Abstract

This paper will discuss a few of the common pitfalls to cost risk estimating, especially those caused by junior estimators lack of familiarity with popular risk tools. None of these tips break new ground, apply fresh theories, or even offer advice that is outside of the "help" section of each tool's operating manual. But nonetheless, these tips will help a new estimator achieve a new level of understanding and insight which will ultimately lead to higher quality and more reliable estimates. The intent is to educate users to understand their project needs and establish a Commercial Off The Shelf (COTS) risk analysis setup that will fulfill requirements and by avoiding pitfalls and following best-practices. Some COTS cost risk tools allows users to sidestep complicated, yet important, steps in order to be user-friendly. Point-and-click features of some models occasionally allow users to blindly make mistakes which reduces an estimate's quality. Ultimately the paper will review 1) what we are doing wrong now by following default settings, 2) what needs to be changed, and 3) what results to expect when cost risk models are implemented correctly.

COTS Risk Models: Basic Advice for Avoiding Pitfalls

Cost estimating is humorously referred to as more of an art than a science, and most estimators agree that risk analysis is the darkest art of them all. Almost always, cost estimators will research figures for months, wrestling with methodologies and carefully documenting results in order to arrive at the most accurate point estimate that they can derive. And then at the last moment, almost as an afterthought (and only if necessary), the estimator will apply a layer of cost risk. For the majority of the cost estimating profession, applying risk is achieved by simply calculating a plus-or-minus 20% boundary to present along with the final cost results, incorporating cost factors in engineering change order (ECO) line items, or worse--hiding management reserve in the estimate.

However, within the past several years, computing horsepower has combined with new cost risk analysis software to create special tools that employ Monte Carlo or Latin Hypercube simulations to help the estimator independently model risk in a very sophisticated fashion. Naturally, there is both a good and a bad side to these new capabilities. Applying these require an amount of insight and training that many estimators do not have. Tools such as Crystal Ball and ACEIT RI$K are very common in an estimator's environment, and with every new release they become easier and more enticing to use. But with ease of use often lies hidden pitfalls. In many cases, the fault rests with the inherent user-friendliness of the models. In an attempt not to be confusing and intimidating, cost risk models streamline processes by rely on graphical user interfaces (GUIs) that allow the user to perform point-and-click risk analysis. Generally, this low "barrier to entry" is beneficial to the profession because it allows estimators to perform sophisticated analysis that lends value to decision makers. However, point-and-click risk analysis using educated guesses is fraught with pitfalls. Choosing distribution shapes, correlating models correctly, including the cost of mitigation, and overcoming the hurdles of default settings are just a few of the issues that require a quick education.

This paper will discuss a few of the common pitfalls to cost risk estimating. None of these tips break new ground, apply fresh theories, or even offer advice that is outside of the "help" section of each tool's operating manual. But nonetheless, these tips will help a new estimator achieve a new level of understanding and insight into performing the art of cost risk estimating which will ultimately lead to higher quality and more reliable estimates.

First Things First – Selecting Distributions

Once an estimator has identified an item that is a good candidate for risk analysis, the first major hurdle is choosing a probability distribution function (PDF) that accurately models the effect of risk on that cost element. The purpose of the PDF is to describe the cost of an element as a range of possible costs rather than a single point estimate. This can be a more accurate way to describe costs since the estimator is predicting the future and almost anything can happen despite how well the cost element is defined. But choosing the right PDF to model costs can be an extremely daunting task, especially when using cost risk software that has so many choices. Figure 1 below shows a sample of the PDFs available to an estimator in the commercial cost risk tool Crystal Ball.

