Being Certain about Uncertainty, Part 1

Andy Prince
NASA/MSFC
256-544-8360
andy.prince@nasa.gov

Introduction

In October 2007, I became manager of the Engineering Cost Office (ECO) at NASA's Marshall Space Flight Center (MSFC). At that time the Constellation Program, a program to return humans to the moon and eventually take us to Mars, was well underway. Performing cost risk analyses to determine the overall Program confidence level had been an integral part of the program management strategy from the very beginning. Every budget cycle and every major program update required that the cost risk analysis be revised to reflect the latest program changes.

As a new manager and a strong believer in the value of doing cost risk analysis, I was eager to see cost risk analyses be performed for our other customers. We conducted training classes in cost risk analysis for the office, shared handbooks and lessons learned, provided one-on-one support to less experienced team members, and encouraged everyone by using the Constellation Program (and the Ares I Project being managed at MSFC) as a shining example of what cost risk (or as we sometimes call it, confidence level) analysis could do.

Despite my best efforts, our office was slow to adopt cost risk analysis as a routine part of doing a good cost estimate. Over the next few years the cycle of training, exhortation, and direct support repeated itself with only modest results. Occasionally I would put my foot down and require one of my staff members to produce a cost risk analysis for whatever estimate he or she was working. The usual response I received was a blank stare followed by a "Well how do I do that?"

Finally I discovered the root of the problem. Being very bright and well-educated people, they could certainly handle the mechanics of doing a cost risk analysis. What proved to be the issue was that they had no idea if the results from their confidence level analysis was any good! They had no yardstick by which to measure the quality of their work. They had no basis by which to judge their cost risk analysis.

The Challenge of Cost Risk Analysis

Below is a chart (Exhibit 1) taken from the paper "Covered in Oil" by Dr. Christian Smart. In this chart Christian captures the cost risk analysis history of one of the Constellation Program's significant projects, from inception to cancellation. Look closely at Exhibit 1. Two conclusions can quickly be drawn from a simple inspection. First, it is obvious that the costs grew rapidly – by almost a factor of two in a little less than three years. Second, the earliest s-curves are steeper than the later s-curves. In other words, the uncertainty in the estimate was *increasing* as uncertainty in the system being developed should have been *decreasing*.

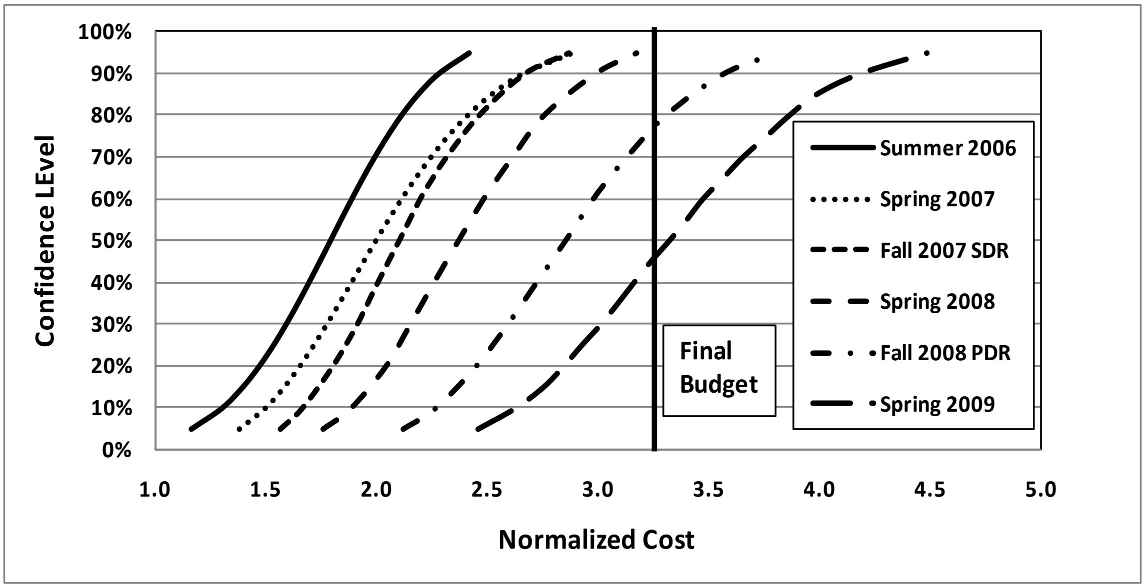

Exhibit 1. History of a Cost Risk Analysis.

Exhibit 1 provides an object lesson in the difficulty of doing good cost risk analyses. We will return to some of those lessons later in this paper.

Cost risk is hard, and it is hard because it is an abstraction. It is an abstract concept built on an assessment of our ability (or lack thereof) to predict the future. Our brains don't like abstract concepts. As Nassim Taleb says, "We harbor a crippling dislike for the abstract." We need something concrete, or at least something to give us a relativistic link to the familiar and understandable. We diligently pursue processes and methods that will provide that link and give us an "explainable" result, a result that "makes sense" given the facts and data. Yet doing a good cost risk analysis is as much art as it is science.

