision for modeling this effect.

One of the most important features of Microsoft Project is its assignment of resources to tasks. Each task consumes resources, and resource-consumption rates (per unit time) are multiplied by task duration to determine task cost, which is then summed over all tasks. RISK+ seems to bypass this entire process. Instead, RISK+ asks the user to supply low, most-likely, and high estimates of task cost, which are then simulated. Thus task costs are simulated directly, rather than being built up from the underlying resources. Microsoft Project contains several useful algorithms for leveling resources across time to avoid bottlenecks. These algorithms seem to be ignored or even nullified by RISK+.

Like Microsoft Project, RISK+ allows each task to be followed by only a single successor task. A possible alternative, for which RISK+ contains no provision, is called *conditional branching*. Under that approach, a task may be followed by one of several successor tasks. The particular task that follows would be determined by a probability mechanism (metaphorically, a roll of the dice) during each Monte Carlo iteration.

RISK+ does not have any mechanism for producing incremental overlay charts, as were demonstrated earlier for both Crystal Ball (Figure IV-2) and @RISK for Microsoft Excel (Figure IV-3). A cumbersome solution would be to first create multiple input data

sets (baseline schedule and costs, baseline plus the first cost shifter, baseline plus the first two cost shifters, baseline plus the first three cost shifters, etc.), then run RISK+ on each data set, and finally export the results to a common Microsoft Excel workbook and manually create an overlay chart. We view the inability to automatically produce overlay charts as a major drawback of RISK+.

3. Distributional Forms

Task cost and task duration are specified in terms of a low, most-likely, and high value. Four distributional forms are allowed for either cost or duration: uniform, triangular, normal, or beta. RISK+ uses the following procedure to fit a normal distribution into the finite interval defined by the user-supplied low and high values. First, they ignore the most-likely value, and estimate the mean instead as the simple midpoint between the low and high values: "Note that the most-likely value of duration or cost is not used—the peak of the distribution is always at the center of the interval specified by the minimum and maximum duration or cost."¹⁴ This procedure may be problematic if the user deliberately supplies a most-likely value much different from the midpoint. Although such a constellation of user input is not well modeled by a normal distribution (which is necessarily symmetric), RISK+ issues no warning messages to that effect. Next, recognizing that over 99% of the probability for a normal distribution falls within the mean ± 3 standard deviations, they treat the difference between high and low values as an estimate of 6σ . Thus they estimate $\sigma = (h - l) / 6$.

Their procedure for fitting a beta distribution is even more suspect. They note that the beta distribution is described by four parameters. Two of these parameters can be immediately identified with the user-supplied low and high endpoints. They then assert that the mode (i.e., the most-likely value) depends upon only the ratio of the two remaining parameters: "It can be shown that the position of the mode depends on the ratio α / β ."¹⁵ This assertion is false. We give the formula for the mode in equation (C-4) of Appendix C. To use their own example, suppose that $l = 0$ and $h = 1$. Parameter values of $\alpha = 2$ and $\beta = 5$ and yield a mode of 0.200, whereas parameter values of $\alpha = 20$ and $\beta = 50$ (which have the same ratio) yield a different mode of 0.279. They set one of the

¹⁴ PMSI, *RISK+: Risk Analysis for Microsoft Project 4.1*, p. 85. If we denote the low and high values as l and h , respectively, the manual gives the formula $\mu = (h - l) / 2$. This formula contains a typographical error; they almost certainly intended to display instead the midpoint formula, $\mu = (h + l) / 2$.

¹⁵ *Ibid.*, p. 85; emphasis in original.

the unknown parameters (either α or β) to the value 6.0, and they determine the other parameter so that the ratio α/β , to paraphrase, "maps into the mode." Unfortunately, as the numerical example demonstrates, the ratio α/β does *not* map uniquely into the mode. Fundamentally, it is impossible to determine a 4-parameter beta distribution from only three pieces of information: the low endpoint, high endpoint, and the mode. Although the restriction that either α or β equals 6.0 provides a fourth piece of information, this restriction is completely unjustified.

We view the limited selection of distributional forms, as well as the tenuous fitting procedures for the normal and beta distributions, as major drawbacks of RISK+.

C. @RISK FOR MICROSOFT PROJECT

1. General Description

@RISK for Microsoft Project is a considerably more sophisticated add-on to Microsoft Project. It is available from the following vendor:

Palisade Corporation
31 Decker Road
Newfield, New York 14867
1-800-432-7475
Fax: 607-277-8001
www.palisade.com

A beta version of @RISK for Microsoft Project (Version 3.5) became available in September 1997 at a price of \$695. Although the beta version contains several bugs, it is adequate to demonstrate the capabilities of this product.

@RISK for Microsoft Project produces simulation statistics and graphics for any cell that the user designates. This software is much more flexible than RISK+, which produces a cost histogram for only the single "preview task." Among the statistics computed by @RISK are the minimum, maximum, mean, standard deviation, and various percentiles. @RISK also produces a variety of graphics, including a histogram and a CDF for any designated cell.

2. Distributional Forms

Like @RISK for Microsoft Excel, @RISK for Microsoft Project features a wide array of some 38 probability distributions. Among these is a triangular distribution, in which the user has the option of specifying either the mode and two endpoints, or the

mode and two percentile points. We verify in Appendix D the @RISK algorithm for determining the low and high endpoints in the latter case.

@RISK also features three versions of the beta distribution. In the first version, the user directly specifies the two shape parameters. In the second version, called *BETASUBJ*, the user specifies low and high endpoints, the mean, and the mode. We verify in Appendix C the @RISK algorithm for determining the two shape parameters. Finally, under the *PERT* distribution, the user specifies the low and high endpoints and the mode, but not the mean. Instead, the mean is estimated using the PERT formula:

$$\mu = \frac{l + 4m + h}{6}. \quad (\text{IV-1})$$

This formula has been the subject of considerable investigation and controversy, which will not be reviewed here.¹⁶ However, if the user is comfortable with the PERT formula, it is available in the software.

@RISK also features two versions of the lognormal distribution. In the first version, called *LOGNORM*, the user directly specifies the mean and standard deviation of the lognormal distribution. The software computes the mean and standard deviation of the associated *normal* distribution:

$$\mu = \ln\left(m^2 / \sqrt{m^2 + s^2}\right), \quad (\text{IV-2})$$

$$\sigma = \sqrt{\ln\left[(m^2 + s^2) / m^2\right]}, \quad (\text{IV-3})$$

where m and s are the mean and standard deviation of the lognormal distribution, and μ and σ are their counterparts for the normal distribution.¹⁷ The software simulates the normal distribution, then applies an anti-logarithmic transformation to recover the lognormal variable. The second version of the function, called *LOGNORM2*, differs only in that the user specifies the mean and standard deviation of the normal (rather than lognormal) distribution. The first version seems generally more useful.

¹⁶ A recent citation is Donald L. Keefer and William A. Verdini, "Better Estimation of PERT Activity Time Parameters," *Management Science*, Vol. 39, No. 9, September 1993, pp. 1086–1091.

¹⁷ Palisade Corporation, *Guide to Using @RISK for Project: Risk Analysis and Simulation Add-In for Microsoft Project*, Newfield, New York, July 1997, p. 207.

3. Correlations

@RISK for Microsoft Project allows for correlation among all input distributions:

During a simulation analysis it is important to account for correlation between input variables. Correlation occurs when the sampling of two or more input distributions are related—for example, when the sampling of one input distribution returns a relatively “high” value, it may be that the sampling of a second input should also return a relatively high value. A good example is the case of one input applied to the *Duration* field of a task named *Construction*, and a second input applied to the *Cost* field of the same task. There may be a distribution for each of these input variables, but the sampling of them should be related to avoid nonsensical results. For example, when a relatively high (or long) *Duration* is sampled, *Costs* should be sampled as relatively high. Conversely, you would expect that when *Duration* is short (i.e., the value sampled is low), *Costs* should be relatively low.¹⁸

@RISK allows for correlation between the cost and duration of the same task, as would be the case if there were a roughly constant burn rate. It also allows for correlation between the costs of different tasks, the durations of different tasks, or indeed the cost of one task and the duration of a different task. Like its counterpart for Microsoft Excel, @RISK for Microsoft Project tests for logical consistency and, if necessary, applies the Davenport and Iman algorithm to adjust the correlation matrix.

4. Incremental Overlay Charts

@RISK for Microsoft Project has the same capability as @RISK for Microsoft Excel to produce incremental overlay charts:

@RISK lets you see the impact of uncertain project parameters on your results. But what if some of the uncertain project parameters are under your control? In this case the value a variable will take is not random, but can be set by you. For example, you might need to choose between a set of possible Start Dates, different possible sequences of tasks, or from a set of possible labor or cost values. To properly analyze your project, you need to run a simulation at each possible value for the “user-controlled” variables and compare the results. A Sensitivity Simulation in @RISK allows you to quickly and easily do this—offering a powerful analysis technique for selecting between available alternatives.¹⁹

¹⁸ *Ibid.*, p. 44.

¹⁹ *Ibid.*, p. 73. A related discussion is contained on p. 174

The overlay capability within @RISK for Microsoft Project is demonstrated in Appendix K. There we construct an example in which a particular project task starts on one date under Scenario #1, but on a later date under Scenario #2. The two start dates are nonstochastic, but different under the two scenarios. We also specify the durations of various other tasks as random variables whose distributions are the same under the two scenarios. We cycle through a series of steps in the software, culminating in a side-by-side display of the CDFs of project completion date under the two scenarios. That display is reproduced here as Figure IV-4.

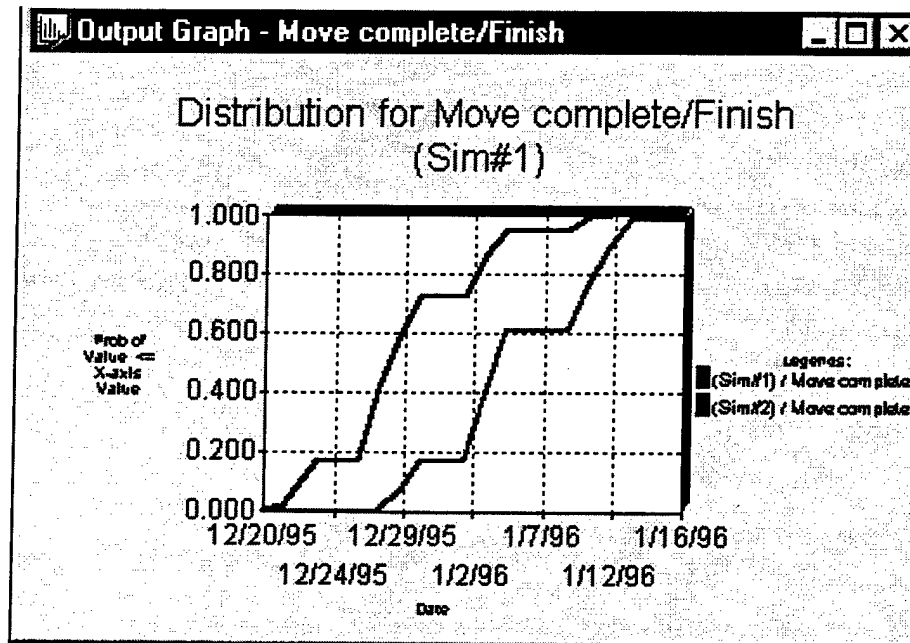


Figure IV-4. @RISK for Microsoft Project Incremental Overlay Chart

Although @RISK's overlay capability is useful, it is somewhat limited. @RISK's *SIMTABLE* function accepts the input values under the various scenarios; in our example, the function call is *SIMTABLE*({10/26/95,11/2/95}). Because both start dates are nonstochastic, the two project schedules tend to be offset by a constant amount. However, the constant offset may be moderated somewhat if the delay in starting the project under Scenario #2 is partially absorbed by other tasks that have slack time relative to the critical path. Thus, we see in Figure IV-4 that the two CDFs are not quite parallel, particularly in the 0.60 to 0.75 probability range.

A more useful capability, not currently available, would be to allow random variables to have different *distributions* (not just different fixed values) under the two scenarios. For example, one might wish to specify the duration of a particular task as a triangular distribution with $(l, m, h) = (1, 2, 3)$ under Scenario #1, but as a triangular distribution with $(l, m, h) = (1, 2, 4)$ under Scenario #2. This specification could represent the situation described earlier in which the j th scenario includes the baseline schedule and costs, plus $j - 1$ incremental cost shifters.

The @RISK for Microsoft Project manual hints at this capability, though no examples are given:

The benefits of Sensitivity Simulation are not limited to evaluating the impacts of user-controlled variables on simulation results. A sensitivity analysis can be run on the probability distributions which describe uncertain variables in your project. *You may wish to repetitively re-run a simulation, each time changing the parameters of one (or several) of the distributions in your project.* After all the individual simulations are complete, you can then compare the results from each.²⁰ [Emphasis added.]

We attempted to use the *SIMTABLE* function in this fashion, specifying *SIMTABLE*(*{TRIANG*(1, 2, 3),*TRIANG*(1, 2, 4) }) or, alternatively, specifying *TRIANG*(1, 2, *SIMTABLE*(*{3, 4}*)). Unfortunately, the software did not interpret these functional calls as we had intended. Instead, the results were unpredictable and ultimately useless.

We view the inability to produce overlay charts with alternative probability distributions as a major drawback of @RISK for Microsoft Project. In this regard, it represents only a minor improvement over RISK+. As with RISK+, a cumbersome manual solution is always available. Specifically, one could create multiple input data sets, run @RISK for Microsoft Project on each one, export the results to a common Microsoft Excel workbook, and manually create an overlay chart.

5. Summary

As we discussed in Chapter III, RACM accounts for schedule effects only to the limited extent of counting the numbers of WBS-element-specific and global parallel paths. Schedule effects would not be modeled with any greater fidelity under variations on RACM that merely substitute Monte Carlo simulation for analytical CDF calculation.

²⁰ *Ibid.*, p. 73.

By contrast, schedule modeling is the *raison d'être* for Microsoft Project. For the narrow purpose of replacing the costs or durations of project tasks with random variables, the @RISK add-on to Microsoft Project is satisfactory.

On the other hand, @RISK for Microsoft Project was designed to model random variables whose distributions are the same under all scenarios. @RISK was not designed to model a situation in which project costs or durations have different distributions under two or more scenarios. Yet the latter is precisely the situation modeled by RACM: alternative distributions that incrementally account for the various sources of risk that drive total project cost. Thus, we face a dilemma in which RACM provides only a perfunctory treatment of schedule effects, yet @RISK for Microsoft Project does not adequately model alternative scenarios.

We conclude that no single tool integrates schedule effects with RACM's scenario structure. Thus, for the immediate future, schedule modeling must remain essentially distinct from scenario-driven modeling of a WBS. The integration of cost and schedule modeling is an important area for additional research.

V. RACM FIELD TEST AT BOEING D&SG

A. DESIGN OF THE FIELD TEST

RACM was developed by a small team of analysts at Lockheed Martin's Sunnyvale plant. An important issue is the portability and applicability of RACM at other contractors' plants. To address this issue, we arranged for a field test at Boeing D&SG in Seattle, Washington.¹ Unfortunately, timing did not permit the use of RACM either on an active proposal-writing effort, or on management of an existing contract. However, we were able to convince Boeing D&SG to retrofit RACM on a proposal that they had just completed. Their experience in attempting to do so is fully documented in Appendix G. This chapter summarizes the most salient points from Boeing's report to IDA and associated dialogue between the two organizations.

Two IDA representatives visited Boeing D&SG in November 1996. We presented a briefing on the RACM concept and software implementation. We portrayed the RACM concept as promising, though we were candid about the limitations of the existing software implementation. With permission from Lockheed Martin, we left copies of the software at Boeing, and we encouraged them to find an application on which to test it. Boeing D&SG assigned their Manager of Statistical Analysis and Simulations, Estimating, as our point of contact.

Boeing attempted to replicate, using RACM, a risk analysis they had just completed on a high-visibility proposal: the pre-Boeing merger Evolved Expendable Launch Vehicle (EELV). In particular, Boeing had modeled life-cycle costs (LCC) for the Engineering and Manufacturing Development (EMD) effort on EELV. We agreed on the following general criteria for evaluating the applicability of RACM to the EELV:

- software ease-of-use,
- ability to generate the RACM input data,

¹ During the period in which we were working with Boeing, it was announced that they would merge with McDonnell Douglas. As an outcome of the merger activity, Boeing D&SG was reorganized as Boeing Information, Space, and Defense Systems (ISDS). We use the designation D&SG throughout this report to mean both D&SG and ISDS.

- numerical comparison between RACM output and the Boeing D&SG internal cost/risk analysis, and
- validity of the RACM assumptions at the Boeing plant.

B. SOFTWARE EASE-OF-USE AND GENERATION OF INPUT DATA

The Boeing team had developed risk estimates at WBS levels 4 and 5. However, there were approximately 100 cost elements at that level, far exceeding the current RACM limitation of 30 cost elements. To facilitate the comparison, Boeing collapsed their WBS structure into 29 cost elements at level 3. Boeing was able to obtain some of the RACM input directly from the element-specific cost distributions that had already been estimated at WBS level 3. In the remaining cases, Boeing had to develop new input from scratch: "For those RACM elements not corresponding to the EELV structure, reasonable estimated values were developed. (The RACM documentation did not provide guidance in this area.)"² The lack of documentation is a common theme throughout Boeing's report: "Additional inputs include, for each labor and non-labor element, Schedule Slip,...Improvement,...Wrap rate growth,...Major Problems,...Schedule Parallelism,...Bid probability, Budget Allocation level, and Management Reserve level. These inputs are not further defined by the RACM documentation."³

In the RFP, the Air Force EELV Project Office provided its own set of cost shifters:⁴

- **Cost Estimating:** Uncertainties in cost due to reliance on available estimating methodologies.
- **Schedule:** Assumes the item being acquired is a developmental item and its schedule meets program goals. The assessment focuses on the adequacy of the time specified for the item relative to schedules for similar items.
- **Supportability:** An evaluation of how well the composite of support considerations necessary to achieve the effective and economical support of a system for its life cycle meets stated quantitative (e.g., Mean Time Between Failures (MTBF)) and qualitative readiness and utilization requirements [*sic*]. This includes integrated logistics support and logistic support resources related Operating and Support (O&S) cost considerations.

² See Appendix G, p. G-9.

³ *Ibid.*, p. G-3.

⁴ See Annex G-2, "EELV Boeing Risk Evaluation Methodology (pre-Boeing merger)."

- **Technology:** Uncertainties to system performance due to reliance on the availability and promise of technology. Technology uncertainty includes the required level of technological sophistication and reflects the current stage of hardware development and testing maturity. Hardware maturity ranges from scientific research, conceptual design, brassboard, breadboard, prototype, to an operational unit.
- **Design and Engineering:** Uncertainties to system performance due to uncertainties and variability in design and engineering process. Design & Engineering uncertainty reflects the degree of difficulty to advance the current state of the art for a given item (e.g. subsystem) to the required (e.g. qualified off-the-shelf item that meets all requirements).
- **Manufacturing:** Uncertainties associated with the production elements used to manufacture the required quantities of an item, within the technical specifications.

The Boeing team had never encountered this particular set of cost shifters prior to receiving the EELV RFP. Although there is some ambiguity in the above definitions, the Boeing team managed to structure their proposal around this set of cost shifters: "The meanings of the cost risk categories, Cost Estimating, Schedule,...and the best and worst case values were left to the engineers making the judgment."⁵

During the RACM retrofit, the Boeing team had great difficulty aligning the EELV cost shifters with the ones hard-wired into RACM. They set the RACM cost improvement and rate increase cost shifters to zero because these effects were already included in their baseline cost estimates. They were unable to back these effects out of the baseline to produce the "bare-bones" baseline expected by RACM. Nor was Boeing able to directly estimate the global schedule or major problem effects. It appears that they just made up numbers to test the latter two features of RACM: "The EELV risk estimating procedure does not have inputs corresponding to the Global Schedule and Major Problem RACM inputs. Due to the intended depth of comparison, an attempt to develop inputs for these RACM risk contributors through the EELV organization was not made. However, to demonstrate RACM's sensitivity to these parameters, values were input."⁶

⁵ See Appendix G, p. G-5.

⁶ *Ibid.*, pp. G-5 to G-6.

C. NUMERICAL COMPARISON AND VALIDITY OF RACM ASSUMPTIONS

Notwithstanding these difficulties, the Boeing team proceeded to make a numerical comparison between their own risk estimates and those produced using RACM. Their comparison is summarized in Table G-3 of Appendix G. Although the numerical results of the Boeing and RACM calculations are somewhat different, the estimates are hardly commensurable and we do not believe there is much to be learned by dissecting the differences. Instead, we believe it is more illuminating to examine the Boeing team's reactions to RACM's assumptions. These reactions are contained in Appendix G, in particular Annex G-3, "Additional Dialogue Between IDA and Boeing D&SG." We summarize the most salient points in this section.

1. Applicability to Proposal Preparation

We asked Boeing to compare the applicability of RACM in two possible domains, proposal preparation and management of existing programs. Boeing responded that the MAIMS effect may invalidate RACM for proposal preparation:

Fundamental to RACM is the concept of the spending of budgeted resources. Internally, within the company, prior to the submittal of a proposal and when one is attempting to establish what the PDF (Probability Density Function) of the program's future cost is, the budget value is not known. *What is needed at this time is an estimate of the PDF that is not conditioned upon a budget value.* If it is true that a project's budget affects its ultimate cost then that effect, well modeled, would certainly be useful in our efforts to both plan and manage a program. RACM's modeling of budget effects is one of its features, that feature makes it more applicable to programs that have budgets.⁷ [Emphasis added.]

We should note that it is a simple matter to back the MAIMS effect out of RACM and still obtain useful results. As Figures II-7 and II-8 illustrate, RACM applies the various cost shifters incrementally. Because the MAIMS effect is the last one applied in the sequence, the CDF of total project cost without this effect is readily available. Thus, RACM could still be used to model the effects of cost improvement, rate increase, and so on, provided the input data could be arrayed according to these categories.

⁷ *Ibid.*, p. G-21.

2. Applicability to Program Management

Notwithstanding the previous quotation, Boeing D&SG denied the existence of the MAIMS effect in their plant. Thus, they argue that RACM may not be valid for management of existing Boeing programs either:

The cost management mechanisms in place at Boeing largely preclude the "money allocated [is] money spent" phenomenon. Budgets are allocated, but at a tight, "motivational" level. The balance-to-go budgets may be reallocated based on the balance-to-go effort remaining to accomplish the plan....Budgets are "held, authorized, monitored, and controlled" at either the level where their use is initiated or often at some one or two levels higher. One effect of this is that the initiator of a cost may not have a budget to "live up to."⁸ [Emphasis added.]

Once again, it would be simple to back the MAIMS effect out of RACM. In this instance, however, little utility remains from using RACM without the MAIMS effect. Recall from Chapter III that, when the MAIMS effect is operative, the optimal budget allocation gives each WBS-element manager the same percentile of his or her respective cost distribution. RACM computes the common percentile for each WBS element that is consistent with a target percentile for total project cost. If the MAIMS effect is turned off, then the equi-percentile budget allocation is no longer necessarily optimal, so there is little theoretical justification behind RACM's calculations.

3. Standing Armies

In Chapter III, we discuss RACM's assumption that the burn rates are constant, so parallel activities continue to accumulate costs at a constant rate while waiting for the slowest activity to finish. This assumption may also be described as level-loaded staffing for the duration of an activity or, more picturesquely, as a "standing army": "[A parallel path] is any one individual time-phase element or serial combination of elements that could cause the WBS element in question to not complete its task within the schedule predicted—cause the holdup of the entire element until it is completed. In the global sense, it is any WBS element which could cause the holdup of the entire program (the 'Standing Army' problem)."⁹

⁸ *Ibid.*, p. G-24.

⁹ See Appendix A, p. A-19.

Boeing D&SG denied the existence of the standing army problem in their plant:

The staffing of a program (contract) is seldom level-loaded, perhaps never, if study contracts are excluded and depending on the definition of "level." Whether this discrepancy is important to the accuracy of the resulting PDF is unknown. However, there is a more fundamental problem. Typically budgets for direct cost are not set and maintained at the set level. Targets are set, below or well below the "contract value," performance is measured, and targets are periodically reset or reallocated commensurate with the performance achieved. In addition, the target value is just that, a target, it is generally not manned to...The idea that a program's cost can be lowered or raised, just by adjusting its budgets, was discovered long ago. That is why we have elaborate mechanisms to produce estimates, budgets, targets, performance measures. This issue could certainly have an effect, if it were not well managed. Much of the cost management effort on a program is expended to manage out the effect that poor budgeting might have.¹⁰ [Emphasis added.]

They continue:

Budget pressure and the ability to move resources to other parts of a project or to other projects or sites allow the organization to operate without "standing-armies." Boeing resources and projects are continuously managed to provide the company with the ability to move resources between requirements. For example the company balance between commercial and DoD work is at least partially intended to provide a continuity of resources despite a time-varying workload, project, and customer mix. Within the D&SG we have a policy of standardizing processes, specifically intended to allow work or resources to be shifted across locations and between projects.¹¹ [Emphasis added.]

4. Distributional Forms

The Boeing team prefers a lognormal distributional form rather than the normal distributional form assumed in RACM:

All of our cost modeling experience strongly suggests that cost PDFs are lognormal...If the cost generating process has multiplicative elements and is driven by approximately normal distributions, an approximately lognormal distribution will result. The cost generating process is at least partially multiplicative. Labor rates, overhead rates, as well as some

¹⁰ See Appendix G, pp. G-22 to G-23.

¹¹ *Ibid.*, p. G-24.

distributed costs factors are typically...more broadly based than the specific program application. They therefore become effectively multiplicative to the specific program.¹²

D. SUMMARY

The following quotation cogently summarizes the mismatch between RACM's modeling assumptions and the reality at Boeing D&SG's Seattle plant:

RACM is an Excel-based tool with a specific view of the causes and quantified effects of cost, schedule, and technical "risk." As such it, in all likelihood, fits into the cost management and estimating infrastructures from which it was created. *RACM does not mesh well with our cost estimating and cost management data sources; estimating procedures, and management requirements nor the management philosophy that I [Manager of Statistical Analysis and Simulations, Estimating] am aware of at the Boeing Company.* This point, of course, could be made about all but the most generic of approaches not invented here, and should not be taken to imply that we should not adjust to the RACM view.¹³ [Emphasis added.]

Boeing also advises against the widespread adoption of RACM until the model has been validated through field testing on an existing contract:

Besides the issues of procedural fit, RACM like any proposed estimating (forecasting) system is intended to guide us in the management of our resources. It would be imprudent to utilize such guidance without some evidence that better results could be expected through its use than otherwise. That is, it must be a valid method. Its risk estimates should be demonstrated to be related to the actual uncertainty. We do not have such a demonstration....A piloting effort should be initiated, deploying RACM on an existing contract to determine its validity and forecasting accuracy. If it is shown to be a good predictor, its forecasts should be attempted to be used in the program management process. The effects should be measured.¹⁴

However, the prospect of further field testing at Boeing D&SG seems imprudent because that organization has already declared the mismatch discussed above.

¹² *Ibid.*, p. G-23.

¹³ *Ibid.*, p. G-9.

¹⁴ *Ibid.*, pp. G-9 to G-10.

VI. CONCLUSIONS

Lockheed Martin Missiles and Space Company developed RACM to help estimate risk margins for proposals that they submit to DoD. They also included a module to help manage risk reserves during program execution.

RACM contains several useful features that all major defense contractors might consider including in their risk estimation and management tools. For example, RACM forces the contractor to account for a set of risk factors that shift the baseline distribution of total program cost. RACM also displays the effects of each risk factor sequentially, building up from the baseline program cost to the final program cost that embodies both the WBS-element-specific and global risk factors.

Our experience with Boeing D&SG suggests that it is not possible to predefine a universal set of risk factors, valid for all programs and contractors. During our field test, Boeing set two of RACM's risk factors (cost improvement and rate increase) to zero because they were unable to back them out of the baseline cost estimates. For two of the other risk factors (global schedule and major problems), it appears that Boeing did not understand the RACM definitions and simply made up numbers. It seems more fruitful to allow contractors to define their own risk factors, based on their accounting system, estimating methods, and management philosophy.

The current RACM software implementation lacks the flexibility to do that, or even to increase the number of WBS elements without considerable effort and possibility for error. One alternative implementation would link RACM to COTS software such as Crystal Ball or @RISK for Microsoft Excel. However, this option would maintain the basic spreadsheet structure with all of its attendant limitations. For example, it would still be difficult to increase the number of WBS elements, and the user would still have to manually create the templates for additional risk factors. Another option would be to rewrite RACM in a structured programming language such as Visual Basic. That option, if pursued efficiently, would eliminate many of the mechanical problems in using RACM. However, it would be difficult to design a Visual Basic implementation flexible enough to accommodate every contractor's accounting system and estimating methods.

Neither the COTS option nor the Visual Basic option would ease the process of eliciting the input from the technical experts. Moreover, the entire RACM concept is based on strong behavioral assumptions such as Money Allocated is Money Spent. Boeing D&SG denied the validity of this assumption at their plant. They also denied the "standing army" problem and the normality assumption for baseline costs.

We conclude that neither RACM nor any other particular tool can be viewed as a "silver bullet" to remove all risk or prevent all cost overruns on defense programs. We recommend instead that contractors be allowed to develop their own tools for risk estimation and management within some broad DoD guidelines. These tools could incorporate some of RACM's best features, such as its sequential view of the various risk factors. DoD might impose other general requirements, such as a set of percentile levels at which cost must be reported. The discussion of risk estimation and management in this report, along with our review of COTS software, should help guide contractors in designing their own tools.

APPENDIX A

LOCKHEED MARTIN RACM STUDY

The Under Secretary of Defense for Acquisition (USDA) was requested in 1991 to examine missile acquisition programs to (1) determine whether they are meeting their cost and schedule goals, (2) determine why some programs fail to meet these goals, and (3) recommend ways to help the acquisition program meet its goals.

To meet the objectives of the review, they selected two different groups of missiles. The first group (12 systems) identified whether current missile systems with five years of production experience had overrun cost or schedule goals. To analyze why such changes occurred, they selected a smaller group of eight missiles at different milestone decision points.

All 12 of the missile systems selected experienced cost and schedule overruns. The unit or total acquisition cost estimates for nine of these systems increased by 20 percent or more. The scheduled completion dates for all 12 systems were extended.

These overruns were attributed to many interrelated factors, some of which are not under DoD's direct control (e.g., changes in threat, congressional direction, etc.). However, optimistic planning assumptions by program officials were a common factor underlying major overruns. The USDA found that "(Government) Program offices often developed cost and schedule estimates that do not adequately reflect the risks associated with the program's design, development, and production." They "found that this is particularly true for technical risk assumptions, which often contribute to cost and schedule overruns."

Technical risk assessments can significantly affect program cost and schedule estimates. Understated program office assessments of technical risks can result in understated cost and schedule estimates. They found that "DoD's independent technical risk assessments were limited and that DoD's prior regulations had provided only limited guidance for such assessments."

The purpose of the Risk Analysis and Cost Management (RACM) process is to successfully manage a program within cost while meeting performance and schedule constraints. This is a process which ideally requires the identification, prior to proposing, of the probability of success (Ps) associated with the cost, schedule, and technical

parameters of the program, during its initial phase. The process should also be used during the conduct of the program to determine the effects on Ps by on-going design, schedule, and budgetary changes within the program. This includes the use of Design to Life-Cycle Cost (DTLCC) and other management tools. Of particular importance is the effect of budget allocation policies during the conduct of the program. This, and how the combining of Cost/Schedule Control Systems Criteria (CSCSC) with the RACM process can be used to better understand the risks associated with the program. These concepts will be discussed in subsequent paragraphs of this paper.

The RACM process combines several subprocesses critical to achieving a possible 10% to 30% reduction in the cost of a program which is analyzed and managed in the standard way. These processes are:

1. A statistical summing of random variables instead of a simple arithmetic summing of single point estimates. For reasons to be explained later, the statistical summing is accomplished using an analytical approach and not a Monte Carlo approach.
2. An analysis of each of the major influencing factors in a program's costs; that is, those factors affecting each WBS element and those factors which are global in nature.
3. Using the knowledge derived from a statistical analysis along with the results of an earned value analysis to manage the program in a way which will achieve the highest probability of success.

The applicability of the process and the model are discussed in the following paragraph. Specifics on operation of the model are discussed in the subsequent paragraphs.

1.0. APPLICABILITY OF RACM METHODOLOGY

The RACM methodology is applicable for any projections of cost which have any uncertainty associated with the estimating process. The more uncertainties there are, the more applicable is the RACM approach. The appropriateness of the RACM process increases as the number of separate sources of uncertainty associated with a project increases. This is important because, if the effect of uncertainty is not recognized, the estimation process will not yield useful answers.

It should be noted that by using a statistical analysis for projecting costs and for cost management, most major programs can realize a 10% to 30% reduction in costs compared with a forecast using simple arithmetic summing.

The RACM approach addresses the following aspects of a program in order to estimate the cost of a program while providing some understanding of the risks associated with the elements of a program, what influences the risks, and how these risks can be mitigated.

1. Size of the program
2. Applicability to each phase in the Life-Cycle Cost (LCC) of the program to be analyzed.
3. The maturity of the program.
4. Use of Historical Data and the similarity of the product or service with other products or services in the data base.
5. How the individual cost elements are combined. (Applicability of Single Point Estimate)
6. The management process. (Cost Management, Integration with Existing Capabilities)
7. The culture of the Government and the Contractor.

A discussion of each of these basic assumptions and considerations is presented below.

1.1. Size and Type of Program

The size and type of the program should be the initial factor used to determine the need for a statistical approach. Size is a function of the number and cost, in dollars, of the WBS elements required to define a program. The greater the number of elements used to define the program and the larger the dollar amount associated with each element, the more applicable is the RACM process. For example, a study contract which is a cost plus contract requires only a simple analysis. However, a contract which requires a great number of WBS elements and large funding can induce considerable errors if a single point estimate approach is used to estimate the cost of the program. Any program requiring the use of "level two" or "level three" elements, as defined in MIL-STD-881, would be considered sufficiently large as to have the potential for large errors if arithmetic summing is used instead of statistical summing. The most easily understood analysis technique would be one that uses Level two elements (usually 11) with selected level three elements. The total number of elements used for analysis should be minimized and, as a general rule, should not exceed number of level 3 elements. However, this decision should be dollar driven. The combination of elements in the WBS should maintain a balance of costs among elements. For example, no single element in a 40

element WBS should represent 20% of the total cost. The ideal would be for each element in a 40 element analysis to represent approximately 2% to 3% of the total cost.

1.2. Applicability to Each Phase in the LCC of a Program

In the life cycle of a new system the major phases of a program are the Development phases, the Procurement phase, the Operations and Maintenance phase, and the Disposal Phase. For the purposes of this paper, only the first three phases of the Life Cycle of a program will be discussed. However, the Disposal phase should not be ignored. For a program dealing with a difficult disposal problem, such as nuclear material, the costs and the uncertainty associated with those costs may be important.

1.2.1. Development costs

Prior to beginning a program and prior to winning a new development program, the contractor must determine what their costs will be using their rate structure, their style of management, etc. The contractor must also determine the risk of any proposed value and the risks associated with the negotiated costs. These estimates can vary greatly, depending upon the number of elements in the Work Breakdown Structure (WBS), the similarity to other programs, the factors affecting each WBS element (contract structure, difficulty of meeting performance requirements, schedule uncertainties, etc.), and the factors affecting the total program (global schedule uncertainties, management policies, etc.). Development programs by their very nature have great uncertainty in the costs associated with meeting performance requirements, the quality of personnel involved, and the schedule time-to-complete.

For example, achieving performance requirements may not be very uncertain if the requirement is a variation on known results (e.g. a minor extension of range). However, if the performance requirement involves new technology (e.g. a major extension of range involving new propellant not yet developed), the uncertainties can be large and important.

1.2.2. Manufacturing

The procurement phase in this discussion will only refer to the manufacture of developed items. This is done to simplify the discussion. This portion of the Life Cycle Cost typically represents approximately 20% to 40% of a fielded weapons system.

In the Manufacturing phase of a program, the degree of uncertainty is dependent upon the maturity of the program, the complexity of the product, the work force, etc. If

the objective is the production of a stable, high volume item such as pills, the uncertainties associated with the cost projections would be minimal. However, the uncertainties associated with most weapons procurements are great because weapons are typically made in low volume and are a new product with an unknown history. This is the type of situation in which a statistical approach would be most useful and would introduce a greater degree of reliability in the estimating process.

1.2.3. Operations and maintenance

The Operations and Maintenance phase should be part of the original analysis as well as any analysis continued during the development of any new system for two reasons.

1. The O&M Phase is usually the most costly
2. The unknowns accompanying an undeveloped system resulting from factors involving such things as training, storage, and quality problems which may be recognized only after deployment, cause each WBS cost element to have large uncertainties.

A life cycle cost model should be developed from the beginning of a program and should include the O&M phase. The model can be continually updated and the effects of design and manufacturing changes made during the development phases can then be evaluated. These updates should include an adjustment to the distributions associated with each element in the WBS.

1.3. Maturity of the Program

During the development of the cost model for the program, many assumptions are made, including the estimates of variable associated with each element in the WBS, each estimate has distributions associated with them. During the conduct of the program, information can be obtained from the CSCSC system that will be able to reduce the size of the distributions and adjust the nominal. As the program continues to mature, the distributions should continue to be adjusted. For example, a subcontractor may have been on a cost plus fixed fee contract during the initial development phase, however, the subcontractor may be put on a firm fixed price contract during the final development phase. Another reduction might occur from an analysis of CSCSC. If, after 50% of the program is completed, the Earned Value shows no significant variance from that projected in the model, the distribution can be significantly reduced.

1.4. Use of Historical Data

Historical data is always used whether explicitly as in parametric models or implicitly as in a model developed from engineering experience. Both rely on what has been done before. However, any historically based model should be used with caution. Because a system was developed in five years at a cost of \$1.0 Billion , does not mean that a “similar” system will take approximately the same length of time or that the costs will be approximately equal. Any use of prior knowledge (historical data) should be questioned and adjusted for:

1. Bias on the high side resulting from unnecessary costs incurred
2. Cultural differences which will affect the costs in the program
3. Technological changes that limit the application of historical experiences.

These seem apparent and are usually considered in any cost analysis. However, there are many hidden changes which the RACM process specifically addresses.