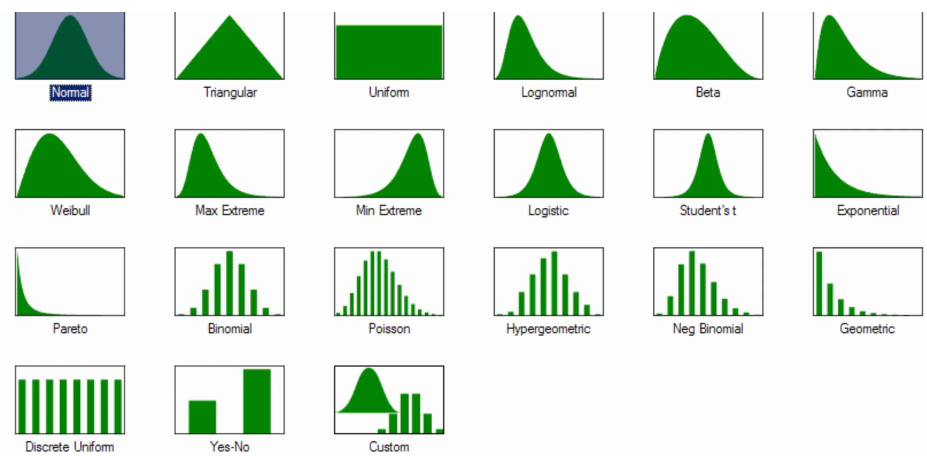

Figure 1: PDFs available in Crystal Ball

Although it appears that there are only 21 choices in Figure 1, the last choice is actually an option that allows the user to create custom distributions—meaning that this list is almost infinite. For a new user, the choices can be overwhelming, and estimators cannot always rely on technical experts tot help pick an option. Each of these PDFs exist because each one has proven to be the most accurate in a particular application, but basic advice does exist for helping the estimator to narrow the choices.

Lognormal distributions should be the default choice as a PDF in the majority of situations where the estimator does not have quantitative reasons to chose another type.[1] The benefits to an inexperienced estimator are numerous. The lognormal distribution is mathematically bounded by zero so there is no chance risk ranges moving below zero and

---

[1] Air Force. Air Force Cost Analysis Agency. Air Force Cost Analysis Agency Cost Risk Handbook DRAFT. Hanscom AFB: n.p., Oct. 2006.

encompassing negative dollars. Also, the nearly asymptotic shape allows for a small chance of extremely high costs, which is always a possibility in a Department of Defense (DoD) program. And most important to the inexperienced practitioner, lognormal distributions require only two data points to define; just the point estimate and the standard deviation or standard error are necessary. In situations where the standard deviation is not available, generic guidelines exist.[2] Despite the benefits of a lognormal distribution, tee triangular distribution remains a popular choice. This PDF requires the estimator to define three points that are fairly straightforward—a low point, the most likely, and the highest point likely. Engineers usually feel comfortable helping an estimator derive these points, and estimators often have enough programmatic insight in order to make accurate educated guesses. Beyond these short paragraphs of advice, it is the estimator's duty to research these PDFs and find a methodology that successfully helps estimate risks for a particular project.

Another pitfall is that most risk tools do not require the estimator to apply risk ranges to any particular cost element within an estimate. SEER, ACEIT and Excel add-ins will compute calculations whether there are risk ranges present or not, so the estimator cannot rely on risk tools to prompt the user to enter risk ranges. Due to the inherent fact that the future is impossible to predict perfectly, the estimator should not be shy about applying PDFs to most of the input variables to account for both cost uncertainty and technical risks that may have cost impacts. It is important to remember that most DoD programs have significant amounts of uncertainty and technical risks, and almost always have suffered from cost overruns.[3] Historically there has not been a problem caused by overly-conservative cost estimates; applying uncertainty to all input variables, as a minimum starting point, will begin to account for the unknown unknowns that lead to cost growth.

Correlation and Default Settings

An important aspect of cost risk analysis that is often overlooked is correlation, and sometimes this is caused by default settings within the risk tool. This is another example of a pitfall where an estimator will perfect his or her model to produce an accurate point estimate, only to ruin the results by applying inaccurate correlation (or worse, no correlation) within the cost risk software. Correlation is a statistical technique used to determine the degree to which variables are related or associated, and indicates the strength and direction of a linear relationship between random variables.[4] The Air Force Cost Risk Handbook has excellent advice on correlation, and begins by describing it as such:

> *An important consideration in risk analysis is to adequately account for the relationship between the cost elements during a risk simulation. This interrelationship between the WBS elements is commonly known as*

---

[2] Air Force, ibid

[3] Chen, Zhihao. "Reduced Parameter Modeling for Cost Estimation Models." Diss. University of Southern California, 2006.

[4] Correlation. 12 Apr. 2007. Wikipedia. 16 June 2007 <http://en.wikipedia.org/wiki/Correlation>.

> *"dependency" or "correlation." For example, if something causes the cost of WBS element A to increases, it may be appropriate that the cost of WBS element B also increases (positive correlation), and perhaps WBS element F should decrease (negative correlation).*[5]

Every element that has an associated level of risk or uncertainty is likely to be correlated to another risky or uncertain element. In many cases, the cost model may already have correlation built in—for example, a weight-based cost estimating relationship (CER). In this example, a CER is entirely dependent on it's input variables, and as they increase or decrease the final results will adjust accordingly. In other words, a weight based CER is very dependent on-- and therefore highly positively correlated to-- weight. As the weight increases, the cost output of that CER will increase at a geometric rate. In other cases, cost estimates rely on different types of estimating techniques other that CERs, so it is important for the estimator add in additional levels of correlation across the estimate.