Cost risk analysis is highly subjective. Numerous decisions must be made to facilitate the analysis that requires a judgment call. Just one example: determining the amount of uncertainty to use for an input variable for a parametric model. We cost estimators can try to estimate the amount of uncertainty, but we are not the experts. We can pass the buck by engaging experts to tell us how much uncertainty to use. But studies show that even experts are often wrong, and that they are just as susceptible to being overly optimistic as the rest of us. We cannot escape using somebody's knowledge and experience in our analysis.

The subjective nature of cost risk analysis is further exacerbated by the lack of consensus on the best method to use. According to the Air Force *Cost Risk and Uncertainty Handbook* a user can select an inputs-based simulation approach, an outputs-based simulation approach, or a scenario-based method (SBM). Each of these methods has its pluses and minuses, and none of them have been proven more effective than the others. They all require a significant amount of analyst's judgment to establish the probability distributions, uncertainty bounds, protect scenarios, and correlation values.

Doing a cost risk analysis requires a basic understanding of probability and statistics. Unfortunately, probability and statistics are foreign to how we think. This is why lotteries and casinos are so successful. I have observed over the years that probability and statistics are especially challenging for most people in upper management and other positions of authority. Several years ago I was doing some cost estimates on a rather important study. I had presented my initial estimates, along with my cost risk analyses, to a group of senior NASA managers, and that presentation had gone well. After two weeks of additional study I presented my updated estimates and risk analyses to the same group of managers. The most senior manager in the group looked at my cost risk analysis and asked me why had the confidence levels not gone up. I thought I must have misunderstood so I asked him what he meant. He responded by saying, "You've studied this for the last two weeks, so shouldn't you be more confident in your estimate?" I was so stunned that I did not know what to say. Fortunately, my boss (Joe Hamaker) was quick on his feet and able to provide a suitable explanation. Otherwise, I might have been fired on the spot!

How we understand and interpret the cost risk abstraction is important to how we tell the story of our analysis. The cost risk analysis probability distribution is derived from analysis, rather than actual data, yet our minds want to treat the results as if they are a true representation of all uncertainty and risk faced by the project, without considering that the analysis itself contains uncertainties and risk. Given that cost risk analysis is usually performed under an organizational imperative to create a product that either confirms the status quo or leads to specific action (is "explainable"), the analyst often has no choice but to find and deliver meaning. But in creating meaning, we may be discounting the real possibility that our abstraction is inadvertently (or perhaps deliberately) limiting uncertainty to make the results more palatable to our masters.

Common Problems with Cost Risk Analysis

The challenges discussed in the previous section are exacerbated by the fact that cost risk analysis is performed by human beings. As discussed in my papers "The Psychology of Cost Estimating" and "The Dangers of Parametrics," we are far from the rational thinking machines that classical economics teaches. We are extremely capable of irrational and extreme behavior. The following paragraphs highlight some of the more common problems we humans face in trying to do credible cost risk analyses.

First of all, we often confuse risk and uncertainty. Risk is a probability of loss. Uncertainty is the indefiniteness of an outcome. In other words, risk is the possibility something might happen, uncertainty is the fuzziness around what will happen. Risk is often expressed in terms of probability and consequences using a 5x5 matrix like the one shown in Exhibit 2.

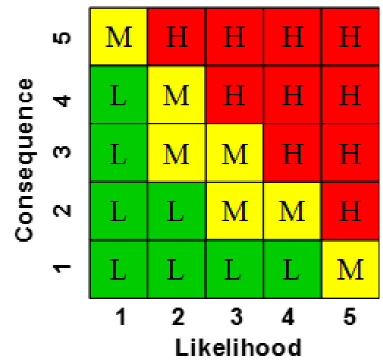

Exhibit 2: 5x5 Risk Assessment Matrix.

The problem I have observed with risk and uncertainty is that people cannot appreciate how uncertain uncertainty can be. Project management types, especially, have a tendency to treat plans as reality except for those pesky risks which can be addressed and managed to a satisfactory resolution. They will acknowledge that uncertainty exists but will whitewash it by using a standard reserve factor or, even worse, a so-called statistical and history-based analysis that eliminates outliers and arbitrarily minimizes the potential for extreme cost growth. This issue will be explored in greater depth in Part II of this paper.

Studies show that when we come to judging probabilities we do a pretty poor job. Most people will assign the probability of an event occurring as either yes, no, or maybe. Which means we see two possible outcomes that have 100% probability of occurring and one that captures everything in between. In their book "Superforecasting," Philip Tetlock and Dan Gardner talk about how one of the traits that super forecasters develop is the ability to see possibilities with much finer granularity. Being able to see the many shades of maybe is difficult. After all, what does the difference between a 60% probability of occurrence versus a 70% really mean? The concept of a probability is an abstraction. Remember, we don't like abstractions.

Humans have the ability to take a messy, confusing, haphazard series of past events and turn them into linear narrative that makes the outcome (meaning the present) seem all but inevitable. We know we do it, and we have even coined a phrase for it: hindsight is twenty/twenty. Nassim Taleb calls this simplification of the past the *narrative fallacy*. In his book "The Black Swan" he has this to say about how the narrative fallacy inhibits our ability to fully appreciate randomness:

> The narrative fallacy addresses our limited ability to look at sequences of facts without weaving an explanation into them, or equivalently, forcing a logical link, *an arrow of relationship*, upon them. Explanations bind facts together. They make them all the more easily remembered; they help them *make more sense*. Where this propensity can go wrong is when it increases our *impression* of understanding.