For example, a bias on the high side is inadvertently realized in most programs by the way in which program costs are estimated and how the budget is allocated during the course of conducting the program. Most cost estimates result from bottom-up estimates adjusted by the program management during the proposal phase. This effect is discussed in section 1.5.

However, there is also a much more subtle effect; the effect of the “self fulfilling prophecy.” This effect results from hidden incentives to spend whatever budget is allocated in order to achieve the most reliable product. Few, if any, incentives exist on most programs which motivate the design engineer or product development team to finish the project ahead of schedule or to reduce costs and risk not meeting all specifications. This results in excess time spans and excess personnel which translates into excess program costs. In fact, the Budgeted Cost of Work Scheduled (BCWS) in a program using CSCSC is often interpreted by people working on the program as a goal which must be met. Underruns are often considered as being as serious as overruns are. Neither may be a serious problem or could be a significant problem. An instant interpretation by program personnel results in a misuse of the intended purpose of the process. However, a significant number of people interpret this as the “budget line” not to be varied from.

There are many more reasons why program costs are higher than they need to be, such as, unwillingness or inability to remove personnel who are no longer needed, inability of the program management to “see” when a program is completed, etc. However, by using the knowledge of the distributions around each element in the WBS

developed in the RACM process, the costs can be minimized and the resulting overall program cost will be less.

The RACM process allows program management to distribute the budget on an equal probability basis to each element in the WBS and to distribute the budget based on an expected value. This is not possible in the management methods employed today. Because budget distributed is usually budget spent, the WBS element managers should not be given more than the "expected value" identified in the RACM process.

1.5. Combining Cost Elements and the Applicability of Single Point Estimates

In combining cost elements in a WBS to find the total cost of a project, it is essential to observe that (1) the total cost is a sum of random variables, and (2) there are multiple influencing factors on each WBS element.

1.5.1. Sum of random variables

One major assumption underlying current bottom-up estimating method is that arithmetically summing WBS cost elements (which have been individually estimated at a high confidence level, e.g., 95%) will result in a program with an *identical* confidence level. This is not the case. Since each WBS element essentially comprises its own "mini-program", the cost estimate for performing the tasks within the element is most appropriately characterized by a statistical distribution of potential values rather than a point estimate. Since each WBS element represents a cost distribution rather than a single cost value, these elements must be combined statistically rather than arithmetically. When analyzed correctly, the standard bottom-up approach yields an increasing confidence level for the total program cost with an attendant increase in total costs. This effect, based on a fundamental error, is unintentional. See Figure A-1.

The graph on the left shows what can happen if all elements in a WBS are estimated at a 95% Ps. As is indicated on the chart, this error will compound with the number of WBS elements being used. One element does not create an error if the distribution is normal. However, with only 25 elements, this example shows almost a doubling of the costs and an increase in the program's Ps to beyond 999 out of 1000.

It should be noted that if a program were contracted at the much increased value, the result might still be an overrun due to the "budget received is budget spent" syndrome. All of this is unintentional on the part of government and contractor personnel and is a "hidden" cost that they are unaware of.

COMBINING COST ELEMENTS

STATISTICAL SUMMING, not SIMPLE ARITHMETIC

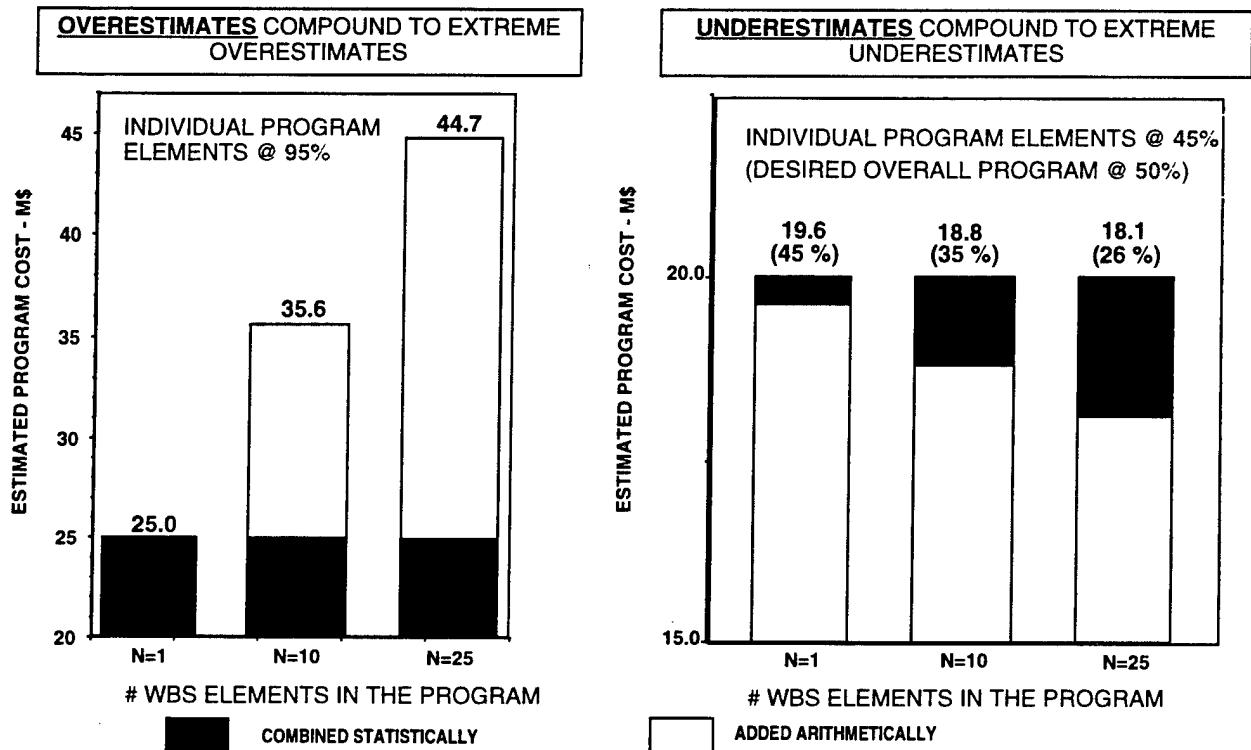


Figure A-1. Comparison Between Arithmetic Summing and Statistical Summing

The graph on the right is an example of the problem which can be encountered if the estimations are too aggressive. With a 45% Ps for each element instead of a 50% Ps, a program cost estimate would be at a Ps of 20%. This is unrealistic.

Compounding the effects shown in Figure A-1 is the problem of overestimation. In typical proposal efforts, individual WBS elements may be over-designed beyond the specifications as a hedge against unknown contingencies. In addition, quite often there is crossover into more than one discipline contributing to over-tasking of certain jobs. This compounds the fundamental error discussed above in the standard bottom-up approach.

1.5.2. Multiple influencing factors

The factors which should be separately considered in any cost forecasting method can be divided into two basic categories, (1) those factors affecting each WBS element, and (2) those factors which are global in nature. The effect of global factors must be

assessed on the total program only. Listed below are those factors which have been identified as the major factors influencing the cost of a program. As such, they should be separately analyzed.

1.5.2.1 Factors affecting each WBS element

1. Performance requirements establish the basic estimate for each element in the WBS
2. Structure of the contract. Is it fixed fee, incentive, firm fixed price, etc. ? This affects the distribution about the estimate.
3. Schedule uncertainties that are identifiable with specific cost elements of the WBS.
4. Effects of improvements using methods such as an Integrated Product Development (IPD) approach.
5. Changes in future parameters which can be anticipated now. e.g. labor rates.
6. Correlation between elements

1.5.2.2. Global factors

1. Global schedule uncertainties affecting all of the elements in the WBS. e.g., a flight failure requiring rescheduling of another test.
2. Potential major design problems discovered during ground or operational flight tests.
3. Management Policies such as the allocation of resources. This is discussed in detail in paragraph 1.6 below.
4. Major changes in business climate and redirection of effort.

1.6. Management Process

The specification and application of management reserves are business policies which can have a dramatic effect on the final outcome of any program. Statistical budget allocation provides management with a technique to:

1. determine the size of the program's reserve, and
2. suggest how the remaining funding should be allocated among the program tasks, so that the project can be completed at a reasonable risk level.

In line with the previous section, looking at the data from Figure A-1 from a different perspective, elements within the WBS should not be allocated with the same probability of success (P_S) as the desired overall program's P_S . This can only be

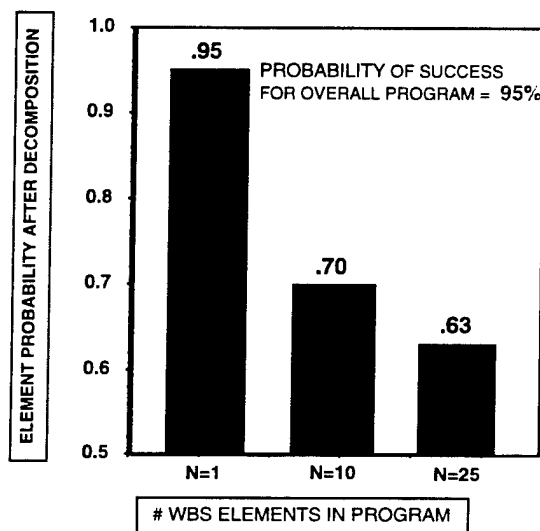
accomplished if the distribution about each element in the WBS is truly normal and the P_s used in arithmetically summing each element in the WBS is equal to 50%. In reality this is not possible. The distribution about each element in the WBS is not normal and identifying the 50% probability of success point is impractical. Most distributions will be skewed because the element manager will normally use all budget allocated in order to provide the best product and to maintain a cadre of knowledgeable personnel.

The combination of budget allocation and the negotiated/proposed contract value is illustrated in Figure A-2. This effect is predicated on the "budget allocated equals budget spent" philosophy. As can be seen, there can be significant changes to the cost of a program depending on the type of management reserve policy invoked.

EFFECTS OF MANAGEMENT PROCESS

STATISTICAL APPORTIONMENT

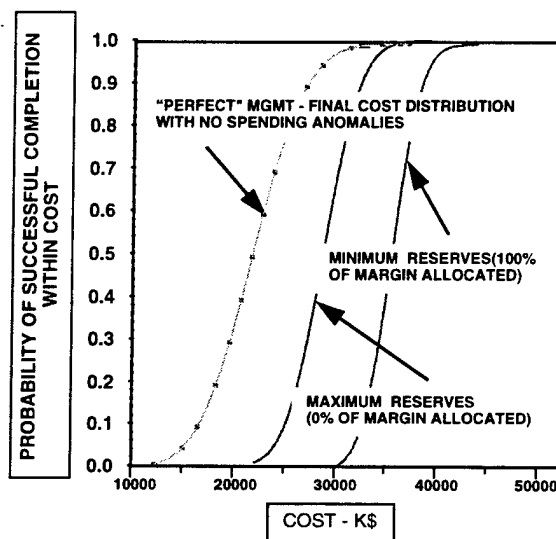
LEVEL SHOULD BE NO GREATER THAN WHAT IS CONSISTENT WITH OVERALL STATISTICAL DISTRIBUTION OF PROBABILITY FOR EACH WBS ELEMENT



ELEMENTS WITHIN THE WBS SHOULD NOT BE ASSIGNED THE SAME P_s AS THE DESIRED PROGRAM P_s . ERRORS MAY COMPOUND INTO EXTREME OVER/UNDER ESTIMATES.

EFFECTS OF ALLOCATION

MANAGEMENT POLICY (RESERVES) IMPACT ON OVERALL (REALIZABLE) COST DISTRIBUTION



- COMBINED EFFECTS OF BUDGET ALLOCATED AND NEGOTIATED/PROPOSED CONTRACT VALUE
- PREDICATED ON "BUDGET RECEIVED IS BUDGET SPENT"

Figure A-2. Effects of Statistical Apportionment and Allocation

1.6.1. Effect on cost management in a disciplined environment

Budgetary risk and an approach for managing the budget is addressed in many government and non-government documents. They typically identify the usual uncertainties as being responsible for risk, i.e., inadequate historical data, inadequate understanding of requirements, unforeseen difficulties, manpower quality/availability, lack of proposal definition, prior experience, etc. This results in an incomplete program forecast model and considerable uncertainty in the level of budget necessary to complete programs. There may never be a complete solution. Uncertainty is a fact of life. The RACM process takes these uncertainties into account.

With a long term goal to reduce program costs and improve business competitiveness, the program management should identify the risks associated with the program. This means that all uncertainties must be identified and considered in developing a cost model for the program. If the cost model is applied during the total course of the program, the program management will be better able to understand how the budget is to be allocated

This requires the additional use of a tool which will allow an understanding of how the program is progressing. This is the CSCSC method of cost management developed by the government. Management must consider the following:

- Rigid budgeting on a project level
- Allowing organizations to duplicate common (overlapping) efforts resulting, in part, from bottom-up program definition and plans.
- Misuse of methodology for monitoring and comparison of progress against program plans (so that project budgets can be adjusted as underpin and overrun situations are identified)
- Not understanding or properly incentivizing project underruns.
- Not having a flexible cadre of performers with the ability to be effectively reassigned as conditions warrant.

At the heart of these considerations are the concepts that program management must 1) accept managed non-conformance with forecast budgets and 2) adjust project budgets so that under-runs can be identified and utilized to support other overrunning projects. Recognizing when projects are ahead of schedule and can be completed below budget could yield reductions in total program costs.

Use of the RACM model provides a conservative estimate of the cost penalty associated with business-as-usual. Since no direct experience with this budget

management approach exists, successful implementation is not automatic. The degree to which these upgraded management practices can be instituted should also be considered when modeling a program. It is similar to the anticipated reductions in IPD. The original expectations for implementing IPD were significant. Some companies did realize the savings anticipated. However, other companies experienced a cost increase. The large variations experienced with IPD are not anticipated with the RACM process.

1.6.2. Integration with existing capabilities

Based on an analysis of 64 completed contracts done by Dr. David Christensen and documented in the 1994 Winter issue of the Acquisition Review Quarterly “the overruns at completion predicted by the contractor and by the government program office were unrealistically optimistic. From as early as the 10 percent completion point through the end of the contracts, the predicted final overruns were less than the current overruns reported on the contracts. Although the estimates supported by the government program offices were less optimistic than the contractors’ estimates, neither was found to be realistic.”

“From the United States Government’s perspective, which must allocate resources between competing products, the performance measurement concept has provided a good method to obtain an accurate assessment of the cost and schedule status of their procurements, on a monthly basis, and it has improved their ability to ascertain the true final costs.”

The key term used in the CSCSC method or measuring contract performance is “earned value.” This method combined with the RACM process can provide a more accurate analysis of each of the following terms by treating each one statistically.

- $BCWS = \text{Budgeted Cost of Work Scheduled} = \text{budgeted value of all work scheduled} = \text{budget or plan}$
- $ACWP = \text{Actual Cost of Work Performed} = \text{Dollars spent} = \text{Cost actuals}$
- $ACWS = \% \text{ of work scheduled} * \text{Dollars budgeted}$
- $BCWP = \text{Budgeted Cost of Work Performed} = \text{Earned Value} = (\% \text{ of work performed} * \text{dollars spent})$
- $CV = \text{Cost Variance} = (\text{Earned Value} - ACWP)$
- $SV = \text{Schedule Variance} = (\text{Earned Value} - BCWS)$
- $CPI = \text{Cost Performance Index} = BCWP/ACWP$
- $IEAC = \text{Independent EAC} = (\text{Budget at Completion} / CPI)$
- $SPI = \text{Schedule Performance Index} = BCWP/BCWS$

Not only should this be more accurate but it should result in fewer variance reports and will provide the program management the ability to determine their risks, and how to best make adjustments to the budget plan.

According to Wayne Abba, government studies indicate that when contract is more than 15% complete (you can't recover.), i.e.

- Overrun at completion will not be less than overrun incurred to-date
- Percent overrun at completion will be greater than percent overrun incurred to-date.

If the RACM process is applied throughout the life of the program, recovery be possible and should not be required as often.

1.7. Cultural—Now and Future Requirements

1.7.1. Incentives

The present incentive structure which exists on most programs are subtle incentives which are incorporated into the present day management styles. There are the unwritten incentives: (1) The need by the design engineers to provide the most reliable product *for the money provided*. (2) The inability and reluctance to off-load and on-load competent engineers from the program. (3) The need to spend the budget provided or "it will not be available next year."

Written incentives include cost, schedule, and performance incentives. However, performance incentives have, in the past, often been large enough to outweigh the cost and schedule incentives. The most subtle incentive that relates to the use of a good risk management process is the misuse of the CSCSC intent. It was intended as a management tool. However, it is often used to provide an absolute budget line from which there can be no variance, i.e. no under-runs or overruns are allowed.

As stated in "Punished by Rewards" by Alfie Kohn 1993 "Not only are incentive systems and pay-for-performance plans pervasive in U.S. companies, but there exists a deep and rarely questioned commitment to the belief that offering people rewards will cause them to do a better job. The evidence, however, suggests that extrinsic motivations in the workplace are not only ineffective but often positively counterproductive. The most familiar reasons proposed to explain this failure deal with relatively minor issues that apply only to specific incentive programs. But several other reasons strike at the heart of the assumptions about motivation that underlie all such programs. The bottom line is that

any approach that offers a reward for better performance is destined to be ineffective. It is simply unfair that employees are held responsible for what are, in reality, systemic factors that are beyond their control.”

Perhaps the *unwritten and misused* incentives should be examined. This can be partially accomplished by changing the culture in several ways. (1) Understanding the “expected” and “most likely” outcome in a program’s cost and schedule structure. (2) Use CSCSC as a tool to manage the program in conjunction with the RACM process. (3) Provide incentives and recognition at the level where costs can best be controlled, i.e. design engineers, integrated product development teams, etc.. (4) Allow companies to make significant profits if they meet all cost, schedule, and performance goals. (5) Include, as part of the cost analysis, how companies will be able to maintain a basic cadre without having them “sit” on the program being proposed. The “basic” cadre, the cadre required during the course of the program, should be identified in the first estimate.

1.8. Analytical Approach Instead of Monte Carlo

Another feature of the RACM process is the use of the analytical approach instead of the Monte Carlo approach used in most cost forecasting models. There are several reasons why the RACM approach is desirable.

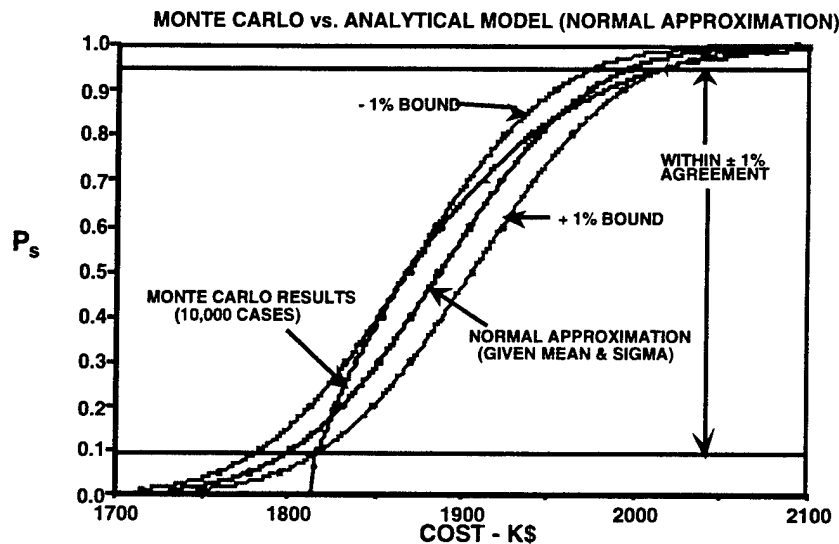
1.8.1. Accuracy (consistency)

The analytical approach has been compared to the Monte Carlo approach for consistency. This does not imply that one method is more accurate than another, but merely indicates that the results of an analytical method are consistent with Monte Carlo. Accuracy is a function of the parameters used to define the function. For this comparison, the same parameters were used for the Monte Carlo and RACM methods. These two approaches provided results that are within 1% of each other and, in the range of normal use (50% Ps to 95% Ps), the agreement is even greater. See Figure A-3.

This level of consistency is surely sufficient. The uncertainty in the inputs for modeling any program will dominate any lack of agreements between the analytical RACM approach and Monte Carlo simulation.

The analytical approach used in RACM is the same approach used to predict accuracy and reliability on the Fleet Ballistic Missile (FBM) program for the past twenty years. The success of the approach on the FBM program is beyond reproach.

SPEED/ACCURACY/ CONSISTENCY CONSIDERATIONS



- WITH ANALYTICAL MODEL DEVELOPED
 - QUICK RESPONSE (SECONDS vs. HOURS)
 - ACCURATE TO WITHIN 1% OF MONTE CARLO

Figure A-3. Accuracy of Monte Carlo Versus Analytical Model

1.8.2. Speed

The RACM model is able to provide an analysis in a matter of seconds. The Monte Carlo models will take an hour or more on some of the large programs. This speed is essential for one very good reason. An understanding of any models forecast requires many “runs” of the model using different inputs. This allows management to understand better what elements in the WBS should be watched and examined further. If each analysis takes an hour or more, there will not be enough time to make all of the runs necessary for a full understanding of the modeling results.

1.9. Visibility of Process

In the very beginning, when this process was first being developed, one of the criterion was to maintain visibility into the various inputs such that management can “see” the results of each input. This is very different from most parametric models. This

visibility has been achieved and is considered by management as one of the strong points. It is not only visible but the process requires management, design engineers, manufacturing engineers, logistics engineers, planners, cost modelers, and other program personnel to work closely together. Furthermore, each discipline involved can see the results of their inputs. This eliminates much of the smoke and mirrors often associated with cost forecasting. It can be time consuming but will result in a better understood result which will ultimately be more accurate.

1.10. Comparison with Other Processes

A search of the literature was completed with the results shown in Annex A. This annex of the report discusses the current state of the art of practical approaches to budget risk analysis. All useful approaches begin with a WBS that treats a program as a set of interconnected elements. Each element of the WBS has an uncertain cost that may be correlated with that of the other elements. If the probability distributions and correlations were known, the probability distribution of the total cost could be found. However, it seems that:

1. the probability distributions of the WBS cost elements are not known;
2. the cost elements are correlated but the amount of correlation is unknown;
3. the cost estimate itself affects the program cost through budgeting, scheduling, and program-approval processes (this suggests that specific modeling of these processes is essential); and
4. the cost of a program is constrained both high and low--the lower bound to generate sufficient profit, the upper bound to avoid cancellation or restructure.

2.0. CONSTRAINTS/CONCERNS

2.1. Cost Elements

2.1.1. How many (What is too many, what is too few?)

The number of cost elements in an individual project varies from project to project. The basic assertion in RACM is that going below WBS level 3 is, in general, unnecessary and potentially misleading since, at too low a level, the estimators are likely to lose sight of the inherent variability associated with the item in question (e.g., - the cost of a nail in estimating the cost of building a house). This viewpoint has led us to believe that provision for 20-30 WBS element items should usually be sufficient in describing the overall project. This is certainly not a fixed requirement and the present physical

limitation of 30 elements reflected in the current RACM mechanization is not intended to be any such restriction, just a convenient size for the format of the demonstration level program that now exists. The actual number of elements is an estimation of a new project is strictly a function of the makeup of that particular WBS breakdown and should always be treated as such when building your cost model.

2.1.2. What is the relative size of each element ?

Similarly, there is certainly no fixed size associated with a proper description of the WBS breakdown. Every project could easily contain unique elements that could influence the final size distribution of the WBS element description used in RACM analysis. Since the combination algorithm is a statistical procedure, in general any element that is below 10% of the relative size of the others already chosen, should be considered a candidate for combination with similar small elements into a new element. This yields a more balanced description of the overall WBS structure. Any element, whether because of its magnitude or variability or interest to program directors, can (and probably should) be included in the model. If later, when the model is complete and sensitivity studies can be performed, an element has minor impact upon the overall answer, it can then be combined with others into a higher level description.

2.1.3. If you combine elements, how do you reduce errors?

As mentioned above, the natural tendency to go to the lowest levels in order to avoid eliminating anything can have an undesired effect: Natural variability that would be recognized at some higher level can easily be missed. Instead of calling for a specific term like "WBS Element Error" for each level which obscures understanding of the whole process, it would be best to follow the recommendation made not to go too deeply into the WBS structure simply because it seems possible to do so. It seems preferable to perform the analysis at WBS levels 3 and 4. As long as the admonition "be careful" is followed, there is no absolute rule that some element should or should not be included. If done carefully, details taken at any specific level should produce the same answer as a model built at another level. Consider the level necessary for desired visibility and the application of the interview process specified later in this document for extraction of "expert" opinion. This along with cost modeling experience should make any specific situation much clearer even if the general "rules" must remain somewhat vague.

2.1.4. How do they recognize uncertainty?

The problem would seem to be in not how to recognize uncertainty, but rather in how to quantify that uncertainty. The general approach is to take a first cut at the uncertainty bounds, even if somewhat crude, and then build a complete model which could then be tested for sensitivities to the assumptions made. This should give a much clearer picture of which inputs are the most critical and which are relatively unimportant in effect and therefore need not be analyzed as deeply as the others. The basic question is as always, "How does one extract expert opinion ?" In addition to the section later in this paper on that specific question, the application of sensitivity analysis using the model itself to clarify such issues should provide the majority of the insights needed. In dealing with a situation where experience and true expert opinion is lacking, other approaches rely on the concept of so-called "fuzzy logic" to convert feelings/verbal confidence statements such as "low", "medium-low", "medium", "medium-high", "high", "extreme" into quantitative values needed in the model. Despite the arbitrariness associated with such a procedure - which RACM does not use - it acknowledges the need for a statistical description of costs. Any statement or admission of uncertainty is better than taking every number as an absolute, unqualified fact and combining them all into a single point estimate of the total with no concept of the possible uncertainty in that final value.

2.1.5. What format should be followed? 881, PDT, Or?

The cost reporting format is not an important issue. If the accepted/familiar format is the WBS 881 structure, then use that structure to build the model and report findings. In a strict PDT environment situation, obviously a PDT structured model and reporting/breakdown system would be useful and provide the visibility required to track and manage the program after the contract is awarded.

2.2. "Schedule" Elements

2.2.1. Define parallel (How to determine number?)

Given a schedule/flow-chart description of the expected time-phased layout of the program in question, the concept of paths or parallel paths within that network should be clear. The approach suggested here is that although all those paths may not be exactly the same length, it is sufficient (and certainly conservative) to assume that they are all of the same length for the RACM analysis. Then the solution can be found by the classical "maximum time to complete across an n-dimensional parallel path network" simulations/algorithms. An example is the familiar PERT chart analysis used in many

cost - schedule analyses. The standard qualifier also applies here: after solving the model, determine the sensitivity of the answer to the particular schedule length & parallel path assumptions made.

2.2.2. Critical parallel paths

As used in this discussion, a "critical path" is any one individual time phase element or serial combination of elements that could cause the WBS element in question to not complete its task within the schedule predicted - cause the holdup of the entire element until it is completed. In the global sense, it is any WBS element which could cause the holdup of the entire program (the "Standing Army" problem).

2.2.3. Where should inputs come from ? Off-line programs or?

Inputs might come from a variety of sources - a PERT type chart is a familiar example. If schedule is an important issue/presumed cost driver in a program, certainly some time-phased flowchart of operational performance must be available and should provide the data required for this phase of RACM descriptors. An off-line program reflecting greater detail and more complicated interactive structure (contingency plans, redundant efforts, etc.) would almost certainly require a Monte-Carlo type analyses and then these results compiled into the present analytical approximation, required by the RACM model, as a simple mean and variance type value. This is an option under consideration.

2.2.4. What would be required in an RFP?

Probably nothing more than what is required today - a PERT chart description of the program. An evaluation of that flowchart as to the number of "critical paths" and possible durations of slippage both within elements and across the global program could be provided to assist the evaluator in determining his inputs for the RACM model.

2.2.5. Algorithm

As mentioned above, the current algorithm is simply the "maximum time to complete across an n-dimensional (equal)path network". The values used in the RACM tables are the results of off-line Monte-Carlo simulations of this process given the assumption all the paths are of equal length and the statistical distribution of time uncertainty for each path follows one of two possible forms. For some cases, it is believed that the amount of time that can be saved during program execution is considerably less than the possible overrun time for that element (usually relating to a "Labor" type

element). In other cases, it is believed that the time can never be shortened - only the possibility for overrun from the nominal description exists (usually perceived as a "Non-Labor/Subcontractor" type element). Since the actual implementation within RACM is to use the empirically derived tabular description of the effects, any other desired form could be equally supported. These were all the forms of slippage description that surfaced from experiences here, and certainly enough to include in this concept demonstration version of RACM.

2.3. Statistical Distributions

2.3.1. Taylor series or exact solution

The only necessary descriptions of each element along the way from "basic" through "rate affected" estimates are the mean and variance (the first two moments) of their distributions. When they are combined at each stage, although it is only for visibility and rough impact comparisons, the Central Limit Theorem is invoked to imply that the overall distribution is adequately approximated by a "Normal" distribution. In this sense, each step of the way is an exact solution of the problem - from calculating the first two moments. The originator of the RACM model used a Taylor Series argument to substantiate the adequacy of describing the distributional form of each element along the way as "Normal" (that the products being considered of the form $x*(1+y)$ were sufficiently described as "Normal" given that x & y were "Normal", etc). At that time, the model did not include the piece-wise linear transformation solution for the "money allocated is money spent" assumption or the "Beta" approximation to represent/display the final results. In the current formulation, I don't believe there is any further need/use in using/referencing the Taylor series expansions for the variables in question. Final adequacy judgments for the overall distributional forms displayed, if deemed necessary, should come from a full blown Monte-Carlo replication of the RACM process (which has been done on example cases as sanity/validity checks).

2.3.2. Does order make a difference? No—Why?

Where order is arbitrary, the individual effects modeled in the RACM stages are multiplicative operations and hence the aggregate effect is independent of the order used. Some of the descriptions are given as additive effects, but their sequence of inclusion/presentation is driven by logic (definition of the term, etc.) and the desired visibility needed to explain to program management the impact (magnitude of the assumption) in transitioning from the previous stage. In the final, global phases of the

estimation process, the truly additive effects of global schedule slips and major problem effects are carefully isolated and removed before the "management effect" transformations are applied and then reapplied afterward to achieve the overall, final cost distributions displayed. See section 6.4 below for additional commentary.

2.3.3. Truncation of normals—how do you get mean and variance?

This is really not a truncation process, but rather a collapsing of the original distribution function from below up to some specified value. The continuous random variable is transformed into a random variable with non-zero "mass" at some point. This represents a specialized case of the piece-wise linear transformation of a Normal distribution function into some (generally analytically undefined) form - a problem which has an exact analytic solution for its first two moments. We can therefore proceed and combine such collapsed distributions to get the first two moments of the sum. A description of the process of interest is as follows : if X_1 is the value in question, all the area under the distribution curve to the left of X_1 is mapped into X_1 precisely - the remaining distribution to the right of X_1 left unaltered. The expressions derived in the paper presented earlier give the exact solution for the mean and variance of the resulting distribution function. The specific case here corresponds to a slope of $a_1=0$ and $a_2=1$ in the more general case presented.

2.3.4. How to compute mean and sigma? (Where used and experience)

The more general case, as depicted in the paper mentioned above, has been recognized and used in several areas of system performance analysis around Lockheed over the past 20 years or so. When it was recognized that this "money allocated is money spent" concept was just a special case of this work/result, it was immediately adopted to assess this "management" phenomenon and allow a display of the impact reserves/no reserves can have on the realization of final cost predictions. This facet of cost estimation is believed to be unique to RACM - there are many statistical cost models available/being used (although mostly Monte-Carlo based) but none of them we are aware of model this (significant) effect.

2.3.5. Check against Johnson and Kotz

The paper given to me from Johnson and Kotz is dealing strictly with the classic truncation problem and hence not directly comparable. In the limited area where they overlap and can be compared, they appear to be totally equivalent (as they should be!). The paper stands by itself and I know of no comparable work in the literature - although

there certainly could be something out there - it is not an approximation and deals with what should be a reasonably common problem. I am just not aware of any other result in print.

2.4. Correlation

2.4.1. Correlation and RACM

As RACM views the world, correlation would be included/modeled at the fundamental "Basic" WBS element level and would represent a relationship between two cost elements. If the original WBS breakdown separated two portions of the same basic task into different items in the RACM model inputs, these elements would likely need a correlation statement describing their combined behavior. It is assumed this will be avoided and that the fundamental WBS element structure chosen will reflect independent efforts. Another potential source of correlation would be if the same person or "expert" made the predictions for two fundamentally distinct elements but he himself is biased in his estimation processes. Although provision was included in the original RACM implementation to handle such an instance, it has been decided that determining the proper correlation coefficients is extremely difficult at this time and the chance for error overshadows the potential gain. It is possible that the "placeholder" for correlation should be left in the proposed future RACM model, but at this time, it has been removed from the current demonstration version.

2.4.2. "Observed" correlation

One of the common reasons given for the necessity for correlation analysis in the cost estimating process is the performance of historical programs. Viewed at the final cost stage, after the possible influence of many external factors (major program problems, management policies, inflation, etc.) it is possible to observe correlations between WBS cost elements - final performance versus initial projected cost. If the model being used does not address these factors directly, then correlation terms would be appropriate and necessary to explain the results why the "whole" did not equal the "sum of the parts". In RACM, it is believed it is better to treat these outside factors directly and hence their effects can be estimated without correlation statements to explain unmodeled (but observed) phenomenon. It also is much easier to explain/validate the final answer if each of these factors can be called out and examined for its own merits. The current list of "influence factors" that RACM addresses could possibly be incomplete at this time - only time and benchmark tests can determine the sufficiency of the model concerns at its

present stage. In any case, the recommendation is to continue to model specific influences directly.

3.0. REQUIREMENTS FOR OPERATING THE RACM MODEL

3.1. Inputs Required by the Model

3.1.1. General guidelines

1. All inputs to RACM are descriptions of individual factors in the realization of final program cost and should be described in a statistical manner. Therefore these inputs should reflect the uncertainty associated with each term - i.e., provide a distribution for the factor, not just a point estimate.
2. In order to avoid "double pork-chopping", at each stage, try to separate inputs so that the term being modeled does not include effects that are to be described later. If such separation is not possible/practical for some reason, then do not repeat the accounting when the later stage asks for some of the same data.
3. Because the nature and physical explanation associated with so-called "labor" and "non-labor" components can be very different, the effects are also separated. After propagation through the appropriate phases, the two considerations are re-combined in the end along with the overall (non WBS-by-WBS element) program level factors to create the final cost curves.
4. The current demonstration version of RACM has a built-in "Menu" system to facilitate the step-by-step process to be followed to complete the necessary input sections.

3.1.2. Current inputs

1. Basic cost estimate—If at all possible, this term should represent the fundamental estimates of "should" cost for the completion of the specific WBS by WBS element tasks. To provide RACM with the desired distribution (estimate reflecting associated uncertainty), the current choice of possible input forms is to provide a high and a low cost estimate for each term along with their associated/estimated probability of success. There are no restrictions on the exact values these estimates must reflect obviously, the Ps for the lower cost should be less than or equal to the Ps for the high estimate. The internal mechanics of the RACM program determine the unique mean and standard deviation of a normal distribution passing through these two points. There are certainly many other possible ways of defining this process - possible future extensions would allow a choice of assumed distributional form (triangular, uniform, beta, user specified, etc.) at this point in the user interface. The

choice of a "normal" distribution seems the most logical since almost any other choice would reflect the consideration of some other influencing factor which is intended to be separated out at this point and entered later in its Proper place. Another consideration that should be noted here is that the distributional form choice is not significant since the only real data operated upon are the mean and standard deviation in proceeding from one step to the next.

2. Schedule uncertainty (within WBS element) - The input asked for is a description of the possible schedule slip within the WBS element in question caused by factors totally within that element. It assumes the value given is a 3-sigma type value (a realistic maximum) with zero representing the nominal or maximum likelihood value. Given some schedule flowchart/time phased description of the element, RACM also requires an input for the equivalent number of critical parallel paths which might cause that slippage to occur. Methods for computing the effects on the mean and standard deviation due to this type of influence factor are discussed elsewhere. Off-line empirical methods have been employed to provide tabulations of possible results for RACM to include. The equations/tabulations use a scale factor (of the original sigma prior to this consideration) formulation and hence reflect a multiplicative process step.
3. Cost Improvement Initiatives - This section or influencing factor is intended to represent a statement of how this program intends to do business differently/better than in the past (better than the business-as-usual approach most likely used as the starting point - the "basic" estimate above). It again is considered element by element throughout the WBS structure developed/modeled. The form of this input is a best estimate/most likely improvement percentage and a min/max uncertainty range (again in the same % format) statement. For example, WBS element X(i) might assert he can improve his previous process by 20% +/-20% for an overall distribution of 0 to 40% improvement where the two extremes are modeled as the 3-sigma points of a normal distribution about the nominal/best estimate improvement point. Done as percentage improvement, etc., this step also reflects a multiplicative process.
4. Rate (inflationary?) influence - Primarily intended/designed to handle a labor (rate) type phenomenon, it can also be a place-holder to accommodate possible effects on the non-labor or subcontractor side. The form of this input is again a percentage estimate effect and a statement of the min/max range variation that might be expected. A nominal (%) is input accompanied by the 3-sigma possible variation (%), so that RACM can define the distribution. Particulars are exactly like the example just given for cost improvement representation, and is a multiplicative process.