Correlation is an important component of cost risk modeling because it has a noticeable impact on final results. Due to the fact that cost growth in one area tends to lead to cost growth in other areas, estimators will usually identify and model positive correlation rather than negative correlation. Applying positive correlation increases the range of outputs because positive correlation has a multiplier effect, which will be discussed later in this paper. Adding this type of correlation to the model will reduce the slope of the resulting cumulative curve, which is a common graph for showing cost risk findings to leadership. Although negative correlation may be less common, but if applied correctly it is equally important at increasing the accuracy of the final product.

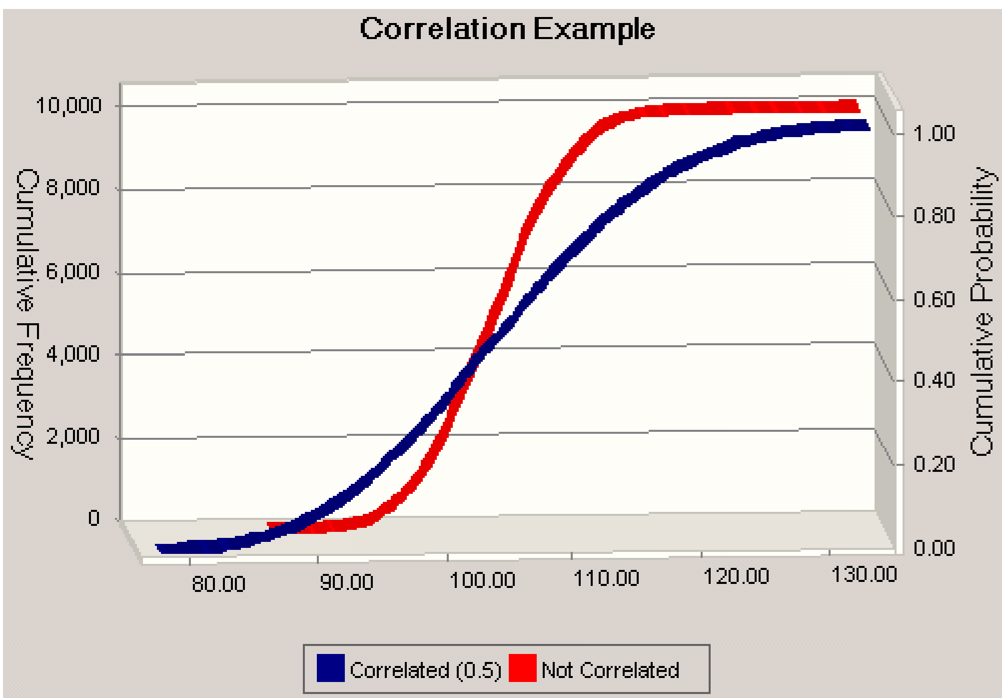

Figure 2: Effect of Correlation

---

[5] Air Force, ibid

Figure 2 above, in addition to prominent papers in the field, have demonstrated the importance of adding correlation to a risk model because of the effect it has on the cumulative curve.[6] Figure 2 is a cumulative curve that shows how added correlation will cause the curve to “flatten,” thus increasing the amount of money a program must expend in order to decrease the likelihood of a cost overrun. The coefficient of variance will increase as well when correlation is added.[7] The pitfall lies in the tools that allow the user avoid this important step in the process. As an example, some parametric estimating tools such as SEER-H and SEER-SEM employ a risk process that will not allow the user to define or manipulate correlation. Others, such as Crystal Ball, have the capability to incorporate correlation into the model, but never require the user to do so. As a result, this estimator’s experience has shown that correlation is usually omitted from simulations run in Crystal Ball. Even though ACEIT uses slightly different methods and terminology, ACEIT RI$K does not require correlation to be defined either. RI$K uses the term “group strength” rather than the analogous term correlation, and defines group strength as “a heuristic or qualitative approach to handling correlation between WBS items.” Under this system, the user defines a dominant cost item (e.g. the prime mission product) and then identifies other cost items that have either a positive or negative association with the dominant row. In effect, this is very analogous to the correlation methodologies incorporated in Crystal Ball but in the case of ACEIT, each line item can only be grouped with one other line item (one to one), whereas in Crystal Ball the user can correlate one line item to every other line item if desired (one to many).

The unfortunate truth of incorporating correlation is that it is often difficult and time-consuming -- especially in models that have scores of line items with associated risk or uncertainty. In models this large, adding correlation often requires a level of insight that estimators do not have, and even in small estimates the process can be frustrating. Because of this complexity, and also due to the model developer’s goals of making tools easier for point-and-click users to operate, this step is often neglected and omitted. There are several methods of determining the proper correlation. The purpose of this paper is not to give the pros and cons of any particular method, but every estimator is encouraged to research these methodologies. In the meantime, in the majority of situations wherein the estimator does not have particular insight into correlation levels, several generic correlation guidelines exist. Figure 3 below shows several of the guidelines as a reference.

