Being Certain About Uncertainty – Part 2
Can We See Extreme Cost Growth Coming?
Andy Prince & Christian Smart

Introduction

In my last paper, Being Certain About Uncertainty – Part 1, I left the reader with a question: Can we see extreme cost growth coming? At the time, I had my reservations. After all, numerous cost growth studies have been done over the years which led me to wonder if 1) could I really contribute anything new or useful to understanding cost growth, and 2) is it possible to identify indicators that would tell us that a project is going to experience extreme cost growth. I quickly realized that this challenge requires more analytical muscle than my few brain cells can muster, so I asked Christian Smart to help.

Christian brings a wealth of knowledge, experience, and analytical horsepower to understanding the causes of extreme cost growth. He has written several award winning papers on cost risk analysis, uncertainty, and the portfolio effect. With Christian's help, I have been able to look at the problem both from both an analytical and a managerial standpoint. His wisdom has provided valuable insights and he has not been hesitant to give me honest feedback on the goodness (or lack thereof) of my own ideas.

In trying to identify the cause (or causes) of extreme cost growth, I quickly realized that despite one's best efforts, we cannot escape the narrative fallacy. For those of you who are new to this subject, the narrative fallacy is our ability to build stories about the past. These stories both simplify events and create clear lines of cause and effect. They lead us to believe that the world is predictable, that it makes sense. Unfortunately, these stories also create the illusion that we truly understand the past, when in fact what we actually understand is our interpretation of the past and the evidence we have gathered to bolster that interpretation.

Therefore, as you read this paper please keep in mind that Christian and I have brought our own stories and our own biases to this subject. Using data and analytical techniques provides a certain amount of objectivity (though many of you understand that even data and analyses are subject to confirmation bias). But at the end of the day we must recognize that we create our own stories of the past, interject our own views of cause and effect, and face our own limitations of understanding. Despite our very human limitations, we hope that you find what we have to say interesting and illuminating.

Defining Extreme Cost Growth

There is no standard definition for extreme cost growth. If costs grow on a given project to the point where the overall portfolio is affected, either by taking resources from other projects or preventing the establishment of new projects, then the cost growth is problematic, but is it

extreme? Congress has imposed on the Department of Defense a policy that directly specifies how DoD must address cost growth. The key elements of that policy are summarized below.

- In order to combat cost growth, Senator Sam Nunn and Representative Dave McCurdy established legislation in the early 1980s that required programs to report on significant cost growth and provided specific guidelines on what is considered a baseline breach.
- A significant breach is 15% growth above the current baseline, or 30% above the original baseline.
- A critical breach is 25% growth above the current baseline, or 50% above the original baseline.
- Critical breaches can result in program cancellation unless the program is restructured and root-cause analysis is conducted on the program's cost growth.

So according to Nunn-McCurdy, cost growth of 25% or more could be considered extreme. For NASA, the policy picture is a little murkier. NASA has no legislative requirement that defines limits for reporting or taking action on cost growth, but; does state in governing documents that projects which exceed 30% cost growth could be rebaselined and a 10% cost growth must be reported to the Office of Management and Budget (OMB).

NASA has a specific analytical process for establishing the baseline on spaceflight projects above $250M life-cycle cost. Called the joint cost schedule confidence level analysis, or more simply the Joint Confidence Level (JCL), this analysis combines cost uncertainties, schedule uncertainties, and risks into a Monte Carlo simulation that produces cost-schedule pairs and their associated confidence level. NASA policy recommends budgeting projects at the 70% JCL, which is a point where both the cost and schedule have a 70% confidence level of being realized. We will talk more about JCLs and their impact on cost growth later in the paper.

For the purposes of this study we define extreme cost growth to be 100% or more above the initial baseline. Since we are using mostly NASA data in our analysis, we established the baseline to be Systems Requirements Review (SRR). We realize that 100% is an arbitrary number. In coming to this conclusion we considered that cost growth in the range of 25% to 50% is not uncommon, nor is it even noteworthy, for high technology space programs. Our stakeholders do not typically make that level of cost growth an issue. Conversely, cost growth of 100% or more can prompt increased scrutiny by third part groups (such as GAO) and lead to Congressional action. Cost growth between 50% and 100% could, under some circumstances, be considered extreme depending upon the collateral impacts. Limiting extreme cost growth to 100% or greater still provided a relatively large (18) and diverse set of projects for examination.

