Risk Analysis in the NASA/Air Force Cost Model

Christian Smart, Ph.D., CCEA

Science Applications International Corporation

675 Discovery Drive

Suite 300

Huntsville, AL 35806

The latest version of the NASA/Air Force Cost Model (NAFCOM) includes a newly developed cost risk estimation capability that provides a seamless and simple interface for accomplishing complete cost risk analyses. It includes technical risk, cost equation uncertainty, and correlation.

Technical risk is provided by allowing the user to set low, most likely, and high values for each CER input. This is combined with estimating uncertainty to produce a probability distribution for cost at the WBS level. The WBS level risk estimates are then aggregated to obtain risk estimates for the subsystem level, stage level, and vehicle level, taking correlation into account. The user can set individual inter- and intra-subsystem correlations to any desired value in the range (-1,1).

NAFCOM risk analysis estimates are calculated using analytic approximation. Comparison tests show that the model produces results that are very similar to those produced by Monte Carlo simulations.

Introduction

Prior to 2004, the NASA/Air Force Cost Model (NAFCOM) provided only point estimates. No risk capability was included. The addition of a probabilistic cost risk analysis module provides the ability to identify, analyze, and quantify risk/uncertainty. It also expands NAFCOM from a static analysis tool to one that provides probabilistic analysis of uncertain results.

In 2003, Science Applications International Corporation (SAIC) began incorporating a cost risk capability in NAFCOM. Experts in the cost risk analysis field, including representatives from the Aerospace Corporation, MCR, MITRE, NASA, and the IPAO participated in the methodology development. A kick-off meeting was held at NASA HQ to discuss requirements and the design of the risk model for NAFCOM, including sampling approaches, incorporation of correlation, and model implementation. Dr. Steve Book worked directly with SAIC to ensure that the best possible approaches were considered for integration into NAFCOM and to consult on the implementation. A version of NAFCOM will full cost-risk capability was released in 2004.

NAFCOM Overview

NAFCOM is a parametric estimating tool for space systems. It uses cost estimating relationships (CERs) that correlate historical costs to mission characteristics to predict new project costs. These CERs are based on historical NASA and Air Force project costs and provides estimates at the subsystem and component levels. NAFCOM has evolved significantly over the past 15 years and includes numerous useful features. The recent addition of a risk capability is only the latest development. See Figure 1 for a timeline of NAFCOM's evolution.