---

[6] Coleman, R. L, S. S. Gupta, R. J. Ayers, G. E. Hartigan. “Cost Risk Estimates Incorporating Functional Correlation, Acquisition Phase Relationships, and Realized Risk.” Society of Cost Estimating and Analysis, SCEA National Conference. 23 June 1997

[7] Covert, Ray. “Correlations in Cost Risk Analysis.” Society of Cost Estimating and Analysis, SCEA National Conference. June 2006.

Strength	Positive	Negative
None	0.00	0.00
Weak	0.25	-0.25
Medium	0.50	-0.50
Strong	0.90	-0.90
Perfect	1.00	-1.00

AFCAA CRH Guidance

# of Elements	Correlation
10	0.333
15	0.214
20	0.158
30	0.103
50	0.061
100	0.03
150	0.020
200	0.015
300	0.010
500	0.006

Covert, MCR

# of Elements	Correlation
5	0.50
10	0.25
20+	0.10

Multiplicative Error Regression Technique

Figure 3: Generic correlation guidelines[8,9]

Although many scholars have argued the validity of one method compared to another, the important takeaway for an inexperienced cost risk estimator is to work within the framework of the tool in question to employ correlation, and never take the easier route of omitting this crucial step.

Iterations and Default Settings

During the initial phases of an Air Force-led Analysis of Alternatives (AoA) the cost team found inconsistencies in the results after performing risk analysis in ACEIT RI$K. The initial phases of the AoA called for the cost team to create independent cost estimates of fourteen technical solutions that had been provided from industry and capture the costs in ACEIT. The solutions were all unique—some proposed solutions included missiles, planes, or submarines. In short order, the team had to assess each of these solutions by providing an independent cost estimate, and assessing the cost impacts of different technical risks was one of the most important aspects of this segment of the study. The team worked with technical experts to determine the Technology Readiness Levels (TRL) of each technology, and then obtained TRL-based cost boundaries from the Air Force Cost Analysis Agency (AFCAA). Once the team had an independent cost estimate, TRL levels of most technologies, and TRL-based risk boundaries, the final step was to apply correlation.

On the first run, the cost team applied a group strength value between 0.3 and 0.5 to every line item and identified the prime mission product (i.e. the key piece of technology such as the plane, missile, or submarine) as the dominant item. After running through several cycles of risk calculations, the results were surprising -- the team found that technologies with highly correlated line item costs did not always have wider distributions, but re-running the simulation sometimes corrected this problem.

Going into the project, the team had operated under the assumption that cost models with "high" correlation would have a wider cost distribution than models with a "low" correlation, all other factors held constant. This effect is caused when one random sample point is above average during an iteration, all other positively correlated sample points are more likely to be above average. As a result, the team expected to see higher

---

[8] Air Force, ibid
[9] Covert, ibid

high and lower low boundaries result in the final rollup distribution. Figure 4 is a notional graphic to illustrate this point.

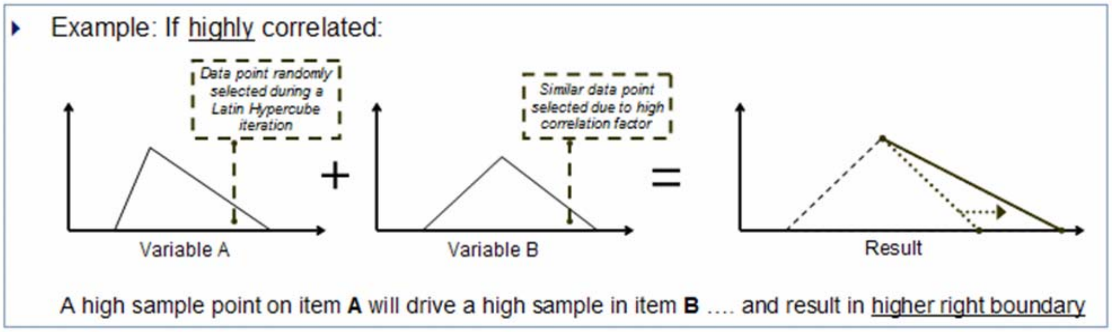

Figure 4: Effect of high correlation on cost outputs

However, Results from ACEIT RI$K runs showed that a model with high correlation did not always have a wider distribution. Costs seemed to go up and down randomly when typical correlation ranges (0.1 – 0.4) were applied. The cost team decided to take an academic approach and run risk simulations on the estimate eleven times—once with zero correlation, again with 0.1 correlation, and again at 0.2, 0.3, 0.4, all the way to 1.0, or full positive correlation. The results of each simulation at the 50% confidence interval are shown in Figure 5.