5. At this point, the element by element influence factors have been completed and the inputs now reflect factors affecting the overall program without specific attachment to individual WBS elements. These are referred to as the "Global" effects. The first of these is the possibility of overall program schedule slips (like the standing army problem) where the input values reflect possible slippage across the entire structure simultaneously and the number of critical parallel paths reflects the expected number of individual WBS elements in the overall description that might cause such a global slippage. The same scale factors are utilized that were used individually in step b) above. The only real difference is that here we are transforming the total program mean and standard deviation, not the individually described element-by-element characteristics. The process is a multiplicative operation, although that consideration really plays no role here. This impact is always maintained separately and presented in the final analysis as a program reserve item for the program manager.
6. The next effect/influence factor considered is referred to as "Major Problems". This implies a consideration for unknown but potential program events that may occur somewhere during the course of the project that might redirect efforts, etc. and in some way cause cost anomalies. Typical examples would be flight test failures, program redirection by the customer, subcontractor difficulties that couldn't be anticipated, etc. Although, at the outset these problems could not be specifically defined, experience with similar projects suggests that such problems can occur and have important consequences. They are described to RACM as the magnitude (cost) of some potential problem and an estimate of the probability of occurrence of that problem or event. The RACM analysis combines all these stated individual events to generate the necessary mean and standard deviation of the discrete event space representing all possible combinations of the scenarios considered. This effect is added in for display purposes at this point, but the effect is also maintained separately so that it can be removed when other factors are considered and then recombined when necessary to reflect the final overall cost estimation. Since this factor is not identified at the individual element level but rather just at the total aggregate "Global" level, the (distribution) dollar value is never considered as a specific WBS level item in allocating resources, etc. This factor would always be considered as a fundamental reserve item for the program manager (as is potential global slippage dollar values).
7. The last influence factor currently considered is the management policy to be employed by program management. That is, RACM addresses the question "How much reserve will be held back during the initial allocation of resources to the WBS elements". This is totally a representation of the concept described as "money allocated is money spent" (the self fulfilling prophecy philosophy). Although the demonstration RACM program now contains options to soften

this principle (a manager will spend at least xx% of what he is allocated, etc.), the principal remains the same. This is a unique consideration of the RACM model - but a fundamental and major player in the overall realizable cost outcome. If a particular program/situation occurs where this assumption is inappropriate, the effect can be removed at the standard input level (zero out the effect). The main usefulness of this term is probably to illustrate to program managers the possible impact their individual reserve policies might have on the final realized cost of their program. The specific RACM inputs used are just the values necessary to describe the management reserve policy intended. The level of success that program management intends to reflect with their cost (bid) estimate (which defines the overall dollars available to the operational program) and the percentage of the possible reserve fund (bid level dollars minus the minimum/nominal estimated program cost) they expect to hand out initially to the individual WBS cost accounts. The current formulation assumes that the minimum reasonable WBS element level allocation is probably the 50% point from the cumulative distribution for the overall WBS identified program dollars. This may seem somewhat arbitrary, but seems correct so that the WBS level projects can design an effort that is appropriate (and nominally should produce the desired product within cost) but will not spend from the reserve fund unless actually required to do so. This is the feature which allows for the "statistical averaging" or "under-runs financing over-runs" concept which may not occur without this scheme or something similar. The mathematical process used to determine this effect was discussed earlier - the individual distributions are collapsed up to the dollars-allocated level using the piece-wise linear transformation algorithm. We could probably just use the Central Limit Theorem at this point so that all these collapsed distributions could be combined and represented as a "Normal" distribution for display purposes, but in extreme cases the Beta distribution definitely provides more freedom to better represent the shape of the resulting cost distribution. The minimum, maximum, mean and sigma data we have available for the (non-symmetric) cost distribution resulting from the overall processes described here are exactly the parameters necessary to generate the Beta curve currently displayed as the final RACM "Cost versus Ps" estimate.

3.2. Sources of the Inputs

The inputs should come from the interview process between the cost analyst who will be running RACM and the individual program managers/cost account managers involved in actually performing the tasks. In practice, we have experienced many variants on this procedure depending upon the experience of the cost analyst, the statistical knowledge of the program personnel, the accumulated knowledge of the interviewers about the specific engineers and program personnel providing the answers, and seemingly

unique situations that arise in every application. Almost surely, the input/interview process will be iterative in nature. This, however, can be very valuable in terms of good understanding of the final product and the confidence that there will not be a lot of missing information in the final estimates. To be sure, the final output is only as good as the quality of the input. But this is true no matter what the cost estimating procedure chosen - certainly not a unique problem with RACM! In fact, because of the possibility for sensitivity analyses that the analytic model provides, it should be considerably easier to identify and resolve potential input errors than with other approaches. Although historical data has quite often been thought of as the best source of information for such analyses, caution should be observed. There are many sources for error buried in most historical data bases. New ways of doing business, the "self-fulfilling prophecy" phenomenon, missing documentation, unknown sources of apparent correlation, and management philosophy followed (what probability of success was associated with that bid & what reserves/fencing of budget were in place, etc.) are only some examples. Historical data may well be the best source available, should always be considered when it exists, but should always be used carefully. The next section discusses methods to be applied in determining the inputs to RACM.

4.0. PROCESS FOR ELICITING THE INPUTS

4.1. Probability Distributions

The fundamental observation guiding design of the RACM process is that modeling uncertainty is the critical issue in understanding the behavior of projects. The RACM process requires that a project be structurally decomposed into a set of (WBS) elements, each of which is characterized by the random variables time-to-complete and cost-to-complete. Therefore, the RACM process requires that these random variables be characterized in some meaningful way so that an appropriate set of inputs to the analysis can be obtained. Clearly, it would be inappropriate to characterize a random variable by a single point estimate of its value; a single value is simply insufficient. Other methodologies have foundered on this very point, since those approaches compound the error of using a single value to characterize a random variable with tending to confuse the mean, median, and maximum likelihood values of a random variable. Thus, point estimates of random variables cannot be coherently interpreted even within a methodology that relies solely upon such estimates. To avoid these kinds of difficulties and errors, RACM assesses directly the probability distribution of each input random variable. The essential reason for assessing probability distributions is to encode

mathematically the uncertainty that an expert assigns to the time-to-complete and cost-to-complete each work-breakdown-structure element. Those uncertainties are used in the RACM process to determine the uncertainty in the entire project.

4.2. Assessment of Uncertainty

There are two general difficulties associated with encoding an expert's perceptions about uncertainty. First, one often attempts to use ordinary language when accounting for uncertainty. Ordinary language can be ambiguous and imprecise. A moment's reflection upon the differences among the following statements should suggest the difficulties attendant upon using ordinary language to describe an uncertain situation. Consider the various interpretations that could be given to (a) "It is most likely to occur"; (b) "It will probably occur"; (c) "It will almost surely occur"; and, (d) "It is likely to occur". RACM eliminates the use of ordinary language in assessment and replaces it with precise probability statements.

A second difficulty is that people are frequently biased when they try to quantify their uncertainty. (Here we adopt a subjectivist view of uncertainty, viz., that uncertainty is a belief held by an expert about the real world.) There are two classes of biases. Motivational biases are present when the expert's statements do not reflect his conscious beliefs. Cognitive biases are present when the expert's conscious beliefs do not reflect his information. RACM attempts to overcome biases by adopting a formal procedure for quantifying an expert's uncertainty.

Aspects of the formal quantification procedure include the following. (a) All relevant information is considered when providing probability assessments. (b) What experts actually think about the variable in question is what is sought by the procedure. (c) Methods that can circumvent or overcome the limits of empirical data are applied in the encoding process. (d) Ambiguity is minimized by expressing expert opinion in mathematical terms. (e) Throughout the process, consistency checking controls for biases.

4.3. Overcoming Biases

Before discussing the quantification procedure itself, we address in more detail the notion of bias. We begin with motivational biases.

Motivation biases can be caused by organizational reward structures. For example, a salesman may bias a sales estimate upward to obtain a large expense account, or downward to assure that his sales exceed his estimate. Similarly, a program manager

may bias cost estimates downward to assure continuation of his program. Often, motivational bias may be indicated by statements such as "I was being conservative in my estimates" or "I am an expert, and experts aren't supposed to be uncertain". The latter statement is a special case of motivational bias, and reflects the expert's preference for the appearance of certainty in an uncertain situation. Motivational bias can be overcome if the causes of such bias--organizational reward structures, expert bias, the need for certainty, etc.--are known and openly discussed during the assessment process.

Cognitive biases result from improper use of subconscious rules of information processing (known as cognitive heuristics). Six cognitive biases are frequently observed. *Adjustment and anchoring bias* results in a tendency to focus on a specific number, such as an initial guess, and not depart from it. *Availability bias* is the tendency to focus on a dramatic or recent event that is easier to recall and impute to it a greater likelihood or representativeness. *Sample bias* results in too much faith in small-sample data and not enough weight on general, abstract information. *Coherence and conjunctive bias* permits a good story to make an unlikely event seem more likely. *Representativeness bias* suggests that an event that is in some sense representative of what is possible is more likely to occur, such as the apparently remarkable occurrence of four heads in a row, compared with any other single mixed sequence of heads and tails when tossing a "fair" coin. *Overconfidence bias* is the tendency to underestimate one's true uncertainty and is often expressed as a firm belief in an incorrect answer. These biases are subconscious and only partially controllable. The encoding process seeks to uncover these biases and thus overcome them.

4.4. Encoding Procedure

The encoding procedure used in RACM is a five-phase procedure, consisting of (1) a *motivating* phase, that is essentially a training session wherein the purpose of the task is explained and the various biases are defined and discussed; (2) a *structuring* phase, in which the variable under assessment is defined and various aspects of the expert's thinking with respect to the variable are explored; (3) a *conditioning* phase that assembles all relevant information to counteract anchoring and availability biases; (4) an *encoding* phase that expresses the expert's understanding numerically by applying various comparison and reference techniques; (5) a *verifying* phase in which the expert is asked to agree to the encoded distribution in terms of his own betting behavior. The entire procedure iterates until the expert is willing to bet his own money according to the odds given by the assessed distribution.

5.0. SENSITIVITY TO ERRORS IN THE INPUTS

5.1. Requirement for Sensitivity Analysis

It seems that it would be almost impossible to accurately produce an overall cost estimate without the ability to do sensitivity studies and iterate through your procedure to scrub the inputs and understand the impact of your assumptions. It is the driving consideration which led to the development of the analytical approximation technique for RACM rather than the more common Monte-Carlo simulations being used elsewhere.

5.2. Standard Procedure

It is strongly recommended that, as a standard procedure, all the inputs and assumptions going into RACM are tested/validated by means of full scale sensitivity runs. A full ensemble of sensitivity results should be generated and utilized even after the final set of nominal values has been examined and agreed upon by both cost analyst and program engineers alike. It is by this means that a set of "partials" is derived and made available for other phases of contract performance enhancement. Separate from cost risk estimation, there is the vital subject of cost risk management or mitigation necessary during the actual performance of the contract. The issue of proper trade-offs between performance issues and cost/risk requires such information. The concept of "Design-to-Cost" is dependent upon this knowledge for optimum success. There is a special display included in the standard RACM outputs to enhance the generation of such sensitivities. It graphically depicts the cost curve "delta" between the current case and some preset nominal, etc.

6.0. RACM METHODOLOGY (QUESTIONS AND ISSUES)

6.1. Modeling the Effects on Cost of Schedule Slippage

6.1.1. Explain development of factor for mapping number of critical paths to cost factor for schedule slippage

As mentioned earlier, this scalar mapping, mean shift and sigma as a function of element sigma and number of critical paths, was developed off-line using 10,000 sample Monte-Carlo simulations of the generic/standardized maximum time to complete across an n-dimensional, equal length, parallel path network situation. The form of the individual elements in the network was varied as described below to provide options to the user in describing the particular situation best applicable to the schedule slippage

problem in question. The resultant tabulation is stored in RACM and utilized by a simple table look-up routine with appropriate rescaling for the individual means and sigmas being considered.

6.1.2. Definitions used for "Critical/Parallel Path"

It is assumed there is some concept of a schedule flowchart, "PERT" chart, or generic time-phased performance/milestone chart associated with each WBS element included in the RACM description as well as a similar, higher level chart for the entire ensemble. The absence of such information would indicate no potential slippage influence factor to be considered and imply an input of zero maximum slippage across a single path through the respective element/task. A "critical" path is any one that might cause the potential slippage on its own/independently of any other. The "number" of critical paths is then just the simple counting of the number of different paths so identified as critical. At the "global" or system level, WBS elements play the roles of paths and the total count of possible elements capable of producing such slippage is the number expected in the RACM input. The conservative approach of assuming all identified paths (element slippages) are of equal length/same approximate duration allows for the standardized off-line generation. Although empirical in nature, it is believed to be adequate at this stage and certainly advantageous since any more detailed analysis would require a Monte-Carlo type solution. In the future, if some special situation arose where more detailed schedule analysis was felt to be necessary (and data was available), this part of the problem might be done in some adjunct, off-line Monte-Carlo simulation and the resulting outputs (mean shift and sigma level details) merely input to RACM much as the "Major Problem" data is now.

6.1.3. Mapping beyond 15 parallel paths

There is no physical limit to the possible number of paths some specific program might have. Being consistent with the present feeling that 20-30 are enough elements to include in the WBS level description, 15 was somewhat arbitrarily chosen as an accordingly appropriate upper limit for the number of those elements/paths which might result in some significant schedule slippage. Going beyond 15 would simply require a few more Monte-Carlo runs for the higher order cases (or possibly some curve fit solution instead of the tabular form now used)

6.1.4. Explain use of triangular distributions for labor input

The non-symmetric, triangular distribution for potential slippage distributions was one of the two options experience suggested were most common/representative. For the labor type elements, some schedule savings might be accomplished, but in general, the potential slippage exceeded the magnitude of the potential savings - a 3 to 1 ratio is utilized in the current demonstration model.

6.1.5. Explain use of half-normal distributions for non-labor input

The half-normal distribution was the other form suggested by experience. History with subcontractors/suppliers indicated that essentially no gains could be expected, only some potential slippage - hence the one sided distribution (normal instead of triangular reflecting the feeling of centrality). Neither of these assumptions is sacred - should really be thought of as two potential assumptions that might be applied to any element/at any level of schedule concern. Any further suggestions/experiences from outside our sphere of knowledge would be greatly appreciated and could be easily incorporated into future versions of RACM.

6.2. Approximating the Distribution of Total Cost with a Beta Density

6.2.1. Explaining procedure for approximating the mean and variance of truncated normal random variables

As mentioned earlier, it is not an approximation (and the problem is not technically one of truncation). The enclosed paper describing the analytic solution for the mean and variance of the distribution after a piece-wise linear transformation of a normal random variable details the equations applied to the concern here - the impact of "money allocated is money spent" or more generally stated, the impact of some implied lower limit on realized cost distributions.

6.2.2. How is beta approximation applied?

Realizing that the "collapsed" distributions from the step above might put a strain on the assumptions for application of the Central Limit Theorem, and observing some of the "actual" results from Monte-Carlo check case runs, another more flexible fitting scheme was sought to represent/present the final results. Especially under the assumption that no reserves might be held back, etc., the final cost curve can be quite non-symmetrical. The thing to be remembered, however, is that this deviation from normality occurs primarily in the lower tail of the overall cost distribution where there is

considerably less interest in the answer (why intentionally run a program at a 10-20% Ps ?) Without implying some theoretical justification for why the "Beta" is the proper form, it was merely chosen because of its reasonable flexibility and recognized ability to use the data already available in RACM (minimum, maximum, mean & sigma) to display a much better fit over the total range of interest. By the way, the output of RACM is in no way intended to be used in the tails - say roughly outside the 10-95% Ps range.

6.3. Incorporating Correlations Among the Cost Elements.

6.3.1. Explain correlation as addressed in the model

It is not currently an influence factor described/utilized by the RACM model. This is because it was always intended only to represent the possible correlation between the fundamental basic cost estimates and has not turned out to be significant (or even definable) for any cases seen so far. It was a place-holder in the original formulation and has since been removed. Every instance brought to our attention so far has actually been a case of some outside influence which is modeled elsewhere in the simulation and wouldn't be applicable (or worse, double counted) if assumed as the correlation asked for in the original RACM input list.

6.3.2. Develop argument that truncation from below of various cost elements induces a correlation

Keep in mind that the correlation referred to would be the observed, historical data derived correlation coefficients and that the term in question in the model was the possible correlation between the fundamental ability of each WBS element to estimate its basic cost to complete - independent of major problems or any outside agent which during the course of the operational program might come in and simultaneously influence the final realized cost of many/all of the elements together. It is the assertion of the RACM formulation that those other agents/influence factors are addressed individually elsewhere in the model and not ever intended to be included as the potential sources for the terms in the RACM model labeled "correlation". By the way, it is not a case of independent, random truncation of the individual element distributions which was intended to infer could result in "observable" correlation in the historical data bases. What was meant to be conveyed (apparently a slip-of-the-tongue/carelessness occurred) was the simultaneous "collapsing" of the distributions which can occur under the self-fulfilling, "allocated implies spent" premise suggested.

6.4. Shifting the Baseline Cost Estimate to Reflect Additional Sources of Uncertainty

6.4.1. Demonstrate the effect on both mean and variance of cost arising from each type of curve shift

Each of these items has been discussed elsewhere in this paper. As to the exact equations used, I would suggest that the respective cells within the RACM spreadsheet be checked. This would serve two purposes - validating the intended operation while simultaneously checking the coding to implement those intentions. This would be greatly appreciated by the author as well as answer the questions posed here.

6.4.2. Argue that the effects are multiplicative (Order is irrelevant)

Some of the effects/influence factors have a natural order associated with them and are included in the presentation sequence accordingly. Global effects come after the final accumulation of the individual WBS element level considerations. Schedule impacts, being converted here to "equivalent" dollars-to-recover factors, are calculated/presented in the sequence before the rate change terms are displayed (Cost Improvement (CI) and "Inflationary" factors). Management policy/reserve consideration impacts are acting only upon the final accumulation of WBS element level factors (right after the "Rate" term), but are displayed at the end since this is (when properly combined with the global level terms) the "final answer". A different presentation order would not produce a different answer, but would only confuse the interpretation process intended in the decision to display all the steps along the way instead of just the "bottom line". Where order is not dictated by logic, etc., the terms are truly multiplicative in mathematical nature and their order of inclusion is arbitrary - CI and then Rate versus Rate and then CI, for example.

7.0. CALCULATION OF THE RESERVES

There seem to be two issues here that are possibly being confused. Holding reserves and the impact that this can have upon the realizable cost of a program from the prospective of "allocated implies-spent" (the collapsing of some/all of the individual WBS element level distributions) is the significant concept being studied. The algorithm chosen for how to spread any given fund of money back across the chosen WBS structure was developed as a management tool to assist program directors with this chore in some hopefully fair and defensible manner. After the fact, from sensitivity type analyses, it has been observed (not proven) that this produces an "optimal" solution as well in terms of the impact on the predicted overall cost curves that results. The fact that RACM now

contains an option to accept any user specified "spread" of resources allows this effect to be directly measured. The observations made reflect a shift in the final cost curve to the right when any non-balanced spread is suggested. Remember, the amount of the reserve is a constraint - dependent on user inputs indicating desired operating point (Ps for your "bid" point) and a simple statement of management policy (how much of the potential reserve do they wish to hand out on day one/ not utilize as actual program reserves) . It is merely an implied choice of dividing the designated funds back across the WBS structure that is at question. Barring absurd decisions (like giving all the money to one element, etc.), this is a significantly second order effect to the overall problem of reserves/"allocated implies spent" impacts. Mathematically optimum or not, the appeal for this breakback scheme is the defensability/saleability of this particular alternative (placing every element at the same implied "risk") when going back to the WBS element level project engineers, etc.

7.1. For Each Element in the Work Breakdown Structure

I believe the question at point here is the actual formulation which accomplishes the breakback displayed in RACM. The first issue is the size of the potential reserves being discussed. The high end of the WBS element unique funding is taken directly from the final WBS curve ("Rate" curve - #4) at the user specified operating point - Ps. As has been discussed elsewhere, a philosophical assumption has been made that no element should be initially funded below its individual 50% Ps level so as to avoid any possible interference in normal planning, etc. A totally "optimum" solution would be to not initially fund anyone - only pay them for work done after the fact. This is certainly not practical and hence the need for the more reasonable, if somewhat arbitrary, decision chosen. Future developments/enlightenment might easily lead to other approaches/lower limits for specific programs being studied. Obviously this would be trivial to incorporate into future versions of RACM as needed. With those groundrules, a potential for a reserve fund has been defined as the simple difference between these two points - the total dollar amount between the bid point and the 50% Ps point, both being taken from the cumulative rate curve (#4) in the sequence normally displayed. The other dollar difference (between this rate "high" point and the overall bid point from the final (curve # 7) is automatically considered to go into program reserves under all circumstances (labeled non-WBS reserves in RACM outputs). These program dollars have no direct known connection with the individual WBS elements/projects and hence their immediate association with program reserves. This then leaves the second issue, how do we

distribute this now identified potential WBS element reserve fund. If the general recommendation to keep maximum reserves is followed, all WBS elements would be initially given the dollar values associated with their individual 50% Ps level and all the potential reserves would be included with the before mentioned non-WBS element reserves into the overall general program reserve fund. If the user input is non-zero, the user wishes to pass out some of his potential reserve capability, then that split is made and his designated percentage of the potential is assigned to the initial allocation fund and only the remaining portion is added to the general program reserves. The assumption of accomplishing this so that each element will be left on an equal "risk" basis simply implies finding some constant "A1" such that the summation of $A1 \cdot \sigma(i)$ across the WBS element structure equals the dollar amount specified. Therefore $A1 = (\text{extra dollars to be distributed}) / \sum(\sigma(i))$ and each element then is assigned $\text{mean}(i) + A1 \cdot \sigma(i)$ in the breakback (mean & sigma being taken at the "Rate" stage of the process - the curve 4 data base).

7.2. During Each Phase of the Program

I may have lost the original significance of this question, but I will address a few issues that come to mind. The current scheme reaches a solution including both the labor and non-labor elements simultaneously and equally - both are assigned the same level of risk in that initial allocation/breakback. This may be altered by the user definable allocation levels that are permitted elsewhere in the current RACM mechanization. It also may be desirable to alter this balance depending on the users definition of labor and non-labor items (and especially the view of their inherent source of variability being modeled). It could also be affected by the phase of the program (development, production, etc.) under consideration. If the recommendation for holding maximum reserves is followed, all of this becomes essentially a moot point since no dollars will have to be distributed - all elements being given their individual 50% Ps values. It also removes the dependence of the final realizable cost curve on the desired bid/operating point chosen - the initial funding and hence lower limit spending levels are unaffected.

8.0. OTHER CONCERNS

8.1. Training—3 levels

Virtually every sophisticated cost model currently available requires extensive training for proper use. In the cases we are aware of (PRICE, WinSight, etc.) this involves

several days of hands on training in formal classes by representatives of the model developer. Implementation of the RACM process will require formal training, as does any model or tool. Because of differing levels of involvement, we believe that training must be developed at 3 levels as described below. It is assumed that each level of training will have appropriate course materials (manuals, tutorial videos, etc.) that are appropriate to the detail of the course.

1. Management Level Training

A key to successful implementation of RACM is understanding of the process and acceptance of the methodology by both contractor and government program management. When the decision is made to implement RACM on a program the first activity should be a familiarization briefing with program management and staff. A top level summary similar to current RACM presentation should provide adequate background and detail on the RACM methodology and how it will be implemented. The intent of this type of briefing is to familiarize Government and contractor management personnel with RACM process and available tools. This would be a one time briefing, two to three hours in duration.

Before proceeding with implementation it is critical that all levels of management understand and accept RACM's assumptions. Without this level of training the output can be misinterpreted by program management and the data may be implemented incorrectly. For example if a program manager followed RACM's budget allocations at the "most likely" point (50/50 P_S) and try to force all CAMs to rigidly adhere to those budget allocations, then the data is being incorrectly used.

2. Cost Risk Model Analyst Training

Each program will require at least one analyst trained in the details of developing and running the RACM model. This analyst will receive detailed training in developing a RACM model, running sensitivity analyses, analyzing output and interpreting the results of a model. This individual must understand how RACM works and have knowledge of how different factors affect the output of the model.

The cost risk model analyst will be responsible for gathering data, setting up the parameters of the model, inputting data, running the initial analysis and the sensitivity analyses on the individual elements and the overall program. This person will have primary responsibility for the development, maintenance and interpretation of the data. Persons most suited for this would likely be budget analysts, schedulers or business oriented program personnel with a good understanding of computers as well as the specifics of the program to be

analyzed. The RACM data will interface with existing cost and schedule management tools.

This person must develop a good understanding of how the model works, so the training will focus on RACM in detail. We envision hands-on training in a classroom environment, 20 to 30 hours in duration, lead by a trained instructor. The courses could be held at contractor facilities, or at central locations where people from several contractors as well as government personnel could attend.

3. Interviewer

An interviewer will have direct contact with technical personnel on a program. This person will be responsible for validating the technical assumptions in an estimate as well as its basis (historical, engineering estimate, CER, parametric estimate, etc.), and determining a high and a low for each estimate and the P_5 for the high and the low estimate.

The interviewer must develop the skills needed to extract expert opinion and point out inconsistencies in the data. This person will need an understanding of the program requirements and the technical disciplines. This can be a joint role for the risk model analyst, but the person's understanding of the program must be broader than just budget and schedule.

A hands-on training program similar to the level of detail in Cost Risk Analyst Training would be most effective. The interviewer must understand all of the requirements of the cost risk analyst because the interviewer will be responsible for validating the data. Formal training in a seminar format would allow the instructors to go through mock interviews to show different types of responses from technical personnel and the best methods of response to these situations.

As with any model, development of good training material is essential. Also, follow-up training should be available to reinforce skills and focus on specifics as people become more familiar with RACM. The feedback from these sessions can provide valuable information for modifications and enhancements in future releases of RACM. Also, as the model is distributed, call in phone support would be helpful for users that have specific problems or questions. This can reinforce training as the users apply RACM to their specific programs.

8.2. Design of an Experiment

The output of RACM is a probability distribution, not a point estimate. Therefore: experimental data that could be used to test the RACM methodology must be collected and interpreted subject to that fact. This is fundamental and bears further discussion. If

the methodology predicted that a project would be accomplished at cost $\$C$ with a standard deviation of $\$s$, then it would be possible to argue that, subject to the usual assumptions, an observed cost greater than $\$C + \$2s$ or less than $\$C - \$2s$ would indicate that the predicted value has been contradicted (with confidence approximately 95%). This is distinctly not the case with the RACM output, since no such point estimate is predicted. Therefore, a similar, simple test is not meaningful in this case.

Nonetheless, it seems possible to construct an experiment using a large, but perhaps not burdensomely so, data set, that could serve to confirm or reject the RACM output. The main idea in constructing such a test is that the RACM output, a probability distribution on project cost and completion time, describes the likelihood of observing at most any particular project cost or completion time. Thus, all costs and completion times can be considered to be members of a family of random variables that is governed by a family of related probability distributions. The relationship that defines the family is that all probability distributions are the results of the RACM methodology.

Each time the RACM methodology is applied, the actual outcome corresponds to a percentile of the output probability distribution. It is straightforward to show that the percentiles of any distribution are themselves uniformly distributed. These two facts suggest that an appropriate statistical test of the RACM distribution is to test whether the observed percentiles come from a uniform distribution. For the Naval Air Systems (NavAir) project, each airplane or group of airplanes produced may be sufficiently uncorrelated with other airplanes or groups so that each outcome may be considered a separate data point. In this way, the RACM methodology may be tested using the NavAir project data.

In general, however, a test of the methodology requires the output of many uncorrelated projects to be expressed as percentiles. Those percentiles are governed by a uniform distribution if the RACM process is providing appropriate probability distributions. The test can be accomplished using elementary goodness-of-fit techniques.

ANNEX: CURRENT APPROACHES TO BUDGET RISK ANALYSIS

1.0 METHODS FOR FINDING TOTAL COST AND TIME-TO-COMPLETE DISTRIBUTIONS

1.1 C-Risk [Smith and Book¹]

This approach quantifies *technical, programmatic, and schedule* risks. These estimates are used to adjust probability distributions of costs of WBS elements. The underlying model is that the cost of each element in the WBS is a normal random variable (mean = best estimate, variance = noise variance) plus a right triangle density for "risk component." These marginal distributions are related by a simple estimate of pairwise correlation coefficients based on assessed values of fraction of "new technology" (P_i) and a "cost growth sensitivity factor" (k) for each WBS element. (The underlying assumption is that all "new" fractions of all WBS elements are perfectly correlated. The cost growth factor (k) is not clearly described.) Simulation is used to find percentiles of total cost.

Comment: This approach is an attempt to model specific factors that are believed to cause cost increases. There is, however, no reason to believe that the sum of a normal and a triangular distribution is a reasonable way to model the cost of an element of the WBS, other than the fact that such a model yields the skewness that is believed to exist in the distribution of costs. The result is gotten by simulation rather than analytically. There is no universally accepted procedure for performing such a simulation.²

1.2 F-Risk [Abramson and Young, Young and Abramson³]

This is an analytic approach, that fits a lognormal distribution to total cost. WBS elements are distributed triangularly determined by expert judgment: the lower bound is

¹ P.L. Smith and Stephen A. Book, "Reducing Subjective Guesswork and Maintaining Traceability When Estimating The 'Risk' Associated With a Cost Estimate," The Aerospace Corporation, Los Angeles, California, 1992.

² Note added by IDA: the statement is false; this problem was solved by Philip M. Lurie and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions," *Management Science*, Vol. 44, No. 2, February 1998.

³ R.L. Abramson and Philip H. Young, "FRISKEM—Formal Risk Evaluation Methodology," The Aerospace Corporation, Los Angeles, California, 1990; Philip H. Young and R.L. Abramson, "Using Risk-Impact Drivers to Form WBS-Element Cost Means and Variances," The Aerospace Corporation, Los Angeles, California, 1991.

the assessed minimum cost to complete the element, the maximum likelihood value is the assessed best estimate, and the upper bound is the assessed maximum cost. Pairwise correlation coefficients are required for the elements of the WBS.

Comment: This approach is similar in spirit to some aspects of the Lockheed approach, which is discussed in more detail below. The objective of the FRISK approach is to avoid the difficulties of simulation using only marginal distributions linked by correlation coefficients. The expected value and the variance of the total cost are determined by the expectations and variances of the marginal distributions of the WBS elements along with the correlation coefficients. Therefore, any distribution that is completely specified by these two parameters will agree up to the second moment with the total cost; and such distributions are uniquely determined by the assessed data. The basis of FRISK is to restrict attention to such random variables. FRISK chooses, in particular, the lognormal, which is a two-parameter distribution. There is no reason to believe that the total cost is a lognormal variable, however, other than the belief that total cost is skewed to the right.

1.3. Monte Carlo Simulation of Total Cost When WBS-Element Costs Are Correlated [Book and Young⁴]

These papers develop a method for simulating the sum of a collection of correlated WBS element marginal distributions. Given the assessed marginals, standardized variates are formed. The correlation matrix is factored to create linear combinations of independent standardized variables that have the correct correlations. Thus, independent simulations of the standardized variables are combined by the factored correlation matrix and added. The sum is a simulation of the actual total cost.

Comment: As noted in the comments to C-Risk, above, there is no generally accepted solution to the problem of simulating a sum of correlated random variables given only marginal distributions and a matrix of correlation coefficients.⁵ This approach is an attempt to solve that general problem. The approach is indeed successful in

⁴ Stephen A. Book and Philip H. Young, "Monte Carlo Generation of Total Cost Distributions When WBS-Element Costs Are Correlated," The Aerospace Corporation, Los Angeles, California, 1990; Stephen A. Book and Philip H. Young, "Monte Carlo Simulation of Project Schedule Duration When Activity Times Are Correlated," The Aerospace Corporation, Los Angeles, California, 1992.

⁵ Note added by IDA: once again, this problem was solved by Philip M. Lurie and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions."

capturing the mean and variance of the total cost correctly. However, it does not express the total cost as a sum of random variables that are distributed by the assessed marginals. Instead, it transforms the samples from the assessed marginals by the square root of the correlation coefficient matrix and adds those transformed samples. Thus, what is being added to obtain the total cost is not a sum of observations from the WBS element distributions. The authors recognize this, of course, and suggest the need for further research.

1.4. Joint Probability Model for Cost and Schedule Uncertainties [Garvey and Taub⁶]

The purpose of this paper is to suggest an approach to modeling the joint behavior of cost-to-complete and time-to-complete a project. The approach fits WBS estimates to a joint lognormal distribution.

Comment: The main idea here is that the estimates of the means and variances of the WBS elements along with the correlation coefficients are sufficient to determine the parameters of the bivariate lognormal. Again, the lognormal is chosen because of the belief that the cost and time are skewed to the right. Further, the bivariate lognormal implies that the marginals are lognormal. This is consistent with the assumptions typically made in modeling the time-to-complete and the cost-to-complete separately (as has been noted above). Yet there is no fundamental basis for this assumption. Indeed, the authors note that the bivariate lognormal assumption failed a goodness-of-fit test on a set of simulated data. That need not be a grave fault, although it is suggestive of a lack of applicability of this model, particularly because it is the asymptotic behavior of the lognormal (the tail of the distribution) that is apparently not consistent with the simulated behavior. The idea of attempting to trade off cost and schedule is of interest and this seems to be the first approach to performing such tradeoffs using WBS cost models.

1.5. ACOP—Analytic Cost Probability Model [Garvey⁷]

This is another approach to determining the total cost probability distribution analytically given a correlated set of WBS element distributions. Each element in the

⁶ Paul R. Garvey and Audrey E. Taub, "A Joint Probability Model for Cost and Schedule Uncertainties," MITRE Corporation, Bedford, Massachusetts, 1992.

⁷ Paul R. Garvey, "A General Analytic Approach to System Cost Uncertainty Analysis," *Cost Analysis and Estimating—Tools and Techniques*, edited by W.R. Greer and D.A. Nussbaum, Springer-Verlag, 1990.

WBS has an uncertain cost vector, where each component of the vector measures a different category of cost risk. Correlation coefficients must be given for all pairs of costs in each category, r_{ikj} , the correlation between Y_{ij} and Y_{kj} , the cost of category j for elements i and k . The expected value and variance of the total cost are determined by the assessed information, as noted in the discussions above. Bounds are provided for the total cost variance.

Comment: This approach is similar to the other analytic methods. The idea is by now familiar: all that can be said with precision, based on the assessments of the WBS elements, is that the mean and variance of the sum is determined. The actual distribution of the sum is unknown. In this paper, Garvey suggests that it may be appropriate to consider the sum to be normal, but "in circumstances where normality does not describe the shape of the [total cost] distribution (e.g., if the [total cost] distribution is skewed) an approximation of the total system cost probability distribution through formal simulation techniques may be warranted" (p. 167). Thus, one is left either to make some very strong assumptions about the shape of the distribution one is attempting to determine or to perform simulation techniques that are not necessarily universally accepted.

1.6. Cost Risk Analysis of the Strategic Defense System [Gupta et al.⁸]

This model assumes that the total project cost is a function of two separate and independent processes, a cost estimating process that generates the random variable C_i for the i th element of the WBS and a schedule/technical process that generates a random variable S_i that acts as a multiplicative correction to the cost C_i . Thus the total project cost is given by $S_i C_i S_i$, the sum of the product of the costs multiplied by the correction factor. The cost estimating process is uncertain because the estimates of cost provided for each element of the WBS are derived from a statistical model, known as a Cost Estimating Relationship (CER). The Cost Estimating Relationship expresses cost as a function of the attributes of a WBS element. These relationships are based on historical data. The schedule/technical process is uncertain because the process modeled is that which must evolve a new design or incorporate a new constraint or respond to some change in the requirements of the system.

⁸ Shishu Gupta, David Olsen, David Hudak and Jennifer Keenan, "Cost Risk Analysis of the Strategic Defense System," TR-9042-2, Revision 1, TASC, Arlington, Virginia, 1992.

The TASC approach is to specify the probability distributions of the random variables C and S for each element of the WBS and apply simulation to yield the probability distribution of the total cost.

Comment: One may question at the outset the notion of modeling the cost estimating process. There may be confusion here because of the weak exposition in this report, yet it seems that what is wanted is the probability distribution on the cost itself, not on the way it is estimated. Even if the correct random variable is being assessed, why match the *mode* of a triangular distribution for the random variable C to the *mean* of a CER? The report provides, in its Appendix C, an argument for doing so that is not convincing; one ought not simply identify means with modes or medians as one wishes with no real justification.

The assessment of technical/schedule risk is based on either historical data (the so-called comparative schedule approach) or direct assessment (the so-called technical assessment approach). Both approaches are arbitrary. The first approach relies on similarity of the current program to past programs and uses ratios to determine the end points of a triangular distribution of mode 1. That is, the ratio of the smallest actual completion time among all similar programs to the point estimate of the scheduled time for the current project is made the lower endpoint of the triangular probability density function. The larger endpoint is the ratio of the largest completion time to the scheduled time. The mode of the distribution is arbitrarily set to 1. There is no convincing justification for this approach. The second approach assesses an arbitrary set of risk scores over six categories of technical risk that are weighted to produce an average score. The average is mapped by two equations into two multipliers that are assumed to be 10 and 90 points of a triangular distribution. There is no theory that suggests why any of this is appropriate; the report cited offers no convincing explanations.

The simulation of the total cost is based on the independence assumptions invoked in this approach. Independence of the WBS elements, whether C or S or their product is surely a questionable assumption.

1.7 Conclusion to Section 1.0

None of these approaches seems to be satisfactory with respect to Lockheed's objectives. What seems particularly remarkable is that not one of these approaches attempts to model the effect of management policies on the cost-to-complete and the time-to-complete a project. It is as if all the approaches share a fundamental, unstated,

assumption that project time and cost is governed by some immutable set of natural laws. What follows from such an (implicit) assumption is that what one must model are the apparent effects of such laws. An alternative to such a relatively passive point of view is that projects are controllable by management policies, and what ought to be modeled is the way policies can change the cost and duration outcomes.

2.0. METHODS FOR FINDING INDIVIDUAL WBS ELEMENT DISTRIBUTIONS

2.1. Risk-Driver Method [Abramson and Book, Young and Abramson⁹]

The purpose of these papers is to provide a method for assessing a triangular distribution for the cost of an individual WBS element. The approach uses the "Maxwell Risk-Driver Assessment Framework." Three parameters must be estimated: H, M, L, the largest, most likely, and least cost values. M, the maximum likelihood estimate, is assessed directly. H, the maximum value, is given by $H = RM$, where R is a risk multiplier that is found by application of the Maxwell Framework. L, the lowest bound, is also assessed directly. The Maxwell Framework identifies categories of risk-drivers (e.g., required technical advancement, complexity, dependencies, etc.) and (typically five) levels of risk (from low to high). Experts assign scores to each category for a particular WBS element. Expert opinion is combined and a representative multiplier is found by an averaging technique.

Comment: This approach attempts to make more analytic and transparent the process of defining triangular WBS element distributions. The essential difficulty with such a method is the actual creation of the multiplier. If experts think that costs are actually multiplicatively affected by these so-called risk-drivers, then the method may be a reasonable assessment approach. One wonders whether the detailed assessments and translation of assessments into multipliers provides a better estimate than simply assessing directly the maximum possible cost.