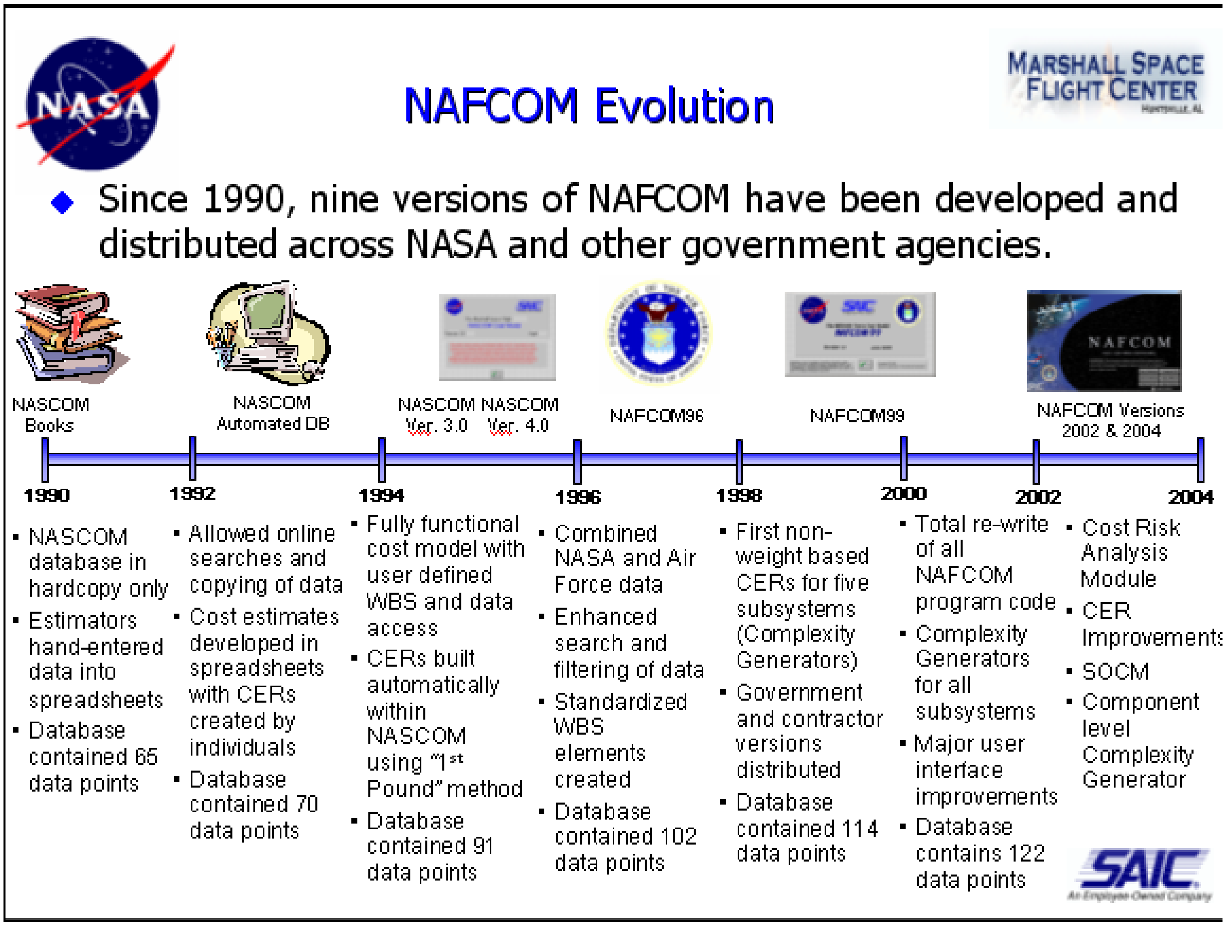

Figure 1. NAFCOM Evolution.

NAFCOM estimates costs at the subsystem level – structures, thermal control, etc. Prior to the development of the risk capability, NAFCOM only provided point estimates at the subsystem level, and the direct roll-up of subsystem point estimates at the stage and vehicle levels.

The user can run NAFCOM in either "Risk On" or "Risk Off" mode. In "Risk Off" mode, point estimates are calculated and aggregated without any consideration of risk or correlation among cost elements. In "Risk On" mode, the user inputs low, most likely, and high values for all CER inputs. In Figure 2 the "Risk On" button is displayed on the toolbar on the left.

Technical Risk

The first step in performing a cost risk analysis in NAFCOM is to define risks for the inputs to the CERs. For each independent variable, the user defines a minimum, most likely, and maximum value. The appropriate CER is used to estimate costs for three cases: all independent variables set to the minimum values; all independent variables set to the most likely values; and all independent variables set to the maximum variables. A triangular distribution is then defined by these three elements. In Figure 2, the user interface for inputting the low, most likely, and high values for each independent variable is displayed for the structures and mechanisms subsystem.

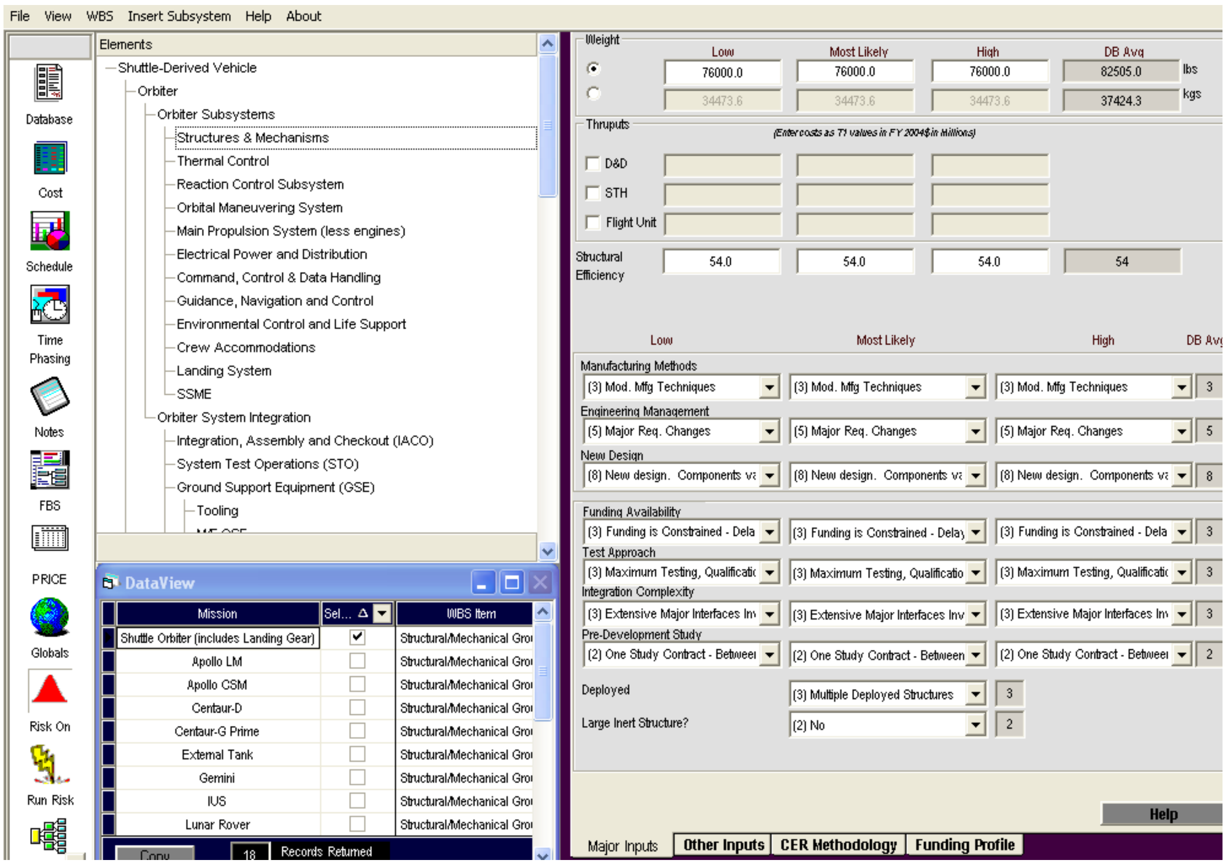

Figure 2. NAFCOM User Interface.