Thus we think we understand because we have simplified the past to make it understandable (some call this processing). But because our simplification can lead to a misunderstanding of what truly happened, we are vulnerable to applying that misunderstanding to the future, to our cost risk analysis. Our perception of the future becomes captive to our understanding of the past, leading to mistakes in our cost risk analysis by giving us a false confidence in our interpretation of how historical projects behaved and thus creating a distorted mirror for analyzing current project data.

It is easy for us to be overconfident and optimistic. It is a basic human trait. Because we think we understand the past, we therefore think we understand the future. In fact, Nate Silver, in his book "The Signal and the Noise," states that "overconfidence is a huge problem in any field in which prediction is involved." Overconfidence hinders our ability to do cost risk analysis because it makes us think we can accurately understand the present (and project into the future) when in fact we don't even accurately understand the past.

For example, if you use prediction intervals to estimate CER uncertainty (as the statisticians tell us to do), then it is possible that under certain conditions you will get extreme values at the upper and lower ends of your s-curve. The upper end extreme values look frightening, and often are frightening to our customers. They might even tell us to make them go away. The temptation to make extreme values go away is strong. After all, how realistic is it to have an s-curve that shows a 95th percentile value that is 3 times greater than the point estimate? If the historical data or some characteristic of the system being analyzed does not support such extremes, we are more than likely going to find a way to truncate the CER error or even ignore it altogether to get an answer that is easier to justify and explain. Such willful ignorance can be dangerous – the statistics are trying to tell us that we are not as certain about our uncertainty as we want to believe.

The final problem that we face when doing a cost risk analysis is known as confirmation bias. Confirmation bias causes us to look for and find data that fits a preconceived opinion or outcome. Confirmation bias causes us to focus on results, rather than achieving results through the analysis process. When an analysis is used to confirm an outcome, rather than to define the outcome, your analysis is no longer objective and fact based. Such an analysis may please the customer, but it may ultimately prove misleading to everyone involved.

The crux of all of these problems associated with risk analysis (and there are several others that I could have chosen to mention) is that they lead towards a notion that I have rarely heard discussed: there is uncertainty in our uncertainty analyses. This uncertainty is driven by our biases and by the subjective judgments needed to facilitate a cost risk analysis. We cannot avoid being biased and judgmental any more than we can avoid breathing. Therefore, we must look for objective ways to measure the quality of our cost risk analyses.

Validating a Cost Risk Analysis

Now we come to the heart of the matter. You have done your cost risk analysis. You have your s-curve. You step back and look at what you have created. How do you know it is any good?

I have identified three ways to validate your cost risk analysis: process; the coefficient of variation; and use of historical data. Process validation is following a systematic approach to determine if you have adequately addressed all sources of uncertainty and correctly calculated the cost risk. Validation processes are defined by the GAO in their "Cost Estimating and Assessment Guide" as well as in cost risk handbooks developed by the Air Force and others. The GAO process is shown in Exhibit 3.

1. Determine the program cost drivers and associated risks.
2. Develop probability distributions to model various types of uncertainty.
3. Account for correlation between cost elements.
4. Perform the uncertainty analysis using a Monte Carlo simulation model.
5. Identify the probability level associated with the point estimate.
6. Recommend sufficient contingency reserves to achieve an acceptable level of confidence.
7. Allocate, phase, and convert a risk-adjusted cost estimate to then-year dollars and identify high-risk elements to help in risk mitigation efforts.

Exhibit 3: Seven steps associated with developing a justifiable s-curve from the GAO Cost Estimating and Assessment Guide, GAO-09-3SP (page 159).

The coefficient of variation (CV) is a standard measure of dispersion for a probability distribution. The CV is calculated as the standard deviation divided by the mean, and is a measure for the amount of uncertainty captured by the probability distribution function (PDF) or s-curve. The CV is a standard output of most Monte Carlo simulation packages (@Risk, Crystal Ball, Argo, etc.). One way to think of the CV is to visualize it as a measure of the flatness of the s-curve. The greater the CV, the more dollars it takes to increase the confidence level percentile. This is illustrated graphically in Exhibit 4.

