Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

Taking a Second Look:
The Potential Pitfalls of Popular Risk Methodologies

SCEA National Conference
New Orleans, LA, June 2007

The Society of Cost Estimating and Analysis

Eric R. Druker
Richard L. Coleman
Christopher J. Leonetti
Peter J. Braxton
Northrop Grumman

Background

Most risk practitioners perform analysis using well-documented methodologies whose origins lie in mathematics. In many situations however, these practitioners come from non-mathematics backgrounds. This can lead to methodologies that may have sound basis being applied incorrectly, albeit innocently, due to a lack of understanding of their underpinnings. The purpose of this paper is to shed light on some of the common mistakes in the execution of risk analysis. It will also try to explain the mathematics behind these mistakes, and the mischief they can cause. As a final note, this paper is not intended to be, nor could it ever be, all-inclusive. It will discuss what seems to be the right mix of common and serious errors in the experience of the writers.

The focus of this paper will be on two areas, risk identification and quantification, and risk modeling. The former will begin with a discussion of continuous vs. discrete risks, below-the-line risks, and combining triangular risks, and will conclude with an important dialogue on how the word "confidence" is meant in terms of risk analysis contrasted with how it is often understood by practitioners and engineers. The risk modeling portion of the paper will focus on potential problems with the two most common methods of consolidating risks, Monte Carlo and Method of Moments. It will then move on to a conversation on the handling of mutually exclusive events, and the dangers of truncating negative values in risk simulations. Lastly, this section will finish with a recommendation on how to present categories of risks to decision makers in a format that is both useful and easy to understand. The paper will conclude with a thought experiment related to the assumption of an underlying log-normal distribution in risk.

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

Discrete vs. Continuous Risks

Although many risk methodologies account for both discrete and continuous risks, some analysts try to squeeze all of their risks into just one of the two categories. There are a couple of reasons why this is desirable. First, it is easier to model risks from the same family of distributions. Second, and probably most important, it is easier to present risks to management when they all come from the same family. Unfortunately, rare is the case that risks can be properly categorized using one family of distributions. Furthermore, improper categorizations cause distortions in risks, usually in their variance, less often in their mean. Using one family of distributions can lead to misguided management decisions brought on by a poor characterization of risk.

Discrete distributions are meant to account for specific events with point estimates for their cost impacts. Examples of these are technical or schedule risks that are due to a specific event. These types of risks are generally characterized as Bernoulli or multi-valued discrete events, described by probabilities and cost impacts. Characterizing a discrete event risk as continuous causes several problems: First, it gives management the impression that they cannot avoid the risk, and second, it can show an opportunity where one does not exist.

Continuous risks, on the other hand, are meant to account for events that will always occur, and have a range of cost impacts. Examples of risks that tend to be continuous include below-the-line risks and estimating risks from estimates built using

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

factors or regression analysis. These risks can be characterized by any number of distributions, some of the most common being: triangular, normal, and log-normal. Characterizing continuous risks as discrete events also causes several problems. It gives management the false idea that they can totally eliminate the risk, and it also tends to leave out information that can show the opportunity side of the risk (in the event that an opportunity exists).

Risk Quantification for Factors and Rates

Another danger area of risk analysis is in the identification and quantification of below-the-line risks such as factors and rates. Generally, one of two errors occurs: either the rate or factor risk is applied to the non-risk-adjusted estimate or a discrete distribution is used to categorize this generally continuous risk. In the first case, only the risk around the rate or factor as it relates to the point estimate will be found. As the point estimate is generally also affected by risk, this will understate the below-the-line risk. Using a discrete distribution to categorize a below-the-line risk can cause another problem by giving the impression that the risk is a specific event. Generally these risks are in relation to the historical variation of the factor or rate that is being used and are not specific events, but continuous distributions. To perform the analysis correctly, the distribution around the rate or factor must be found and applied to the risk adjusted distribution of the cost that it is to be applied to.

Sometimes, when developing a risk distribution for a portion of an estimate, analysts will collect information on distributions at a lower level, and then roll them up to obtain the risk distribution for the level at which they are performing their analysis. A

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

common mistake often happens when analysts try to do this with triangular distributions. Some analysts simply take the lower level triangular distribution parameters (minimum, mode, maximum) and add them together to get a top level distribution. The problem in this is that unlike means, percentiles do not add across distributions. Doing this and assuming a top level triangular distribution incorrectly adds weight to the tails of the top level distribution. The lower level distributions must be run through a simulation to obtain the upper level distribution.