CER Risk

The NAFCOM CERs have the power equation form: $Y = aX_1^{b1}X_2^{b2}...X_n^{bn}$**.** The CERs are calculated using transformed ordinary least squares (OLS). The natural logarithms of the dependent variable and independent variables are calculated, and then OLS is applied to the transformed data. That is, OLS is applied to the logarithmic transformed ("log") model: $ln(Y) = ln(a)+b_1ln(X_1)+b_2ln(X_2)+...+b_nln(X_n)$.

The coefficients calculated by ordinary least squares, $\boldsymbol{b_1, b_2, \ldots, b_n}$, are used in the model $\boldsymbol{Y = aX_1^{b1}X_2^{b2}\ldots X_n^{bn}}$, and the transformation exp(ln(log model a-value)) = a-value is performed to yield the corresponding coefficient for the power equation.

The estimation error for the log model is a normal distribution, with mean equal to zero and standard deviation equal to the standard error of the model. The log model standard error is:

$$SE = \sqrt{\frac{\sum_{i=1}^{n}(\ln(Y_i) - \ln(\hat{Y}_i))^2}{n - k - 2}} ,$$

where n = number of data points and k = number of independent variables.

Note this is *not* the standard error of the CER. However, the CER standard error can be calculated from this standard error. A random variable *X* is said to be lognormally distributed if *ln(X)* is normally distributed. Since the errors of the log model are normally distributed and the log model is simply a logarithmic transformation of the power equation, then the error for the power equation model is lognormally distributed, with mean and standard deviation:

$$\text{Mean} = \mu = e^{P + \frac{1}{2}Q^2}$$

$$\text{Standard Deviation} = \sigma = e^{P + \frac{1}{2}Q^2}\sqrt{e^{Q^2} - 1}$$

where P and Q are the mean and standard deviation of the corresponding normal distribution.

Since OLS was used to perform the regression, it follows that P = 0 and Q = SE. Therefore the estimation error of the CER is lognormally distributed with mean and standard deviation:

$$\text{Mean} = \mu = e^{\frac{1}{2}SE^2}$$

$$\text{Standard Deviation} = \sigma = e^{\frac{1}{2}SE^2}\sqrt{e^{SE^2} - 1} .$$

Combining CER and Technical Risk

As mentioned in the previous section, the error model is assumed to be multiplicative, so in order to combine the technical risk distribution with the CER risk distribution, the two distributions are multiplied. For computational ease, a lognormal distribution is fitted to the first two moments (mean and standard deviation) of the technical risk triangular distribution before the multiplication, allowing for a closed form solution for the resulting distribution, which is lognormal. This process is illustrated in Figure 3.