A Few Lessons from History

Those of you familiar with our work have seen bar graphs like the one shown in Exhibit 1 in our previous papers and presentations. Despite changes in data composition (some old data removed, some new data added), the general shape is enduring. Looking like a lognormal distribution, the data shows a few success stories (cost growth <= zero) and a few disasters

(cost growth >= 100%). In fact, in this data set the number of the former is almost exactly equal to the number of the latter: 17 versus 18.
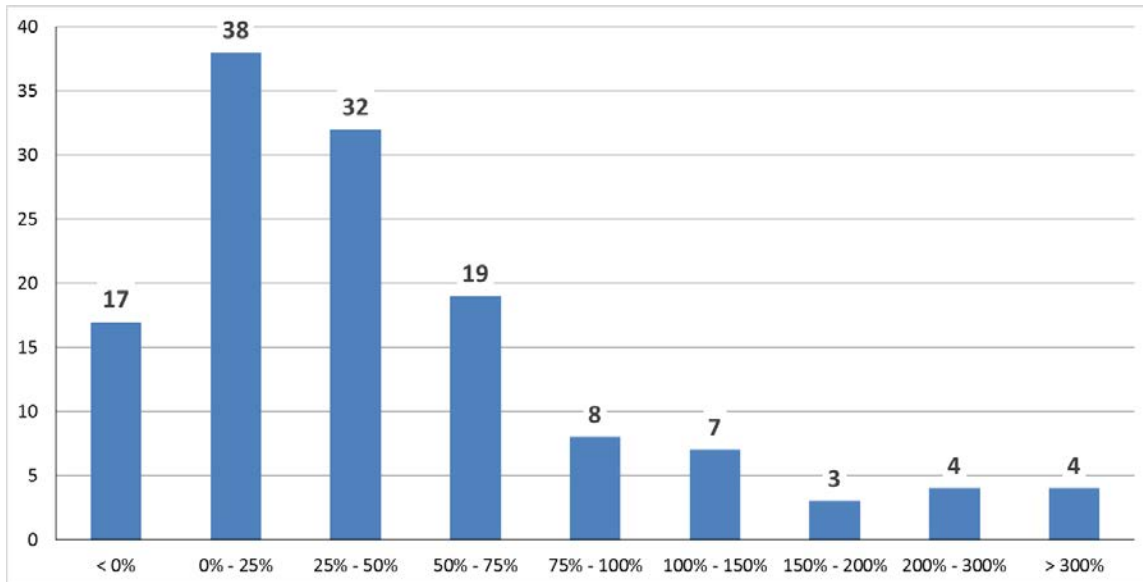

Exhibit 1. Historical Cost Growth Bar Chart.

From our historical data set we can draw some common statistics. The mean cost growth is 56.2% with a standard deviation of 82.5%. The median is 35.1%. The best performing project had a cost underrun of -26.8%; the worst had a cost growth of 498.3%.

Visually, the bar chart can be somewhat misleading. Because the sizes of the bins change (from 25% to 50% to 100%) the bar chart understates the size of the tail. Plotting the frequency against the average percent growth, as shown in Exhibit 2, gives a truer representation of the possibility of extreme outcomes.
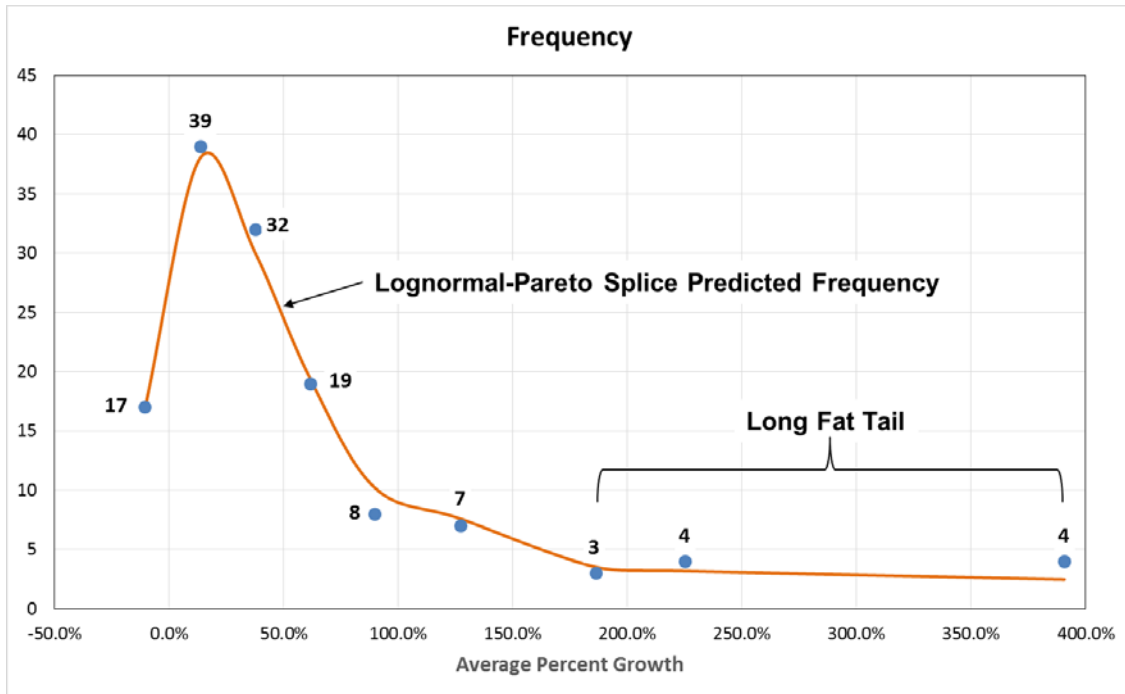
Exhibit 2. Frequency Graph of Historical Data.