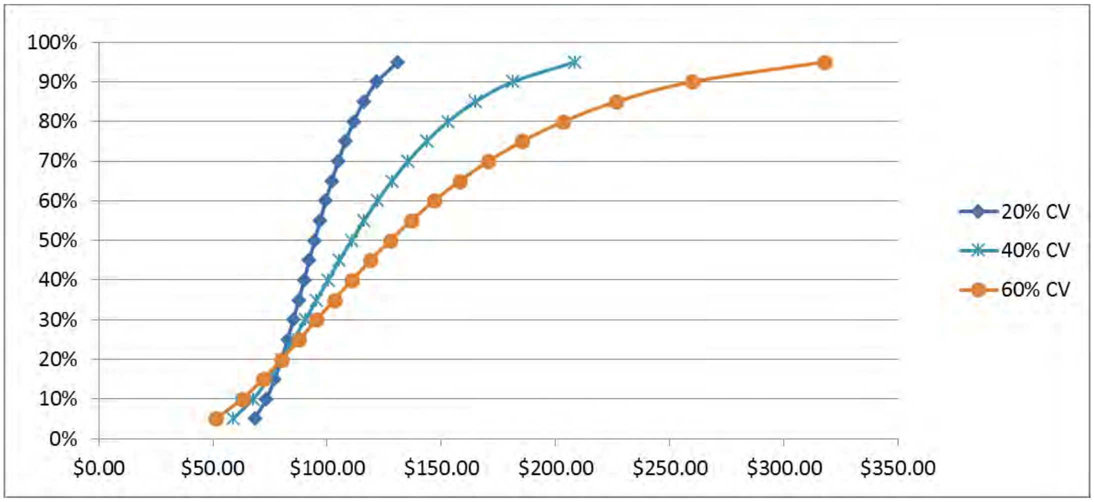

Exhibit 4: Cost Risk S-Curves with Different Coefficients of Variation.

The value of using CV as a measure of goodness is readily apparent upon examination of the curves in Exhibit 4. For each curve the point estimate is $80. The 70th percentile value on the 20% CV curve is about $105, or about 31% greater. The 70th percentile value on the 40% CV curve is about $135, which is almost 69% greater than the point estimate. Therefore, a cost risk analysis that yields a 40% CV s-curve is telling you there is more uncertainty in the estimate than a curve with a smaller CV.

Research into various handbooks and Wikipedia turned up very few recommendations for CV values. The Air Force Cost Risk and Uncertainty Analysis Handbook gave the following recommendation:

> … early in the project 35-45% is typical for space systems and software intensive projects; 25-35% is typical for aircraft and similar complexity hardware; and 10-20% is typical of large electronic system procurements.

The Joint Cost Schedule Risk Uncertainty Handbook, a joint publication of NASA and various branches of the Department of Defense, has a table of CV's based on the cost growth experience of the Naval Center for Cost Analysis (NCCA).

Using historical data to validate your cost risk analysis requires that you have historical project cost growth data. Fortunately, NASA has a rich set of project data on cost overruns. Exhibit 5 shows a bar chart based on the experience of 158 NASA projects.

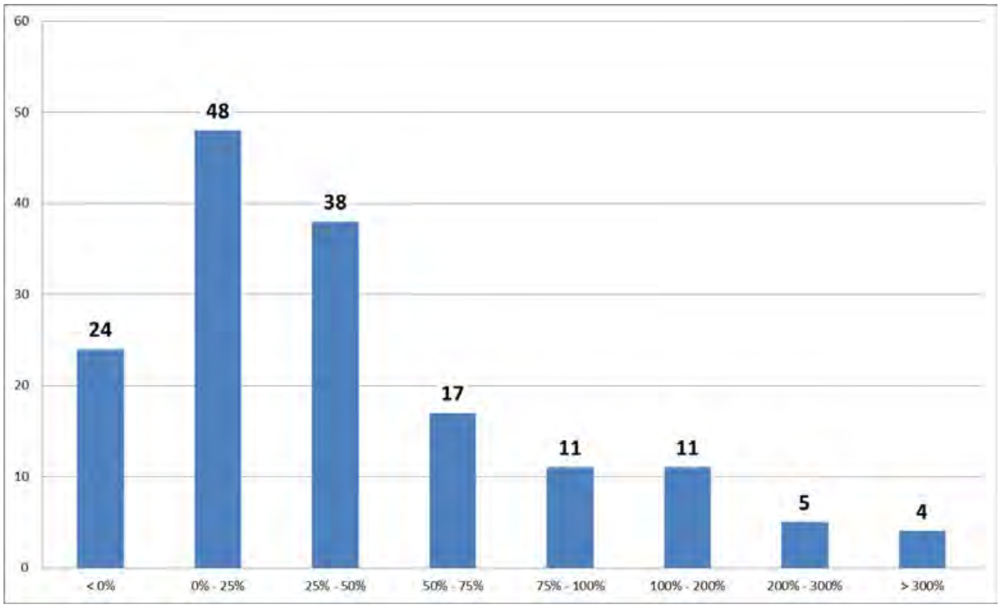

Exhibit 5: NASA Cost Growth History.

Out of these 158 projects, 24 (or 15.2%) have experienced no cost growth. However, 19 projects (or about 12.0%) have experienced cost growth of more than 100%. Therefore the chance a project will experience no cost growth is only slightly better than the chance it will experience extreme cost growth.

To help us use this data to evaluate a cost risk analysis, I have used the data to calculate a PDF, which is shown in Exhibit 6.