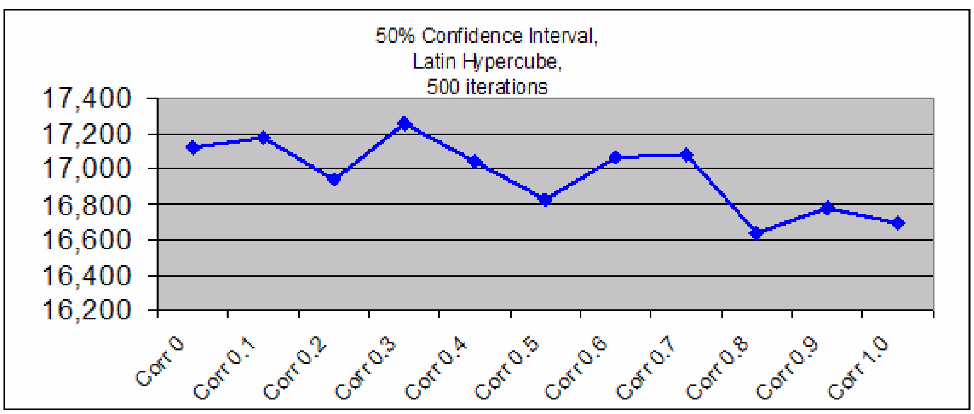

Figure 5: Correlation cost excursion results

Figure 5 shows a level of unpredictable variability that illustrates the cost team's original concerns. Namely, assigning a particular correlation to cost line items will not lead to a particular result. In back-to-back Latin Hypercube simulations using the exact same data set, very different results appeared on each run. (It is important to note that the cost team was not trying to engineer a particular number as a result—rather the team was just trying to explain to leadership why multiple simulations and changes in correlation yielded such small, yet capricious, effects on the final result).

As seen in Figure 6 below, costs across different correlations did not follow a linear pattern and as a result the 50% confidence interval costs were always different. In Figure 3, the lowest observation at the 50% confidence interval was $16,787M, and the highest was $17,086M. That equates to a $299M dollar variance! Ultimately, the estimator cannot compare multiple alternatives with this level of uncertainty present.

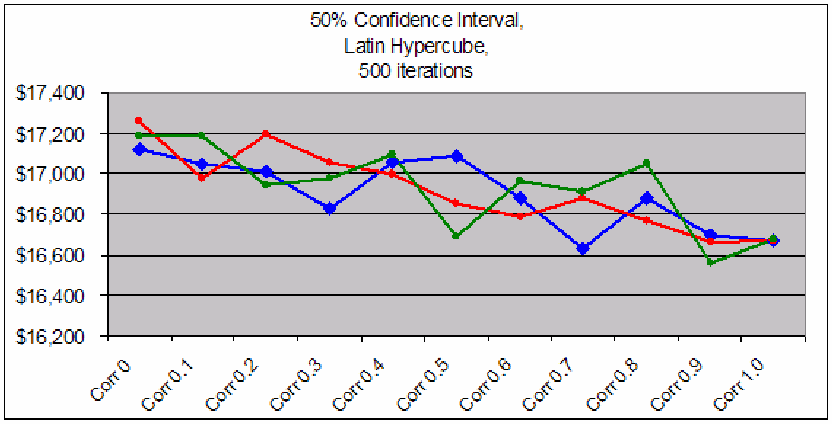

Figure 6: Output of multiple successive 500 iterations of risk simulations

To find a solution, the cost team turned to Crystal Ball, which operates in a more transparent manor and was more familiar to the team. It only took a handful of simulations before the answer became clear: The non-linear change was caused by the default ACEIT risk settings of only 500 iterations. As seen in Figure 7 below, 500 iterations is not enough to produce a smooth distribution. Optimally, running 250,000 simulations produces a very smooth distribution. Please note that ACEIT is not alone in this default pitfall-- the Crystal Ball default is only 1000 iterations, which yields similarly low levels of fidelity.