With such a long, fat tail fitting a probability distribution proved to be a challenge. In the end, Christian grafted a Pareto distribution onto a lognormal distribution to model the general trend. The Pareto distribution has a heavier tail than a lognormal distribution and is used to model extreme variations. The logic and derivation behind the lognormal-Pareto splice is given in Appendix A.

For the cost engineer, there are a few other important lessons that can be drawn from the historical data. First, as explained in previous papers by the authors, analysts can develop enhanced Scenario-Based Risk (eSBR) or calibration models using historical data. Such models can provide a more realistic picture of project risk, can be used to perform a risk analysis in the absence of other information, and be used to validate (or perhaps invalidate) risk analyses developed by others.

Second, using historical data, the analyst can look for relationships that could provide insight into a specific project's potential for cost growth. Interestingly, we found no correlation between the estimated cost and the cost growth percentage or the actual cost and percentage cost growth. However, the data indicated a high correlation (cc = 0.91) between the cost estimate and the amount of cost growth and an equally high correlation between the actual cost and the amount of cost growth. In fact, the relationship between estimated cost and actual cost can easily be predicted with a simple power function as shown in Exhibit 3.
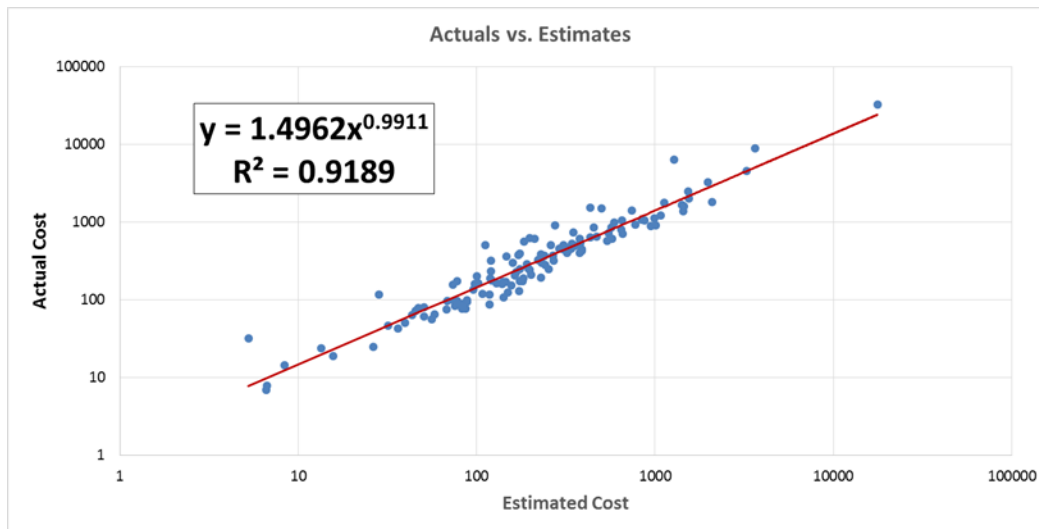
Exhibit 3. Relationship between Estimated and Actual Cost.

The equation in Exhibit 3 demonstrates how easy the job of a NASA manager could be. All they need to do is add 50% to any cost estimate coming from a project and their portfolio would average out to around zero percent cost growth. However, managers often struggle with empirical information. In the next section we explore the differences in how the project manager views the world versus the view of the cost analyst.

The Cost Estimator and the Project Manager

Key to understanding why extreme cost growth occurs is understanding the difference in how cost estimators and analysts view the world and how project managers view the world. Parametricians and other data driven analysts know the importance of historical data and how to effectively use that data to make cost estimates. Cost estimators also know that performing cost estimates and analyses requires significant judgment. For most of us, that judgment is developed through experience, and sometimes experience can be a hard teacher.