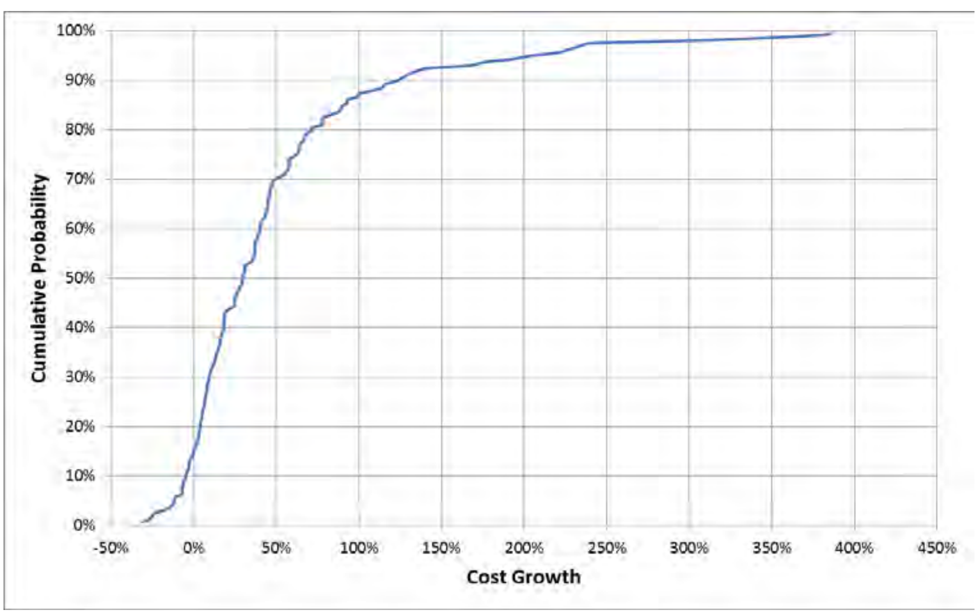

Exhibit 6: Historical Data Cost Growth Probability Distribution Function.

Several conclusions can be drawn from the graphs in Exhibits 5 and 6. First, a hypothetical project's cost estimate is historically at a 15% confidence level (CL). This means that 85% of all NASA projects will, on average, experience cost growth. Of those projects that experience cost growth, the median increase is about 30% (50% CL value). NASA typically budgets at the 70% CL. The historical data tells us to expect about 50% cost growth to achieve a 70% CL. Interestingly, up to about the 70% CL the slope of the line is greater than one. However, after 70% CL the slope of the line is less than one, and begins to go asymptotic as it approaches the 100% CL.

One approach to using the historical cost growth data is to fit a probability distribution. Exhibit 7 shows the result of fitting a lognormal distribution to the PDF in Exhibit 6.

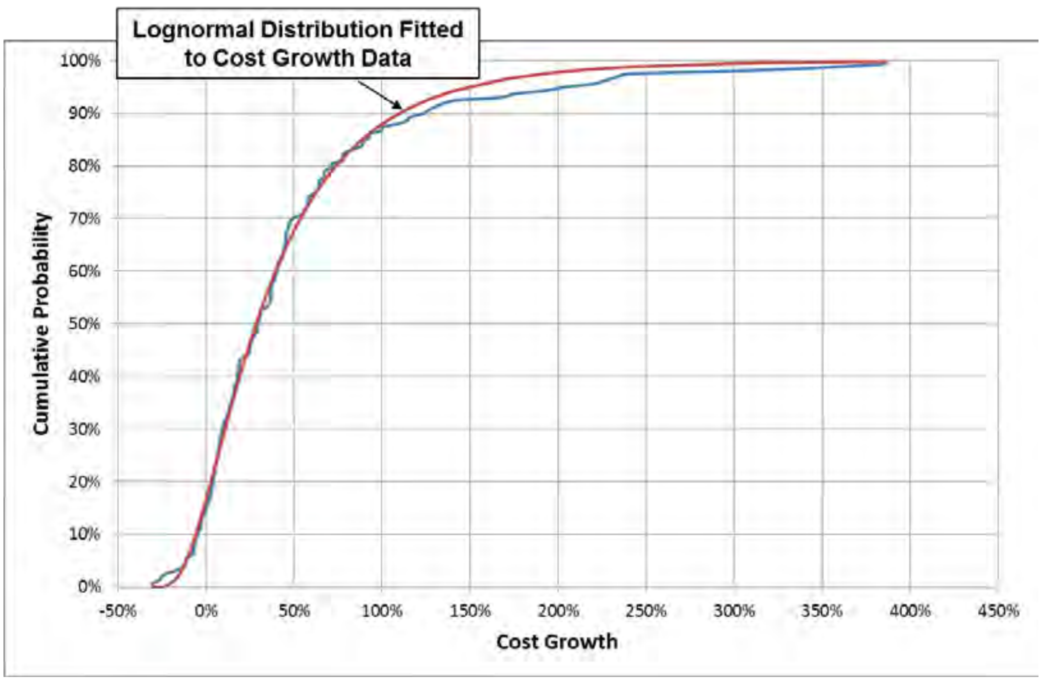

Exhibit 7: Lognormal Distribution Fitted to the Historical Cost Growth Data.