⁹ R.L. Abramson and Stephen A. Book, "A Quantification Structure for Assessing Risk-Impact Drivers," The Aerospace Corporation, Los Angeles, California, 1990; Philip H. Young and R.L. Abramson, "Using Risk-Impact Drivers to Form WBS-Element Cost Means and Variances," The Aerospace Corporation, Los Angeles, California, 1991.

2.2. Cost Risk Analysis of the Strategic Defense System [Gupta et al.¹⁰]

The approach to individual WBS elements was noted above. The main idea here is that the cost of each WBS element is the product of two random variables, the cost estimate and the schedule/technical multiplier. The cost estimate is derived from a cost estimating relationship, while the schedule/technical multiplier is assessed using either historical data or directly assessed using a set of technical risk multipliers.

Comment: This approach was discussed above. Very little justification was given for this approach in the paper cited.

2.3 Conclusion to Section 2.0

There is no single accepted technique for the assessment of the cost-to-complete individual WBS elements. Most researchers believe that WBS element costs are skewed to the right; that triangular or beta distributions are reasonable models; and that the parameters of these model distributions can be assessed by experts or found in historical data. It is also believed that individual WBS element costs are correlated and those correlations can be assessed by experts. In order to do this appropriately and accurately, the experts must understand that they are providing either parameters of probability distributions or fractiles of those distributions. This is not what actually occurs in practice. In practice, experts provide answers to questions posed by analysts. Those answers are then *interpreted* as statements about the probability distribution of the WBS element cost. The interpretation process can easily introduce errors and inaccuracies.

¹⁰ Shishu Gupta, David Olsen, David Hudak and Jennifer Keenan, "Cost Risk Analysis of the Strategic Defense System."

APPENDIX B

PROOF OF OPTIMALITY OF THE RACM BUDGET ALLOCATION

This appendix demonstrates the optimality of an equi-percentile budget allocation, as implemented in RACM. At the start of program execution, upper management allocates to each WBS-element manager a fraction of his or her total budget. Under the MAIMS principle, the WBS-element manager will spend *at least* this initial allocation. This phenomenon causes shifts in the density function and the CDF of cost.

Let $f(x)$ denote the density function and $F(x)$ the CDF of cost, in a particular WBS element, prior to budget allocation. Let $f^*(x)$ and $F^*(x)$ denote the corresponding functions after budget allocation. If we let θ denote the budget level, the two sets of functions are related in the following way:

$$\begin{aligned} f^*(x) &= 0, & F^*(x) &= 0 & \text{if } x < \theta; \\ f^*(x) &= F(\theta), & F^*(x) &= F(\theta) & \text{if } x = \theta; \\ f^*(x) &= f(x), & F^*(x) &= F(x) & \text{if } x > \theta. \end{aligned} \tag{B-1}$$

CALCULATION OF EXPECTED COST

We are interested in the mean of the transformed distribution, which we denote as $\phi(\theta)$. Defining the survivor function as $S(x) = 1 - F(x)$:

$$\phi(\theta) = [1 - S(\theta)] \times \theta + S(\theta) \times E(X | X > \theta). \tag{B-2}$$

The final term in equation (B-2) may be computed as follows:

$$E(X | X > \theta) = \frac{\int_{\theta}^{\infty} x f(x) dx}{\int_{\theta}^{\infty} f(x) dx} = \frac{\int_{\theta}^{\infty} x f(x) dx}{S(\theta)}. \tag{B-3}$$

The numerator of equation (B-3) may be integrated by parts, yielding the following result:¹

$$E(X | X > \theta) = \left[\theta S(\theta) + \int_{\theta}^{\infty} S(x) dx \right] / S(\theta) = \theta + \left[\int_{\theta}^{\infty} S(x) dx / S(\theta) \right]. \quad (\text{B-4})$$

Substituting equation (B-4) into equation (B-2) yields a convenient expression for the mean of the transformed distribution:²

$$\phi(\theta) = \theta + \int_{\theta}^{\infty} S(x) dx. \quad (\text{B-5})$$

We may compute the first two derivatives of $\phi(\theta)$ with respect to the budget level, θ :

$$\begin{aligned} \phi'(\theta) &= 1 - S(\theta) = F(\theta), \\ \phi''(\theta) &= F'(\theta) = f(\theta) > 0. \end{aligned}$$

By virtue of the strictly positive second derivative, the mean is a strictly convex function of the budget level, θ . We note that a convex function always lies above its tangent:³

$$\phi(\theta + h) > \phi(\theta) + h\phi'(\theta) = \phi(\theta) + hF(\theta) \text{ for all } h \neq 0. \quad (\text{B-6})$$

BUDGET ALLOCATION THEOREM

We have n WBS elements, and we begin from an equi-percentile budget allocation (*prior* to transformation):

$$F_1(\theta_1) = F_2(\theta_2) = \dots = F_n(\theta_n) = \bar{F} \quad (\text{B-7})$$

¹ Use the transformation: $u = -x$ and $v = S(x)$. Note that, for positive random variables [$F(0) = 0$] with finite mean, $0 \leq xS(x) = \int_x^{\infty} xf(y)dy \leq \int_x^{\infty} yf(y)dy \leq \int_0^{\infty} yf(y)dy = \mu < \infty$. It follows that $\lim_{x \rightarrow \infty} xS(x) = 0$.

² When applied to the normal distribution, this equation yields the same result obtained using a different methodology by Ronald A. Welch, "Piece-Wise Linear Transformation of a Gaussian Random Variable," Lockheed Corporation, Sunnyvale, California, 15 July 1969. The result for normal distributions is also contained in G. S. Maddala, *Limited-Dependent and Qualitative Variables in Econometrics*, Cambridge: Cambridge University Press, 1983, pp. 158-160. Welch also derives the variance of a transformed normal variable. We have verified the latter result using methods developed in Norman L. Johnson and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions-1*, chapter 13.

³ J. M. Ortega and W. C. Rheinboldt, *Iterative Solution of Nonlinear Equations in Several Variables*, New York: Academic Press, 1970, pp. 84-85.

where \bar{F} is the common percentile that arises from the various budget levels $\{\theta_1, \dots, \theta_n\}$. Starting from this position, we reallocate the budgets in the following way: the budget for the i th WBS element is increased by the amount h_i , where not all $\{h_i\}$ are equal to zero. The increments h_i may be either positive or negative and, in fact, they sum to zero:

$$\sum_{i=1}^n h_i = 0. \quad (\text{B-8})$$

The respective means of the transformed distributions are now given by $\phi_i(\theta_i + h_i)$ for $i = 1, \dots, n$. We now show that the mean of total cost is always larger after perturbation from the equi-percentile position.

Theorem: $\sum_{i=1}^n \phi_i(\theta_i + h_i) > \sum_{i=1}^n \phi_i(\theta_i).$

Proof: By equations (B-6) through (B-8),

$$\sum_{i=1}^n \phi_i(\theta_i + h_i) > \sum_{i=1}^n [\phi_i(\theta_i) + h_i F_i(\theta_i)] = \sum_{i=1}^n [\phi_i(\theta_i) + h_i \bar{F}] = \sum_{i=1}^n \phi_i(\theta_i) + \bar{F} \sum_{i=1}^n h_i = \sum_{i=1}^n \phi_i(\theta_i).$$

Intuitively, the theorem follows from the convexity demonstrated in equation (B-6). The positive increments h_i to certain WBS elements lead to greater increases in the mean cost than do the negative increments to the remaining WBS elements.

APPENDIX C

PROPERTIES OF THE BETA DISTRIBUTION

BASIC PROPERTIES

The beta distribution is specified by its low and high endpoints, l and h , and two additional shape parameters, which we denote as $\alpha > 0$ and $\beta > 0$. The density function is given by:

$$f(y) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)(h-l)^{\alpha+\beta-1}} (y-l)^{\alpha-1} (h-y)^{\beta-1}, \quad (\text{C-1})$$

for $l \leq y \leq h$.

The mean of the beta distribution depends on α and β only through their ratio:

$$\mu = l + (h-l) \frac{\alpha}{(\alpha + \beta)} = l + \left[(h-l) / \left\{ 1 + (\beta/\alpha) \right\} \right]. \quad (\text{C-2})$$

The variance of the beta distribution is equal to:

$$\sigma^2 = (h-l)^2 \frac{\alpha\beta}{(\alpha + \beta)^2 (\alpha + \beta + 1)}. \quad (\text{C-3})$$

The beta distribution has a mode (i.e., a most-likely value) only if $\alpha > 1$ and $\beta > 1$. However, the mode depends on the individual values of α and β , not just their ratio:¹

$$m = l + (h-l) \frac{\alpha - 1}{(\alpha + \beta - 2)}. \quad (\text{C-4})$$

¹ Norman L. Johnson and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions-2*, chapter 24, p. 41; or Philip M. Lurie, Matthew S. Goldberg, and Mitchell S. Robinson, "A Handbook of Cost-Risk Analysis Methods," Paper P-2734, Institute for Defense Analyses, April 1993, p. 9.

The beta distribution has an anti-mode (i.e., a *least*-likely value) if $\alpha < 1$ and $\beta < 1$. The beta distribution is J-shaped if $0 < \beta < 1$ and $\alpha > 1$, and it is backwards J-shaped if $0 < \alpha < 1$ and $\beta > 1$.

PARAMETER ESTIMATION

The “method of moments” may be used to estimate the parameters of the beta distribution. If the user supplies the low and high endpoints, then the parameters l and h are known immediately. Suppose that the user also supplies the mean and variance. Then by inverting equations (C-2) and (C-3), the following solutions for α and β are obtained:²

$$\alpha = \frac{(\mu - l)^2 (h - \mu)}{\sigma^2 (h - l)} - \frac{(\mu - l)}{(h - l)}, \quad (\text{C-5})$$

$$\beta = \frac{\alpha (h - \mu)}{(\mu - l)}. \quad (\text{C-6})$$

Alternatively, if the user supplies the mode (rather than the mean) and variance, the remaining parameters α and β may be determined by inverting equations (C-3) and (C-4). However, there is no explicit closed-form solution in this case. Instead, the equations must be inverted numerically using Newton’s or some other method.

A final case is relevant to the *BETASUBJ* function found in both the @RISK for Microsoft Excel and @RISK for Microsoft Project software. In this case, the user supplies the low and high endpoints, as well as the mean and the mode (but not the variance). The parameters α and β may be determined by inverting equations (C-2) and (C-4). The following closed-form solution is available:³

$$\alpha = \frac{(\mu - l)(2m - l - h)}{(m - \mu)(h - l)}, \quad (\text{C-7})$$

² Norman L. Johnson and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions*–2, chapter 24, p. 44. The slightly more compact solution shown here is taken from Philip M. Lurie, Matthew S. Goldberg, and Mitchell S. Robinson, “A Handbook of Cost-Risk Analysis Methods,” p. 9.

³ Palisade Corporation, *Guide to Using @RISK: Risk Analysis and Simulation Add-In for Microsoft Excel or Lotus 1-2-3*, Newfield, New York, July 1997, p. 221; *Guide to Using @RISK for Project: Risk Analysis and Simulation Add-In for Microsoft Project*, Newfield, New York, July 1997, p. 185.

$$\beta = \frac{\alpha (h - \mu)}{(\mu - l)} . \quad (C-8)$$

Note that the ratio of β to α is the same in equations (C-6) and (C-8). This result follows because, for the beta distribution, the mean directly maps into the parameter ratio. Thus, for the two cases in which the mean is assumed to be known, the parameter ratio is determined as well.

APPENDIX D

PROPERTIES OF THE TRIANGULAR DISTRIBUTION

BASIC PROPERTIES

The triangular distribution is fully specified by its low and high endpoints, denoted l and h , and its mode, denoted m .

The mean and variance of the triangular distribution are:

$$\mu = \frac{l + m + h}{3} \quad (\text{D-1})$$

and

$$\sigma^2 = \frac{(h-l)^2 + (m-h)(m-l)}{18}. \quad (\text{D-2})$$

The CDF of the triangular distribution is:

$$F(x) = \begin{cases} (x-l)^2 / [(h-l)(m-l)] & \text{for } l < x < m, \\ 1 - \{(h-x)^2 / [(h-l)(h-m)]\} & \text{for } m < x < h. \end{cases} \quad (\text{D-3})$$

The height of the triangle at the mode is $2/(h-l)$.

A triangular distribution is said to be *right-skewed* if $(h-m) > (m-l)$ or, equivalently, $m < (l+h)/2$. Under this condition, the area to the left of the mode is less than one-half:

$$F(m) = \frac{(m-l)}{(h-l)} = \frac{(m-l)}{[(h-m) + (m-l)]} < \frac{(m-l)}{[(m-l) + (m-l)]} = 0.5. \quad (\text{D-4})$$

Thus the median, denoted x_5 , exceeds the mode:

$$F(m) < 0.5 = F(x_5) \rightarrow m < x_5. \quad (\text{D-5})$$

A right-skewed triangular distribution is depicted in Figure D-1.




Figure D-1. Right-Skewed Triangular Distribution

PARAMETER ESTIMATION

User input does not always consist of the canonical parameters l , h , and m . Table D-1 shows the five cases considered, depending upon which of these parameters are provided by the user. In the table, x_1 is a percentile point such that the area to the left is p_1 , and x_2 is another percentile point such that the area to the left is p_2 (e.g., the 10th and 90th percentiles).

Table D-1. Cases for Estimation of Triangular Distribution

Case Number	Known Parameters	Unknown Parameters
Known endpoints:		
I	l, h, μ	m
II	$l, h, x_{.5}$	m
Unknown endpoints:		
III	x_1, p_1, x_2, p_2, m	l, h
IV	x_1, p_1, x_2, p_2, μ	l, h, m
V	$x_1, p_1, x_2, p_2, x_{.5}$	l, h, m

Case I: Endpoints and Mean

We can solve for the mode given the low and high endpoints and either the mean or the median. First, suppose we are given l , h , and μ . Then from equation (D-1):

$$m = 3\mu - l - h. \quad (D-6)$$

Case II: Endpoints and Median

Next, suppose we are given l , h , and x_5 (i.e., we are given the median rather than the mean). Suppose also that the distribution is right-skewed, so that $m < x_5$. The median satisfies the following equation:

$$F(x_5) = 1 - \{(h - x_5)^2 / [(h - l)(h - m)]\} = 0.5. \quad (D-7)$$

Inverting equation (D-7) yields a solution for the mode in terms of the median:

$$m = h - [2 \times (h - x_5)^2 / (h - l)]. \quad (D-8)$$

Finally, we note that a different solution would obtain if the distribution were left-skewed, so that $m > x_5$. The median would satisfy the following equation:

$$F(x_5) = (x_5 - l)^2 / [(h - l)(m - l)] = 0.5. \quad (D-9)$$

Inverting equation (D-9) yields a solution for the mode in terms of the median:

$$m = l + [2 \times (x_5 - l)^2 / (h - l)]. \quad (D-10)$$

Case III: Two Percentile Points and Mode

The triangular distribution may also be specified by two percentile points and its mode. We denote the percentile points as (x_1, p_1) and (x_2, p_2) . Then from equation (D-3), the following two equations must hold:

$$p_1 = (x_1 - l)^2 / [(h - l)(m - l)] \quad (D-11)$$

and

$$p_2 = 1 - \{(h - x_2)^2 / [(h - l)(h - m)]\}. \quad (D-12)$$

These equations may be expanded as follows:

$$0 = f_1(l, h) = -x_1^2 + (2x_1 - mp_1)l + mp_1h - p_1lh - (1 - p_1)l^2 \quad (\text{D-13})$$

and

$$0 = f_2(l, h) = -x_2^2 + (1 - p_2)ml + [2x_2 - m(1 - p_2)]h - (1 - p_2)lh - p_2h^2. \quad (\text{D-14})$$

The mode, m , is assumed to be known. Equations (D-13) and (D-14) may be solved for the two unknowns, l and h , using Newton's method.¹ Given estimates (l^k, h^k) at iteration k , the following update is performed:

$$\begin{pmatrix} l^{k+1} \\ h^{k+1} \end{pmatrix} = \begin{pmatrix} l^k \\ h^k \end{pmatrix} - \begin{pmatrix} (2x_1 - mp_1) - p_1h^k - 2(1 - p_1)l^k & (m - l^k)p_1 \\ (m - h^k)(1 - p_2) & [2x_2 - m(1 - p_2)] - (1 - p_2)l^k - 2p_2h^k \end{pmatrix}^{-1} \begin{pmatrix} f_1(l^k, h^k) \\ f_2(l^k, h^k) \end{pmatrix}$$

Because equations (D-13) and (D-14) are quadratic, there is some danger of converging to an extraneous solution, e.g., one in which $x_1 < l$ or $h < x_2$. Extraneous solutions may be avoided by a judicious choice of starting values. We recommend the starting values $l^0 = 2x_1 - m$ and $h^0 = 2x_2 - m$. Additional convergence guarantees are available using quasi-Newton methods, though these should not normally be necessary.²

Example

Consider a triangular distribution with the parameters shown in Table D-2. Using Newton's method, convergence to three decimal places was achieved after a mere four iterations.

¹ William H. Press, Brian P. Flannery, Saul A. Teukolsky, and Willam T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge: Cambridge University Press, second edition, 1992, chapter 9.6.

² *Ibid.*, chapter 9.7; or J. E. Dennis and Robert B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Englewood Cliffs, N.J.: Prentice-Hall, 1983, chapter 6.

Table D-2. Parameters of Triangular Distribution

Parameter		Value
Problem data:	10% point	0.0
	mode	1.0
	90% point	3.0
Starting values:	low endpoint	-1.0
	high endpoint	5.0
Problem solution:	low endpoint	-1.052
	high endpoint	4.343

Using @RISK Software

By default, the @RISK software requires the user to specify the triangular distribution in terms of its low and high endpoints and its mode—in our notation, $\text{RiskTriang}(l, m, h)$. Optionally, @RISK allows the user to specify the triangular distribution in terms of two percentile points and the mode. To do so, the user must declare the random variable either as $\text{RiskTrigen}(x_1, m, x_2, p_1, p_2)$ or as $\text{RiskTrigenMulti}(x_1, m, x_2, p_1, p_2)$.³ The @RISK algorithm for estimating l and h is described in Annex D-1. @RISK does not immediately return its estimates of l and h . However, these estimates may be recovered by simulating a large sample (e.g., 10,000 iterations) from each input cell, and simply observing the empirical low and high values. We have verified through several examples that @RISK's solution essentially agrees with our own (i.e., the one obtained using Newton's method).

Using Crystal Ball Software

By default, the Crystal Ball software requires the user to specify the triangular distribution in terms of its low and high endpoints and its mode. Optionally, Crystal Ball allows the user to specify the triangular distribution in terms of two percentile points.⁴ Unfortunately, Crystal Ball then sets the mode equal to the midpoint of the two percentile points, $m = (x_1 + x_2)/2$. Thus, Crystal Ball's feature is of no value for skewed triangular

³ *Guide to Using @RISK: Risk Analysis and Simulation Add-In for Microsoft Excel or Lotus 1-2-3*, Palisade Corporation, Newfield, New York, July 1997, p.176. According to Palisade Corporation technical support: "The 'multi' feature for the @RISK functions allows the use of @RISK functions in multi-way what-if analysis in TopRank (another decision tool product). There is no difference in how these functions are sampled."

⁴ *Crystal Ball: Forecasting and Risk Analysis for Spreadsheet Users, Version 4.0*, Decisioneering, Inc., Denver, Colorado, March 1996, p. 64: "You can use the 'Parms' or Alternative Parameter menu found in the upper-right corner of the dialog box to specify a triangular distribution using different combinations of parameters." See also pp. 118-119.

distributions. In Annex D-2, we show a series of screens to illustrate Crystal Ball's behavior in this situation. We have, however, verified through several examples that Crystal Ball's solution agrees with our own (i.e., the one obtained using Newton's method) in the special case of symmetric triangular distributions.

Case IV: Two Percentile Points and Mean

Yet another option is to specify two percentile points and the mean. We now have three unknown parameters, l , h , and m . We must find a corresponding set of three equations. We can repeat equations (D-13) and (D-14) except that the mode, m , is added to the list of unknowns:

$$0 = f_1(l, h, m) = -x_1^2 + (2x_1 - mp_1)l + mp_1h - p_1lh - (1 - p_1)l^2 \quad (\text{D-15})$$

and

$$0 = f_2(l, h, m) = -x_2^2 + (1 - p_2)ml + [2x_2 - m(1 - p_2)]h - (1 - p_2)lh - p_2h^2. \quad (\text{D-16})$$

The third equation uses information on the mean. It suffices to use equation (D-6), which is repeated here for convenience:

$$0 = f_3(l, h, m) = l + h + m - 3\mu. \quad (\text{D-17})$$

A three-dimensional version of Newton's method can be applied to solve for l , h , and m using equations (D-15) through (D-17). Given estimates (l^k, h^k, m^k) at iteration k , the following update is performed:

$$\begin{pmatrix} l^{k+1} \\ h^{k+1} \\ m^{k+1} \end{pmatrix} = \begin{pmatrix} l^k \\ h^k \\ m^k \end{pmatrix} - \begin{pmatrix} (2x_1 - mp_1) - p_1h^k - 2(1 - p_1)l^k & (m - l^k)p_1 & (h - l)p_1 \\ (m - h^k)(1 - p_2) & [2x_2 - m(1 - p_2)] - (1 - p_2)l^k - 2p_2h^k & (l - h)(1 - p_2) \\ 1 & 1 & 1 \end{pmatrix}^{-1} \begin{pmatrix} f_1(l^k, h^k, m^k) \\ f_2(l^k, h^k, m^k) \\ f_3(l^k, h^k, m^k) \end{pmatrix}$$

Because equations (D-15) and (D-16) are quadratic, there is again a danger of converging to an extraneous solution. We recommend the starting values $l^0 = 2x_1 - \mu$, $h^0 = 2x_2 - \mu$, and $m^0 = \mu$.

Case V: Three Percentiles

The final option we considered is to specify three percentiles. The specific case we explored has the median as one of the three percentiles, although this is not absolutely necessary. The two tail percentiles are represented as (x_1, p_1) and (x_2, p_2) , respectively, and the median is similarly represented as $(x_5, .5)$. We assume that one percentile falls into each tail, so $l < x_1 < x_5 < x_2 < h$. We also assume that the tail probabilities are equal, so that $p_1 + p_2 = 1$ (e.g., the 10th and 90th percentiles).

We must again find a corresponding set of three equations. Equations (D-15) and (D-16) are still valid for the two tail percentiles. The third equation uses information on the median. The distribution will be right-skewed if $(x_2 - x_5) > (x_5 - x_1)$ or, equivalently, $x_5 < (x_1 + x_2)/2$. Under this condition it suffices to use equation (D-7), which is repeated here for convenience:

$$0 = f_4(l, h, m) = \{(h - x_5)^2 / [(h - l)(h - m)]\} - 0.5. \quad (D-18)$$

If the distribution is left-skewed, we use equation (D-9) instead:

$$0 = f_5(l, h, m) = (x_5 - l)^2 / [(h - l)(m - l)] - 0.5. \quad (D-19)$$

A three-dimensional version of Newton's method can be applied to solve for l , h , and m using equations (D-15), (D-16), and either (D-18) or (D-19). We leave it to the interested reader to develop the appropriate update equation.

TRANSFORMED DISTRIBUTION

This section considers a triangular distribution that is transformed in the manner described in Appendix B, so the area to the left of a threshold value θ is replaced by a mass point at θ . From equation (B-5), the mean of the transformed distribution may be written in terms of the survivor function, $S(x) = 1 - F(x)$:

$$\phi(\theta) = \theta + \int_{\theta}^h S(x) dx. \quad (D-20)$$

This expression may be evaluated using the triangular CDF as presented previously in equation (D-3). For $\theta \geq m$, the result is:

$$\phi(\theta) = \theta + \frac{(h - \theta)^3}{3(h - l)(h - m)}. \quad (\text{D-21})$$

For $\theta < m$, the result is:

$$\phi(\theta) = \mu + \frac{(\theta - l)^3}{3(h - l)(m - l)}. \quad (\text{D-22})$$

ANNEX D-1: @RISK ALGORITHM FOR CASE III

The triangular distribution may be specified by its mode, m , and two percentile points, (x_1, p_1) and (x_2, p_2) . The unknown parameters are the low and high endpoints, l and h . We introduce an additional unknown parameter, the height of the triangle at its mode, $\lambda = 2/(h-l)$.

Given a tentative estimate of λ , it is possible to solve for endpoints l and h such that the area to the left of x_1 is p_1 , and the area to the left of x_2 is p_2 . However, these estimates may require further adjustment if the resulting total area under the triangle differs from 1.0. To estimate l and h , we first express equations (D-11) and (D-12) in terms of the new parameter λ :

$$0 = (x_1 - l)^2 - 2(m - l)p_1/\lambda \quad (\text{D-23})$$

and

$$0 = (h - x_2)^2 - 2(h - m)(1 - p_2)/\lambda. \quad (\text{D-24})$$

Now introduce the transformations $y_1 = (x_1 - l) > 0$ and $y_2 = (h - x_2) > 0$, so equations (D-23) and (D-24) become, respectively:

$$0 = \lambda y_1^2 - 2p_1 y_1 - 2(m - x_1)p_1 \quad (\text{D-25})$$

and

$$0 = \lambda y_2^2 - 2(1 - p_2)y_2 - 2(x_2 - m)(1 - p_2). \quad (\text{D-26})$$

Equations (D-25) and (D-26) may be solved by the quadratic formula. To ensure positive solutions for y_1 and y_2 , we choose the larger of the two quadratic roots in each case:

$$y_1 = \frac{2p_1 + \sqrt{4p_1^2 + 8\lambda(m - x_1)p_1}}{2\lambda} \quad (\text{D-27})$$

and

$$y_2 = \frac{2(1-p_2) + \sqrt{4(1-p_2)^2 + 8\lambda(x_2-m)(1-p_2)}}{2\lambda} \quad (D-28)$$

We may now solve for l and h :

$$l = \frac{\lambda x_1 - p_1 - \sqrt{p_1^2 + 2\lambda(m-x_1)p_1}}{\lambda} \quad (D-29)$$

and

$$h = \frac{\lambda x_2 + (1-p_2) + \sqrt{(1-p_2)^2 + 2\lambda(x_2-m)(1-p_2)}}{\lambda} \quad (D-30)$$

Given these estimates of l and h , the resulting total area under the triangle is simply $A = \lambda(h-l)/2$. If this value is within an allowable tolerance of 1.0, then the iterative scheme is completed. However, the first iteration is certain to yield a total area greater than 1.0. To bracket the solution, we successively halve the estimate of λ until the implied estimates of l and h (via equations (D-29) and (D-30)) yield a total area less than 1.0. We then interpolate between the two estimates of λ , as illustrated in Figure D-2.




Figure D-2. Iterative Solution for λ Using @RISK

The desired solution is indicated as λ^* in Figure D-2. The initial estimate, λ_1 , yields a total area greater than 1.0. Successive halving of λ_1 eventually produces a second estimate, λ_2 , with total area less than 1.0. The secant connecting these two estimates is depicted as a dashed line. The equation for the secant is:

$$g(\lambda) = \frac{\lambda_2 f(\lambda_1) - \lambda_1 f(\lambda_2) + [f(\lambda_2) - f(\lambda_1)]\lambda}{\lambda_2 - \lambda_1}. \quad (\text{D-31})$$

This secant achieves height 1.0 at the value λ_3 , which is computed as follows:⁵

$$\lambda_3 = \frac{\lambda_2[1 - f(\lambda_1)] - \lambda_1[1 - f(\lambda_2)]}{[f(\lambda_2) - f(\lambda_1)]}, \quad (\text{D-32})$$

or, in incremental form

$$\lambda_3 = \lambda_2 + \frac{[\lambda_2 - \lambda_1][1 - f(\lambda_2)]}{[f(\lambda_2) - f(\lambda_1)]}. \quad (\text{D-33})$$

The algorithm then replaces either λ_1 or λ_2 with the updated estimate, λ_3 . The total area under the triangle, $A = \lambda(h-l)/2$, is less than 1.0 when evaluated at λ_3 . In order to continue bracketing the solution at λ^* , λ_1 is retained and λ_2 is discarded in favor of λ_3 . Then the iterative scheme continues and we interpolate between λ_1 and λ_3 .

⁵ This method is known as *regula falsi* (false position). See Peter Henrici, *Elements of Numerical Analysis*, New York: John Wiley, 1964, p. 87; or Stephen Kellison, *Fundamentals of Numerical Analysis*, Homewood, Ill.: Richard D. Irwin, 1975, p. 259.

ANNEX D-2: CRYSTAL BALL ALGORITHM FOR CASE III

Figures D-3 through D-7 show a series of screens that result from an attempt to use Crystal Ball when specifying only two percentile points. As can be seen in the final screen, there is no provision for setting the mode other than as the midpoint of the two percentile points.

The figures are explained below.

- Figure D-3: Click the leftmost button on the toolbar, "Define Assumption," to bring up the input screen.
- Figure D-4: Specify the triangular distributional family.
- Figure D-5: View the default parameter screen (i.e., minimum, likeliest, maximum). Click "Parms" to view the optional parameter screen.
- Figure D-6: Specify the 10%, 90% percentile option.
- Figure D-7: View the resulting parameter screen.

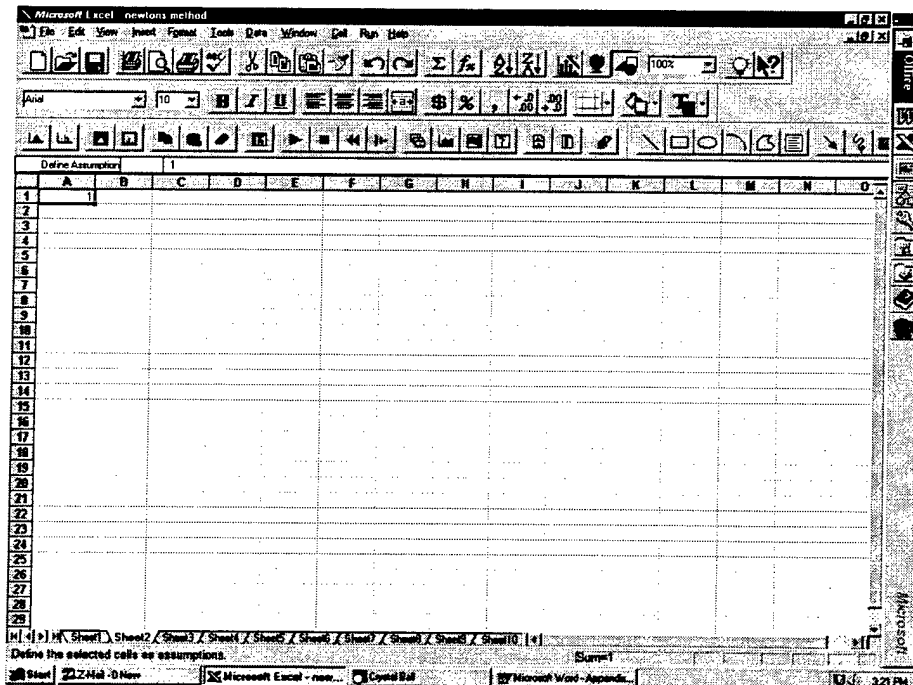


Figure D-3. Crystal Ball: Define Assumptions

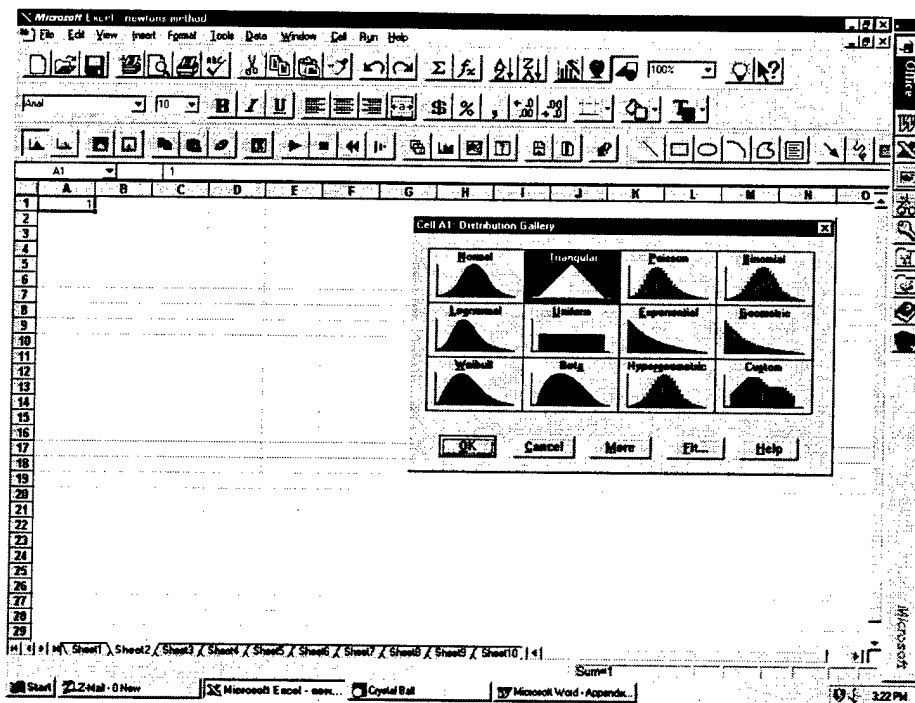


Figure D-4. Crystal Ball: Specify Triangular Distributional Family

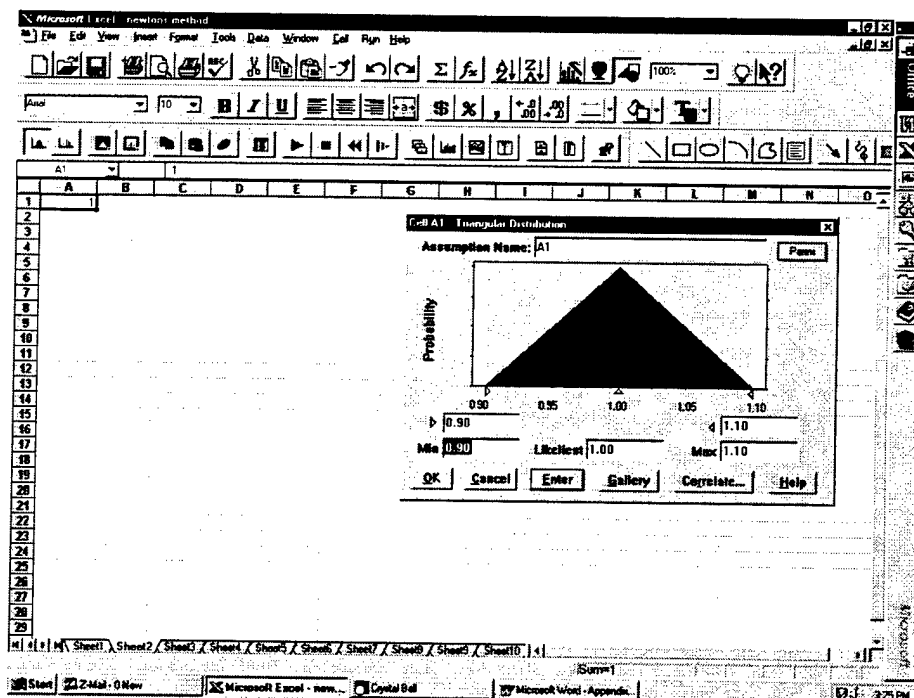


Figure D-5. Crystal Ball: View Default Parameter Screen

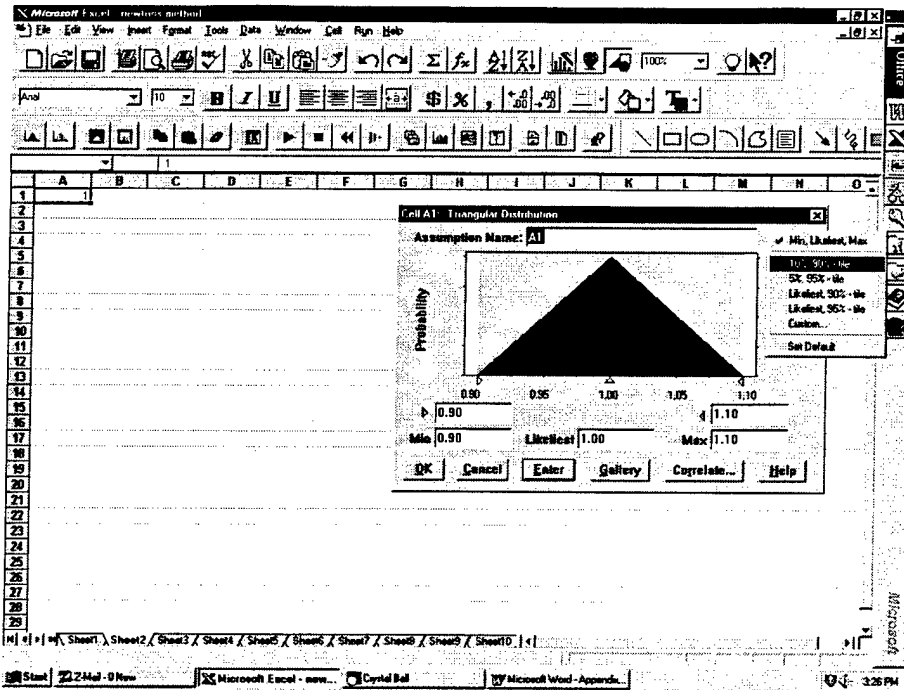


Figure D-6. Crystal Ball: Specify Percentile Option

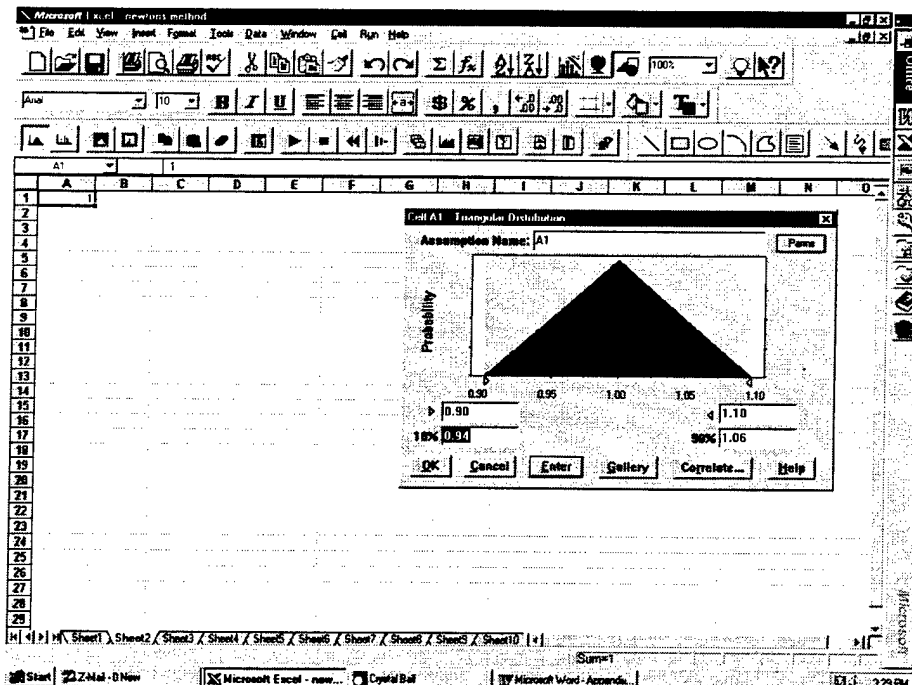


Figure D-7. Crystal Ball: View Percentile Parameter Screen

APPENDIX E

USING CORRELATIONS WITH CRYSTAL BALL

This appendix demonstrates the use of correlations within the Crystal Ball add-in software to Microsoft Excel. Specifically, it demonstrates Crystal Ball for Windows, Version 4.0 for Windows 95/Excel 7.0, released March 1996. Crystal Ball is available from the following vendor:

Decisioneering Inc.
1515 Arapahoe Street, Suite 1311
Denver, Colorado 80202
1-800-289-2550
Fax: 303-534-4818
www.decisioneering.com

As of July 1997, the prices for Version 4.0 were \$495 for new users and \$149 for upgrades from a previous version.