Because our data is often messy and we know that predicting the future is challenging, good estimators learn to be comfortable with uncertainty and to see the world probabilistically. Developing credible estimates requires that we use all the information at our disposal, taking into account historical experience as well as the unique characteristics of the system we are estimating. At the end of the day, we estimators know that what we put forth will be scrutinized by everyone in our management chain. Therefore, we focus on being able to explain, defend, and support every decision, every bit of data, and every judgment, so that our results will be seen as credible.

The project manager has a very different mindset than the cost estimator. While we cost estimators see our estimate as the end product, the project manager sees our work merely as a means to an end. Project managers are results driven. They have a job to do, and that job is to

successfully complete the project. Thus, to the project manager, the cost estimator is there to help them get the job done.

Project managers see the world deterministically. In a deterministic view of the world, plans can be made, problems uncovered, and actions taken which lead to a successful result. Yes, the good ones do recognize uncertainty, but they usually believe that uncertainty can be bounded in the same way a safety factor can be put on a bridge design. Their deterministic view of the world leads to a problem known as the *planning fallacy*. Daniel Kahneman and Amos Tversky describe the planning fallacy as "plans and forecasts that are unrealistically close to best-case scenarios." In other words, the plans are success oriented.

Because the continuation of government projects is dependent upon the support of key stakeholders, including politicians, good project managers know that they need a consistent message and that they must build and maintain relationships. The consistent message is important for creating a sense of realism, excitement, and inevitability about a project. Such messages are necessary when building advocacy among customers, stakeholders, and senior agency leaders. A good project manager is a good salesperson for their project. But sometimes the sales message does not agree with the analysis results.

There is a natural tension between the cost estimator and the project manager. The cost estimator is on a quest for knowledge and understanding. The project manager is focused on achievement.

It is in the melding of both points of view that success can be found. Used correctly, the analysis community functions as a governor, keeping the results-focused optimism of the project manager in check. As we will see, it is when an organization loses this tension between project advocacy and good estimating that the risk of extreme cost growth increases.

History for Managers

History for project managers is the same history that drives cost analysts. Only the explanations have changed to protect the guilty. Whereas the cost engineer sees trends and generalities, the project manager sees lessons learned. Take a look at Exhibit 4, a nice summary of findings from previous NASA cost growth studies (see Appendix B for a list of the studies used in building the table).

Cost Growth Reasons	1970s	1980s	1990s	2000s
Inadequate definitions prior to agency budget decision and to external commitments	X	X	X	X
Optimistic Cost Estimates/Estimating Errors	X	X	X	X
Inability to execute initial schedule baseline	X	X	X	X
Inadequate risk assessments	X	X	X	X
Higher technical complexity of projects than anticipated	X	X	X	X
Changes in Scope (Design/Content)	X	X	X	X
Inadequate assessment of impacts of schedule changes on cost		X	X	X
Annual Funding instability			X	X
Eroding in-house technical expertise			X	X
Poor tracking of contractor requirements against plans			X	X
Launch Vehicle			X	
Reserve Position adequacy		X		X
Lack of Probabilistic estimating		X		X
"Go as you can afford" Approach				X
Lack of formal document for recording key technical, schedule and programmatic assumptions (CARD)**				X

** CADRe has since been implemented as a requirement of NPR 7120.5

Exhibit 4. Summary of NASA Cost Growth Studies.

Notice the top four reasons: inadequate project definition; optimistic cost estimates; unexecutable schedules; and inadequate risk assessments. These are all management failures. Yes, you can blame the cost and schedule communities for being optimistic. But those of us who have been in the business a few years know it is management that makes the final decision on cost and schedule, and management is rarely pessimistic.

It is interesting to note that the number of reasons is increasing over time. We are not sure if this is because NASA is getting better at identifying reasons for cost growth, or if we just fail to learn.

Nevertheless, project managers like studies like these. Why? Because they explain cause and effect and lead to specific actions that can implemented. Specific actions that address specific problems increase our confidence that we are doing the right things to ensure success, which helps with messaging and advocacy. Good project managers do not see history in terms of warnings to be heeded, but rather as pointing out potential challenges that can be overcome (see planning fallacy above).

History is an Illusion

Obviously, history itself is not an illusion. What is an illusion is our understanding of history. The illusion begins with cause and effect. We all believe in cause and effect. We us it every day to make decisions and for planning. What we don't realize is that what we believe is cause and what we believe is effect are influenced by a number of biases. These biases lead to a distortion in our understanding which makes actions and outcomes seem more direct and linear than they really are. For example, the representative bias, which causes us to value information that is more easily recalled, makes us susceptible to the influence of a good story, especially a story that plays to our stereotypes.