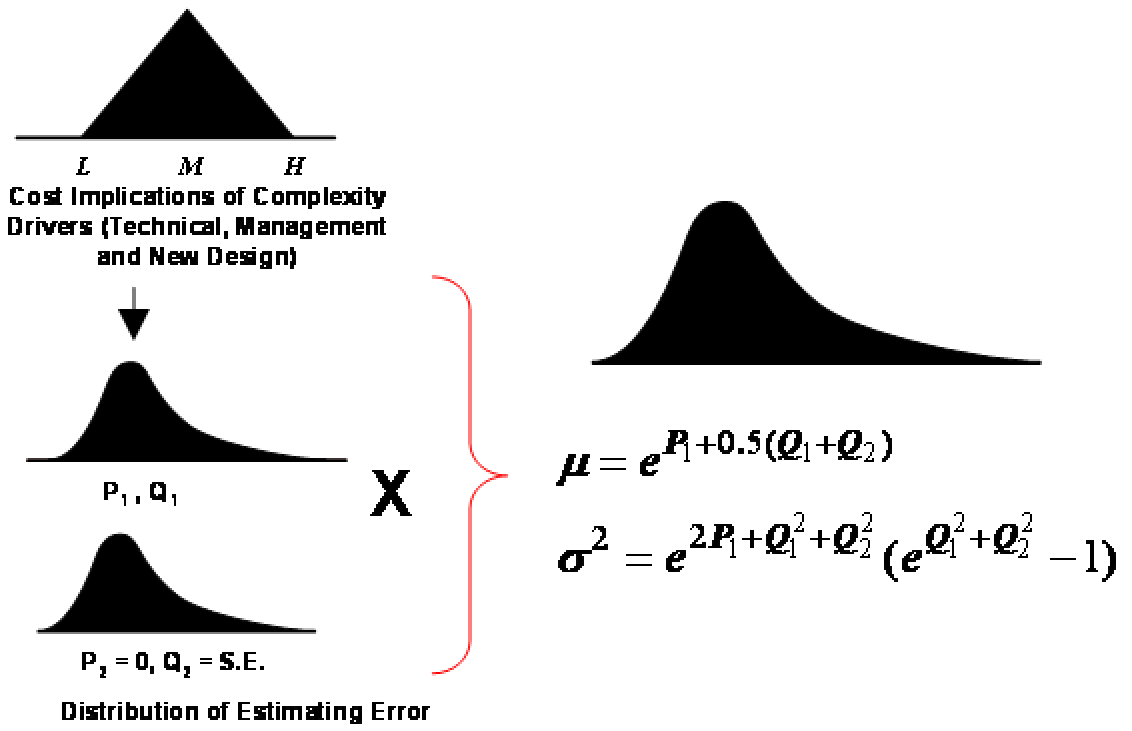

Figure 3. Combining Technical and CER Risk.

Correlation

Correlation between pairs of WBS-element cost distributions is a significant contributor to overall system cost uncertainty (Book, 1999). Default values in NAFCOM are 0.2 for inter-subsystem correlations, 0.5 for intra-subsystem correlations, and 1.0 for systems integration and similar items. The user can review the correlation matrix and set individual correlations to any value between -1 and +1. The only potential issue with allowing unfettered access to the correlation matrix is that it can lead to an inconsistent correlation matrix. The current version of NAFCOM does not check for consistency, but this enhancement will be implemented this year.

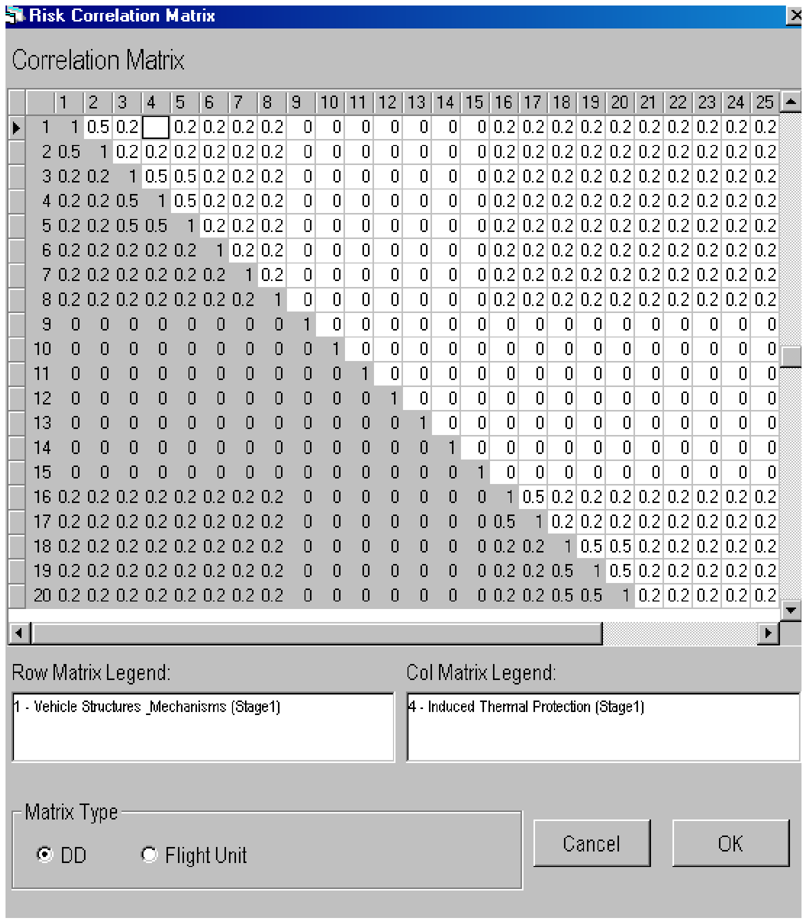

Figure 4. Example of a Default Correlation Matrix.

Aggregating Risk Results

One of the first methodology choices that had to be made was whether to use Monte Carlo simulation or analytic approximation to aggregate the WBS-level risk results. SAIC chose to use analytic approximation because we wanted a method that is computationally simple, calculates the correct top-level means and standard deviations, is faster than Monte Carlo, and allows full access to the correlation matrix. NAFCOM implements the method of moments methodology to aggregate subsystem-level mean and variances. Separately for DDT&E, flight unit, and production costs, NAFCOM sums subsystem-level and variance, taking correlation into account. Specifically, for an n-element WBS in a NAFCOM cost risk estimate, let μ_i and σ_i represent the mean and standard deviation for the lognormal distribution that represents WBS element i, i = 1,…n., and ρ_{ij} represent the correlation between any two WBS element pairs i and j, i = 1,…,n and j = 1,…,n. Then the total cost means and standard deviations are calculated as:

$$\text{Total Cost Mean} = \sum_{k=1}^{n} \mu_k$$

$$\text{Total Cost Standard Deviation} = \sqrt{\sum_{k=1}^{n} \sigma_k^{\ 2} + 2\sum_{k=2}^{n}\sum_{j=1}^{k-1} \rho_{jk}\sigma_j\sigma_k}$$

See Figure 5 for an illustration of this approach.