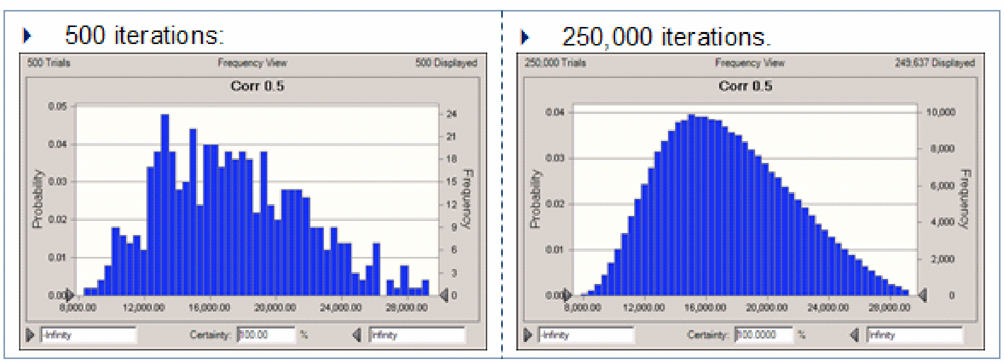

Figure 7: Iteration comparison

Because this 500 iteration distribution on the left is so scattered, one can assume another simulation run under the same scenario might yield a <u>different</u> 50% value. However, despite the fact that the 250,000 iteration chart on the right chose samples using the same Latin Hypercube methodology, the fact that there were so many more iterations yields a smooth curve that very closely resembles the input assumptions that were used to create the distribution. Due to the evenness of the distribution, one can assume another simulation will yield basically the <u>same</u> 50% value. Therefore, an analyst that is comparing two close alternatives can feel confident that using more iterations increases the likelihood that the correct alternative will be chosen every time.

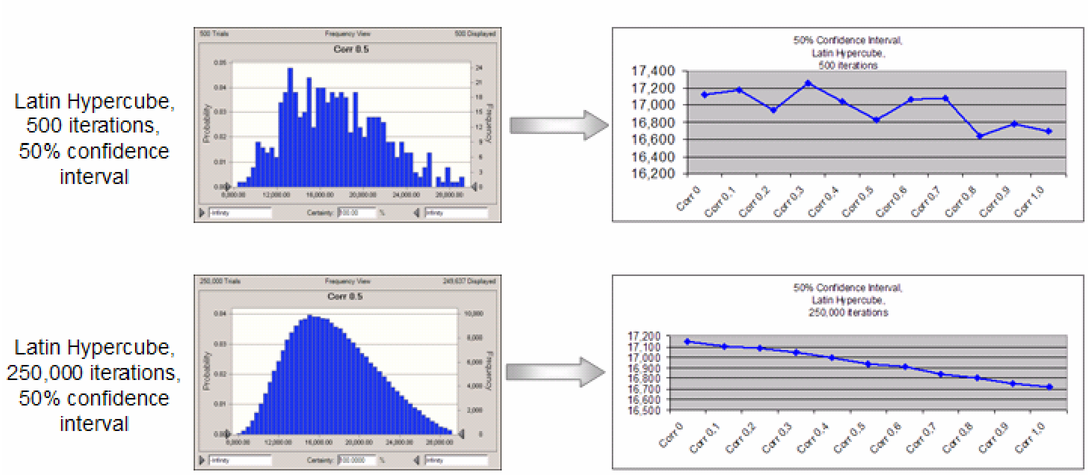

Figure 8: Effects of multiple simulations on correlation.

In order to tie this back to the original problem of correlation not having the desired effect on overall cost, Figure 8 illustrates that running too few simulations was indeed the cause of the non-linear results, and that running more simulations will result in linear outputs across a range of correlation. But Figure 8 only shows the 50% confidence interval, and indicates that higher correlation causes costs to decrease when measured at the 50% confidence interval. However, across the spectrum from 0% confidence to 100% confidence, cost boundaries widen as higher levels of correlation are added. Figure 9 below shows a unique view of the effects of correlation on cost ranges.