Note that the cause and effect we speak of in the previous paragraph is not direct, physical cause and effect (I was hungry, so I bought lunch). Rather, the type of cause and effect that we

are talking about, the type which is subject to biases and oversimplification, is the cause and effect we create in the absence of knowledge and understanding. The cause and effect we create to explain things like the outcomes of political races, macro-economic behavior, and cost growth on government projects.

Heuristics are a codification of our incomplete understanding of cause and effect. Heuristics that are intuitive, simple, and easily understood become the most powerful. Despite their flawed foundation, many heuristics are useful. For example, most people would agree that system complexity and cost are positively correlated. Such a heuristic can have value when trying to provide a relative ranking in cost between two systems. The problem arises when people treat a heuristic as fact rather than as a rule of thumb.

Ultimately, as discussed in Part 1 of this paper, history fails us because of the Narrative Fallacy. For those of you who did not read "Being Certain about Uncertainty – Part 1," the Narrative Fallacy, as defined by Nassim Taleb, is our ability to assign a linear narrative to history, to make events seem almost inevitable. We simplify the past to make it more understandable and memorable. What our simplification ignores is the impact of random events.

We struggle with randomness. If we believe that randomness plays a significant role in our lives then we are accepting that we have limited control over our future. Such acceptance is hard. It goes directly against our illusion of control, an illusion that is necessary for us to function. Daniel Kahneman, in his book *Thinking, Fast and Slow*, gets right to the heart of why this is so difficult.

> The illusion that one has understood the past feeds the further illusion that one can predict and control the future. These illusions are comforting. They reduce the anxiety we would experience if we allowed ourselves to fully acknowledge the uncertainties of existence. (*Thinking, Fast and Slow*; pages 204-205)

How does this set us up for failure? The illusion that we understand history leads to an overly simplistic understanding of the past. Because we believe we understand the past, we now believe that (within limits) the future is deterministic. But to achieve a desired end result, we must plan to take the proper actions based on the correct set of initial conditions. To create the perception that a project can be accomplished within a desired cost and schedule, a good project manager constructs a plausible scenario (plan) based on what is known today. This plausible scenario becomes the basis for selling the project.

For example, to save money a project manager may look for ways to use more heritage hardware. A less ethical project manager might even go so far as to browbeat the poor cost engineer into assuming more heritage. But let's assume our project manager is ethical. Even ethical project managers are going to work hard to solve the cost problem. Sometimes the solution leads to what we have termed "over-specification."

An over-specified estimate is one which is based on a significant number of cost-saving assumptions. These assumptions are critical to project success. Over-specification also feeds the illusion that we are reducing uncertainty, that we can determine our future. But over-specification comes with a very high price. The greater the number of assumptions, the more risk we have of one or more of those assumptions failing to hold. This is basic probability: the greater the number of events required for success, the greater the probability that at least one event will fail.

The Triumph of Randomness

In the prolog of his book *Antifragile*, Nassim Taleb defines three types of responses to exposure to negative events: fragile; robust; and antifragile. Obviously, a fragile response is one that is harmed by a negative event. A robust response is neither harmed nor helped. An antifragile response is one that responds positively (it benefits) from harm or volatility. For each of these responses Taleb defines systems, people, and institutions that are typical of the response. He puts high technology projects in the fragile category.

High technology projects are complex and fragile, small things can have large negative consequences. To quote Taleb

> Complex Systems are full of interdependencies – hard to detect – and nonlinear responses. … Man-made complex systems tend to develop cascades and runaway chains of reactions that decrease, even eliminate, predictability and cause outsized events. (*Antifragile*, page 7)

If you have studied the history of even a few high technology development projects (yes, I know what I just said about history) you will see that what Taleb says rings true. It is extremely difficult to foresee, with any reasonable reliability, what will cause cost growth or how extreme that cost growth will be.   If we could, it would be relatively easy to identify and correct. When we over-specify the conditions for a cost estimate, we think we are addressing and minimizing the possibility of cost growth, when in fact the opposite is true. We have increased the complexity and interdependencies.

Thus randomness can defeat our best laid plans. An empirical study of NASA and DoD history shows that somewhere between 12% and 15% of all projects will experience cost growth of more than 100%, yet how many risk analyses have been done that show that level of potential growth in the tail? We invariably provide input into our cost risk analyses based on the triangular distribution, as if we can truly bound everything we believe we know and understand. Yet the future continues to confound us with events and outcomes that we neither expect nor are adequately prepared for.

To illustrate the dominance of randomness over specificity, I offer up the example of the Hubble Space Telescope. The Hubble Space Telescope (HST) is an engineering marvel and a great scientific achievement. It has literally rewritten textbooks. Yet, as an example of how to

manage a large, complex project, it was a failure. In constant year dollars the development cost of HST grew 274%. Originally planned to launch in 1983, it did not actually make it into space until 1990. Once in space, it was found to have a flawed primary mirror (due to a small fleck of reflective material being lost from an alignment rod) that had to be corrected.