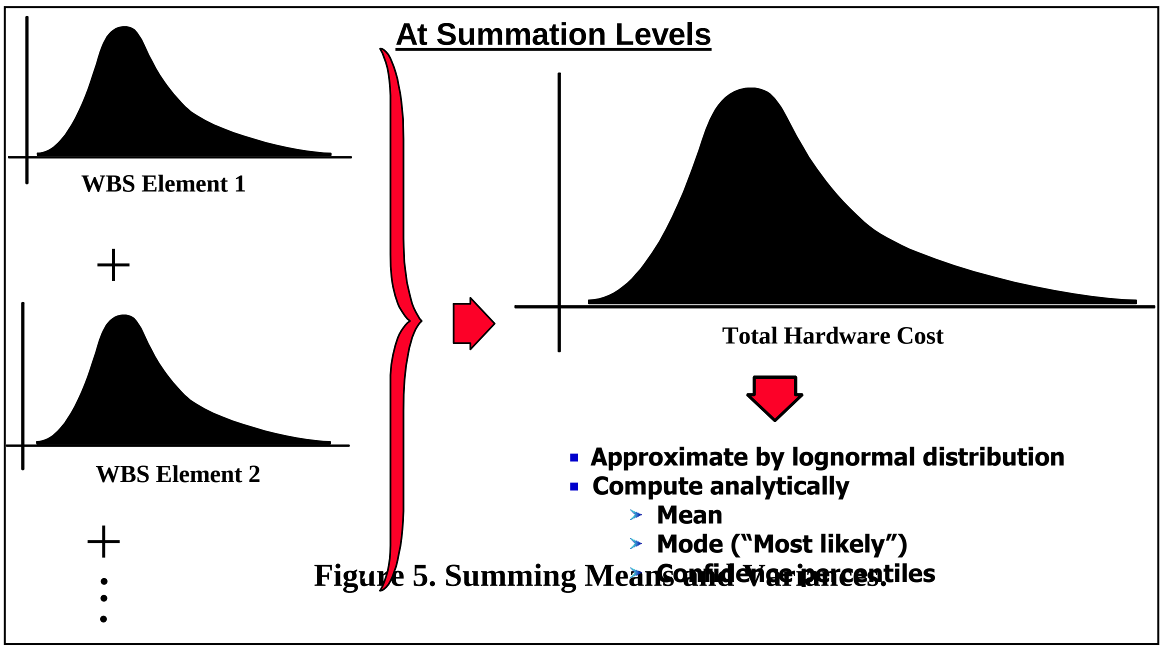

Figure 5. Summing Means and Variances.

NAFCOM offers the analyst the opportunity to approximate the total-cost distribution by both normal and lognormal distributions. According to statistical theory, the normal distribution should provide a better approximation to a statistical sum of triangular distributions than would the lognormal distribution under three circumstances: there is a large number of WBS elements, so that the Central Limit Theorem of statistics applies; the triangular distributions are not very skewed, so that convergence of their sum to the (symmetric) normal distribution does not require very many WBS elements; or there is little or no correlation between WBS elements, so that each WBS element contributes fully to the statistical sum, thereby achieving acceptable convergence with a smaller number of elements. The normal approximation has been recommended by (Simpson and Grant, 2001).

NAFCOM Risk Outputs

The NAFCOM risk module outputs consist of statistics and graphs. For DDT&E, Flight Unit, Production, and Total Cost, NAFCOM prints a table consisting of the mean, median, mode, standard deviation, the 5th and 95th percentiles, and every 10th percentile from the 10th to the 90th percentile. The graphical display consists of the probability density function (PDF) and the cumulative distribution function (CDF) for the lognormal (or normal) distribution for the DDT&E, Flight Unit, Production, and Total Cost. See Figure 6 for an example PDF output graph..