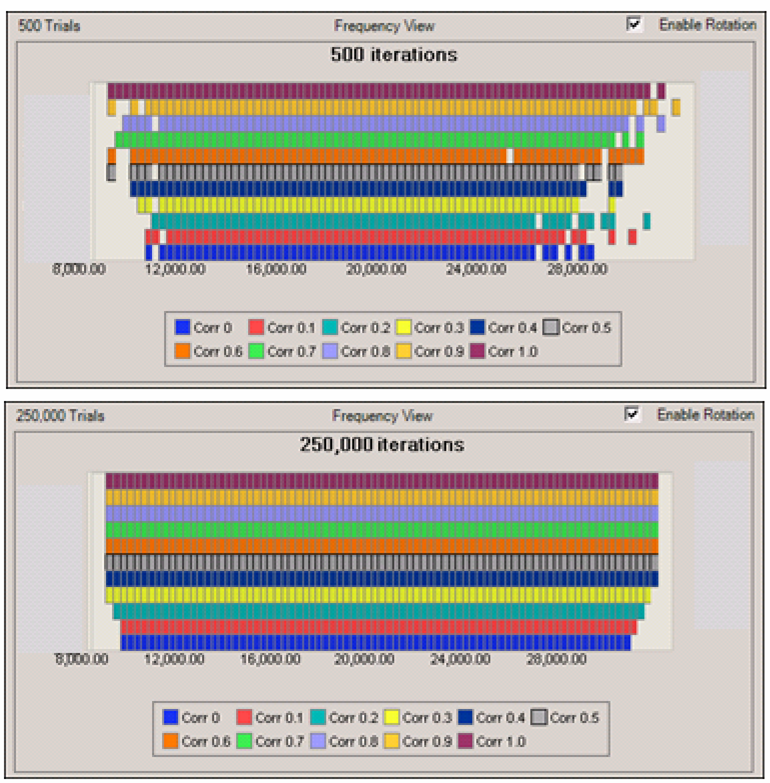

Figure 9: Effect of correlation on upper and lower cost ranges

In Figure 9 above, a sample cost model was set up and run with either 500 or 250,000 iterations using the Latin Hypercube sampling methodology. Each model was run ten times, once for 1.0 correlation between cost elements, once for 0.9, once for 0.8, once for 0.7, all the way down to 0.0 correlation (Figure 9 does not show the effects of negative correlation). The charts are meant to validate two concepts—that highly correlated models have wider cost distributions, and that running a model with too few iterations is apt to produce inconsistent results. Demonstrating the first point, it is clear that highly correlated results near the top of each chart do have wider distributions than uncorrelated models. Setting the correlation in this example to the 1.0 between elements, which is the highest level of positive correlation, creates the widest cost distribution (from approximately $9M to $32M given the sample's risk assumptions), while a correlation of 0.0, which indicates there is no correlation between cost risk items, creates the most narrow distribution (approximately $11M to $29M given the sample's risk assumptions). Next, demonstrating the second point that models with less correlation have inconsistent results, only running 500 iterations (Figure 9, top) can create results that lack uniformity. This lack of uniformity is demonstrated by disconnects in the sampling, as illustrated when a histogram block appears to the right or left of the main body of the distribution. When more iterations are simulated, even when using Latin Hypercube methodologies, the increase creates uniform distributions with repeatable results. As shown in the bottom half of Figure 9, there are no gaps in the results. Ultimately, in the AoA, cost estimators saw that 500 iterations were not enough to create dependable results for decision making. Something needed to change.

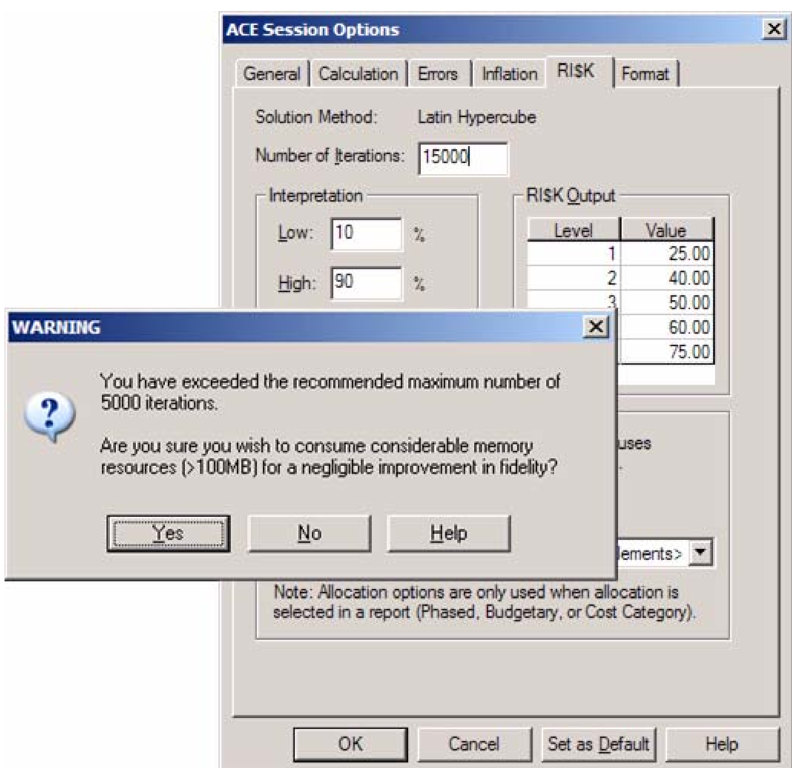

Figure 10: Changing ACEIT defaults

Leaving the default settings at 500 iterations is an acceptable practice while the cost model is still under construction, because fewer iterations take less time to compute. But when entering the final stages of the cost modeling process when results may be used for decision making, the estimator should change the default settings. Within ACEIT, Crystal Ball, or similar risk tools, increase the number of iterations to at least 5000 (the max recommended by ACEIT), and preferably even more since most modern computers can handle tens of thousands of iterations with aplomb. Many of the examples in this paper were created using 250,000 iterations, which resulted in almost perfect uniformity of results in only a few minutes of calculating. Despite the ACEIT warnings, the improvement in fidelity may be more than "negligible!"