As in most cases of extreme cost growth, management attempted to discover the causes. What they found was that a number of assumptions about the HST development were not realized. In some cases decisions made to reduce cost actually led to cost increases. The findings by NASA management are summarized below.

- HST had a complex management interface, with two lead NASA centers (Marshall Space Flight Center for the spacecraft, telescope, and overall system integration; Goddard Space Flight Center for the instruments and ground system) and two prime contractors (Lockheed for the spacecraft and Perkin Elmer for the telescope).
- Assumed use of existing and standard space flight hardware did not materialize.
- The original estimate did not include sufficient spares.
- It was assumed that because HST was launched on the Space Shuttle, there would be large mass margins and that these margins would translate into cost savings. In reality, weight growth exceeded the available margin and the design had to be light weighted.
- The telescope was sold as design-to-cost with performance allowed to vary. However, performance was held constant and cost increased.
- Telescope contamination requirements increased.
- Historical data indicating high costs for optics, fine guidance sensors, and optical structures was removed from the CERs.

Particularly troubling is the last finding: removal of relevant data from the CERs. A willingness to explain away the past as irrelevant to the future is a manifestation of the narrative fallacy. We believe we understand the past well enough so as not to make the same mistakes. Obviously NASA did not. You would think that having endured the management and technical failures of HST, NASA management would approach the development of the next large space telescope with a measure of humility and circumspection. How wrong you would be.

Take a look at Exhibit 5. Exhibit 5 compares the costs of HST and the James Webb Space Telescope (JWST) with their initial estimates.
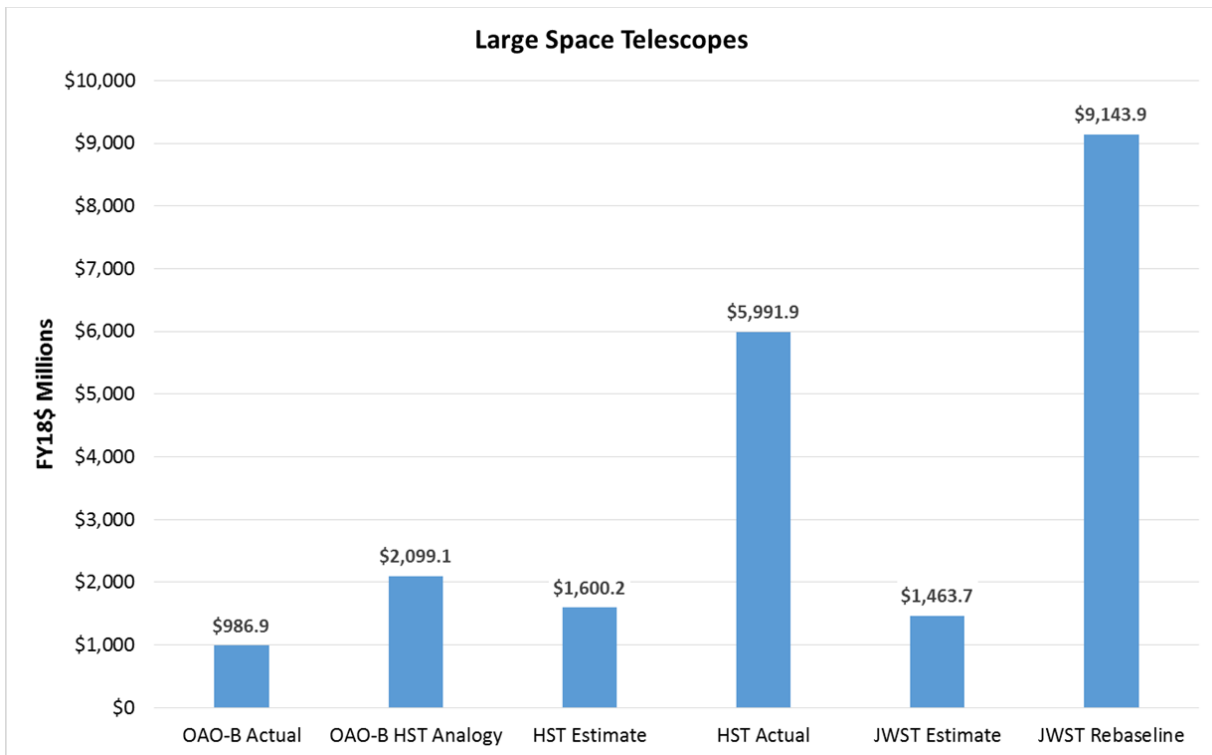
Exhibit 5. A History of Large Space Telescope Cost Estimates and Actuals.