The remainder of this appendix provides a series of instructions and associated screen captures to illustrate the use of correlations. In particular, a logically inconsistent correlation matrix is first input, and the software responds by calculating a “nearby” logically consistent alternative.

Figure E-1: The screen upon opening Crystal Ball.

Figure E-2: Click the leftmost button on the toolbar, "Define Assumption," to bring up the input screen. Alternatively, select "Define Assumption" from the "Cell" pull-down menu. Note that a number must first be typed into the cell; attempting to define an assumption in an empty cell generates an error message.

Figure E-3: Specify a distributional family for the selected cell.

Figure E-4: Specify the parameters (in this case, mean and standard deviation) for the selected cell. Repeat this and the two preceding steps for all additional input cells. Alternatively, first select a *range* of input cells, then specify assumptions for all of them (the input screens will cycle through them sequentially).

Figure E-5: Select a range of input cells, and return to the "Define Assumption" screen. Click the "Correlate" button to generate the screen shown in the figure. Enter a cell (B1) to be correlated with the current cell (A1), and also enter the value of the correlation coefficient. Repeat for all cells correlated with the current cell, and for any additional cells so that a correlation coefficient has been entered for *every distinct pair* of input cells.

Figure E-6: Enter the simulation output variable by algebraic formula, in this case, summation. Click the second button from the left on the toolbar, "Define Forecast" (alternatively, select "Define Forecast" from the "Cell" pull-down menu).

Figure E-7. Click the right-arrow button, "Start Simulation" (alternatively, select it from the "Run" pull-down menu).

Figure E-8: Receive a warning regarding the inconsistent correlation matrix. Click "Adjust Coefficients Permanently" to initiate the adjustment procedure, and *retain a record of adjusted correlations*.

Figure E-9: View Help screen that describes the options for proceeding.

Figure E-10: After the adjustment is performed and the simulation is run, view the simulation results.

Figure E-11: Recover the adjusted correlations. To do so, select the range of correlated input cells, and return to the "Define Assumption" screen. The correlations indicated are the outcome of the adjustment procedure.




Figure E-1. Crystal Ball: Opening Screen

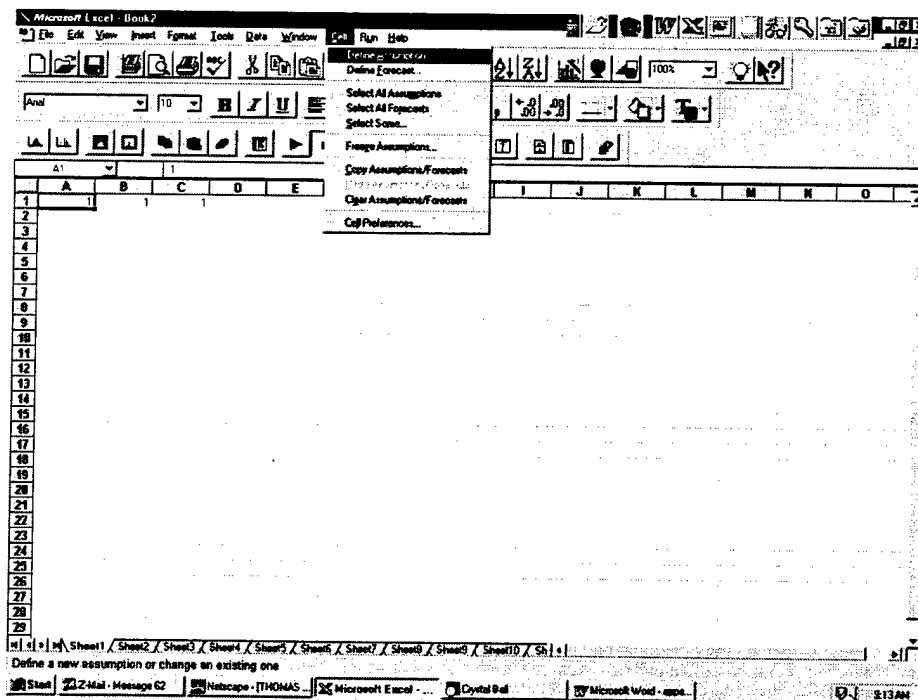


Figure E-2. Crystal Ball: Define Assumptions

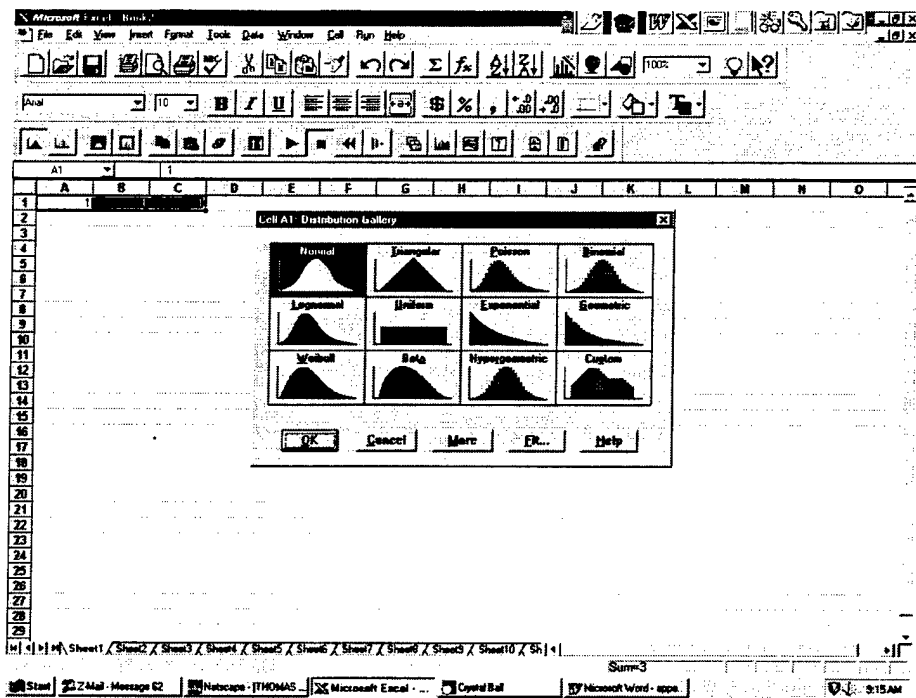


Figure E-3. Crystal Ball: Specify Distributional Family

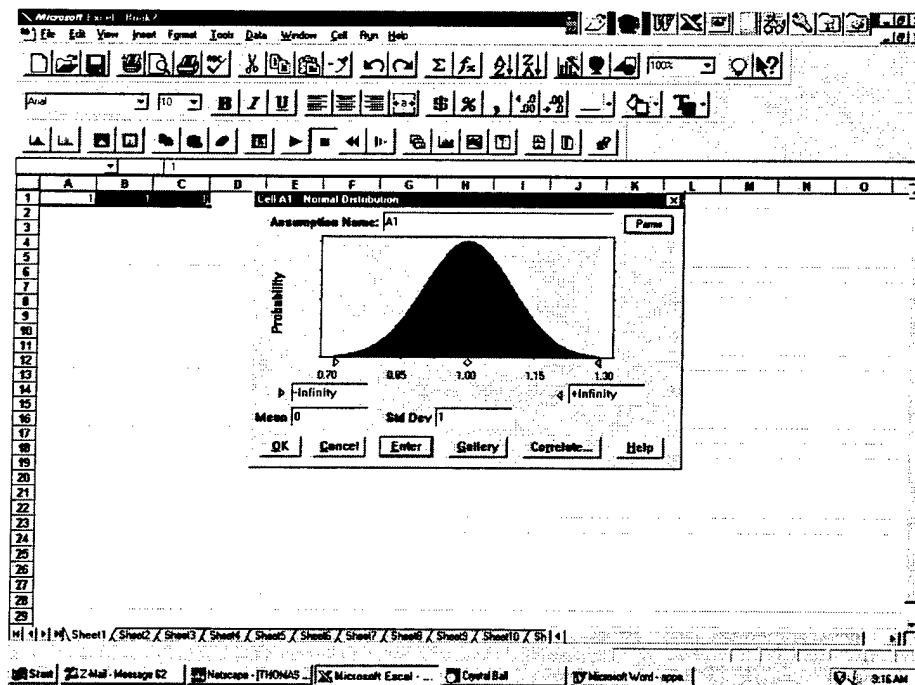


Figure E-4. Crystal Ball: Specify Distributional Parameters

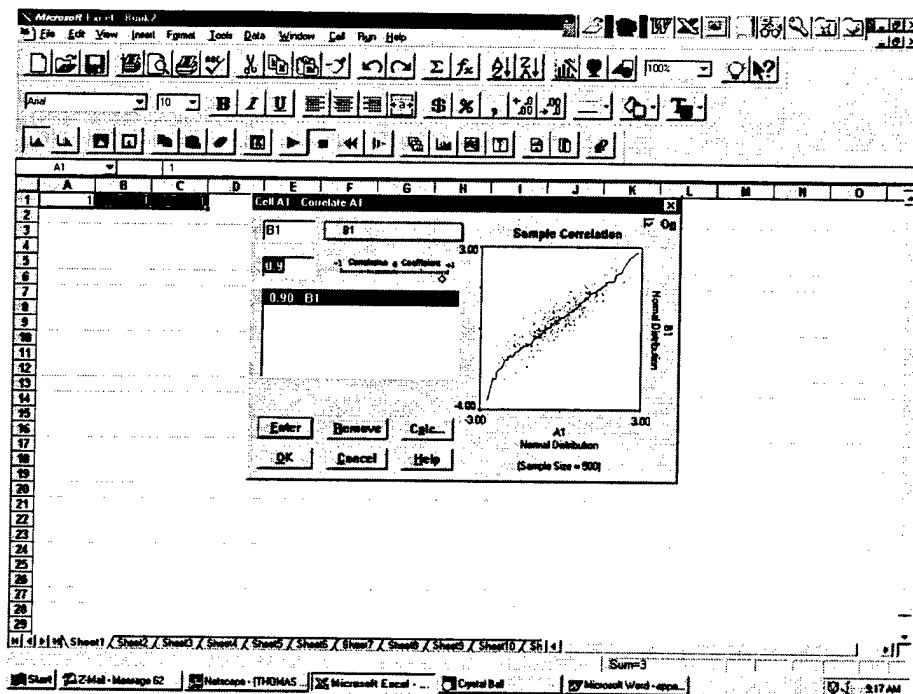


Figure E-5. Crystal Ball: Input Correlations

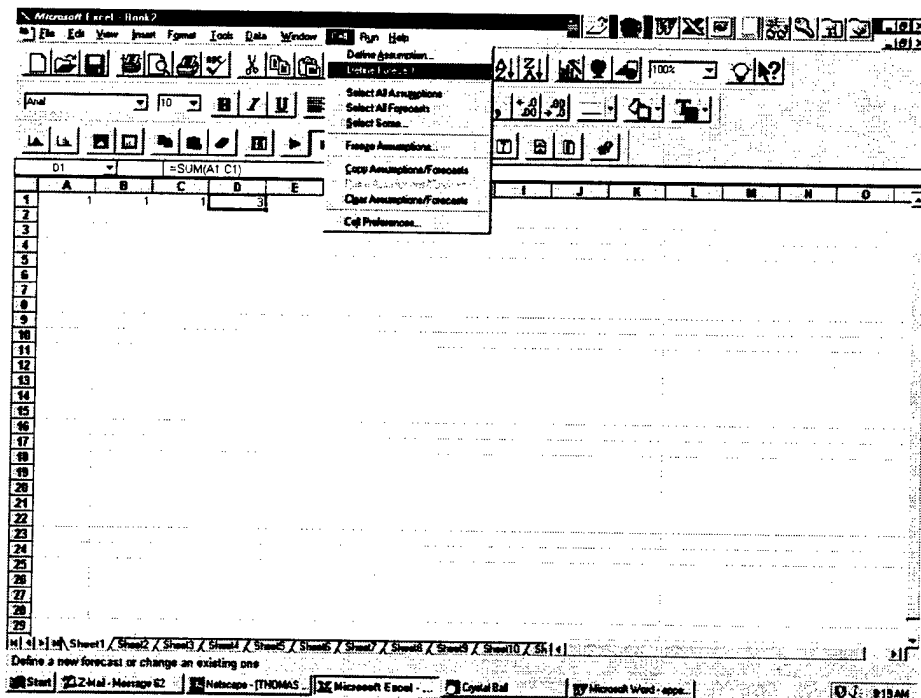


Figure E-6. Crystal Ball: Define Forecasts

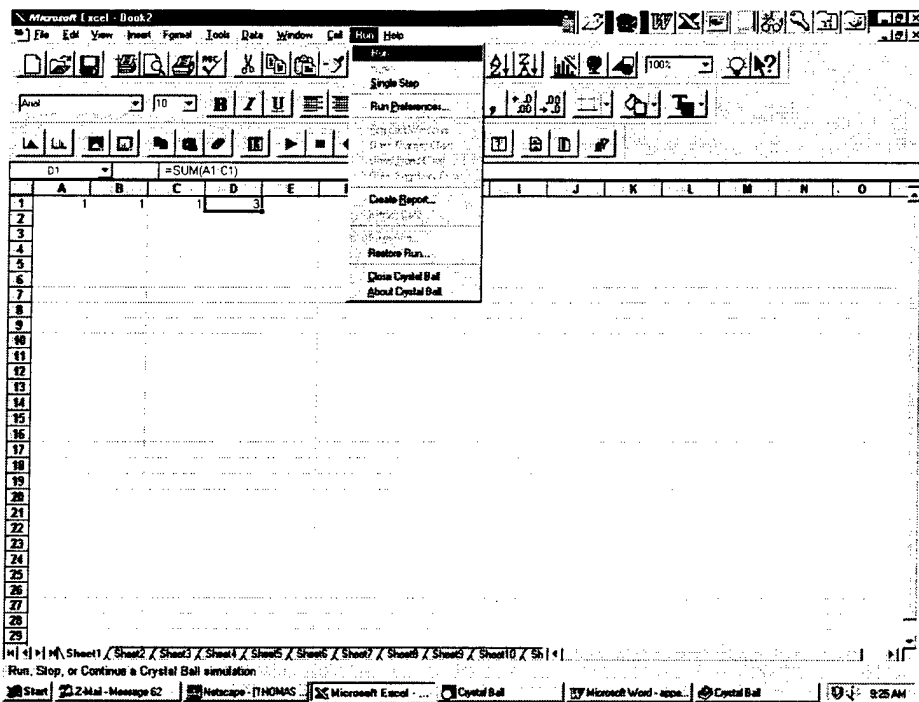


Figure E-7. Crystal Ball: Run Simulation

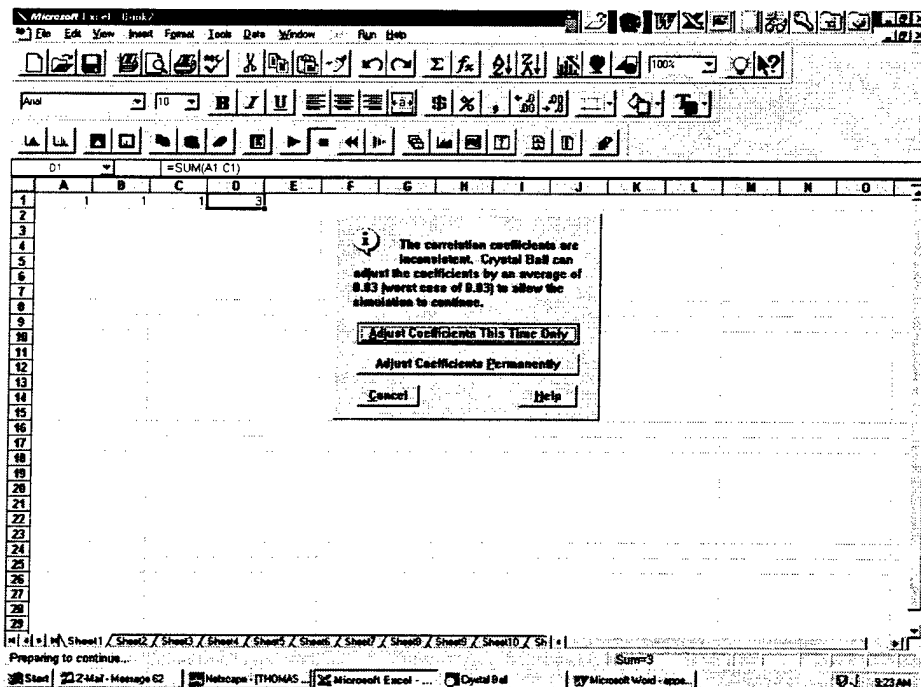


Figure E-8. Crystal Ball: Receive Warning Message, Inconsistent Correlation Matrix

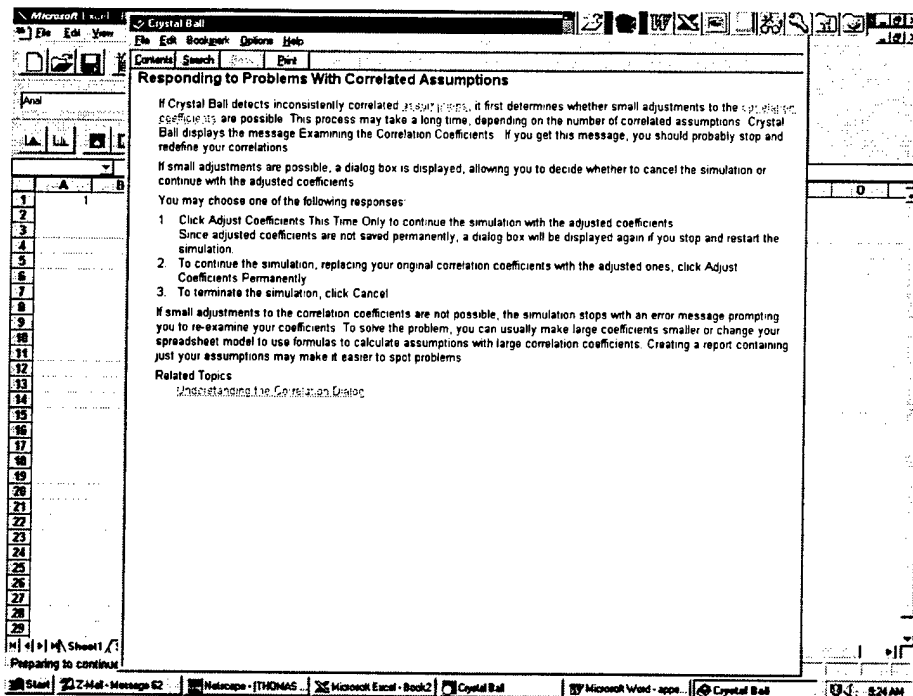


Figure E-9. Crystal Ball: Help Screen on Inconsistent Correlation Matrix

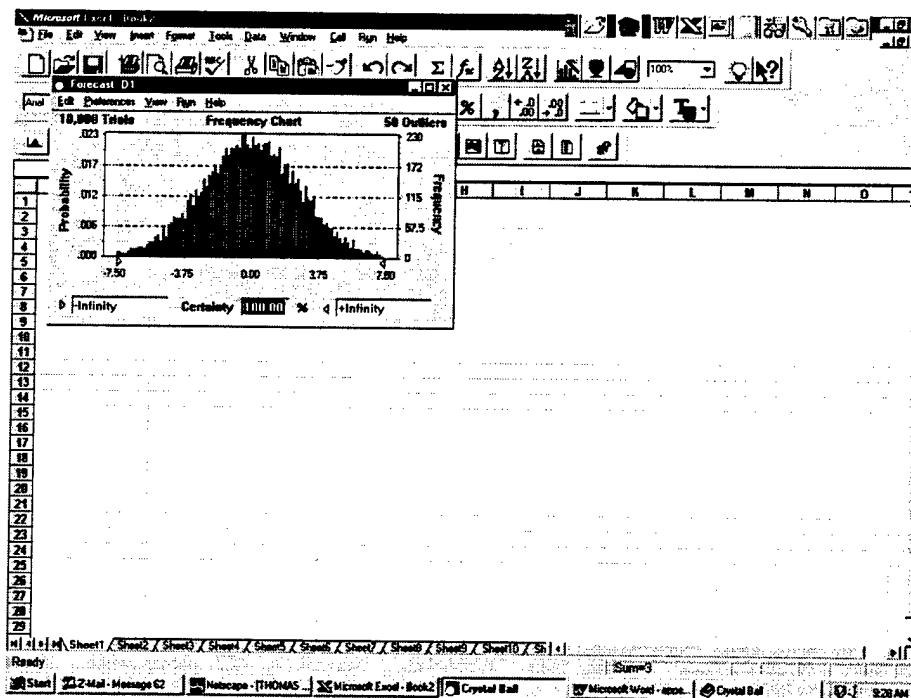


Figure E-10. Crystal Ball: View Simulation Results

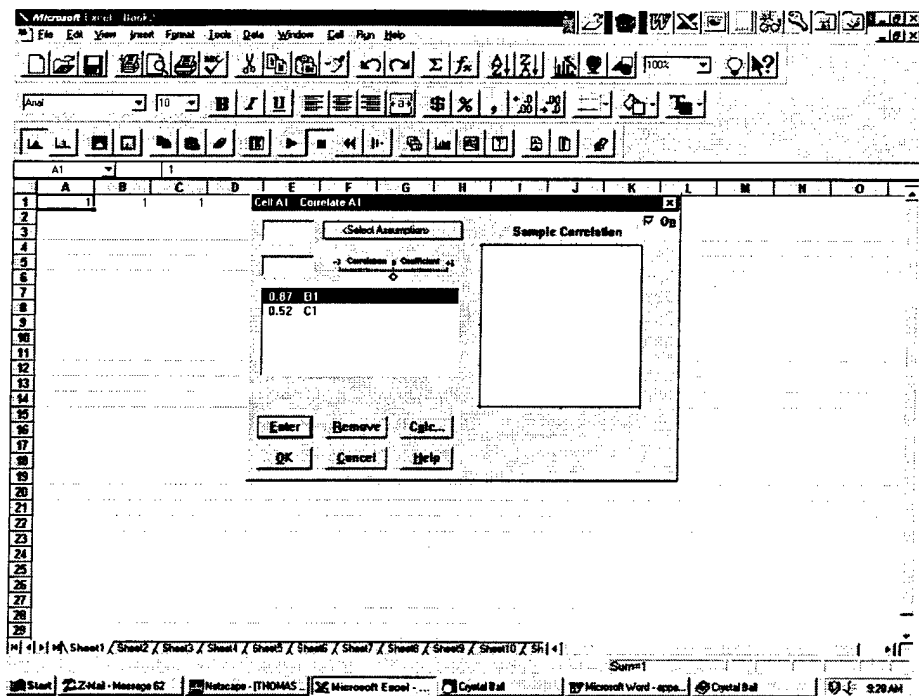


Figure E-11. Crystal Ball: Recover the Correlation Matrix

APPENDIX F

USING CORRELATIONS WITH @RISK

This appendix demonstrates the use of correlations within the @RISK add-in software to Microsoft Excel. Specifically, it demonstrates @RISK for Microsoft Excel for Windows, Version 3.5 for Windows 95/Excel 7.0, released July 1997. @RISK is available from the following vendor:

Palisade Corporation
31 Decker Road
Newfield, New York 14867
1-800-432-7475
Fax: 607-277-8001
www.palisade.com

As of July 1997, the price for Version 3.5e was \$395.

The remainder of this appendix provides a series of instructions and associated screen captures to illustrate the use of correlations. In particular, a logically inconsistent correlation matrix is first input, and the software responds by calculating a “nearby” logically consistent alternative.

Figure F-1: The screen upon opening @RISK.

Figure F-2: Enter the simulation input variables, in this case, three standard normal variables.

Figure F-3: Enter the simulation output variable by algebraic formula, in this case, summation. Click the fourth button from the left on the toolbar, "Add the selected cells as @RISK outputs."

Figure F-4: Click the rightmost button on the toolbar, "Show main @RISK window." Then click the "Correlate" button to generate the screen shown in the figure. Enter the correlations into the matrix, then click "OK." Click "Hide" to return to the spreadsheet.

Figure F-5: Click the second button from the right on the toolbar, "Run Simulation."

Figure F-6: Receive a warning regarding the inconsistent correlation matrix. Click "OK" to initiate the correction procedure.

Figure F-7: After the correction is performed and the simulation is run, view the simulation results.

Figure F-8: Attempt to recover the corrected correlations by toggling to the @RISK window, clicking the "List" button, and clicking the "Correlate" button. The corrected correlations are not shown, only the initial (inconsistent) correlations.

Figure F-9: The corrected correlations may be found in the file Windows\Temp\Corrmat.mtx. This file can be opened by Excel.

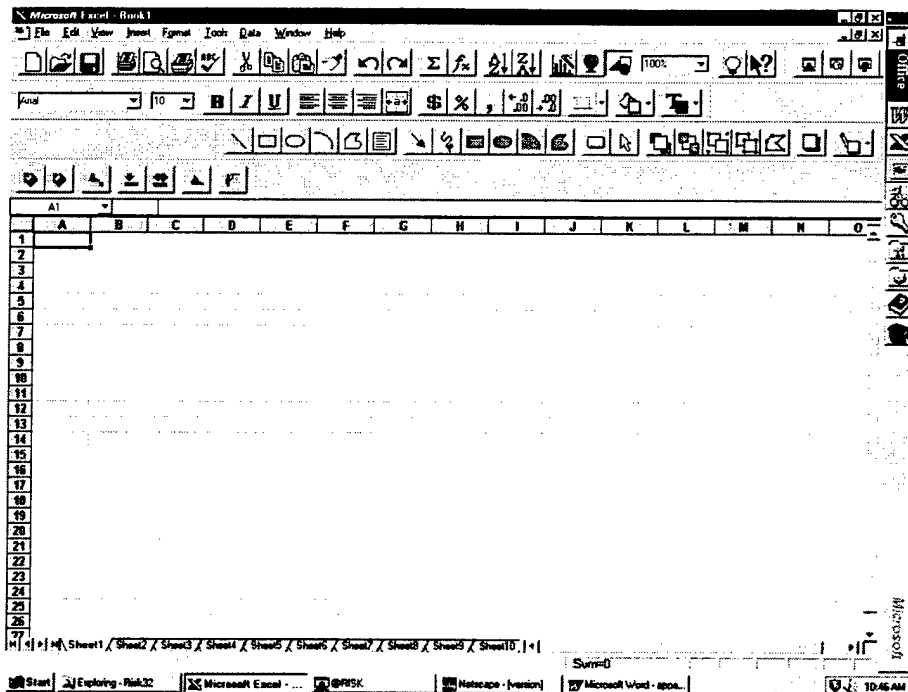


Figure F-1. @RISK: Opening Screen




Figure F-2. @RISK: Enter Input Variables

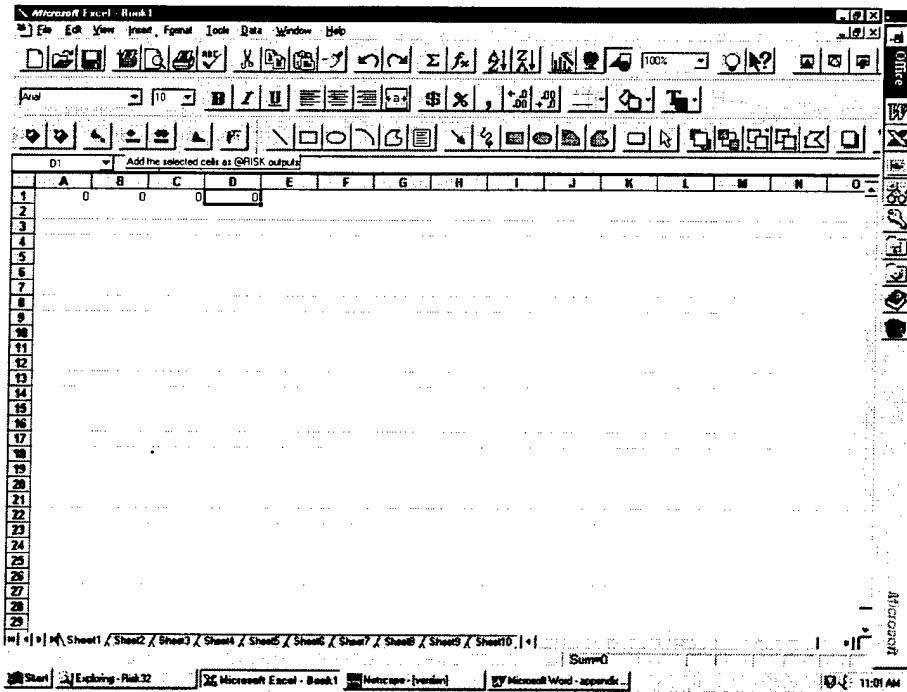


Figure F-3. @RISK: Enter Output Variables

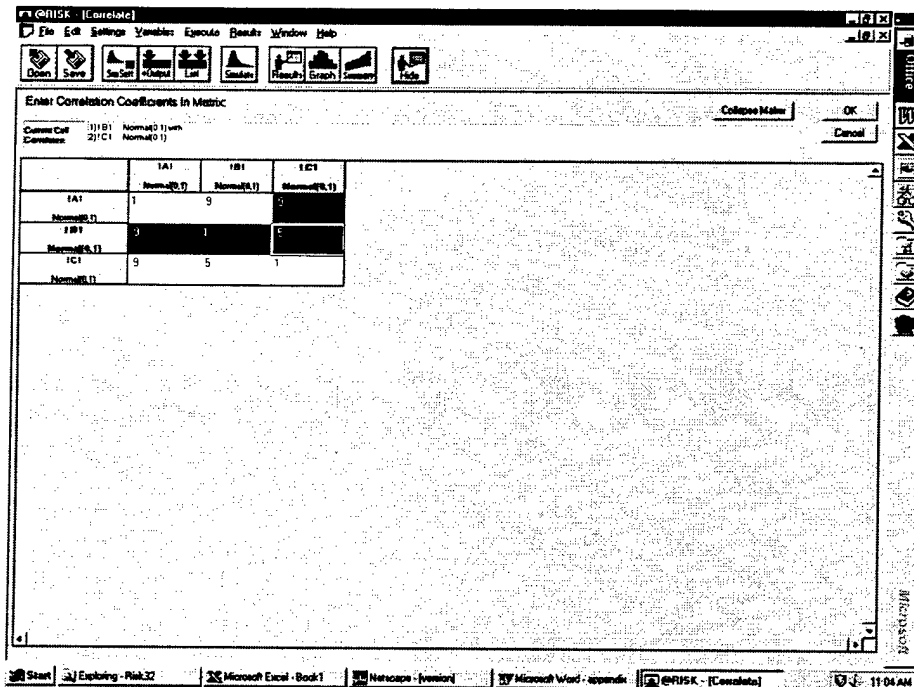


Figure F-4. @RISK: Input Correlations




Figure F-5. @RISK: Run Simulation

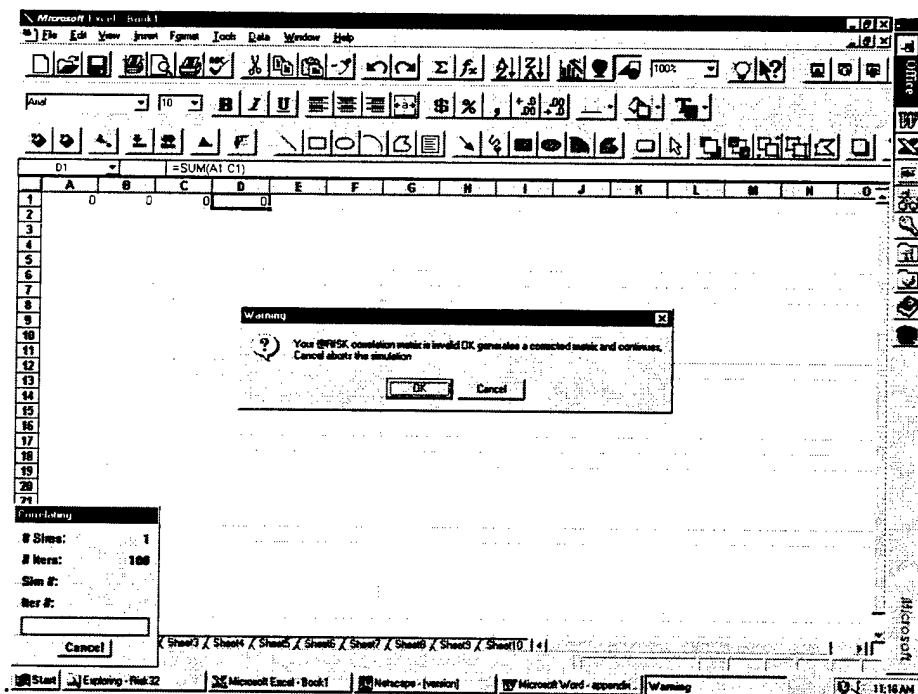


Figure F-6. @RISK: Receive Warning Message, Inconsistent Correlation Matrix

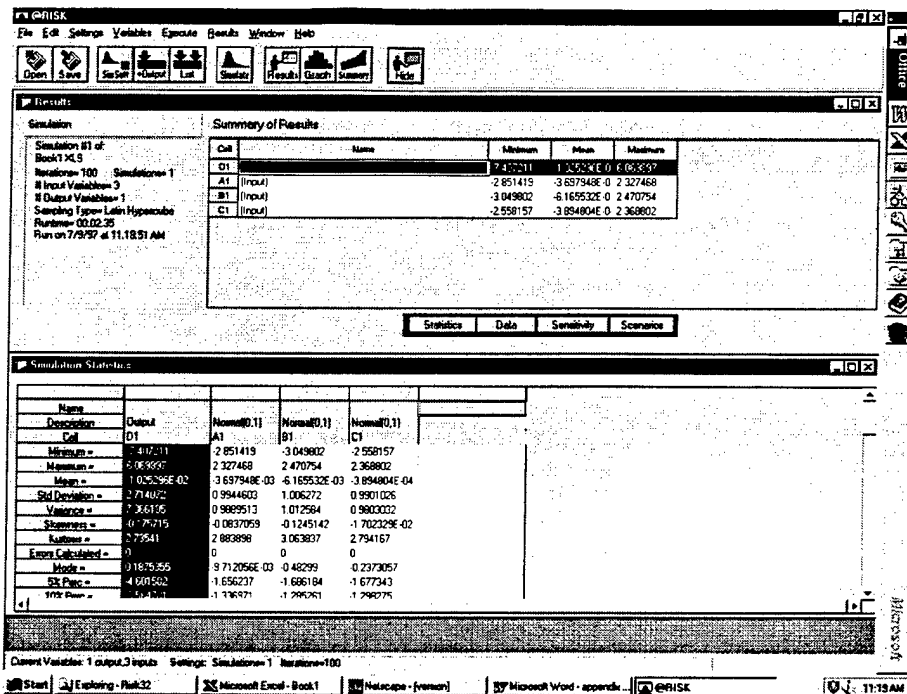


Figure F-7. @RISK: View Simulation Results

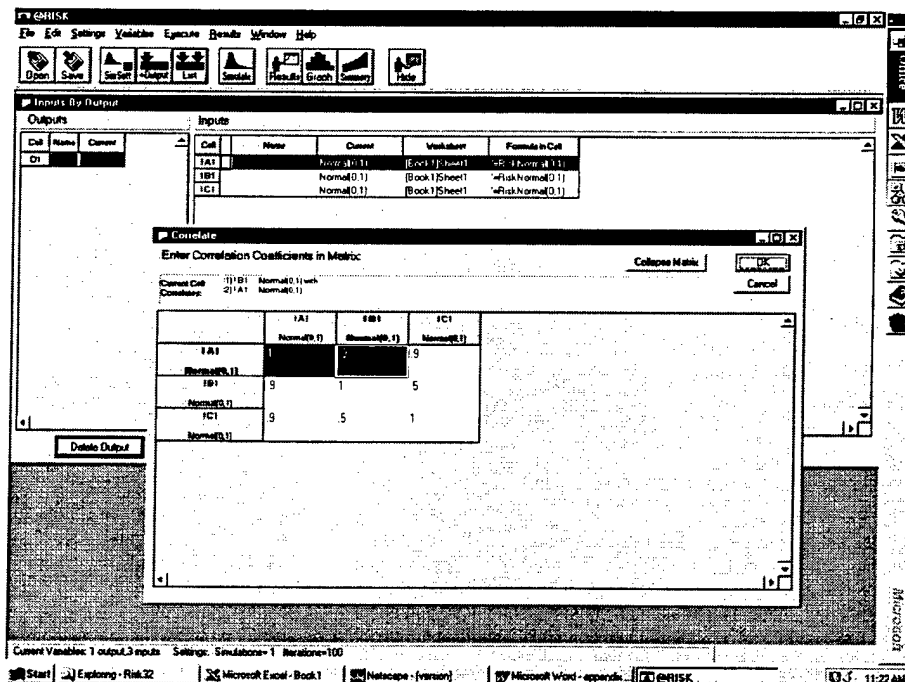


Figure F-8. @RISK: Attempt to Recover the Correlation Matrix

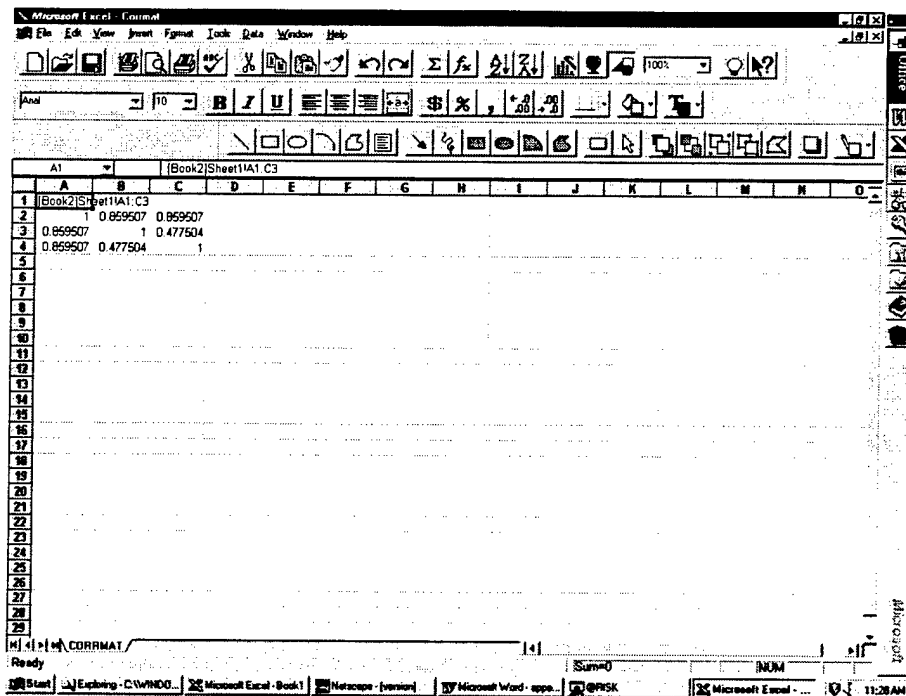


Figure F-9. @RISK: Recover the Correlation Matrix

APPENDIX G

BOEING D&SG REVIEW OF RACM AND ITS APPLICATION TO COST MANAGEMENT, PROJECT PLANNING, AND PROPOSALS

ABSTRACT

Background

Wayne Abba, OUSD (A&T), and Dr. Matthew S. Goldberg, IDA, presented to Boeing D&SG an overview of RACM, a Lockheed Martin risk evaluation tool. We were asked to comment on its utility in applications at the D&SG. Shortly before the IDA presentation on RACM an EELV proposal had been submitted. Its risk submittal was organized by WBS and at its level 3 summaries had data corresponding to some of the RACM input requirements. Boeing D&SG Finance offered to make a comparison of RACM to the EELV submittal and comment further on the system.