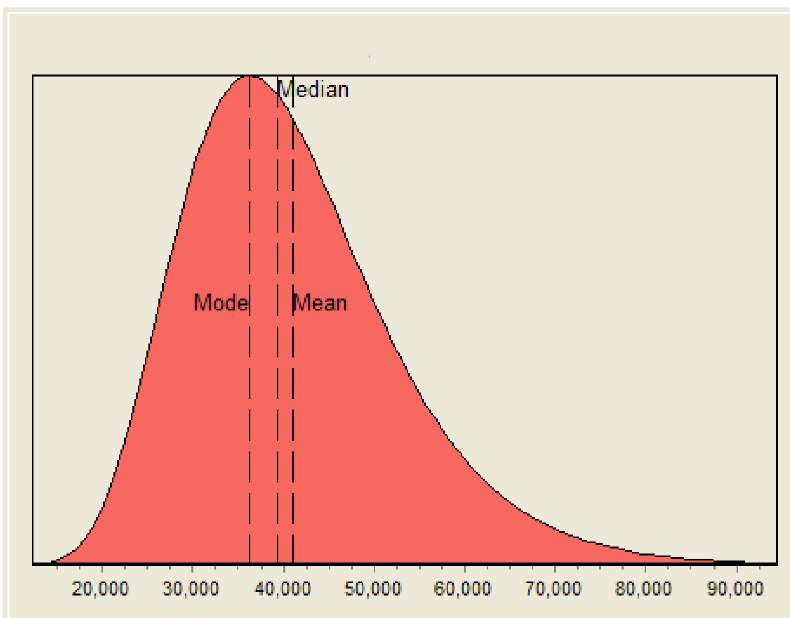

Figure 6. Lognormal PDF for the Total Cost Distribution.

Monte-Carlo and Analytic Approximation Comparisons

There is no consensus in the cost analysis community on whether Monte Carlo simulation or analytic approximation is more accurate. Recent studies by Tecolote (Smith, 2004) and MCR (Alexander et al., 2004) have shown that the two methods provide similar results. In November, 2004, SAIC collaborated with Dr. Steve Book and Erik Burgess of MCR to perform comparison tests between NAFCOM's analytic approximation technique and Monte Carlo simulation. To achieve this, SAIC created a series of Excel spreadsheets that replicate NAFCOM's cost risk engine. These spreadsheets were used as a basis for Monte Carlo simulations using @Risk, a Monte Carlo simulation tool and Excel add-in. Excel and @Risk were used to perform the same risk calculations in NAFCOM using Monte Carlo simulation to combine CER and technical risk to obtain WBS-level risk distributions and to aggregate WBS-level costs to obtain top-level results. For this comparison, all the @Risk simulations used 10,000 trials. One caveat is that @Risk uses rank correlation, which means that the Monte Carlo results will have smaller tails than if product-moment correlation were used (Garvey, 1999).

The results are extremely close. Even for the worst case, the 5th and 95th percentiles are within 10% of one another. Also, the tails for the Monte Carlo simulations are slightly thinner than for analytic approximation, as expected because of the rank correlation issue. Had product-moment correlation been used in the simulations, the results would likely have been even closer. The results of these tests add evidence to other studies cited above that Monte Carlo simulation and analytic approximation provide similar results.

SAIC and MCR analyzed four test cases. The first test case comprises four subsystems: structures; thermal control; electric power, distribution, and control; and command, control, and data handling, with no systems level costs. In all four cases both top-level DDT&E and flight unit costs were analyzed. The results for both are similar and so we present here only the results for the DDT&E comparisons. The results of the simulation and the analytic approximation for case 1 are displayed in Figure 7. In Figure 7, the histogram represents the Monte Carlo simulation results and the green line graph represents the analytic approximation results, fit to a lognormal distribution. Note that in Figure 7, the simulation results are slightly more peaked than the analytic approximation. This is to be expected, since as we have mentioned, use of rank correlation results in thinner tails. Figure 8 displays percentiles comparisons for DDT&E costs.