The early HST estimates were informed by an earlier telescope mission called OAO-B (OAO for Orbiting Astronomical Observatory, B because in NASA nomenclature spacecraft in a series are designated using letters on the ground, and numbers once they are launched into space). OAO-B suffered a launch failure and never reached orbit. However, its development costs were well documented and used to develop an early analogy estimate for HST. Of course management, being smarter because of history, believed that they could do better, and went forward with an estimate for HST that was over $400 million less than the analogy estimate and well short of the actual development cost.

It took only a few short years for management to forget the lessons of HST. When JWST was initially formulated in the late 1990's the development cost was assumed to be only $1.464 billion, *or over $100M less than NASA's commitment to Congress for HST*! This was despite the fact that JWST is larger and more complex that HST. What could possibly justify such an estimate?

In explaining what happened, I (Andy Prince) must expose my own culpability in creating the mess that was to come. The program that was to become JWST began formulation in 1996. Then called the Next Generation Space Telescope (NGST), it was to be a follow-on to the Hubble Space Telescope (HST). I was invited to attend a gathering of scientists, engineers, managers, and programmatic professionals where we were told, in no uncertain terms, that our job was to show how NGST could be design, built, and launched for $1B dollars. Notice that our job was *not* to develop an estimate of the cost for NGST. Rather, our job was to justify an estimate

11

being handed down from the highest levels of NASA. So we did our job. We produced a report showing how it could be done and even included the phrase "…our costs are realistic and conservative." We did this despite knowing that HST had cost over $3.3B to develop in 1996 dollars.

Why Extreme Cost Growth Occurs

To this point we have looked at the empirical history of cost growth and seen that extreme cost growth is real, and that it occurs with some regularity. We have examined the psychology of why management fails to grasp the possibility of extreme outcomes or believes that such outcomes can be managed away. We have discussed the fragility of high-technology projects and how easily they are upset by random events. And we have seen how management is slow to learn.

Thus we are now ready to explain why extreme cost growth occurs and what can be done to address it. The key component in the failure of any organization to prevent extreme cost growth is the lack of independent cost estimates and assessments being presented to and discussed by senior management. A second and equally important component is a lack of an independent technical review, whose results are reported to and discussed by the senior decision makers.

When only one point of view is presented, that of the project manager, senior decision makers are deprived of information and put into a position of thinking they must make a yes or no decision. Decision-makers believe that their options are limited. When we are put into a situation where we believe our choices appear to be limited, we are *narrow framing* the problem.

Humans by nature prefer problems that are simple. We are also risk-averse. If you have ever dealt with a competent car salesperson, you have experienced narrow framing first-hand. A good car salesperson will lead you to see the decision, to buy or not to buy, as a loss if you do not buy and a win if you do. They do this by using human psychology to get you to eliminate other alternatives (keeping your old car) and playing on your fear of loss (you will miss a great opportunity). By the way, most car salespeople don't take psychology classes to learn how to do this. Their techniques are based on thousands of years of trial and error by salesmen and saleswomen. That is why they are so effective.

Now, we are not trying to equate project managers with car salespeople. But it is true that when the decision makers are only told one side of the story, it is easier for management to accept the project's position. To do otherwise would be to turn down the possibility of success by accepting the sure failure (by saying no). When presented with more than one point of view, decision makers are forced to consider other possible outcomes. Incorporating an independent cost risk analysis, or as well we see, a joint cost schedule risk analysis, into the discussion goes even further, by giving decision-makers a sense of the risk involved in accepting the project's position and some understanding into how that risk, if considered too high, can be mitigated.

12

To effectively counter our tendency to narrow frame, senior management must create a culture that values and expects to see alternative points of view. A management culture that is not willing to entertain independent assessments of their projects is probably a management culture more focused on selling projects than on successfully executing projects.

One final note: independent cost estimates and assessments will not eliminate extreme cost growth. But such practices can reduce the possibility. The alternative, failure to perform independent assessments of high-technology projects, guarantees that extreme cost growth will occur.

Some Observations from DoD

The Department of Defense has a management culture that creates its own challenges to effective project management. Some of these challenges are shared by NASA project managers. One that can affect both organizations (though perhaps to different degrees) is learned helplessness.

Learned helplessness starts with the situation that many government project managers find themselves in: they are responsible for the project's success but they are not the ultimate authority. In other words, some decisions (usually the big ones like major requirements changes) must be approved by a higher authority. This puts the project manager in the position of having to justify and defend important changes that he or she believes are necessary for project success.   Learned helplessness is exacerbated when project managers have to live with suboptimal decisions imposed upon them by their boss or by a key stakeholder, such as Congress.