Approach

The Risk Analysis Cost Management (RACM) computerized risk model was developed to help Program Management quantify and evaluate cost and schedule risks associated with a proposed plan or existing contract. We ran a comparison analysis of the RACM model to a recently generated risk evaluation on a high-visibility proposal, the pre-Boeing merger Evolved Expendable Launch Vehicle (EELV). This analysis was completed to establish a "benchmark" position and allow a side-by-side comparison of the two approaches. The EELV proposal team developed a risk position based upon Integrated Product Team (IPT) judgments of risk at WBS levels 4 and 5. The RACM model was loaded with the level 3 EELV risk outputs to produce the RACM comparison.

Results

The RACM model will suggest a higher cost estimate (~5+%) would be necessary to achieve a level of risk equal to that proposed on the pre-Boeing merger EELV submittal. In addition, a broader range of possible outcomes was developed by the RACM model than was developed by the pre-Boeing merger EELV proposal team.

Table G-1. Summary of RACM and EELV Estimates

Confidence Level	RACM Base	EELV Proposal
50/50	\$1,182	\$1,124

Concerns

1) The cornerstone of RACM logic is the assumption that a budget allocated is a budget consumed. There is zero probability of an underrun, expenditures will continue to accrue until budgets are exhausted.

2) RACM currently accommodates 30 Work Breakdown Structure (WBS) elements which would be inadequate for proposal evaluations, for program budget tracking and other Cost Management/Control activities. (It should be noted that this shortfall can be remedied with some "tailoring" and reprogramming.)

3) Evidence that RACM predicts the variation experienced between cost estimates and realized cost is lacking. Such evidence is needed before resource decisions are based upon RACM estimates.

Recommendations

1) RACM seems best used in the areas relating to existing contracts; program management, forecasting [Management Estimates at Completion¹ (MEACs)] and cost management tools rather than in support on new business proposals

2) To further research in the field, we should foster a continuing relationship with IDA relating to the RACM model and others they may wish to evaluate. A piloting effort should be initiated, deploying RACM on an existing contract so that more in-depth evaluation can be made. The evaluation should include the development of statistics describing RACM's performance as predictor of the variation between estimates and realized cost.

¹ The MEAC is technically a lump-sum contract completion estimate for a cost category, e.g., the total contract value, a sub-element's value. It is made up of the accumulated actual cost to a current point in time plus an estimate of the balance to go. The estimate of the balance to go is typically made by the budget responsible organization. The "M" implies that it has been approved by the budget responsible management. It usually has an underlying time series form and as such it can be converted into a new budget easily. It is a mechanism by which cost performance is incorporated into budgets. A related estimate is the Estimate at Completion (EAC). It is the same as the MEAC except it is not approved by the budget responsible management and it is usually developed independently of the budgeted organizations, for example by Finance.

OVERVIEW OF RACM

In the implementation provided to us:

- RACM is on an EXCEL spread sheet,
- it has a limit of about 30 WBS elements,
- it has two cost element categories, labor and non labor.

The inputs for each WBS labor cost element are the low-probability equivalent heads estimate and its probability, the high-probability equivalent heads and its probability, along with the duration of those heads for the project. For each non-labor cost element the inputs are the low-probability dollar estimate and its probability, and the high-probability dollar estimate and its probability, along with the duration of those dollars expenditures for the project. For example:

Table G-2. Example of RACM Input Data for a WBS Element

Equivalent People or Non-Labor Dollars		Duration	Average Wrap Rate	Probability	
<u>Low</u>	<u>High</u>			<u>Low</u>	<u>High</u>
10	15	36	100	.16	.84

These inputs are used to develop Normal Probability Density Functions. However the model's internal structure effectively utilizes only the upper cost or schedule tail in many instances.

Addition inputs include, for each labor and non-labor element:

1. Schedule Slip at a 3- σ probability level (99.87%),
2. Improvement percent and its probability (potential cost improvement from estimate basis),
3. Wrap rate growth and its probability;

As global parameters:

4. Milestone cost estimates and their probabilities,
5. Global project attributes, Major Problems and Schedule Parallelism,
6. Bid probability, Budget Allocation level, and Management Reserve level.

These inputs are not further defined by RACM documentation.

It incorporates a view of cost behavior based on the idea that budgeted costs are at least expended, and they may be overrun due to two causes:

1. The costs exceeding the budgeted amount (controlled by the density function input),
2. The simulated schedule flow exceeds the planned duration. The duration expansion has two causes:
 - a). the schedule duration probability density function extends the time expenditures and thereby the cost,
 - b). the schedule network parallelism simulates the continuation on all subtasks until the entire task is completed.

INTERNAL POTENTIAL USES OF RISK ESTIMATES

Guiding the Cost Management and Project Planning Processes

RACM estimates lower cost when the budgets are set lower. The concept “Money Allocated Is Money Spent” is incorporated into RACM. For example, if budgets are set at the lowest possible cost, cost and schedule variation would be funded from management reserved funds not as a planned expenditure.

RACM estimates lower cost when the schedule networks are sequential rather than parallel.

Meeting Requirements for Cost Risk Estimates in Proposals

RACM estimates are sensitive to cost and schedule uncertainties and could be used to weigh the merits of program plans and designs with varying uncertainties. For example, if the risks of two alternative designs differ significantly, these risk differences could influence the design choice.

COMPARISON TO EELV RISK APPROACH

To compare the pre-Boeing merger EELV and RACM approaches, the level 3 WBS estimates from the EELV Engineering and Manufacturing Development (EMD) Life-Cycle Cost (LCC) estimate were input to RACM. These 29 EELV WBS elements were Monte Carlo summaries of the 101 EELV LCC WBS elements.

In both approaches inputs describing the costs density functions are utilized.

In the EELV case, those density function inputs are broken out for each WBS into six causative cost-risk contributors: Cost Estimating, Schedule, Supportability, Technology, Design and Engineering, and Manufacturing. Like RACM, the EELV cost risk input categories were not further defined. The risk categories are each input as a triangle distribution using the distribution mode and the best and worst "realistic" endpoint values. Mathematical adjustments were made where inputs were not available for modes but were available for means or medians. Judgments were solicited from experts (engineers), in each of these categories (IPT/WBS), for the category's relative impact on cost and the amount of risk, ranging from very low to very high, for each category by WBS. A table, supplied by the Air Force Program Office, translating the very low to very high risk assessments into the best and worst "realistic" endpoint values was used to quantify the judgments; the table is attached to this appendix.² (Neither the use of this table nor any aspect of this procedure was required by the Air Force.) A summary WBS distribution was generated for these six categories for each WBS element. The total (life cycle cost) distribution is calculated as the sum of the WBS components. No non-linear interactions between WBS elements were incorporated in the summation, that is, the mean of the total equaled the sum of the WBS means. In RACM this is not the case, the mean of the total is greater than the sum of the means. The EELV approach allowed for the incorporation of correlations between WBS elements; however, in EELV implementation the cross correlations were set to zero. The meanings of the cost risk categories, Cost Estimating, Schedule,..., and the best and worst case values were left to the engineers making the judgments. Also, in contrast to RACM symmetric distributions, the EELV triangle distributions were typically skewed toward higher costs.

THE NUMERIC RESULTS FOR THE EELV EMD PHASE

The RACM values were generated using the WBS level 3 (29 elements) outputs of the EELV risk calculation as an input to RACM. Improvement and Rate RACM risk inputs were set to zero for comparability. The EELV risk estimating procedure does not have inputs corresponding to the Global Schedule and Major Problem RACM inputs. Due to the intended depth of comparison, an attempt to develop inputs for these RACM risk

² Note added by IDA: The Air Force table appears to be a modification of a table found in Shishu Gupta, David Olsen, David Hudak and Jennifer Keenan, "Cost Risk Analysis of the Strategic Defense System." Related work is found in Fred Biery, David Hudak and Shishu Gupta, "Improving Cost Risk Analyses," *Journal of Cost Analysis*, Spring 1994, pp. 57-86, especially Table 3.

contributors through the EELV organization was not made. However, to demonstrate RACM's sensitivity to these parameters, values were input.

Attached is a tabular summary of the RACM output (relative dollars in thousands). It contains incremental cumulative distributions due to each sequentially additive RACM risk component. An examination of the RACM data for the Basic estimate shows that the risk (70% – 50%) estimate was \$14.5, an amount not differing much from the EELV difference of \$10.9. The additional RACM variation is largely due to the Major Problem input (judgmental). Also of note is that the adjustment at the 50% level, \$172.1, comprises the bulk of the change in estimate.

Table G-3. Decomposition of RACM and EELV Estimates

	50%	70%	Difference	Percent
RACM Base	\$1181.5	\$1196.0	\$14.5	1.23
RACM Schedule	\$1228.4	\$1243.5	\$15.1	1.23
RACM Global Schedule	\$1281.9	\$1300.9	\$19.0	1.48
RACM Major Problem	\$1296.4	\$1327.1	\$30.7	2.37
EELV	\$1124.3	\$1135.2	\$10.9	0.97
RACM Major Problem less EELV	\$172.1	\$191.9	\$19.9	

The RACM approach differs from that used in EELV:

- Its distributions are normal. Non-linear interactions occur between elements during the summation process.
- The budget non-underrun effect.
- The extension of all individual organization's cost for parallel activities until the entire activity is completed.

CONCLUSIONS

The RACM limitation of 30 WBS elements would be constraining in many applications, both in the proposal and cost management domains. However the RACM EXCEL spreadsheet is small, and it could be expanded with a comparatively small programming effort. Also, probably, for most on-going applications some tailoring of RACM to the specific program requirements would be indicated.

RACM forecasts cost effects not estimated by the standard linear summations common to most risk solutions. To the extent that these forecasts describe reality, they provide information that would guide us towards producing products at a lower cost.

The input data required by RACM is probably more easily obtainable than that developed for EELV.

RACM has its greatest value in the ongoing planning and management of projects.

Attached is the numeric output of RACM. It shows the RACM risk estimates cumulatively due to its several risk input categories.

Table G-4. Sequential Distribution Functions of RACM Risk Estimates

	Basic		Schedule		Improvement		Rate		Global Schedule		Major Problem	
	Mean	Sigma	Mean	Sigma	Mean	Sigma	Mean	Sigma	Mean	Sigma	Mean	Sigma
	867.7	25.3	904.7	26.6	904.7	26.6	904.7	26.6	958.3	34.3	972.8	57.4
Labor:	313.9	10.8	323.7	11.1	323.7	11.1	323.7	11.1	323.7	11.1	323.7	11.1
Non-Labor:	1181.5	27.5	1228.4	28.8	1228.4	28.8	1228.4	28.8	1281.9	36.1	1296.4	58.4
Total:	Combined Basic (Cum. Dist.)		Combined Schedule (Cum. Dist.)		Combined Improvement (Cum. Dist.)		Combined Rate (Cum. Dist.)		Combined Schedule (Cum. Dist.)		Combined Major Problem (Cum. Dist.)	
	Prob(S)	Cost-K\$	Prob(S)	Cost-K\$	Prob(S)	Cost-K\$	Prob(S)	Cost-K\$	Prob(S)	Cost-K\$	Prob(S)	Cost-K\$
	0.01	1117.6	0.01	1161.4	0.01	1161.4	0.01	1161.4	0.01	1198.1	0.01	1160.5
	0.05	1136.4	0.05	1181.0	0.05	1181.0	0.05	1181.0	0.05	1222.6	0.05	1200.3
	0.1	1146.3	0.1	1191.5	0.1	1191.5	0.1	1191.5	0.1	1235.7	0.1	1221.6
	0.2	1158.4	0.2	1204.1	0.2	1204.1	0.2	1204.1	0.2	1251.6	0.2	1247.3
	0.3	1167.1	0.3	1213.3	0.3	1213.3	0.3	1213.3	0.3	1263.0	0.3	1265.8
	0.4	1174.6	0.4	1221.1	0.4	1221.1	0.4	1221.1	0.4	1272.8	0.4	1281.6
	0.5	1181.5	0.5	1228.4	0.5	1228.4	0.5	1228.4	0.5	1281.9	0.5	1296.4
	0.6	1188.5	0.6	1235.7	0.6	1235.7	0.6	1235.7	0.6	1291.1	0.6	1311.3
	0.7	1196.0	0.7	1243.5	0.7	1243.5	0.7	1243.5	0.7	1300.9	0.7	1327.1
	0.8	1204.7	0.8	1252.6	0.8	1252.6	0.8	1252.6	0.8	1312.3	0.8	1345.6
	0.9	1216.8	0.9	1265.3	0.9	1265.3	0.9	1265.3	0.9	1328.2	0.9	1371.3
	0.95	1226.7	0.95	1275.7	0.95	1275.7	0.95	1275.7	0.95	1341.3	0.95	1392.6
	0.99	1245.5	0.99	1295.3	0.99	1295.3	0.99	1295.3	0.99	1365.8	0.99	1432.4
	0.995	1252.3	0.995	1302.5	0.995	1302.5	0.995	1302.5	0.995	1374.8	0.995	1446.9
	0.999	1266.4	0.999	1317.3	0.999	1317.3	0.999	1317.3	0.999	1393.4	0.999	1477.0
	0.9999	1283.7	0.9999	1335.4	0.9999	1335.4	0.9999	1335.4	0.9999	1416.1	0.9999	1513.8

The data contained in Table G-4 are plotted in Figure G-1 and Figure G-2.³

³ Only the table, not the figures, appeared in Boeing's original report. The figures were added by IDA.




Figure G-1. Crystal Ball Graph of D&SG Data




Figure G-2. @RISK Graph of D&SG Data

ANNEX G-1: TRANSMITTAL LETTER FROM BOEING D&SG TO IDA

November 20, 1997

Subject: Overview, IDA request for evaluation of Lockheed Martin's RACM

Wayne Abba, OUSD (A&T), and Dr. Matthew S. Goldberg, IDA, presented to us an overview of RACM, a Lockheed Martin risk evaluation tool. We were asked to comment on its utility in applications at the D&SG.

Description of Analysis Approach

Shortly before the IDA presentation on RACM the EELV proposal had been submitted. It contained a risk estimate by WBS similar to the RACM WBS input structure. The EELV approach also had as internal values, some of the data required by RACM for its estimates. Where the EELV approach had inputs corresponding to RACM's the EELV probabilistic estimates were utilized. For those RACM elements not corresponding to the EELV structure, reasonable estimated values were developed. (The RACM documentation did not provide guidance in this area.) A stepwise set of RACM risk estimates were produced and compared with the EELV estimate. Each step introduced an incremental RACM risk.

Description of Evaluation Process

In addition to generating comparative output, Dr. Goldberg and Dr. Charles A. Weber carried on a dialog with us to gain an understanding of the EELV approach, how it related to RACM, how the RACM approach would fit into Boeing D&SG procedures and methods, and our opinion of it as a useful risk estimation procedure. A subset of this dialog, deemed significant by Dr. Goldberg, is included in the following document.

Evaluation Summary

RACM is an Excel-based tool with a specific view of the causes and quantified effects of cost, schedule, and technical "risk." As such it, in all likelihood, fits into the cost management and estimating infrastructures from which it was created. RACM does not mesh well with our cost estimating and cost management data sources; estimating procedures, and management requirements nor the management philosophy that I am aware of at the Boeing Company. This point, of course, could be made about all but the most generic of approaches not invented here, and should not be taken to imply that we should not adjust to the RACM view.

Besides the issues of procedural fit, RACM like any proposed estimating (forecasting) system is intended to guide us in the management of our resources. It would

be imprudent to utilize such guidance without some evidence that better results could be expected through its use than otherwise. That is, it must be a valid method. Its risk estimates should be demonstrated to be related to the actual uncertainty. We do not have such a demonstration.

Recommendation

Due to its dependency on budgeting data, RACM seems best suited to existing contracts; program management, forecasting, MEACs, and cost management tools rather than in support on new business proposals. A piloting effort should be initiated, deploying RACM on an existing contract to determine its validity and forecasting accuracy. If it is shown to be a good predictor, its forecasts should be attempted to be used in the program management process. The effects should be measured.

Dr. Goldberg has an assignment to issue a report to DoD on RACM. (Lockheed/Martin under contract to IDA is contributing some 50 pages to his report.) He has asked for our permission to include in his report portions of the dialog carried on between us. The following document, with changes as noted, is what he has deemed the most useful and relevant of our communications. I recommend that permission be granted. He also must attribute these communication to Boeing. I suggest that they be attributed to the Manager of Statistical Analysis and Simulations, Boeing Information, Space, & Defense Systems (ISDS), Estimating.

Manager of Statistical Analysis and Simulations, Boeing ISDS, Estimating.
[Original signed]

ANNEX G-2: EELV BOEING RISK EVALUATION METHODOLOGY (PRE-BOEING MERGER)

- Disciplined analytical approach using Monte Carlo simulation to bound the point estimates and establish confidence levels of the estimate.
- Incorporates cost risk impact of six (6) elements:
 - ✦ Estimating Methodology,
 - ✦ Design & Engineering,
 - ✦ Program Schedule,
 - ✦ Manufacturing,
 - ✦ Technology,
 - ✦ Supportability (Logistics).

Each element assigned weighting (judgmental) for impact to total cost of the particular sub-system/program element under consideration (total of weightings = 100%).

- Risk assessment at Element or Sub-system Level.
- Seven (7) risk ratings from Very Low to Very High with error (+/-) bands established for each ranking:

Table G-5. Risk Ratings from EELV Program Office

Risk Rating	Error Band	
	Low	High
Very Low	5%	10%
Low	7%	15%
Moderately Low	10%	25%
Moderate	15%	50%
Moderately High	25%	75%
High	35%	125%
Very High	50%	200%

- Document basis of risk rankings based on defined criteria (following pages).

NON-RECURRING

Input Form **PBS:** _____ **Evaluator:** _____ **Date:** _____

Risk Ranking	Cost	Schedule	Technology	Design/Engineering	Manufacturing	Supportability
Relative Weight ⁽¹⁾						

⁽¹⁾ Expressed as a percentage, must total to 100%. Weighting is intended to reflect relative impact of risk category on projected cost (e.g., a cost estimate for program management/system engineering may be more impacted by schedule changes/slips than by technology or design & engineering risks, etc.).

Risk Rating Rationale

Risk Category	Basis of Ranking
Cost	
Schedule	
Technology	
Design & Engineering	
Manufacturing	
Supportability	

RECURRING

Input Form **PBS:** _____ **Evaluator:** _____ **Date:** _____

Risk Ranking	Cost	Schedule	Technology	Design/Engineering	Manufacturing	Supportability
Relative Weight ⁽¹⁾						

⁽¹⁾ Expressed as a percentage, must total to 100%. Weighting is intended to reflect relative impact of risk category on projected cost (e.g., a cost estimate for program management/system engineering may be more impacted by schedule changes/slips than by technology or design & engineering risks, etc.).

Risk Rating Rationale

Risk Category	Basis of Ranking
Cost	
Schedule	
Technology	
Design & Engineering	
Manufacturing	
Supportability	

COST: Uncertainties in cost due to reliance on available estimating methodologies.

- | | |
|------------------|---|
| Very low: | Estimate based on current actual production program costs, vendor firm quotes, catalog pricing for identical hardware/equipment/material. |
| Low: | Estimate based on analogy to actual current hardware/program costs, vendor quotes. |
| Moderately low: | Estimate based on analogy to actual historic hardware/program(s), material cost history. |
| Moderate: | Estimate based on calibrated pricing model. |
| Moderately high: | Estimate based on detailed functional 'bottom up' estimate. |
| High: | Estimate based on detailed engineering 'bottom up' estimate. |
| Very high: | Estimate based on informal engineering/management estimate. |

Note: The level of program definition must also be considered in determining the risk rating. The less detailed the program/product definition, the greater will be the risk associated with ANY of the described estimating methods.

SCHEDULE: Assumes the item being acquired is a developmental item and its schedule meets program goals. The assessment focuses on the adequacy of the time specified for the item relative to schedules for similar items.

- | | |
|------------------|---|
| Very low: | Schedule is much longer than is typical for similar items. (At the system level, there are at least 69 months between program start and launch.) |
| Low: | Schedule is longer than is typical for similar items. (At the system level, there are at least 58 months between program start and launch.) |
| Moderately low: | Schedule provides development time. (At the system level, there are at least 51 months between program start and launch.) |
| Moderate: | Schedule is achievable. (At the system level, there are at least 46 months between program start and launch.) |
| Moderately high: | Schedule is challenging. (At the system level, there are at least 40 months between program start and launch.) |
| High: | Schedule is very challenging. (At the system level, there are at least 34 months between program start and launch.) |
| Very high: | Schedule represents an unprecedented situation in terms of time. (At the system level, there are less than 32 months between program start and launch.) |

Note: Data from unmanned space vehicle programs (including NASA, military and commercial) range from 26 to 88 months for the period between Authority to Proceed (program start) and first launch.

TECHNOLOGY: Uncertainties to system performance due to reliance on the availability and promise of technology. Technology uncertainty includes the required level of technological sophistication and reflects the current stage of hardware development and testing maturity. Hardware maturity ranges from scientific research, conceptual design, brassboard, breadboard, prototype, to an operational unit.

Very low:	Hardware is currently operational and deployed.
Low:	Hardware is in limited production and has passed all acceptance tests.
Moderately low:	Prototype is currently in qualification test, but has passed performance requirements.
Moderate:	A brassboard example has been fabricated and passed performance and applicable development tests.
Moderately high:	Critical functions/characteristics have been demonstrated by a brassboard example.
High:	Conceptual design has been formulated and tested for performance and qualification considerations.
Very high:	Scientific research is required and ongoing.

Note: The two categories of 'Technology' and 'Design & Engineering' include some overlap, since both involve the level of maturity of an item. The Technology risk category primarily focuses on the hardware INDEPENDENT of how it will be used. The Design & Engineering category primarily focuses on hardware implementation partially independent of the inherent level of technological readiness. For example, a qualified component may still require modification necessitated by form, fit and function changes and specialized modifications (i.e., radiation shielding, vibration damping, etc.) unique to the system.

DESIGN & ENGINEERING: Uncertainties to system performance due to uncertainties and variability in design and engineering process. Design & Engineering uncertainty reflects the degree of difficulty to advance the current state of the art for a given item (e.g. subsystem) to the required (e.g. qualified off-the-shelf item that meets all requirements).

Very low:	Qualified off-the-shelf item that meets all requirements.
Low:	Off-the-shelf items that meet requirements, but need qualification.
Moderately low:	Design effort required using standard, existing components within their original specification levels.
Moderate:	Design effort required using standard, existing components beyond their original accepted specification levels.
Moderately high:	Moderate engineering development is required using existing design knowledge.
High:	Major engineering development is required using existing design knowledge.
Very high:	New or breakthrough advance in design capability is required.

Note: The two categories of 'Technology' and 'Design & Engineering' include some overlap, since both involve the level of maturity of an item. The Technology risk category primarily focuses on the hardware INDEPENDENT of how it will be used. The Design & Engineering category primarily focuses on hardware implementation partially independent of the inherent level of technological readiness. For example, a qualified component may still require modification necessitated by form, fit and function changes and specialized modifications (i.e., radiation shielding, vibration damping, etc.) unique to the system.

MANUFACTURING: Uncertainties associated with the production elements used to manufacture the required quantities of an item, within the technical specifications.

- Very low: An existing process meets key attributes (C,Y,T, TP and PC). Test equipment is operational. Sufficiently trained personnel exist on production lines. Multiple, adequate domestic sources for material exist. Production equipment currently manufacturing the product.
- Low: Modification of an existing integrated process to meet key attributes (C, Y, T, TP and PC). Test equipment is in limited use producing the item. Multiple domestic sources exist, but for an inadequate amount of material; the balance of required material is available from multiple foreign sources. Sufficient personnel exist but not currently being used in producing the item. Suitable production equipment exists, but insufficient quantity on hand.
- Moderately low: Integrated process is a combination of demonstrated processes and one key attribute (C, Y, T, TP or PC) exceeds the norm for these processes. Insufficient amount of test equipment/facilities are on hand but available from outside sources. Sufficient amount of personnel exist, but require minor training. Single domestic source exists, but for an inadequate amount of material, the balance of material is available from a single foreign source. The set of production equipment requires minor modification.
- Moderate: Integrated process is a combination of demonstrated processes and two or more attributes (C, Y, T, TP or PC) exceeds the norm for these processes. Current test equipment requires minor modification. Multiple, adequate foreign sources for material exist. Sufficient amount of personnel are available but at a sub-optimum mix. The set of production equipment requires moderate modification.

- Moderately high:** Integrated process is a combination of demonstrated processes and all relevant attributes (C, Y, T, TP or PC) are within the state-of-the-art. Some custom designed test equipment is required to be developed. A single, adequate domestic source for material exists. Insufficient amount of moderately skilled personnel are available, with no immediate replacement. The set of production equipment is available, but requires significant modification.
- High:** No comparable process, but all of the requirements for all relevant attributes (C, Y, T, TP and PC) are expected to be within the state-of-the-art. Majority of test equipment is required to be developed. A single, adequate foreign source for material exists. Insufficient amount of highly skilled personnel are available, with no immediate replacement. Some production equipment needs to be designed.
- Very high:** No comparable process and one or more key attributes (C, Y, T, TP or PC) exceed the state-of-the-art. Suitable test procedures, or test equipment/facilities have not been designed. No defined source of material. Insufficient amount of expertly skilled personnel are available, with no immediate replacement. Suitable production equipment and facilities are unavailable.

Note: Y = Yield; T = Tolerance and/or Precision; TP = throughput; C = Complexity; PC = Process Controls. To achieve a rating of Very Low, ALL criteria must be met. For ratings of Low to Very High, the most restrictive criteria that would result in the highest rating is used to determine the rating.

SUPPORTABILITY: An evaluation of how well the composite of support considerations necessary to achieve the effective and economical support of a system for its life cycle meets stated quantitative (e.g., Mean Time Between Failures (MTBF)) and qualitative readiness and utilization requirements. This includes integrated logistics support and logistic support resources related Operating and Support (O&S) cost considerations.

- | | |
|------------------|--|
| Very low: | Mature, well defined and known support requirements, and an in-place operations and support capability that is demonstrated and satisfactory. |
| Low: | Stable existing logistics support system with some new aspects utilizing an existing logistics support system which has undergone major modifications. |
| Moderately low: | New logistics support design utilizing no part of any existing logistics support system. Initial unproven support system in place but lacking substantial actual history. |
| Moderate: | Logistics support analysis completed resulting in firm definitions and requirements for the logistics support system design. |
| Moderately high: | Logistics support analysis initiated. Enough definition to conduct preliminary O&S cost trade studies. |
| High: | Logistics support system design is conceptual in nature and supportability requirements for the system, subsystems and components are not adequately defined to start planning the logistics support system. |
| Very high: | No conceptual logistics support system. No activity to date associated with any formal logistics support analysis. |

ANNEX G-3: ADDITIONAL DIALOGUE BETWEEN IDA AND BOEING D&SG

- A. Why does Boeing believe that RACM is applicable more to management of existing programs than to proposal preparation for new programs? Would RACM be too demanding for the latter purpose, with regard to either historical data or expert judgment? Can you conceive of any modifications to RACM that would render it more amenable to handling new project proposals, or does the RACM modeling approach inherently run counter to this type of analysis?

Fundamental to RACM is the concept of the spending of budgeted resources. Internally, within the company, prior to the submittal of a proposal and when one is attempting to establish what the PDF (Probability Density Function) of the program's future cost is, the budget value is not known. What is needed at this time is an estimate of the PDF that is not conditioned upon a budget value. If it is true that a project's budget affects its ultimate cost then that effect, well modeled, would certainly be useful in our efforts to both plan and manage a program. RACM's modeling of budget effects is one of its features, that feature makes it more applicable to programs that have budgets. This should not be taken to mean that it might not be valuable at the early stages of a program.

Some RACM inputs are, although easy to input, not organized as is our available data. In these cases, of course, RACM inputs could be estimated from our data, but with a loss of fidelity. For example:

A project manning plan would typically be available during the development of a design or study proposal submittal. Such a plan might not exist for estimates submitted in a proposal that included estimates of follow-on contracts. Certainly during the production phases of a program the manpower levels would vary widely.

Similarly, a project's network description may well be known for the immediate contract being proposed, but is unlikely to be known for follow-on contracts except as a historical precedent.

The cost improvement percentages inputs could be well suited to some phenomena, but would be an awkward way of dealing with, for example, the learning curve.

Rate increases: Their uncertainty in time is typically not a constant. The amount of uncertainty is usually related to the time distance into the future of the proposal expenditures. Wrap rate changes between WBS elements are, in my experience, highly correlated, one might have to, for example, inflate the rate uncertainties for each WBS in order that their treatment as independent variables generate the approximately correct total effect—not a desirable solution.⁴

RACM combines features that generate an estimate of variation, the uncertainty parameters, and elements that change the estimate's central tendency, e.g., the number of expected problems and plan parallelism. Because of this RACM appears to be intended to be an add-on estimating model that generates PDFs, not strictly a "risk" estimating model...As an estimating model RACM can only be viewed as an add-on, it does not, and certainly was not intended to include the bulk of the issues involved in cost estimating. But as an add-on, its value is related to its context, the statistical properties of the estimate to which it is applied. The input uncertainty parameters to RACM allow some of these estimate variances to be passed into RACM for summarization. In the calculation of PDFs the vast majority of the work is in estimating the variance parameters, not in their summarization, operations that RACM and the EELV risk models are largely limited to.

B. Please elaborate on Boeings concerns regarding: i) RACM's assumption of level-loaded staffing throughout the duration of a project; ii) RACM's use of normal (vs. lognormal) cost distributions; iii) the ability of users to estimate extreme tail events (e.g., the 90th or 95th percentile of cost).

i) RACM's assumption of level-loaded staffing throughout the duration of a project

The staffing of a program (contract) is seldom level loaded, perhaps never, if study contracts are excluded and depending on the definition of "level". Whether this discrepancy is important to the accuracy of the resulting PDF is unknown. However, there is a more fundamental problem. Typically budgets for direct cost are not set and maintained at the set level. Targets are set, below or well below the "contract value", performance is measured, and targets are periodically reset or reallocated commensurate with the performance achieved. In addition, the target value is just that, a target, it is

⁴ Note added by IDA: the idea is to inflate the variances to compensate for lack of explicit treatment of correlations. We have made a similar observation.

generally not manned to. During program's execution, manning tends to be related, in direct areas, to schedule and technical performance. Program specific indirect areas typically are budgeted as a ratio to program direct costs and workload performance. More broadly based indirect costs are multiplicative, i.e., they are ratio allocated via a rate.

The idea that a program's cost can be lowered or raised, just by adjusting its budgets, was discovered long ago. That is why we have elaborate mechanisms to produce estimates, budgets, targets, performance measures. This issue could certainly have an effect, if it were not well managed. Much of the cost management effort on a program is expended to manage out the effect that poor budgeting might have.

ii) RACM's use of normal (vs. lognormal) cost distributions

All of our cost modeling experience strongly suggests that cost PDFs are lognormal. I know of two possible mechanisms that generate such outcomes in linear systems. First, if the driving phenomena are exponential, the expanding sum will pass through a gamma distribution to become approximately lognormal. Second, if the cost generating process has multiplicative elements and is driven by approximately normal distributions, an approximately lognormal distribution will result. The cost generating process is at least partially multiplicative. Labor rates, overhead rates, as well as some distributed costs factors are typically a phenomena more broadly based than the specific program application. They therefore become effectively multiplicative to the specific program.

iii) the ability of users to estimate extreme tail events (e.g., the 90th or 95th percentile of cost).

If the users are deriving the extreme tail events statistics from historic data there will be relatively few occurrences of these low probability events, perhaps none in there relevant historic experience or data. For example, it is not unusual for a cost estimating relationship to have fewer than 10 data points. The accuracy of a 90th percentile estimate is necessarily low in this cases like this, unless the distribution form is assumed or known. A judgmental assessment may not assume the correct distribution. More likely, in my opinion, a judgmental estimate would treat the data, implicitly, as non-parametric, if so there would be insufficient data to determine the 90th percentile with any accuracy.

Again, the likelihood that a person will have experienced the extreme tails of a distribution for a class of major program events is necessarily low. Even if experienced, for him to know that an experienced event was extreme he must have knowledge of the remainder of the distribution. Depending on the class of events, a person, in their working life, may never experience enough events to have that knowledge.

- C. What is Boeing's opinion of RACM's assumption that "money allocated is money spent," so that initial WBS budget allocations (during execution) are never underrun. What about the "standing-army problem," whereby parallel activities continue to burn money at a constant rate while waiting for the slowest activity to complete.

The cost management mechanisms in place at Boeing largely preclude the "money allocated money spent" phenomenon. Budgets are allocated, but at a tight, "motivational," level. The balance-to-go budgets may be reallocated based on the balance-to-go effort remaining to accomplish the plan. In my experience, budgets are "held, authorized, monitored, and controlled," at either the level where their use is initiated or often at some one or two levels higher. One effect of this is that the initiator of a cost may not have a budget to "live up to." For example: the Control Account Manager (CAM), typically an IPT leader, is usually the lowest level of budget control, while the members of the IPT initiate its expenditures. For these budgets the money isn't actually allocated to the people who spend it.

Budget pressure and the ability to move resources to other parts of a project or to other projects or sites allow the organization to operate without "standing-armies". Boeing resources and projects are continuously managed to provide the company with the ability to move resources between requirements. For example the company balance between commercial and DoD work is at least partially intended to provide a continuity of resources despite a time varying workload, project, and customer mix. Within the D&SG we have a policy of standardizing processes, specifically intended to allow work or resources to be shifted across locations and between projects.

- D. How well do either RACM or the EELV approach interface with D&SG accounting systems? Specifically, comment on the adequacy of the two models' treatment of indirect costs and overhead rates.

The pre-Boeing merger EELV model did not treat or deal with any rate or overhead issue. There was no relationship to the accounting system other than the WBS. RACM has rate uncertainty and growth inputs by WBS. In my experience rates are highly correlated between WBS elements. The RACM assumption of independence would lead to misleading results.

E. Lockheed-Martin believes that RACM can be used during program execution, feeding actual costs into the model and estimating cost-to-completion. Does Boeing believe that RACM can be used in this fashion? Would there be any utility from doing so? Does Boeing use their own models in this fashion (i.e., are the same models used up-front and during follow-up, or are they distinct models?).

Yes. At the present time in the Puget Sound D&SG Program Financial Controls and the Business Planning and Control groups, the organizations reporting cost and schedule status in D&SG, do not deal with issues of uncertainty. The balance to go estimates and plans are developed typically by the functionally responsible organizations as part of the budgeting, targeting and performance measurement process. The estimating organization very largely does not produce balance to go estimates, it does contract change estimating and follow-on contract estimating.

Would there be any utility from doing so?

If by this question you are asking whether RACM is a good mechanism for cost-to-completion estimating as it normally is done, that is, without an uncertainty component, then answer is that programs at Boeing have far better methods for producing such estimates. Conversely, if you are asking whether PDF information would be useful, I believe that answer is probably yes. However, I am not aware of any organizations in my local area (Puget Sound D&SG) that at the present time generates a PDF output in their cost-to-completion estimating process.

Does Boeing use their own models in this fashion (i.e., are the same models used up-front and during follow-up, or are they distinct models?)