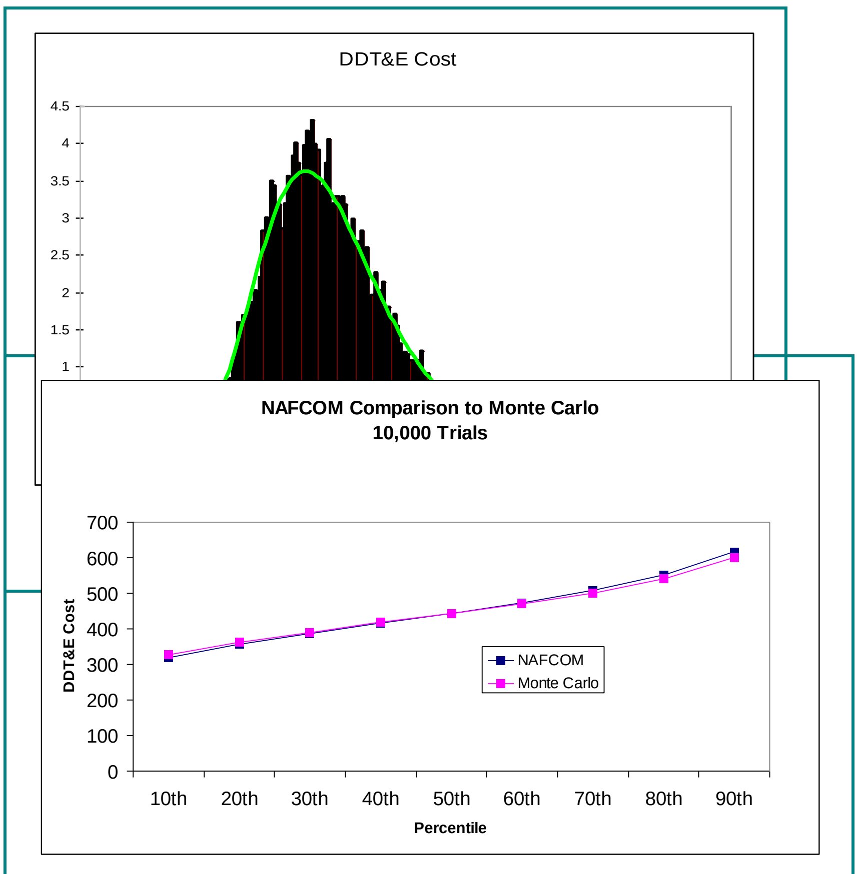

Figure 8. Test Case 1, DDT&E Percentile Comparisons

The second test case includes four subsystems: structures; thermal control; electric power, distribution, and control; and command, control, and data handling, with systems engineering costs included. The results of the two simulations are displayed in Figures 9 and 10.

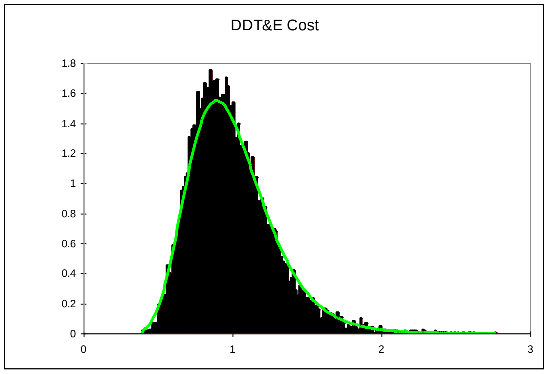

Figure 9. Test Case 2, DDT&E

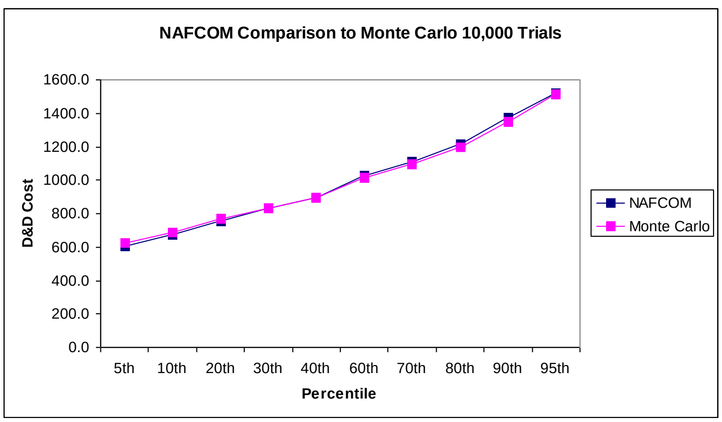

Figure 10. Test Case 2, DDT&E Percentile Comparisons

The third test case includes seven subsystems – structures; thermal control; electric power, distribution, and control; command, control, and data handling; guidance, navigation, and control; reaction control and auxiliary propulsion; and main propulsion (less engines); with systems engineering costs. The results of the two approaches for this case are displayed in Figures 11 and 12.

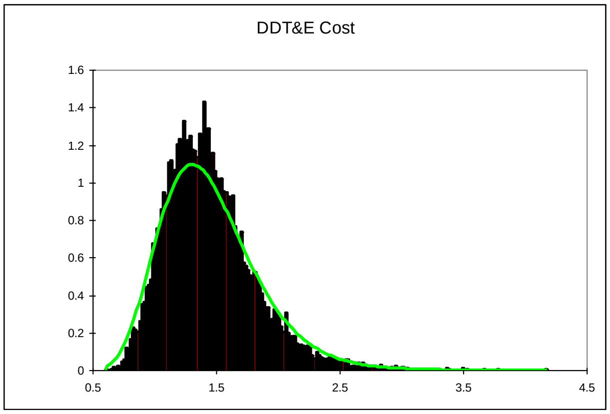

Figure 11. Test Case 3, DDT&E

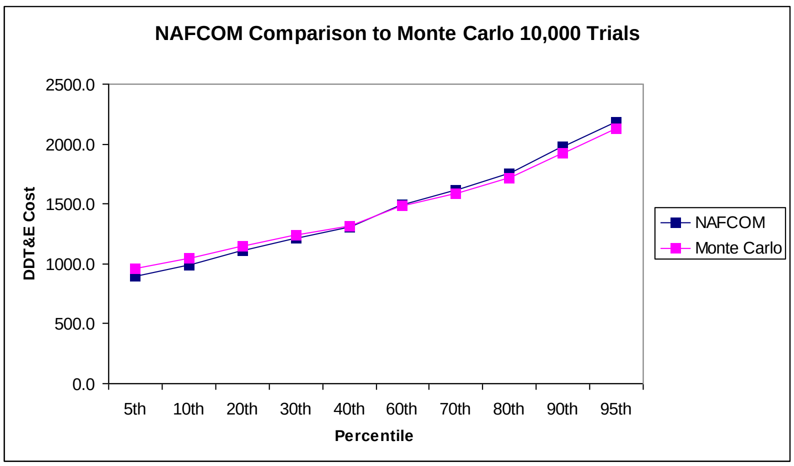

Figure 12. Test Case 3, DDT&E Percentile Comparisons

The fourth test case is the same as the third case, except that the CER inputs vary more widely. That is, the percentage differences between the low inputs and the high inputs vary more in test case 4 than in test case 3. The results from this test case are summarized in Figures 13 and 14. One interesting result of this test case is that perhaps because of the widely skewed inputs, the normal distribution is clearly not a good fit of the simulation results, but the lognormal compares closely.