Confidence in Risk Models

Some methodologies rely on an input of "confidence" in order to ultimately produce a distribution around the point estimate. The mistake that can occur in these types of methodologies lies in a breakdown of understanding somewhere in the chain between methodology developer and cost estimator. What these models are generally looking for is "confidence" defined as: What is the probability that the actual costs incurred for this program will fall at or under the point estimate? Sometimes, this is misunderstood by the estimator to mean: What is the probability that the actual costs incurred for this program will fall on or close to my point estimate? Adding another layer to the problem, sometimes interviews are conducted to ascertain the confidence in an estimate, when the confidence is already known. When estimates are made using data-driven approaches including regressions, parametric, or EVM for example, the confidence level of the estimate is almost always 50%. The exception to this is when the estimate was intentionally developed at a level higher than 50%, in which case the confidence can be derived from the data as well.

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

There are three problems in using the approach of specifying confidence as an input that make it inherently dangerous: First, it requires both the risk analyst and the estimator being interviewed to have a considerable level of statistical sophistication. Second, in the case where the risk analysis is being performed by an independent observer and the BOEs were written to a target, it requires them to look deeper than the BOEs to obtain true confidence. This is because the desired confidence should come from the method used to develop the target cost, not the justification used to support it. Lastly, in cases where actual risk items do not constitute a large percentage of the total estimate, these “confidences in the estimate” can drive the entire analysis!

The impact of this misunderstanding of the word “confidence” on the results of the analysis can be *substantial***.** Whereas adding risks to a point estimate *derives* the distribution of an estimate, specifying confidence in an estimate *assumes* that the estimate is distributed, with whatever distribution the methodology decides to use. Overstating confidence and assuming a distribution around the point estimate will lead to poor analysis that gives decision makers the impression that they should be lowering their cost targets to reach a lower confidence level.

Combining (Rolling Up) Risks

Now that some of the common mistakes in identifying and quantifying risks have been discussed, it’s time to look at how risks are compiled into results for presentation to management. There are two main ways of calculating the combined effects of a large number or risks: a Monte Carlo simulation and a Method of Moments model. Both

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

methods work equally well when applied correctly, what follows will be a quick summary of how each method works and the pros and cons of each.

Monte Carlo simulations arrive at the distribution of the combined risks by simulating multiple, independent "runs of the contract" and portraying the range of outcomes. The upsides to this method are that it is the most common approach and therefore will be understood by the largest general audience; it is also far more intuitive and makes fewer assumptions than Method of Moments. The downsides are that it is very difficult to apply correlation correctly and the output correlation matrix will rarely match the input correlation when multiple families of distributions are used. It can also be time consuming and require somewhat heavy computing power due to the fact that (generally) thousands of runs are needed to converge to the actual distribution.

Method of Moments, on the other hand, arrives at the distribution of the combined effects of risks by relying on the Central Limit Theorem (C.L.T). The C. L. T. asserts that a sufficiently large number of risks will eventually combine to a parent distribution (generally normal) whose moments match the combined moments of the child distributions. The pros of Methods of Moments are that, because it assumes one underlying family of distributions (again, generally normal), it is very easy to apply correlation. Also, no simulation, and therefore less computing power, is needed because the mean, standard deviation, and percentiles of the overall distribution are deterministic. The cons are that this method is very non-intuitive (why would I model a Bernoulli risk as Normal?) and requires considerable statistical sophistication to understand. The

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

method also makes several potentially dangerous assumptions. First, assuming normality assumes no skew in the overall distribution. Second, this method assumes that all the risks converge to the C.L.T. C.L.T. assumes there are many distributions all of which are independent and identically distributed. This is often not the case with risk registers.

Method of Moments has another downside. One very dangerous situation that can arise when using this technique occurs when there is a risk (or series of risks) that skew(s) the distribution. This happens when the risk register (when represented as random variables) does not satisfy the Lyapunov condition in regards to the third moments of the child distributions. In cases like this, Method of Moments will produce analysis with inaccurate total percentiles of risk. This calls the viability of Method of Moments into question as a risk tool for a couple reasons. First, without a math background, risk practitioners will be unaware that this mistake has occurred. Second, this mistake can not be caught without running a Monte Carlo simulation on the risk register and comparing the outputs to Method of Moments. At this point, it is reasonable to raise the question as to if Method of Moments is a viable risk tool in the first place.

Truncating Negative Values of Risk Distributions

The next mistake this paper will address is the somewhat common practice of the removal (truncating to zero) of negative values when they occur in risk simulations. This does make intuitive sense (how could there be a negative cost?), however it shifts the mean away from the intended distribution. The mean, arguably the most important result, will be unavoidably higher than what actually exists in the data due to the additional weight placed on zero. All percentiles will be adversely affected as well.