The lognormal distribution function models the cost growth data extremely well. The functional form of the lognormal PDF used in this case is a three-parameter model, rather than the standard two. The third parameter is an offset to account for the negative cost growth values captured by the actual data. The parameters for the fitted PDF are a mean of 74.4%, standard deviation of 56.6%, an offset of -30.7%, and a CV of 76.1%. We have used our lognormal PDF with the enhanced Scenario Based Method (eSBM – Garvey, Flynn, Braxton, and Lee) to develop a spreadsheet model. The spreadsheet model can be used to do an eSBM risk analysis or as a tool for validating a risk analysis performed using other methods.

So in summary, when validating a risk analysis you should check your process, check your CV, and compare your results to historical experience. When checking your process, make sure you are accounting for correlation. Check to make sure all significant sources of uncertainty are adequately addressed. Be aware that optimism and overconfidence will cause us to believe that things are more certain than they really are. If you are using triangular distributions you may be artificially limiting the potential for cost growth.

The CV provides a simple way to determine the amount of uncertainty in your cost risk analysis. You can compare the CV of your s-curve to a CV derived from historical cost growth data to see how much uncertainty your analysis has captured compared to past projects. Historical cost growth data can also be used to do more specific comparisons, such as by type of system (the NASA data shown in Exhibit 5 can be subdivided into earth orbiting spacecraft, launch vehicle stages, human-rated vehicles, etc.). PDF's can be fitted to historical data and used to make simple models for validating. Finally, techniques such as eSBM can be used to develop alternative cost risk analyses for comparison to more traditional approaches (and vice versa).

Your cost risk analysis should be a logical outcome of all the evidence, and all that evidence must be fairly weighted. Putting too much emphasis on the project's story is probably going to result in an s-curve that is too steep (as measured by the CV) and insufficient to address all uncertainty. But if your analysis is not supported by facts and data, then key stakeholders might discount the value of your work and fail to appreciate the potential for cost growth.

Extreme Cost Growth

Look again at the bar chart in Exhibit 5 or the PDF in Exhibit 6. Notice that most projects experience cost growth of less than 50%, but; a significant number of projects have cost growth greater than 75%. In fact, 31 out of the 158 projects (almost 20%) have more than 75% cost growth. Cost growth of 75% or 100% or greater is extreme, and has a disruptive effect on a government agency's fiscal management. Being able to foresee the possibility of extreme cost growth with sufficient time to take corrective action would be a major improvement to government financial management.

As Nassim Taleb says, again from his book *The Black Swan*, "The inability to predict outliers implies the inability to predict the course of history…" and right now we struggle to predict outliers. To recover from a problem, you must first realize that you have a problem. It is easy to believe that our processes and analyses make extreme cost growth unlikely, but history does not support such a belief. Take a look at Exhibit 8. Exhibit 8 is a power function fitted to the cost growth frequency data from Exhibit 5, converted to probabilities, to yield a model to estimate the possibility of cost growth. The projects that underran or had no cost growth have been excluded.

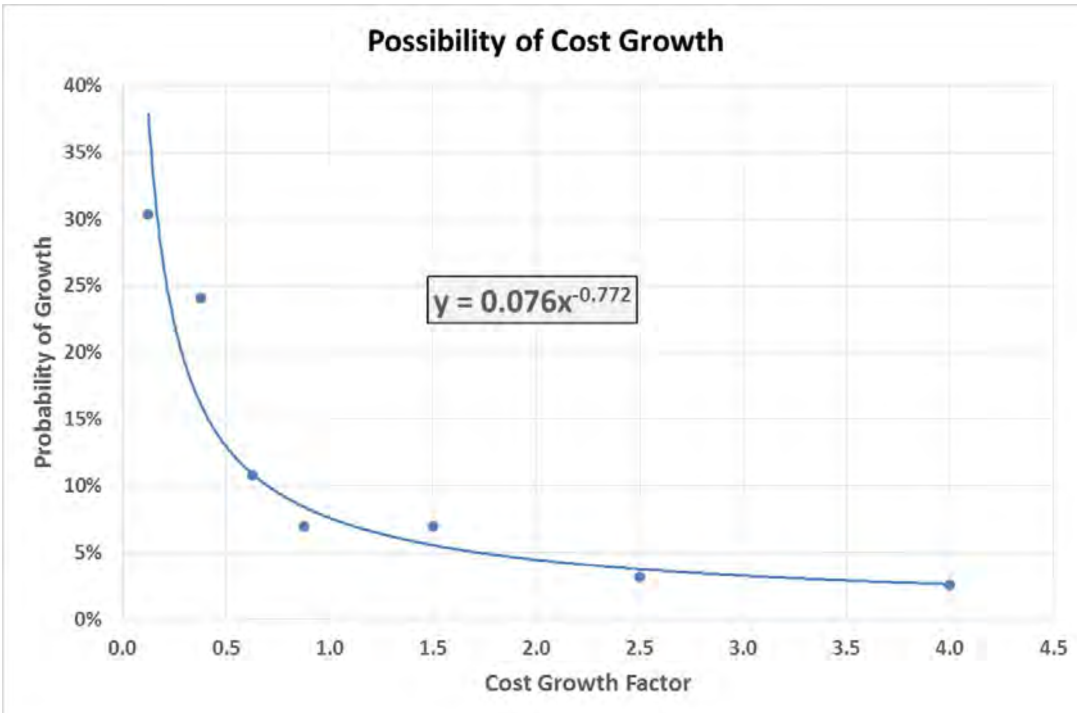

Exhibit 8: Probability of Cost Growth.