Historically there have been programs/divisions where a single model has to some degree dealt with both new program and balance to complete estimates. The two problems have very different conditionalities. As one moves through a program's phases the dominant predictors of future cost change. Initially it is the product design and schedule. As these become realized, cost and schedule performance realizes much of the

uncertainty in those initial estimates. For example, once the first unit has been produced, the second unit's cost is more closely related to it than it is to the product's weight.

F. Does Boeing feel that they would have benefited from a training course on RACM? On the EELV Program Office approach? Is there any standard training within Boeing on risk analysis, or proposal preparation more generally?

Yes, there is standard training for proposal preparation. It includes some information on risk analysis. I would not recommend that either approach be utilized for internal applications. If a customer finds the approaches to be helpful then, of course, appropriate resources should be expended. Before either approach is implemented it should be demonstrated to be valid, i.e., predict the future, and useful in managing our resources.

APPENDIX H

BOOK AND YOUNG'S BUDGET ALLOCATION ALGORITHM

This appendix summarizes Book and Young's budget allocation algorithm.¹ Their algorithm is an alternative to the one contained in RACM.

For the j th WBS element ($j = 1, \dots, n$), Book and Young define the *need* as the difference between some high percentile of the cost distribution (in their examples, they use the 70th percentile) and the most likely estimate. They also truncate the need at zero:

$$\begin{aligned} \text{Need}_j &= (70\text{th percentile cost}) - (\text{most likely cost}) \text{ if } > 0, \\ &= 0 \text{ otherwise.} \end{aligned} \quad (\text{H-1})$$

Book and Young define the *need base* as a quadratic function of the needs for the various WBS elements:

$$\text{Base} = \sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \text{Need}_i \text{Need}_j = \sum_{i=1}^n (\text{Need}_i)^2 + 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n \rho_{ij} \text{Need}_i \text{Need}_j, \quad (\text{H-2})$$

where ρ_{ij} is the correlation between the costs of the i th and j th WBS elements. Book and Young indicate that definition (H-2) was motivated by the somewhat analogous formula for the variance of a sum (the square of our equation (III-15)).

Book and Young then identify total *risk dollars* on the project as the difference between some high percentile of the distribution of total cost and its most likely estimate.

¹ To our knowledge, this algorithm was first published in Philip H. Young, "FRISK: Formal Risk Assessment of System Cost Estimates," The Aerospace Corporation, Los Angeles, California, February 1992. The algorithm has been presented at a variety of conferences over the ensuing years, starting with Stephen A. Book, "Recent Developments in Cost Risk," at the 4th Annual IDA Cost Research Symposium, May 1992. Minor updates to the algorithm have been presented more recently at the Institute for Operations Research and the Management Sciences (INFORMS) conference in Atlanta, Georgia, November 1996; and the Military Operations Research Society (MORS) conference in Quantico, Virginia, June 1997.

They generally choose the same percentile of the total-cost distribution (e.g., the 70th) as in the definition of WBS-element need:

$$\begin{aligned} \text{Project risk dollars} &= (70\text{th percentile total cost}) - (\text{most-likely total cost}) \text{ if } > 0, \\ &= 0 \text{ otherwise.} \end{aligned} \quad (\text{H-3})$$

Note that that the 70th percentile of total cost is *not* generally equal to the sum of the 70th percentiles of the individual WBS elements. Nor is the most-likely total cost equal to the sum of the most-likely estimates for the individual WBS elements. Thus equation (H-3) is *not* the sum of the WBS-element needs.

Finally, Book and Young allocate to the i th WBS element the following *share* of project risk dollars:

$$\begin{aligned} \text{Share}_i &= \sum_{j=1}^n \rho_{ij} \text{Need}_i \text{Need}_j / \text{Need base} \\ &= \sum_{j=1}^n \rho_{ij} \text{Need}_i \text{Need}_j / \sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \text{Need}_i \text{Need}_j \\ &= \text{Need}_i \times \sum_{j=1}^n \rho_{ij} \text{Need}_j / \sum_{i=1}^n \sum_{j=1}^n \rho_{ij} \text{Need}_i \text{Need}_j. \end{aligned} \quad (\text{H-4})$$

If the correlations are non-negative, the shares will be non-negative as well. Also, it is obvious from equation (H-4) that $\sum_{i=1}^n \text{Share}_i = 1.0$.

Although the shares defined in equation (H-4) certainly exhaust the risk dollars available on the project, Book and Young have not demonstrated any sense in which their allocation scheme is optimal. By contrast, the optimality of RACM's algorithm was explored in Appendix B.

APPENDIX I

USING OVERLAY CHARTS WITH CRYSTAL BALL

This appendix demonstrates the use of overlay charts with the Crystal Ball add-in software to Microsoft Excel. According to the manual:

After completing a simulation with multiple related forecasts, you can use Crystal Ball's overlay chart feature to view the relative characteristics of those forecasts on one chart. The overlay chart superimposes the frequency data from selected forecasts in one location so that you can compare differences or similarities that otherwise might not be apparent. There is no limit to the number of forecasts you can view at one time on the overlay chart.¹

The remainder of this appendix provides a series of instructions and associated screen captures to illustrate the use of overlay charts. The specific example is taken from the Boeing D&SG data discussed in Appendix G.

Figure I-1: Designate multiple output variables using the "Define Forecast" command.

Figure I-2: Click the right-arrow button, "Start Simulation."

Figure I-3: Select "Open Overlay Chart" from the "Run" pull-down menu.

Figure I-4: Click the "Add Forecasts" button.

Figure I-5: Select the (multiple) output variables for inclusion in the overlay chart.

Figure I-6: View the overlay chart.

¹ Decisioneering, Inc., *Crystal Ball: Forecasting and Risk Analysis for Spreadsheet Users, Version 4.0*, Denver, Colorado, March 1996, p. 174.

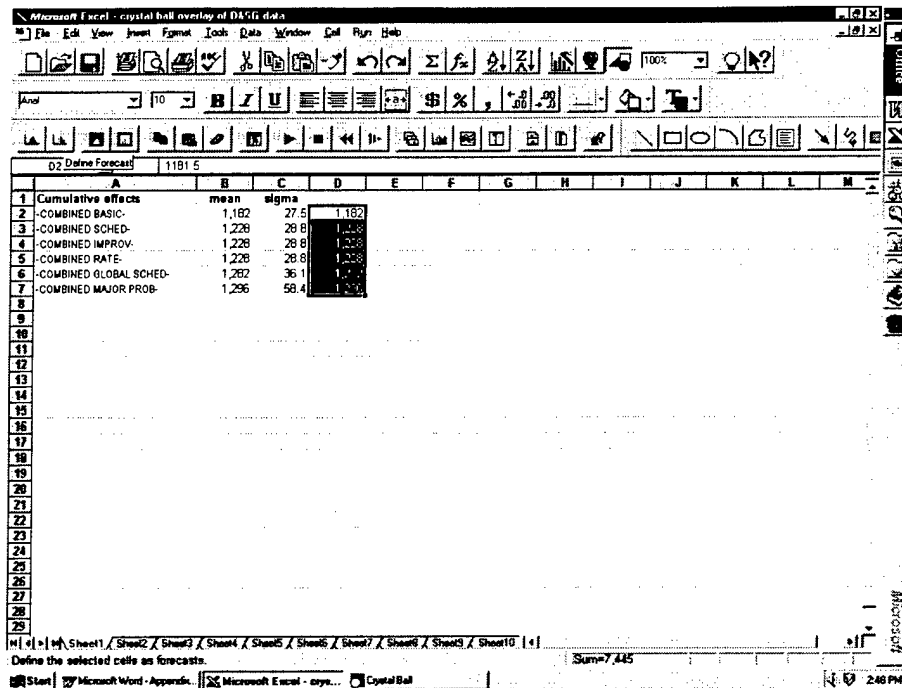


Figure I-1. Crystal Ball: Define Multiple Output Variables

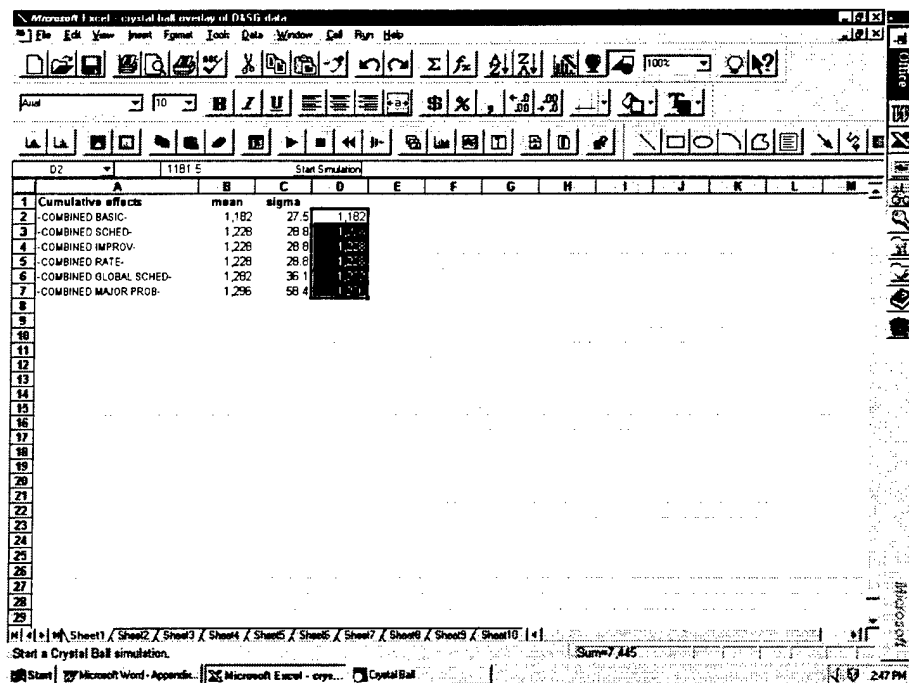


Figure I-2. Crystal Ball: Run Simulation

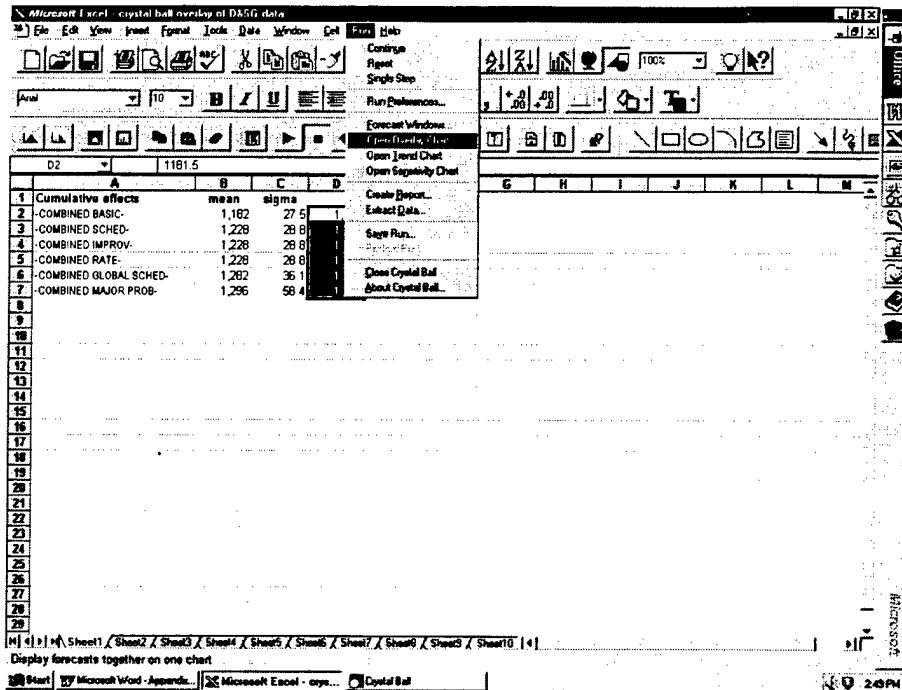


Figure I-3. Crystal Ball: Open Overlay Chart

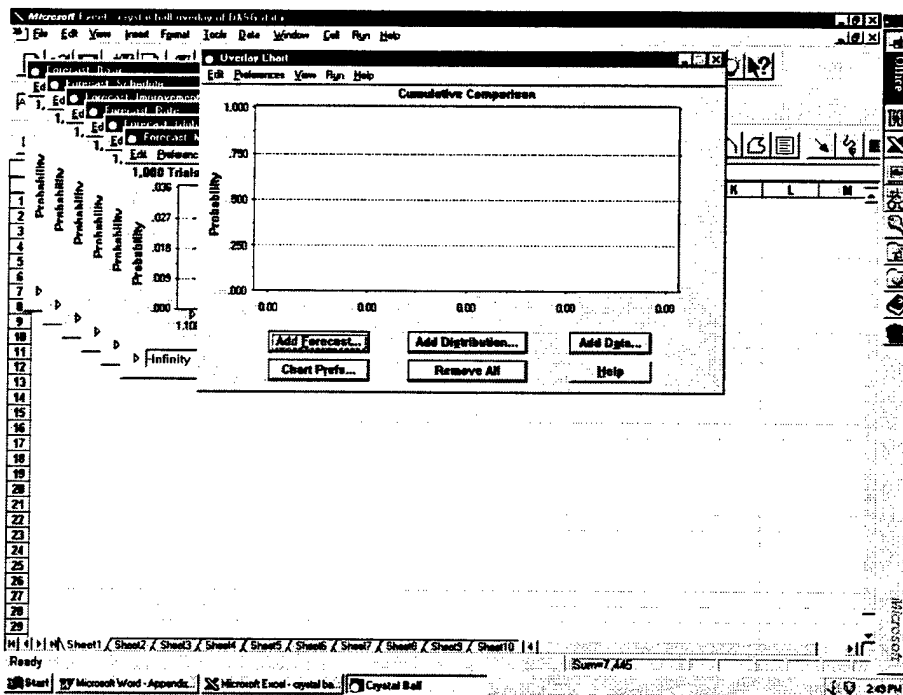


Figure I-4. Crystal Ball: Add Forecasts

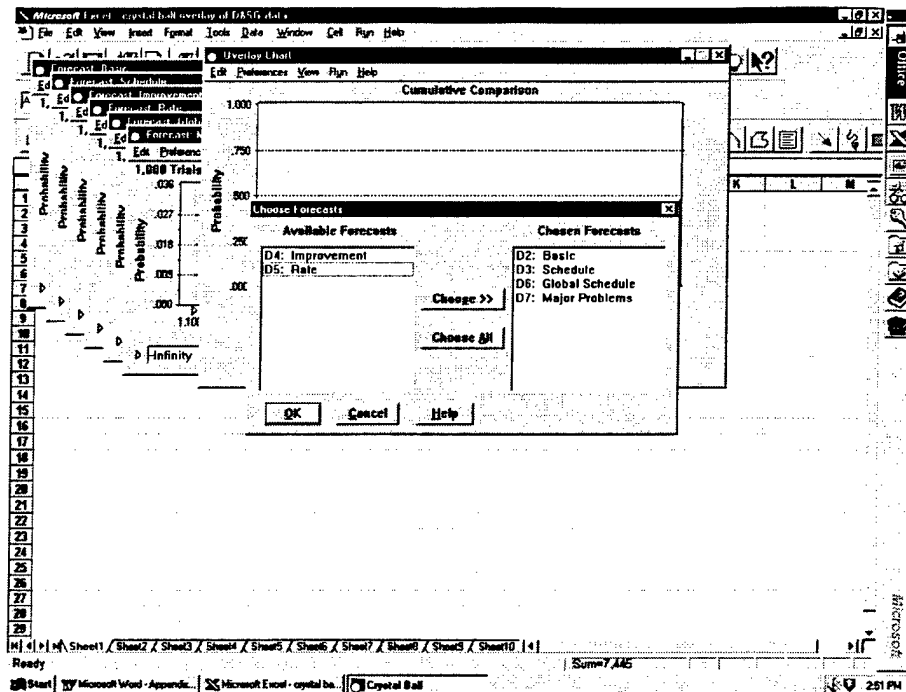


Figure I-5. Crystal Ball: Select the Multiple Output Variables

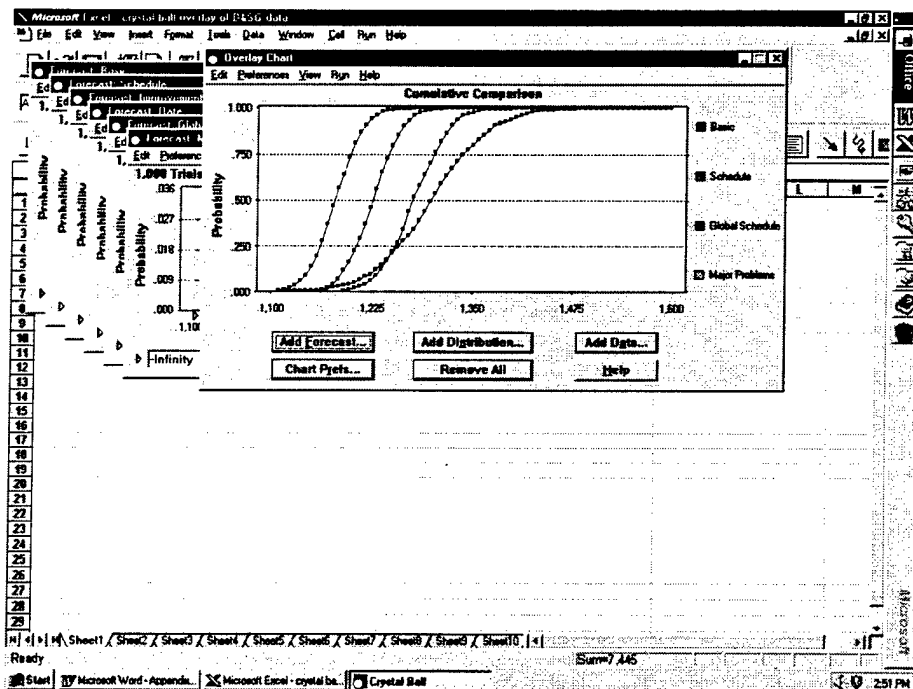


Figure I-6. Crystal Ball: View Overlay Chart

APPENDIX J

USING OVERLAY CHARTS WITH @RISK

This appendix demonstrates the use of overlay charts with the @RISK add-in software to Microsoft Excel. According to the manual:

[@RISK 3.5] allows multiple simulation outputs to be overlaid on the same graph in histogram or cumulative format....Overlays are especially useful for comparing output distributions in cumulative format. This comparison shows which outputs have higher probabilities at various points across the X-scale range. Comparisons may also be made for outputs in different simulations when multiple simulations are run.¹

The remainder of this appendix provides a series of instructions and associated screen captures to illustrate the use of overlay charts. The specific example is taken from the Boeing D&SG data discussed in Appendix G.

Figure J-1: Designate multiple output variables using the fourth button from the left on the toolbar, "Add the selected cells as @RISK outputs."

Figure J-2: Click the second button from the right on the toolbar, "Run Simulation."

Figure J-3: View the simulation results. Click the "Graph" button to graph the first output variable.

Figure J-4: Under the "Results" menu, cascade to "Graph" and then "Format..." Alternatively, right-click the graph window and select "Format..."

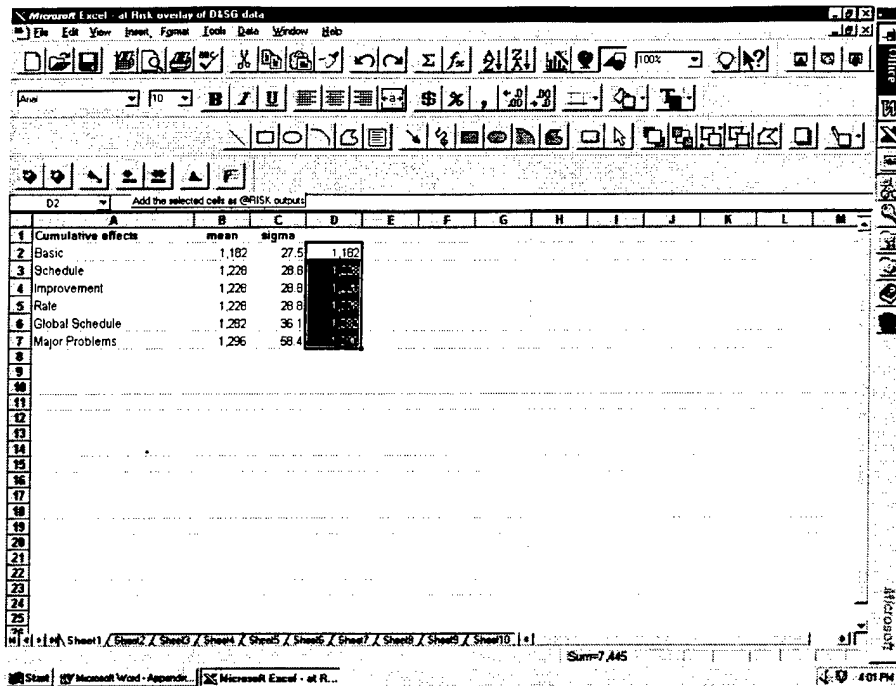
Figure J-5: Change the graph type to "Cumulative ascending," and display using "Fitted curve." Click the "Patterns" tab.

Figure J-6. Change the pattern to "Outline." Click the "Variables to Graph" tab.

Figure J-7: Using the <shift> or <control> keys, select the multiple output variables for inclusion in the overlay chart.

Figure J-8: View the overlay chart.

¹ Palisade Corporation, *Upgrade Guide for @RISK 3.5: Risk Analysis and Simulation Add-In for Microsoft Excel*, Newfield, New York, September 1996, pp. 10–11.



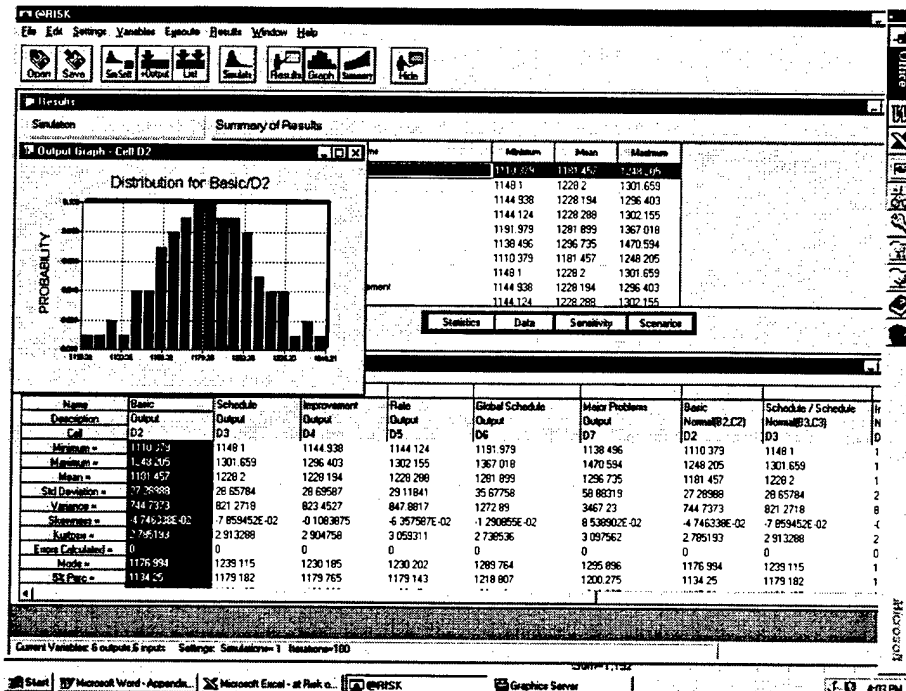


Figure J-3. @RISK: Graph the First Output Variable

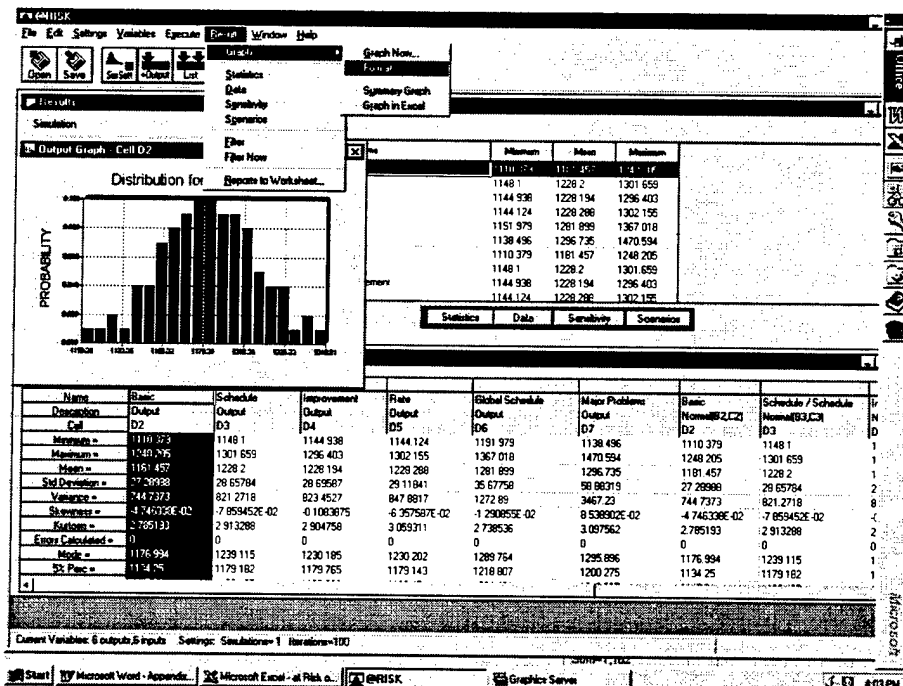


Figure J-4. @RISK: Format the Output Graph

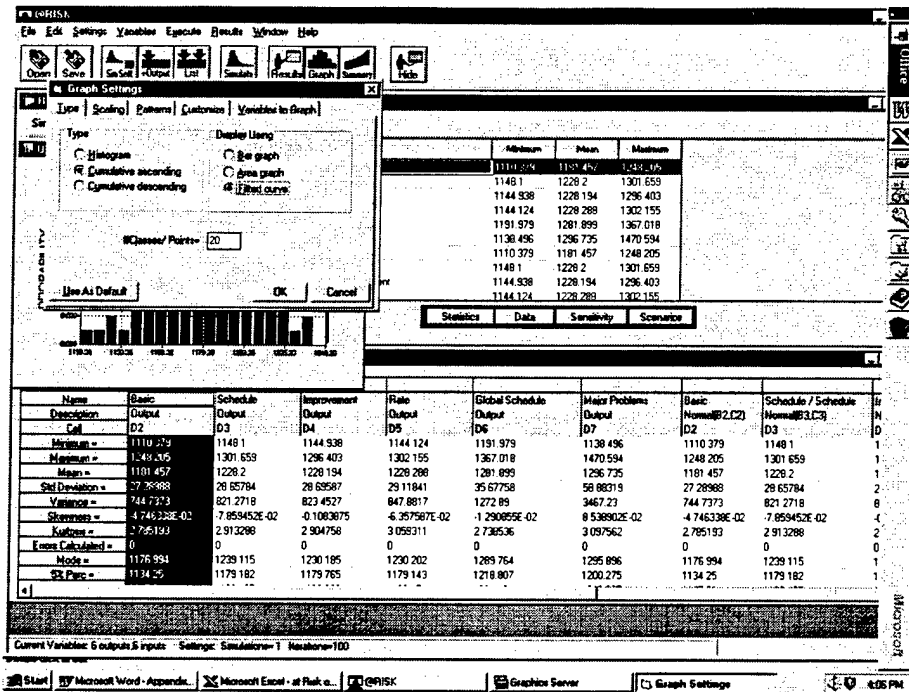


Figure J-5. @RISK: Change Output Graph Type

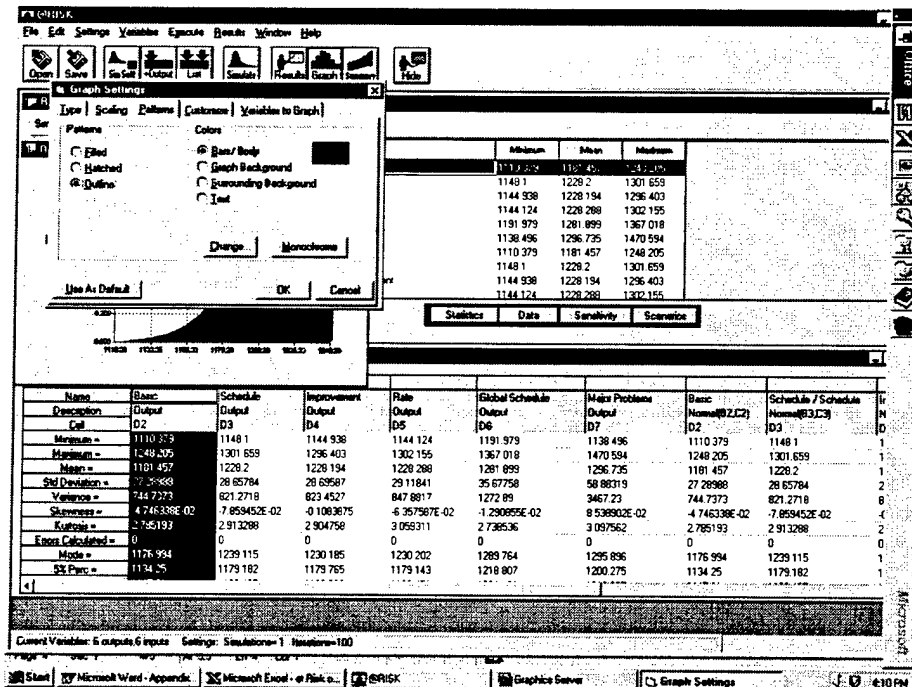


Figure J-6. @RISK: Change Output Graph Pattern

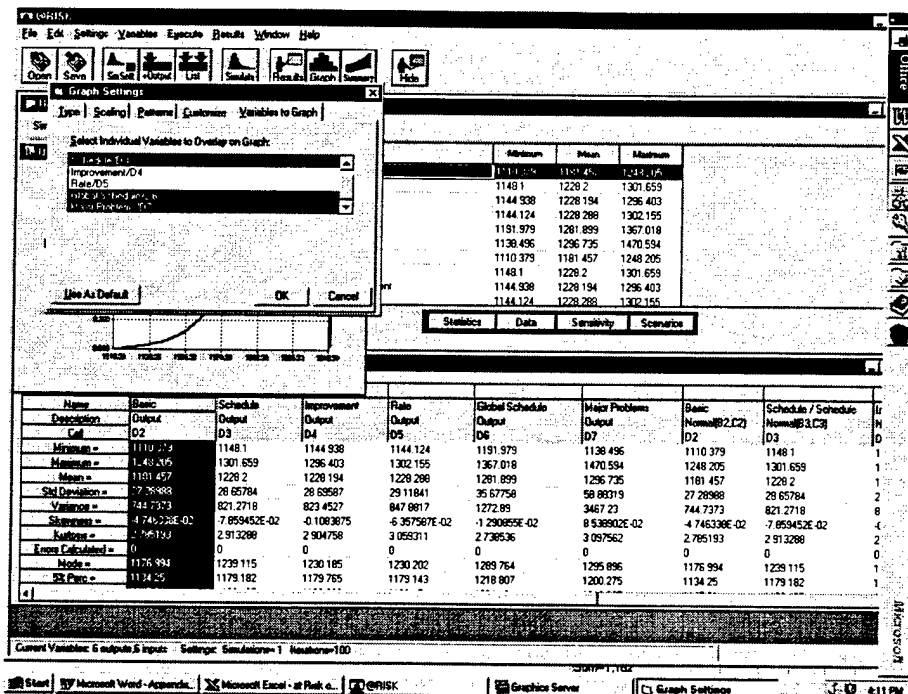


Figure J-7. @RISK: Select the Multiple Output Variables

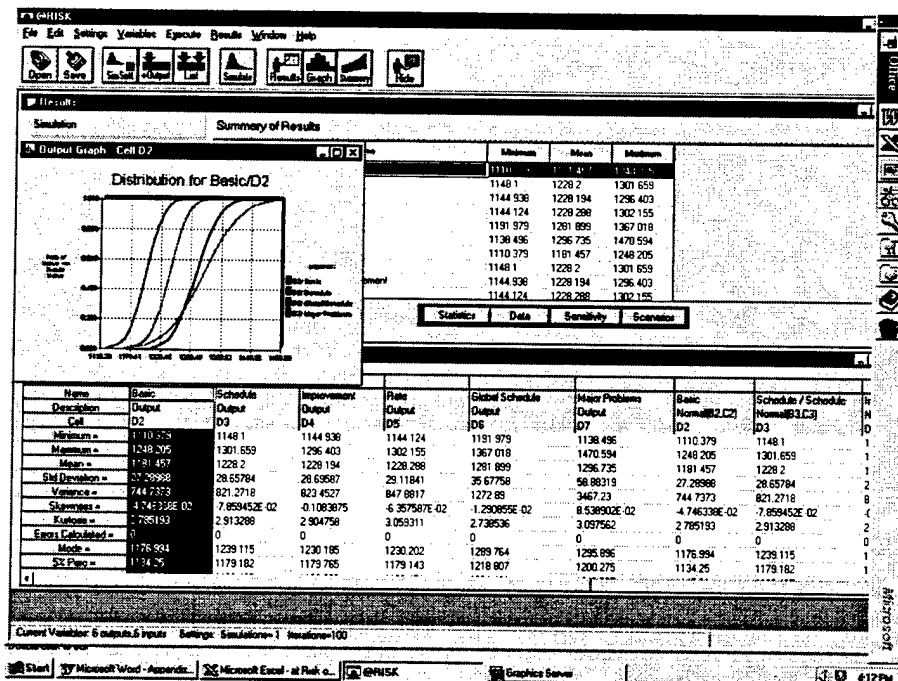


Figure J-8. @RISK: View Overlay Chart

APPENDIX K

OVERLAY CHARTS IN @RISK FOR MICROSOFT PROJECT

This appendix demonstrates the use of overlay charts in @RISK for Microsoft Project. This product is available from the following vendor:

Palisade Corporation
31 Decker Road
Newfield, New York 14867
1-800-432-7475
Fax: 607-277-8001
www.palisade.com

A beta version of @RISK for Microsoft Project (Version 3.5e) became available in September 1997 at a price of \$695. Although the beta version contains several bugs, it is adequate to demonstrate the capabilities of this product.

Figure K-1: The screen upon opening an existing Microsoft Project file under @RISK.

Figure K-2: Use the *SIMTABLE* function to specify a variable that will take on a different, deterministic value under each of the two scenarios. In this instance, the task "Write Proposal" will start on 10/26/95 under Scenario #1, and on 11/2/95 under Scenario #2.

Figure K-3: Specify random variables whose distributions are the same under the two scenarios.

Figure K-4: Highlight an output cell, and click the "Add @RISK output" button on the toolbar.

Figure K-5: Click the "Simulation Settings" button, and set the number of simulations (i.e., scenarios) equal to 2.

Figure K-6: Click the "Run Simulation" button, and view the output. Click the "Merge Sim#s" button to enable graphing of two scenario outputs on a common, overlay chart.

Figure K-7: Select one of the output variables from the merged list, and click the "Graph" button.

Figure K-8: Under the “Results” menu, cascade to “Graph” and then “Format...” Alternatively, right-click the graph window and select “Format...”

Figure K-9: Change the graph type to “Cumulative ascending,” and display using “Fitted curve.” Click the “Patterns” tab.

Figure K-10. Change the pattern to “Outline.” Click the “Variables to Graph” tab.

Figure K-11: Using the <shift> or <control> keys, select the multiple output variables for inclusion in the overlay chart.

Figure K-12: View the overlay chart.