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

Risks With Multiple Possible Outcomes

Another common mis-practice occurs when a risk practitioner faces a situation where there are two outcomes for a risk (these are most commonly discrete events). Most of the time, these are meant to be mutually exclusive events. Sometimes, when the analyst places this risk into the simulation, they will enter them as two separate, and unrelated line items. This has the impact of allowing *both* events to occur in the simulation. In other words, instead of having three possible events (a risk, an opportunity, or nothing), there are now four (a risk, an opportunity, nothing, *and* both the risk and opportunity). Although this does not cause a shift in the *mean*, it does change the standard deviation, and thus the *shape of the total risk distribution*.

Representation of Categories of Risk and Display of Risks

One of the most frequent topics of discussion between risk practitioners and their management is in regards to how categories of risks are represented. The hurdle in this presentation lies in the fact that the subcategories of risk will never sum to the total. Several methodologies contain processes for adjusting the results by category so that they sum to the total. It is the belief of the authors that in general, management does understand this fact; but, giving decision makers some of the basic tools needed to understand our analysis increases its usefulness to them. Instead of explaining the math behind this feature of risk analysis, it is worth trying to relate it to real life.

Suppose I have one die, and roll it once. What is the probability of getting a one? Since the die has six sides, and (one would hope) the probability of landing on any given side is equal, the probability of getting a one (or any other number) is 1/6. Now, let’s take this one step further. Suppose I have two dice, and I roll each of them once, what is the probability of the sum of my two rolls equaling two? The only way this event can occur is if I was to roll a 1 twice. Yet, since probabilities do not add, but rather multiply, the probability of this event occurring is (1/6) x (1/6) = 1/36. It is quickly noted that this is not the sum of the "subcategories" (or rolling a one with one die). When the logic in this description is passed to risk assessment (in terms of cost risk, a "1" would be having an opportunity occur). It is easy to see that the probability of every opportunity occurring (and causing the best case scenario) is generally very low. This same logic applies to percentiles as well as the extrema described above.

Now that a method of explaining to management the nature of subcategories of risk has been illustrated, a way of presenting the information is needed. Risk analysis is generally only a piece of the puzzle when decision makers receive a presentation on a program. This generally leads to the risk assessment results being compressed onto a couple of slides. It is therefore critical that we present the information in a way that is both compressed and evocative.

	Point Estimate	20th %	**50th %**	80th %	**Risk %**	**Risk $**
Labor	$ 100,000	$ 101,144	$ 104,046	$ 108,072	4.0%	$ 4,046
Material	$ 25,000	$ 26,144	$ 29,046	$ 33,072	16.2%	$ 4,046
Total	**$ 125,000**	**$ 129,616**	**$ 133,990**	**$ 138,768**	**7.2%**	**$ 8,990**

Table 1: Display of Risk Results

The above table shows how categories can be presented along with the bottom line. The point estimate is included for reference, along with the $20^{th}/50^{th}/80^{th}$ percentiles. Risk dollars and percentages (based on the 50^{th} percentile) are shown off to the right. This allows decision makers to see the risks from both important perspectives.

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

The Assumption of the Log Normality Distribution of Risks

The last topic in this paper relates to the assumption of log-normality in risk distributions and is meant purely as a thought experiment. Many studies have asserted the CGF distribution across many DoD programs to be distributed log-normally, an example is Arena and Younossi[1]. A paper by Summerville and Coleman[2] presented a risk approach that recommended applying a normal distribution with a mean and standard deviation based on a weighted average risk score based on several objective measures. Could it be possible that the log-normal distribution described in the Arena and Younossi paper is due to the risk scores from the Summerville and Coleman paper being distributed log-normally? This would give the illusion of an underlying log-normal distribution when the actual distribution is normal with a mean and standard deviation dependant on the technical score. This is not to say that it would be right to drop the umbrella log-normal assumption that is being used in many methods, especially when the technical score is unknown.

Conclusion

One of the biggest issues involving risk analysis lies in the fact that it is impossible to catch all mistakes just by looking at percentiles or an S-Curve (which is often all that management sees). Catching mistakes requires looking at not just the models and their outputs, but the methods used to produce the inputs. It is known that "garbage in equals garbage out" but it is generally overlooked that good data into bad

---

[1] Arena, Mark, Obaid Younossi, and et. al.. Impossible Certainty: Cost Risk Analysis for Air Force Systems. Santa Monica: RAND Corporation, 2006

[2] "*Cost and Schedule Risk CE V*" Coleman, Summerville and Dameron, TASC Inc., June 2002

Presented at the 2007 ISPA/SCEA Joint Annual International Conference and Workshop - www.iceaaonline.com

methods also equals garbage out. Due to the mathematical sophistication required to catch many of these mistakes, the authors advocate the vetting of all risk analysis performed within an organization with someone (or some group) who understands both the process and the math behind it. Normally, a few days to a week is all that is needed to catch problems like the ones discussed in this paper. Once problems have been caught they can generally be fixed in order to present the most accurate information available to management. This grants them the ability to make well-informed situations, which in the end, is really what risk analysis is for.