Exhibit 8 tells us that the potential for extreme cost growth has no upper bound. The probability does diminish rapidly starting at about 100%, but even as far out as 400% (a cost growth factor of 4), there is still a 2.5% chance that a NASA project could experience that level of cost growth. 2.5% is small, but it is not zero, and the possibility of less extreme cost growth in the range of 100% to 150% is more than 5%. This is not theoretical. Some recent examples of extreme cost growth include the James Webb Space Telescope (JWST), the Ares I launch vehicle (cancelled), the A-3 test stand (built for Ares I), and the F-35 fighter. Had I taken the time to look, I am sure I would have found more examples from other government agencies.

What is surprising about extreme cost growth is that it continues to happen. The examples cited in the previous paragraph are all relatively recent, and two are ongoing programs. Extreme cost growth is not limited to government. It also happens in industry. It happens despite being a problem known and studied for at least the last 40 years. In the REDSTAR (Resource Data Storage and Retrieval) Library there are 1,127 individual documents - studies, surveys, assessments, and recommendations on how to prevent extreme cost growth. Yet extreme cost growth continues to be a problem routinely highlighted by the GAO and Inspectors General. So the question must be asked: Why aren't we doing better? Why has our study of cost growth, which is really a study of history, failed us?

When we study project histories we can read documents, talk to key participants, and review financial reports. However, no matter how hard we try, we cannot fully grasp all the complexities and details buried within those documents and reports. Even if we worked on the project, we would only have insight from the advantage afforded to us by our position. We would not know what it was like to work in a different position, or at a different level in the organization.

In the words of Nassim Taleb, "history is opaque." He has identified three problems with how we understand history that he calls his *"triplet of opacity."*

a. The illusion of understanding, or how everyone thinks he or she knows what is going on in a world that is more complicated (and random) than they realize;
b. The retrospective distortion, or how we can assess matters only after the fact, as if we are looking in a rearview mirror (history seems clearer and more organized in history books than in empirical reality); and
c. The overvaluation of factual information and the handicap of authoritative and learned people, particularly when they create categories – when they "Platonify." (Nassim Taleb, The Black Swan, page 8)

Parts a and b refer to the Narrative Fallacy discussed earlier. Part a is also making the point that we create simplified models of the world to explain how the world works (which can be called heuristics). The problem with heuristics, as discussed by Douglas Hubbard in his book *How to Measure Anything,* is that some are reasonable but some are wrong. And I believe that both Hubbard and Taleb would agree that even reasonable heuristics ignore the complications and randomness of the real world.

In the final part of the triplet, Taleb uses the term "Platonify" to describe how we humans will build entire disciplines (like perhaps, cost risk analysis) around observed facts and imposed logic. We take what we see, which is very messy, use categorization and logical structures to explain the messiness, then create disciplines that promulgate and solidify our categorization and logic. Because the structures and logic define the discipline, they can become impediments to true understanding, but more importantly; they do become impediments to the humbling realization that we understand the world far less than we think we do.

Onward to the Future

If we seek to do better than all the previous cost growth studies, none of which seem to have had any significant or lasting impact, we must learn why the approaches we keep using fail and look for different ways to address the problem. One thing we have learned from studies of human behavior is that we primed are from birth to find causality – we see events as linked, such that the occurrence of one event explains the occurrence of another. Therefore, when we study history, especially when we study history with the intent of finding causes to explain a specific outcome, we find causes. And as long as those causes make logical sense, we believe them.

Douglas Hubbard calls seeing what we want to see the Expectancy Bias. The danger of seeking, finding, and explaining causes (impolitely called "jumping to conclusions") is that we underestimate the randomness that contributes to outcomes.

We can also fail to see what the data is not telling us. For example, if you read cost growth studies on major government acquisitions, almost all of them will cite requirements instability or lack of requirements definition as a source of cost growth. Yet I have never seen a study that looked at requirements instability in projects that had no or only modest cost growth. Therefore, I believe we need to bring a more holistic approach to the study of extreme cost growth.

The approach we use must create a division between knowing what happened and explaining what happened. Our approach must seek knowledge without prejudice, focused solely on the knowing while attempting to identify and eliminate biases. We must be observers of the past without judging the past.

But how do we do this? I don't have any easy answers. One approach would be to seek to understand the larger environment in which the project is executed, including both projects with extreme cost growth and projects without. Taking a larger, more holistic point of view might allow us to put those programs and projects that experienced extreme cost growth in a larger and more complete context. Such a point of view has the potential to alter our worldview, to let us see what may previously have been unseen.

Another approach, which could be used in conjunction with a study of the larger environment, is to have two or three analysts prepare a project history *independent* of each other. By doing independent studies, personal biases could be identified and minimized. Plus, having several people familiar with a program or project would allow for different points of view to be brought to the table when the time comes to begin looking at causes.