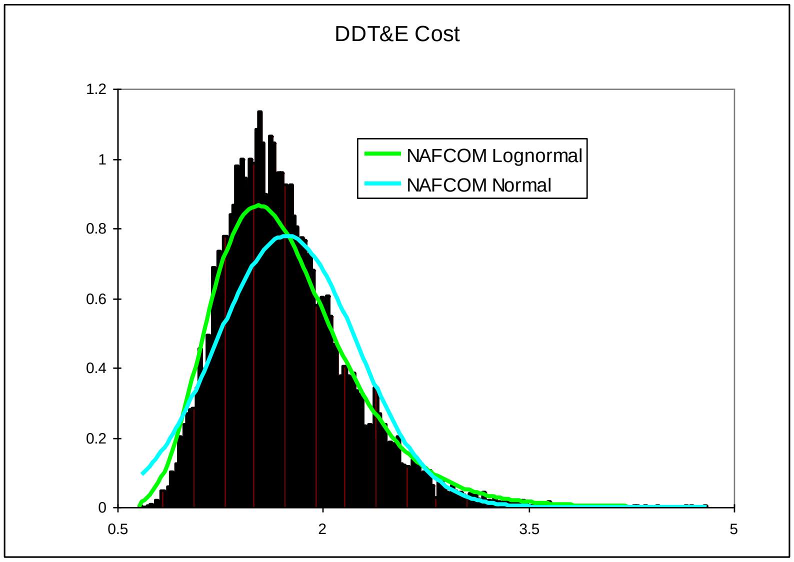

Figure 13. Test Case 4, DDT&E

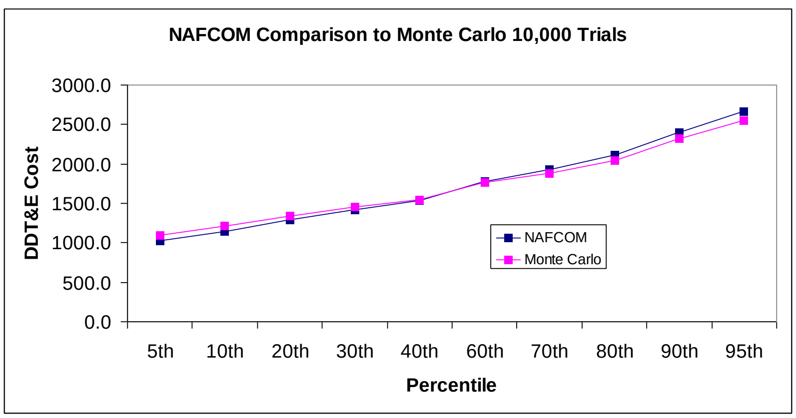

Figure 14. Test Case 4, DDT&E Percentile Comparisons

References

Alexander, A., Book, S.A., Burgess, E. L., and Maltz, I., "Cost Risk Model Assessment Report," Electronics Systems Center, Global Air Traffic Operations/ Mobility Command and Control Financial Management Division Technical Report, 2004.

Book, S.A., "Why Correlation Matters in Cost Estimating," Advanced Training Session, 32nd Annual DOD Cost Analysis Symposium, Williamsburg, VA, 1999.

Garvey, P.R., "Do Not Use Rank Correlation in Cost Risk Analysis," 32nd Annual Department of Defense Cost Analysis Symposium, Williamsburg, VA, 1999.

Garvey, P.R., *Probability Methods for Cost Uncertainty Analysis - A Systems Engineering Perspective*, New York, Marcel Dekker, 2000.

Simpson, W.P., and Grant, K.P. "An Investigation of the Accuracy of Heuristic Methods for Cost Uncertainty Analysis," *The Journal of Cost Analysis & Management,* Winter 2001.

Smith, Alfred, "Cost Risk Analysis for the Masses," *2004 Conference of the Society of Cost Estimating and Analysis,* Manhattan Beach, California, 2004.

Biography

Dr. Christian Smart is employed as a senior cost analyst with the Science Applications International Corporation. He is a SCEA Certified Cost Estimator/Analyst and served as President of the Greater Alabama Chapter of SCEA during the 2004-2005 program year. Dr. Smart's paper "Process-Based Modeling" was awarded best paper in the applications track at the 2004 annual ISPA conference.

Dr. Smart earned bachelor's degrees in economics and mathematics from Jacksonville State University, and a Ph.D. in applied mathematics from the University of Alabama in Huntsville.