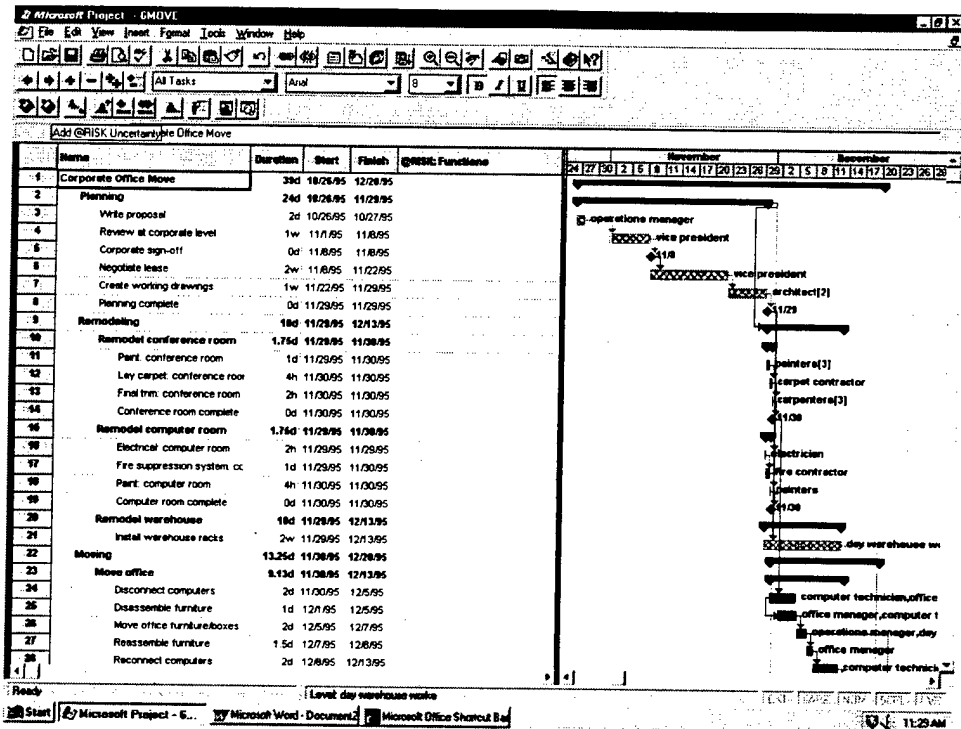


Figure K-1. @RISK for Microsoft Project: Opening Screen

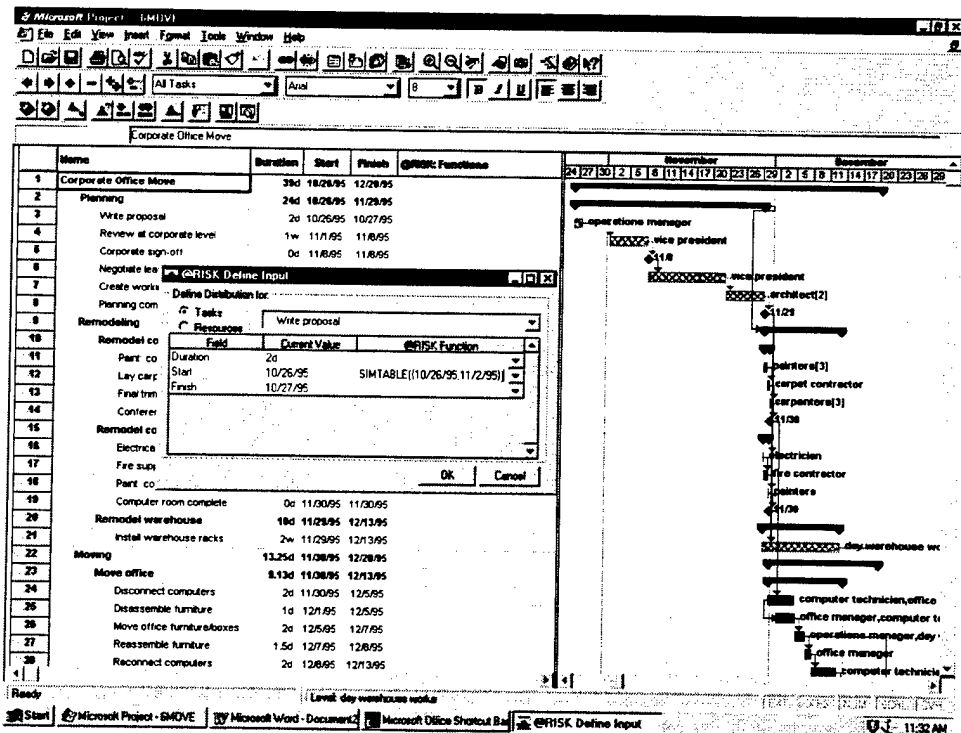


Figure K-2. @RISK for Microsoft Project: Specify a Scenario Variable

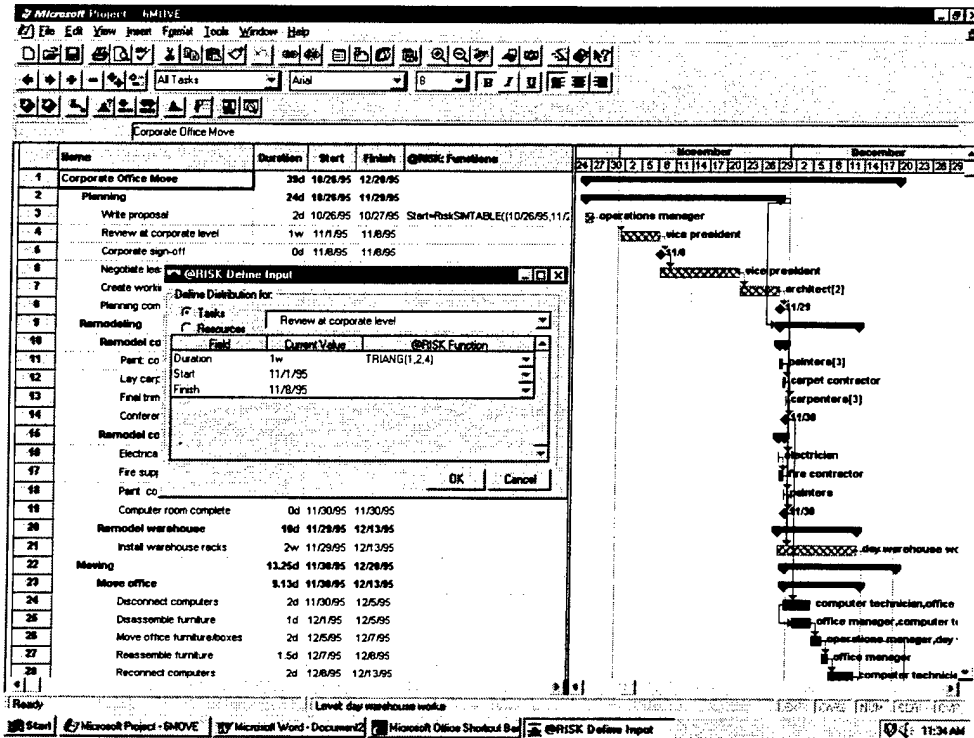


Figure K-3. @RISK for Microsoft Project: Specify a Random Variable

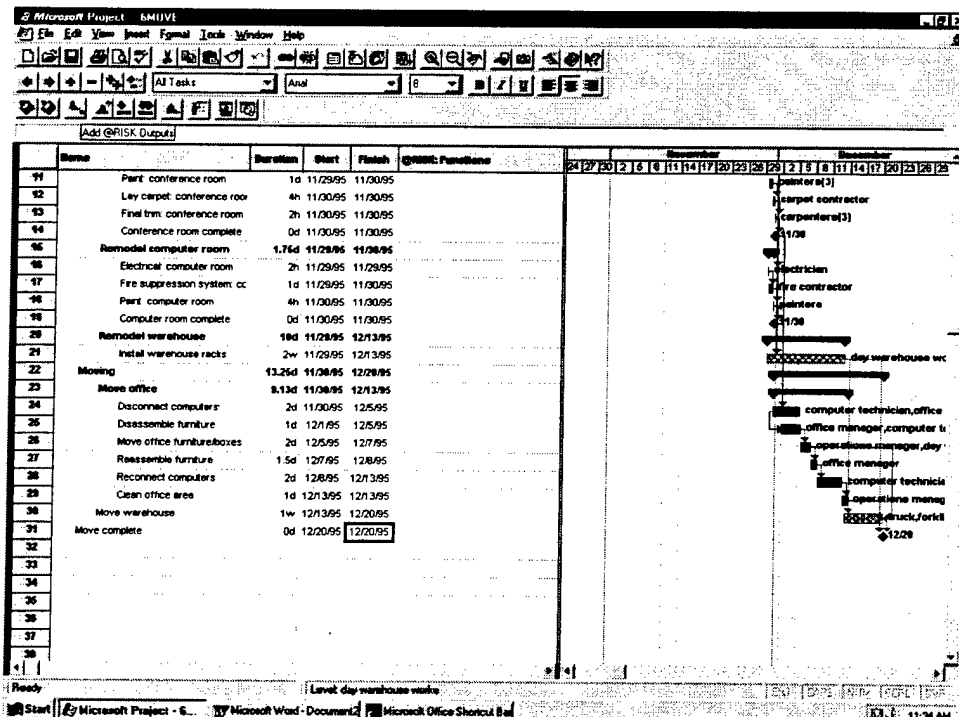


Figure K-4. @RISK for Microsoft Project: Enter an Output Variable

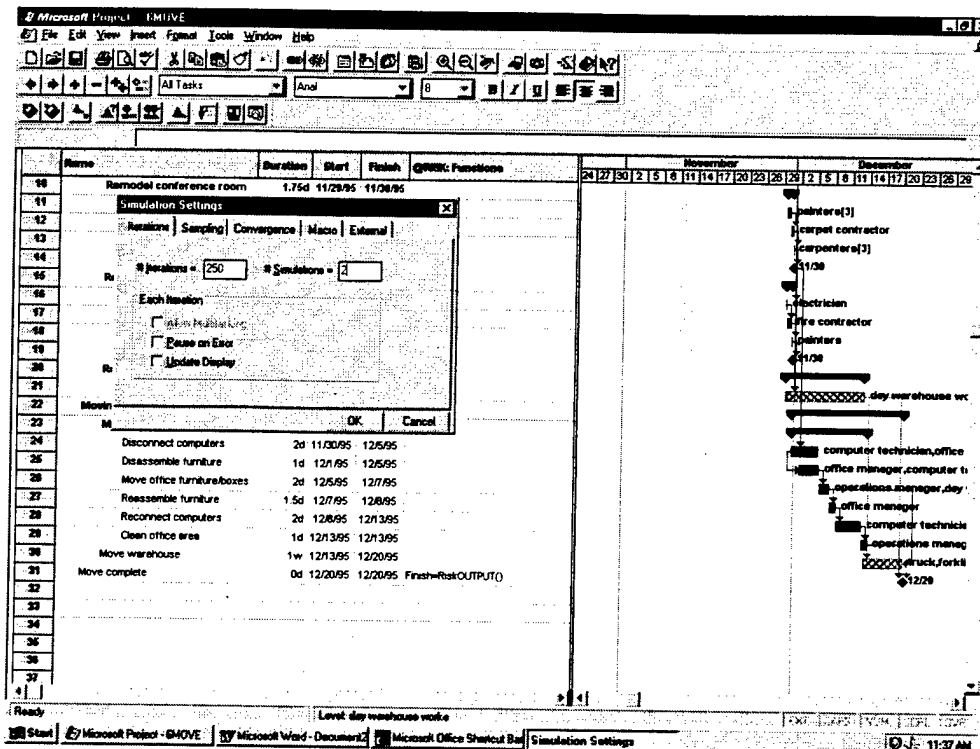


Figure K-5. @RISK for Microsoft Project: Set the Number of Simulations

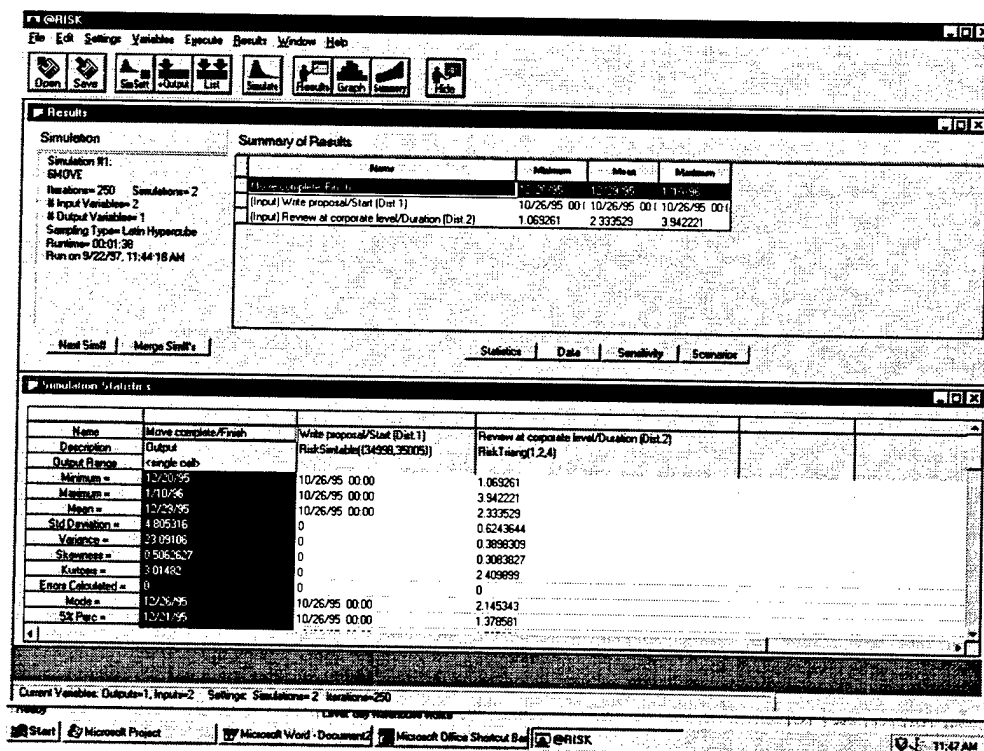


Figure K-6. @RISK for Microsoft Project: View Simulation Results

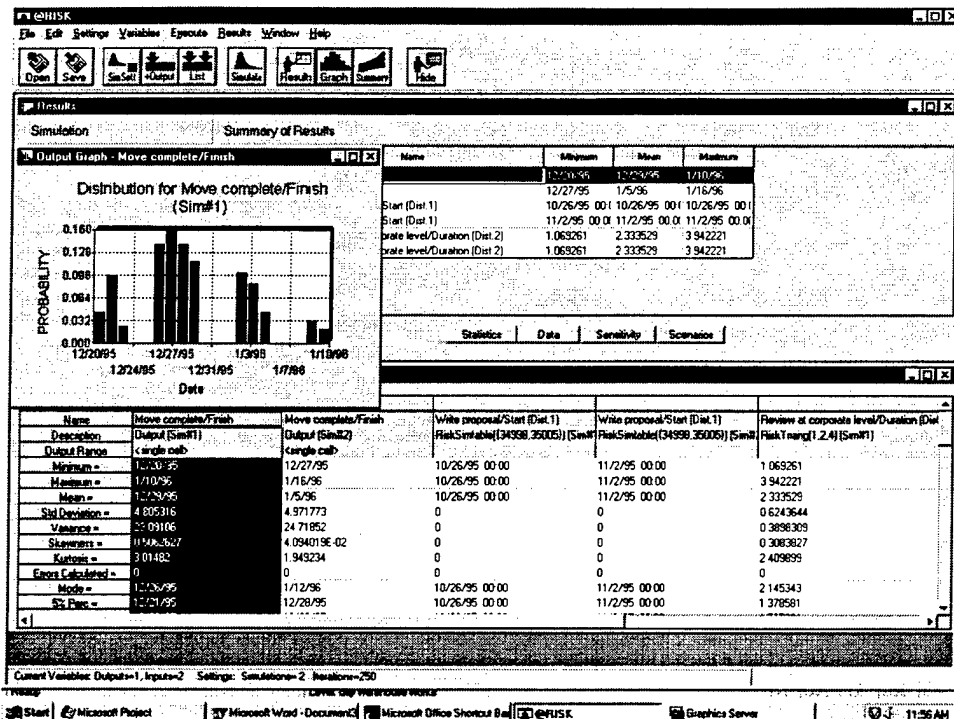


Figure K-7. @RISK for Microsoft Project: Graph the First Output Variable

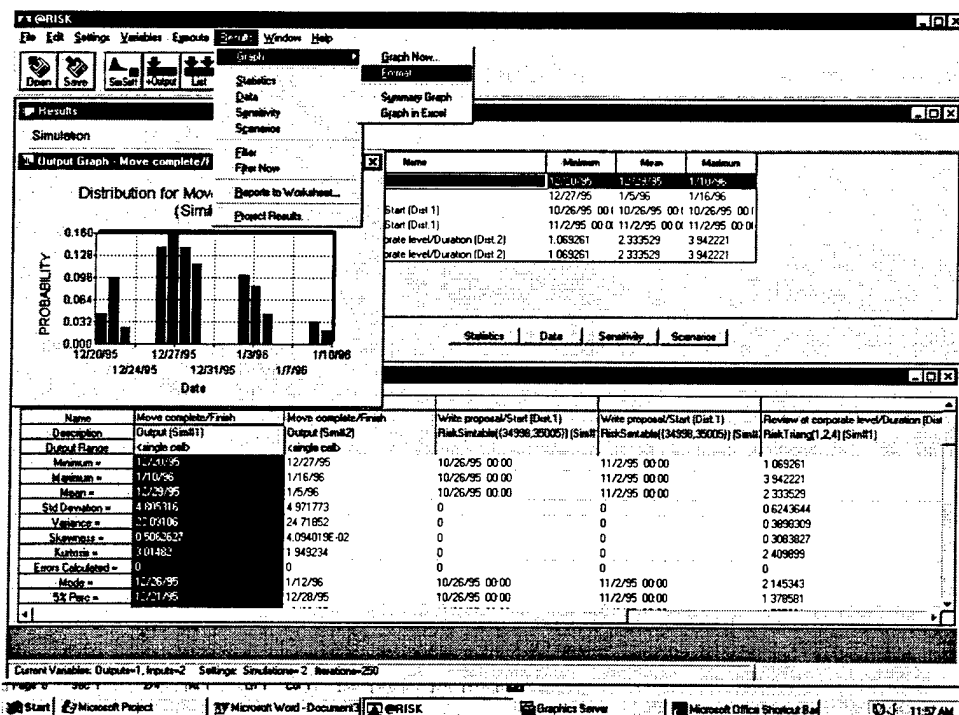


Figure K-8. @RISK for Microsoft Project: Format the Output Graph

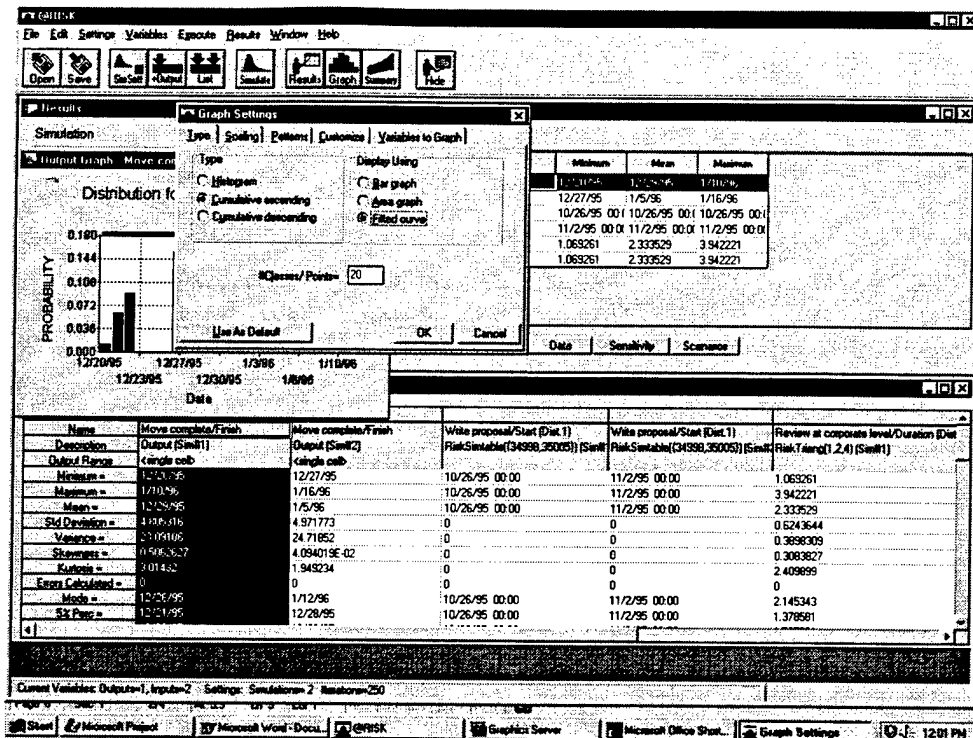


Figure K-9. @RISK for Microsoft Project: Change Output Graph Type

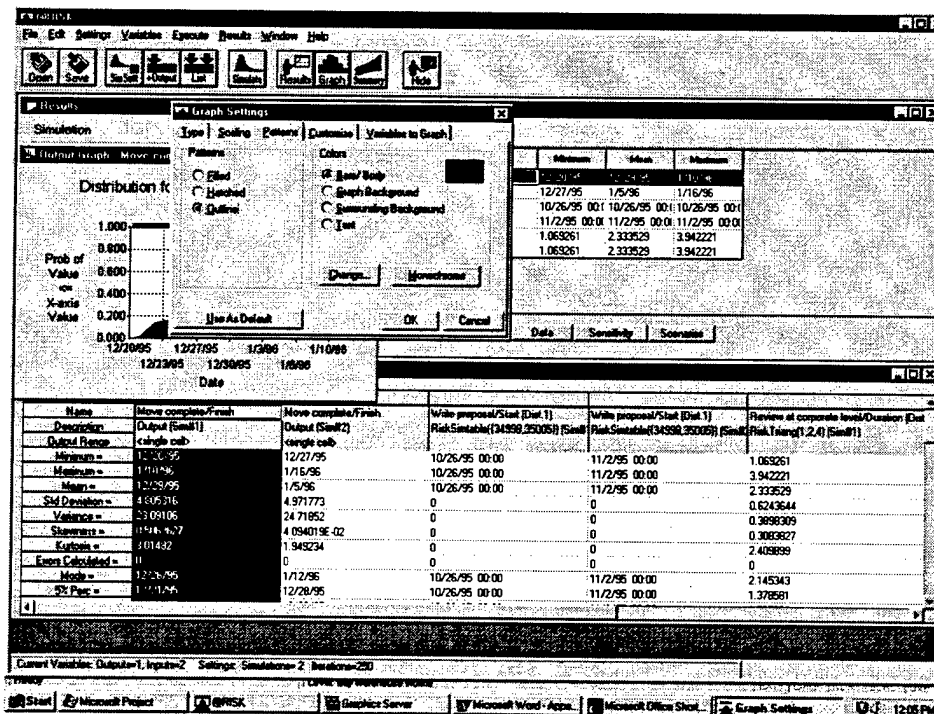


Figure K-10. @RISK for Microsoft Project: Change Output Graph Pattern

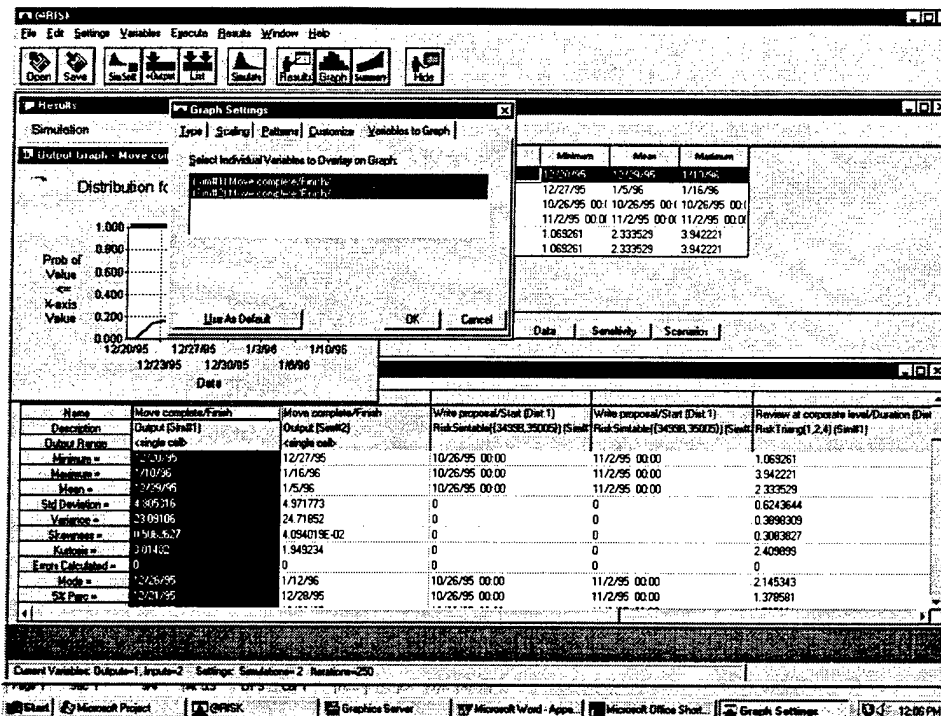


Figure K-11. @RISK for Microsoft Project: Select Multiple Output Variables

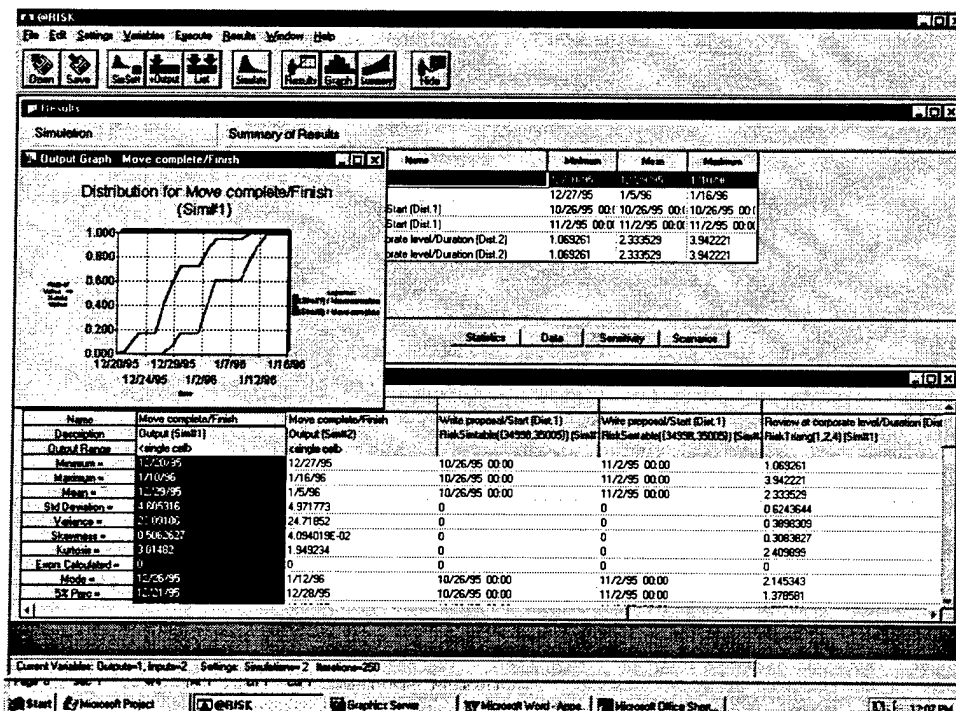


Figure K-12. @RISK for Microsoft Project: View Overlay Chart

ABBREVIATIONS

ACWP	Actual Cost of Work Performed
ACWS	Actual Cost of Work Scheduled
AFMC	Air Force Materiel Command
API	Acquisition Program Integration
BCWP	Budgeted Cost of Work Performed
BCWS	Budgeted Cost of Work Scheduled
BFL	Basic Factory Labor
BMDO	Ballistic Missile Defense Organization
BTU	Basic Task Unit
CAIG	Cost Analysis Improvement Group
CAM	Control Account Manager
CDF	Cumulative Distribution Function
CER	Cost Estimating Relationship
CI	Cost Improvement
CLT	Central Limit Theorem
COTS	Commercial Off-the-Shelf
CPI	Cost Performance Index
CPR	Cost Performance Report
C/SCSC	Cost/Schedule Control Systems Criteria
C/SSR	Cost/Schedule Status Report
CV	Cost Variance
D&SG	Defense and Space Group
DARPA	Defense Advanced Research Projects Agency
DoD	Department of Defense
DTC	Design to Cost
DTIC	Defense Technical Information Center
DTLCC	Design to Life-Cycle Cost
EAC	Estimate at Completion
EDI	Electronic Data Interchange
EDIS	Estimating Data and Information Services
EELV	Evolved Expendable Launch Vehicle
EMD	Engineering and Manufacturing Development

FBM	Fleet Ballistic Missile
FRISK	Formal Risk
FTE	Full-Time Equivalent
G&C	Guidance and Control
GUI	Graphical User Interface
IA&T	Integration, Assembly and Testing
IDA	Institute for Defense Analyses
IEAC	Independent Estimate at Completion
iid	independently and identically distributed
INFORMS	Institute for Operations Research and the Management Sciences
IPD	Integrated Product Development
IPT	Integrated Product Team
ISDS	Information, Space, and Defense Systems
LCC	Life-Cycle Cost
MAIMS	Money Allocated Is Money Spent
MEAC	Management Estimate at Completion
MIL-STD	Military Standard
MORS	Military Operations Research Society
MTBF	Mean Time Between Failures
NASA	National Aeronautics and Space Administration
NavAir	Naval Air Systems
O&M	Operations and Maintenance
O&S	Operating and Support
OBS	Organizational Breakdown Structure
OSD	Office of the Secretary of Defense
PA	Performance Analyzer
PDF	Probability Density Function
PERT	Program Evaluation and Review Technique
PMSI	Program Management Solutions, Inc.
PRICE-H	Parametric Review of Information for Costing and Evaluation of Hardware
PRICE-HL	Parametric Review of Information for Costing and Evaluation of Hardware-Life Cycle

Ps	Probability of Success
RACM	Risk Analysis and Cost Management
RFP	Request for Proposals
RSS	Root Sum-of-Squares
SCC	Statistical Cost Control
SPI	Schedule Performance Index
SV	Schedule Variance
S/W	Software
USDA	Under Secretary of Defense for Acquisition
WBS	Work Breakdown Structure

BIBLIOGRAPHY

- Abramson, R.L. and Stephen A. Book, "A Quantification Structure for Assessing Risk-Impact Drivers," The Aerospace Corporation, Los Angeles, California, September 1990.
- Abramson, R.L. and Philip H. Young, "FRISKEM—Formal Risk Evaluation Methodology," The Aerospace Corporation, Los Angeles, California, October 1990; also appeared in *Journal of Cost Analysis*, Spring 1997, pp. 29-38.
- Biery, Fred, David Hudak and Shishu Gupta, "Improving Cost Risk Analyses," *Journal of Cost Analysis*, Spring 1994, pp. 57-86.
- Book, Stephen A., "Recent Developments in Cost Risk," The Aerospace Corporation, Los Angeles, California, May 1992.
- Book, Stephen A., "Do Not Sum 'Most-Likely' Cost Estimates," The Aerospace Corporation, Los Angeles, California, May 1995.
- Book, Stephen A., O.F. Blackshire and Philip H. Young, "Validation Report on *RISK+* Risk Modeling Software for Microsoft Project 4.0 (U)," The Aerospace Corporation, Los Angeles, California, October 1995, Contractor-Proprietary.
- Book, Stephen A. and Philip H. Young, "Monte Carlo Generation of Total Cost Distributions When WBS-Element Costs Are Correlated," The Aerospace Corporation, Los Angeles, California, 1990.
- Book, Stephen A. and Philip H. Young, "Monte Carlo Simulation of Project Schedule Duration When Activity Times Are Correlated," The Aerospace Corporation, Los Angeles, California, 1992.
- Cost Management Systems, Inc., "Performance Analyzer V1.0 for Windows," Vienna, Virginia, 1995.
- Davenport, James M. and Ronald L. Iman, "An Iterative Algorithm to Produce a Positive Definite Correlation Matrix from an Approximate Correlation Matrix," Technical Report SAND-81-1376, Sandia National Laboratories, 1981.
- Decisioneering, Inc., *Crystal Ball: Forecasting and Risk Analysis for Spreadsheet Users, Version 4.0*, Denver, Colorado, March 1996.
- Dennis, J.E. and Robert B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Englewood Cliffs, New Jersey: Prentice-Hall, 1983.
- Elmaghraby, S.E., "The Estimation of Some Network Parameters in the PERT Model of Activity Networks: Review and Critique," *Advances in Project Scheduling*, edited by R. Slowinski and J. Weglarz, Elsevier, 1989.
- Feller, William, *An Introduction to Probability Theory and Its Applications*, Volume II, New York: John Wiley, 1971.
- Gallagher, Mark A. and David A. Lee, "Final Cost Estimates for Research and Development Programs Conditioned on Realized Costs," *Military Operations Research*, Vol. 2, No. 2, 1996, pp. 51-65.

- Garvey, Paul R., "A General Analytic Approach to System Cost Uncertainty Analysis," *Cost Analysis and Estimating—Tools and Techniques*, edited by W.R. Greer and D.A. Nussbaum, New York: Springer-Verlag, June 1990, pp. 161-181.
- Garvey, Paul R. and Audrey E. Taub, "A Joint Probability Model for Cost and Schedule Uncertainties," MITRE Corporation, Bedford, Massachusetts, September 1992.
- Goldberg, Matthew S. and Thomas P. Frazier, "Employment and Utilization of Engineers Among Defense Contractors," *Naval Research Logistics*, Vol. 41, No. 7, December 1994, pp. 853-874.
- Goldberg, Matthew S., "Some Fallacies in Cost-Risk Analysis," *PHALANX: The Bulletin of Military Operations Research*, September 1996.
- Goodman, Leo, "On the Exact Variance of Products," *Journal of the American Statistical Association*, Vol. 55, 1960, pp. 708-713.
- Gordon, Creaghe, "Risk Analysis and Cost Management (RACM): A Cost/Schedule Management Approach Using Statistical Cost Control (SCC)," Los Gatos, California, June 1997.
- Gupta, Shishu, David Olsen, David Hudak and Jennifer Keenan, "Cost Risk Analysis of the Strategic Defense System," TR-9042-2, Revision 1, TASC, Arlington, Virginia, 1992.
- Henrici, Peter, *Elements of Numerical Analysis*, New York: John Wiley, 1964.
- Hillier, Frederick S. and Gerald J. Lieberman, *Introduction to Operations Research*, fifth edition, New York: McGraw Hill, 1990.
- Iman, Ronald L. and W.J. Conover, "A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables," *Communications in Statistics—Simulation and Computation*, Vol. B11, No. 3, 1982, pp. 311-334.
- Iman, Ronald L., W.J. Conover and D.K. Zeigler, "Latin Hypercube Sampling (A Program User's Guide)," Technical Report SAND-79-1473, Sandia National Laboratories, 1979.
- Johnson, Norman L. and Samuel Kotz, *Distributions in Statistics, Continuous Univariate Distributions*, Volumes 1 and 2, New York: John Wiley, 1970.
- Kazanowski, A.D., "A Quantitative Methodology for Estimating Total System Cost Risk," in *Management of Risk and Uncertainty in Systems Acquisition: Proceedings of the 1983 Defense Risk and Uncertainty Workshop*, Defense Systems Management College, Fort Belvoir, Virginia, July 1983, pp. 135-163.
- Keefer, Donald L. and William A. Verdini, "Better Estimation of PERT Activity Time Parameters," *Management Science*, Vol. 39, No. 9, September 1993, pp. 1086-1091.
- Kellison, Stephen, *Fundamentals of Numerical Analysis*, Homewood, Illinois: Richard D. Irwin, 1975.

- Lee, David A. and John A. Dukovich, "Using the Rayleigh Analyzer," AT701C1, Logistics Management Institute, March 1998.
- Lootsma, F.A., "Stochastic and Fuzzy PERT," *European Journal of Operational Research*, Vol. 43, 1989, pp. 174-183.
- Lurie, Philip M. and Matthew S. Goldberg, "An Approximate Method for Sampling Correlated Random Variables from Partially-Specified Distributions," *Management Science*, Vol. 44, No. 2, February 1998, pp. 203-218.
- Lurie, Philip M., Matthew S. Goldberg, and Mitchell S. Robinson, "A Handbook of Cost-Risk Analysis Methods," Paper P-2734, Institute for Defense Analyses, April 1993.
- Maddala, G.S., *Limited-Dependent and Qualitative Variables in Econometrics*, Cambridge: Cambridge University Press, 1983.
- Merkhofer, M.L., "Quantifying Judgmental Uncertainty," *IEEE Transactions on Systems, Man and Cybernetics*, SMC-7, 1987, pp. 741-752.
- Morgan, M. Granger and Max Henrion, *Uncertainty: A Guide to Dealing With Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge: Cambridge University Press, 1990.
- Ortega, J.M. and W.C. Rheinboldt, *Iterative Solution of Nonlinear Equations in Several Variables*, New York: Academic Press, 1970.
- Palisade Corporation, *Guide to Using @RISK: Risk Analysis and Simulation Add-In for Microsoft Excel or Lotus 1-2-3*, Newfield, New York, July 1997.
- Palisade Corporation, *Guide to Using @RISK for Project: Risk Analysis and Simulation Add-In for Microsoft Project*, Newfield, New York, July 1997.
- Papoulis, Athanasios, *Probability, Random Variables, and Stochastic Processes*, New York: McGraw-Hill, 1965.
- Press, William H., Brian P. Flannery, Saul A. Teukolsky and William T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge: Cambridge University Press, second edition, 1992.
- Program Management Solutions, Inc., *RISK+: Risk Analysis for Microsoft Project 4.1*, Redondo Beach, California, September 1996.
- Seiler, F., "Error Propagation for Large Errors," *Risk Analysis*, Vol. 7, 1987, pp. 509-518.
- Smith, P.L. and Stephen A. Book, "Reducing Subjective Guesswork and Maintaining Traceability When Estimating the 'Risk' Associated With a Cost Estimate," The Aerospace Corporation, Los Angeles, California, 1992.
- Spetzler, C.S. and C.A. Stael von Holstein, "Probability Encoding in Decision Analysis," *Management Science*, Vol. 29, 1983, pp. 151-173.
- U.S. Department of Defense, "Military Standard: Work Breakdown Structures for Defense Materiel Items," MIL-STD-881B, 25 March 1993.

- U.S. General Accounting Office, "Major Acquisitions: Significant Changes Underway in DoD's Earned Value Management Process," NSIAD-97-108, Washington, D.C., May 5, 1997.
- Welch, Ronald A., "Piece-Wise Linear Transformation of a Gaussian Random Variable," Lockheed Corporation, Sunnyvale, California, 15 July 1969.
- Young, Philip H., "FRISK: Formal Risk Assessment of System Cost Estimates," The Aerospace Corporation, Los Angeles, California, February 1992.
- Young, Philip H. and R.L. Abramson, "Using Risk-Impact Drivers to Form WBS-Element Cost Means and Variances," The Aerospace Corporation, Los Angeles, California, 1991.

UNCLASSIFIED

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22204-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE August 1998	3. REPORT TYPE AND DATES COVERED Final Report, Dec 94-Aug 98	
4. TITLE AND SUBTITLE "Evaluation of the Risk Analysis and Cost Management (RACM) Model"			5. FUNDING NUMBERS DASW01 94 C 0054 Task Order T-G7-1353	
6. AUTHOR(S) Matthew S. Goldberg and Charles A. Weber				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 1801 N. Beauregard Street Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Paper P-3388	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Mr. Wayne F. Abba Office of the Director, Acquisition Program Integration Room 3E1025, The Pentagon Washington, DC 20301			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12B. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) IDA evaluated the RACM concept and computer software developed by Lockheed Martin Missiles and Space Company. RACM was designed to help defense contractors estimate risk margins for proposals they submit to DoD. RACM also contains a module for managing risk reserves during program execution. RACM is based on a particular set of assumptions regarding contractor behavior. Given these assumptions, IDA determined that RACM computes rapidly and essentially accurately. Many of RACM's best features could be easily duplicated in a commercial software environment, which would also open up the possibility of substituting Monte Carlo simulation for analytical approximation. IDA conducted a field test at Boeing Defense and Space Group to determine whether RACM is portable to another defense contractor's site. The Boeing team had difficulty both eliciting the input required by RACM and mapping its cost and risk drivers into those required by RACM. Although RACM may have proved useful at Lockheed Martin's Sunnyvale site, it does not, in its current form, appear to be generally applicable throughout the defense industry.				
14. SUBJECT TERMS Risk Analysis; Cost Management; Computer Programs; Contractors; Department of Defense			15. NUMBER OF PAGES 276	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT SAR	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102**UNCLASSIFIED**