In Summary

There are several key points that I hope my readers take with them. First, doing a good cost risk analysis is hard. Every project is unique, with its own special challenges and opportunities. Give your analysis the attention it deserves. Please don't make the mistake of reducing hard (that which is difficult) to a formulaic exercise.

Second, the coefficient of variation (CV) is a good yardstick for checking to see if you have included enough uncertainly in your s-curve. However, you have to evaluate the CV in the context of past historical cost growth experience. Otherwise, using the CV to judge your result is engaging in guesswork.

Third, check to make sure you are not ignoring key sources of uncertainty. These sources include CER error, assumptions regarding inheritance and technology readiness, and the maturity of the technical design. Sensitivity analyses are an excellent way to understand the level of uncertainty incorporated in your key assumptions. As always, history can provide a guide.

Fourth, be a realist about the possibility of extreme cost growth. Project managers and other senior leaders will be loath to acknowledge this possibility despite evidence to the contrary. Being loathed does not make our jobs any easier, but it is our job to speak truth to those who must make important decisions. We may face the difficult choice of either disappointing a customer today, or satisfying a customer at the expense of making our organization look foolish down the road. Be careful with this one.

Finally, always remember that the less you and anyone else know about the system, the greater the uncertainty. We cannot help but build a better story the less we know, and when you are part of a group a herd mentality takes over, driven by optimism and overconfidence. In this setting empirical evidence of past failures is minimized or forgotten. To overcome this one, we must rely on objective evidence over impressions. Standing apart from the group is hard. Make sure you have the data to support your position.

Into the Void

The story of being certain about uncertainty does not end with this paper. Until we have some way to see extreme cost growth coming, we will always be operating with one eye closed, impervious to the possibility that our analysis is misinforming management and that someday it could blow up in our faces. During the next several months our office will be researching this issue, applying the principles I discussed earlier in an analytical framework that might, but might not, lead us to an answer.

Stay tuned for Part 2.

Bibliography

Ariely, Dan, *Predictably Irrational*, Revised and Expanded Edition, New York: Harper Perennial, 2009

Aschwanden, Christie, "Your Brain is Primed to Reach False Conclusions." *fivethirtyeight.* February 17, 2015. <http://fivethirtyeight.com/features/your-brain-is-primed-to-reach-false-conclusions/>

Garvey, Paul R.; Flynn, Brian; Braxton, Peter; and Lee, Richard, "Enhanced Scenario-Based Method for Cost Risk Analysis: Theory, Application, and Implementation," *Journal of Cost Analysis and Parametrics,* Vol. 5, Issue No. 2, 2012: 98-142.

Hubbard, Douglas W., *How to Measure Anything*, New Jersey: John Wiley & Sons, 2010

Kahneman, Daniel, *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux, 2011

Levitt, Steven D. and Dubner, Stephen J., *Freakonomics, a Rogue Economist Explores the Hidden Side of Everything*, New York: Harper Perennial, 2009

Mlodinow, Leonard, *The Drunkards Walk: How Randomness Rules Our Lives*, New York: Pantheon Books, 2008

Mooney, Chris, "The Science of Why We Don't Believe Science." *Mother Jones.* May/June 2011. http://www.motherjones.com/politics/2011/03/denial-science-chris-mooney

National Aeronautics and Space Administration and the Department of Defense, *Joint Cost Schedule Risk and Uncertainty Handbook,* April 2013

Nuzzo, Regina, "How scientists fool themselves – and how they can stop." Nature. October 7, 2015. http://www.nature.com/news/how-scientists-fool-themselves-and-how-they-can-stop/

Prince, Frank, "The Dangers of Parametrics," *Proceedings of the 2016 International Cost Estimating and Analysis Association Professional Development and Training Workshop,* Atlanta, June, 2016

Prince, Frank, "The Psychology of Cost Estimating," *Proceedings of the 2015 International Cost Estimating and Analysis Association Professional Development and Training Workshop*, San Diego, June, 2015

Siegel, Eric, "The One Rule Every Data Scientist (and Manager) Should Know By Heart," GovExec.com, December 21, 2015. <http://www.govexec.com/technology/2015/12/oneruleeverydatascientistandmanagershouldknowheart/124803/print/>

Silver, Nate, *The Signal and the Noise: Why most Predictions Fail but some Don't*, New York: The Penguin Press, 2012

Smart, Christian, "Covered in Oil, Realism in Cost Risk Analysis," *Journal of Cost Analysis and Parametrics,* Vol. 8, Issue No. 3, 2015: 186-205.

Smart, Christian, "Here, There Be Dragons: Considering the Right Tail in Risk Management," *Journal of Cost Analysis and Parametrics,* Vol. 5, Issue No. 2, 2012: 64-86.

Taleb, Nassim Nicholas, *The Black Swan: The Impact of the Highly Improbable,* New York: Random House, 2010

Tetlock, Philip E., and Gardner, Dan, *Superforecasting: The Art and Science of Prediction,* New York, Crown Publishing Group, 2015.

U.S. Air Force, *United States Air Force Cost Risk and Uncertainty Analysis Handbook,* April 2007

U.S. Government Accountability Office, *GAO Cost Estimating and Assessment Guide,* March 2009