A Final Pitfall – Cost of Mitigation

Another pitfall that snares many estimators involves the costs of mitigation of risks. Mitigation is the set of actions a project undertakes in order to prevent a risk from occurring, and mitigation costs are the costs related to these actions. Mitigation actions are inherent to every risk management process, and are arguably the most important aspect of the process because the goal of every risk management process is to prevent risks from occurring, and the mitigation plans are the actions that will accomplish this goal. Despite the importance of these plans, the estimator is usually focused on estimating all of the items described in a requirements document or program description was created prior to the identification of risks. Risk tools, which are not much more than glorified calculators, do not have the intuition to prompt the user to include these costs either. The effect is that the cost of mitigations plans is omitted from the estimate.

Every risk process identifies mitigation actions, but few processes incorporate a feedback loop to require the estimator to include the costs of mitigation in the original estimate.[10] For the purposes of this paper, it is not particularly important to note which risk management processes are being employed because so many different risk management processes exist. In fact, risk analysis processes are like snowflakes—no two are ever the same. Even in organizations where standardized processes exist, risk managers and cost estimators will tailor processes to meet the needs of the specific project. And yet in many cases there is no feedback loop for the estimator to capture the cost impacts of these actions, and very few popular cost risk models remind the estimator to include these costs. This is a significant oversight, because while the estimator is likely to be directed to reduce the amount of risk in an estimate after mitigation plans are developed, the estimator is rarely told to increase the point estimate due to these mitigations plans.

In order to avoid this pitfall, the advice is somewhat elementary. The best way to be a great cost risk estimator is to become a functional member of the project instead of just a manipulator of software. This is perhaps the most important advice for avoiding any of the pitfalls mentioned here because so many problems can be solved through awareness and teamwork. Unfortunately, cost estimators are secluded individuals who sit in back rooms and basements, and do not interface with the scheduling team, risk team,

[10] Clark, R. K. "True Risk Cost: Including Mitigation in Upfront Cost Risk Analysis." Society of Cost Estimating and Analysis, SCEA National Conference. June 2006.

or systems engineering team on a daily basis. Many major programs do not include cost estimator in risk assessment meetings, even though the technical risks and mitigation plans that the team identifies all have cost and schedule impacts. There is no magical formula, at least not yet, in cost risk tools that can estimate the costs of mitigation plans; therefore the only solution is for the estimator to attend meetings and be observant. It is difficult to understate the importance of an estimator's reliance on his or her team instead of a piece of software. Popular cost risk tools take a lot of time to set up correctly -- before pointing and clicking oneself into a dead end, make sure that the cost team is on the same page as the risk team, and to a lesser degree the scheduling team. As great as the tools are, sometimes you have to do the work yourself!

Avoiding Cost Risk Pitfalls

Cost risk estimating is a minefield, and this paper only touched on a few of the common pitfalls to cost risk estimating. None of these tips were meant to offer groundbreaking advice, but these tips will help a new estimator achieve a new level of understanding and insight into performing the art of cost risk estimating. Hopefully, this will ultimately lead to higher quality and more reliable estimates.

Estimators who are ambitious enough to employ new cost software should be cognizant of the particulars of different tools in order not to unwittingly sabotage estimates that have been researched for months. Applying these software tools require an amount of insight and training that many estimators do not have, and in many cases, the fault lies with the inherent user-friendliness of the models. The software is beneficial to the profession because it allows estimators to perform sophisticated analysis, but conversely the point-and-click risk analysis using educated guesses is fraught with pitfalls. Remember the preceding advice when choosing distribution shapes, correlating models correctly, overcoming the hurdles of default settings, and including the cost of mitigation. And always keep in mind that the estimator needs to be cognizant of the mechanics of the project, and sometimes the best idea is to push the tool aside. These are just a few of the issues that require a quick education. Doing so will reduce the perception that cost risk analysis is a dark art, and increase the probability of defensible results.

Speaker Biography

John Teal is an associate with Booz Allen Hamilton in Colorado Springs and works in the areas of cost analysis, life cycle cost estimating, business case analysis, cost risk analysis, and economic analysis. His four years of cost and finance experience have been on space and aircraft programs for the Air Force, intelligence agencies, DoD joint programs, and NASA. He is currently pursuing an MBA at Regis University and has a Bachelor of Science degree from the United States Military Academy at West Point. Mr. Teal is a Certified Cost Estimator/Analyst (CCE/A) and a Project Management Professional (PMP). He is a member of the Pikes Peak SCEA chapter where he serves as the Program Chair. He can be reached at teal_john@bah.com or 719-387-3757.