This loss of control creates a psychic stress (in what is already a very stressful job) that can make the project manager feel helpless. Once a project manager begins to feel helpless, they lose their motivation, their sense of responsibility. Demotivation and a loss of responsibility leads to degraded performance. Degraded performance increases the likelihood of cost growth.

Another DoD culture challenge to good project management is the high turnover of uniformed project personnel. These military officers are typically only in the job for two or three years. Thus, poor decisions by these short-term project managers often will not bear fruit until they have moved to their next assignment. Since they are not affected by the consequences of their decisions, they lack what Taleb calls *skin in the game*.

Is the Joint Cost Schedule Confidence Level the Solution?

The philosopher and mathematician Gottfried Leibniz is noted for a variety of contributions including the co-discovery of Calculus and invention of binary numbers. However, his optimism, particularly his phrase that we live in the "best of all possible worlds," has been heavily lampooned, including by Voltaire and the singer/songwriter/actor Kris Kristofferson. Modern program managers however seem to have adopted this philosophy wholeheartedly, especially

when it comes to cost. This is one of the primary reasons that cost growth is a ubiquitous phenomenon in government programs.

A recent claim that joint confidence level analysis is reducing cost growth may be proving effective at corralling cost growth, according to a recent study of 13 missions. As you can see in Exhibit 5 below, the average overrun for these 13 missions from Systems Requirements Review (SRR) was 23.9%, with a range from -28.7% to +81.1%.   Average cost growth of -2.7% from the Agency Baseline Commitment (ABC), which is the milestone when the JCL is performed is even better.

Mission	Data				Cost Growth		
	SRR	ABC	Actuals		SRR to ABC	SRR to Actuals	ABC to Actuals
NuSTAR	$96.2	$109.9	$104.0		14.2%	8.1%	-5.4%
Landsat 8	$382.1	$587.6	$395.7		53.8%	3.6%	-32.7%
IRIS	$86.2	$140.7	$156.0		63.3%	81.1%	10.9%
LADEE	$117.9	$168.2	$188.3		42.6%	59.6%	11.9%
MAVEN	$488.7	$567.2	$467.9		16.1%	-4.2%	-17.5%
GPM	$660.2	$555.2	$470.5		-15.9%	-28.7%	-15.3%
OCO-2	$225.2	$249.0	$304.6		10.6%	35.3%	22.3%
SMAP	$412.0	$485.7	$469.9		17.9%	14.0%	-3.2%
MMS	$741.0	$857.3	$962.3		15.7%	29.9%	12.2%
Astro-H	$30.0	$44.9	$51.0		49.9%	70.1%	13.5%
OSIRIS-Rex	$515.7	$778.6	$648.7		51.0%	25.8%	-16.7%
CYGNSS	$125.0	$152.8	$90.1		22.2%	-27.9%	-41.0%
SAGE-III	$56.8	$64.6	$81.6		13.7%	43.7%	26.3%
Average					27.3%	23.9%	-2.7%

Exhibit 5: Cost Performance for JCL Projects.

This is a small sample, so a legitimate question is whether or not this was just good luck? Small samples are notorious for not being representative of the underlying population and are exceptionally prone to the influence of noise. As a result, our initial hypothesis was that this sample is too small to be conclusive. In order to test this hypothesis, we compared this recent success to the larger set of cost growth data from Exhibit 1. This larger data set is more representative of typical government experience, with mean growth equal to 56.2 % and a range from -26.8% to 498.3%.

In order to test this hypothesis, we conducted a means comparison test, assuming that cost growth is lognormally distributed. The test was run using the cost growth data from SRR since this is what most closely matches the data in Exhibit 1. This results in a statistically significant difference. In log-space, the distributions are both normal. The difference of the two distributions is also normal. The combined mean is 0.3040, and the combined standard deviation of the means is 0.1786. Using a t-test with 12 degrees of freedom and the hypothesis that the difference is zero yielded a p-value equal to 5.7%. If we adopt the 5% Fisherian threshold, we cannot reject the hypothesis that the two sample means are not different. This is

a strict criterion aimed at rejecting a true hypothesis, since there is less than a 6% chance of seeing differences this extreme if our null hypothesis is correct.

To take another look at the data, we used the larger sample of data as an empirical distribution. We took samples of 13 data points without replacement from the 132 data points. We ran 1,000,000 iterations. The range of average cost growth for the samples was -0.8% to 168.1%. The overall mean was 52.0%, similar to the overall mean. There were 2 samples with mean less than 0% and 369 with mean less than 10%. See the Exhibit 6 for a histogram of these sampled values.
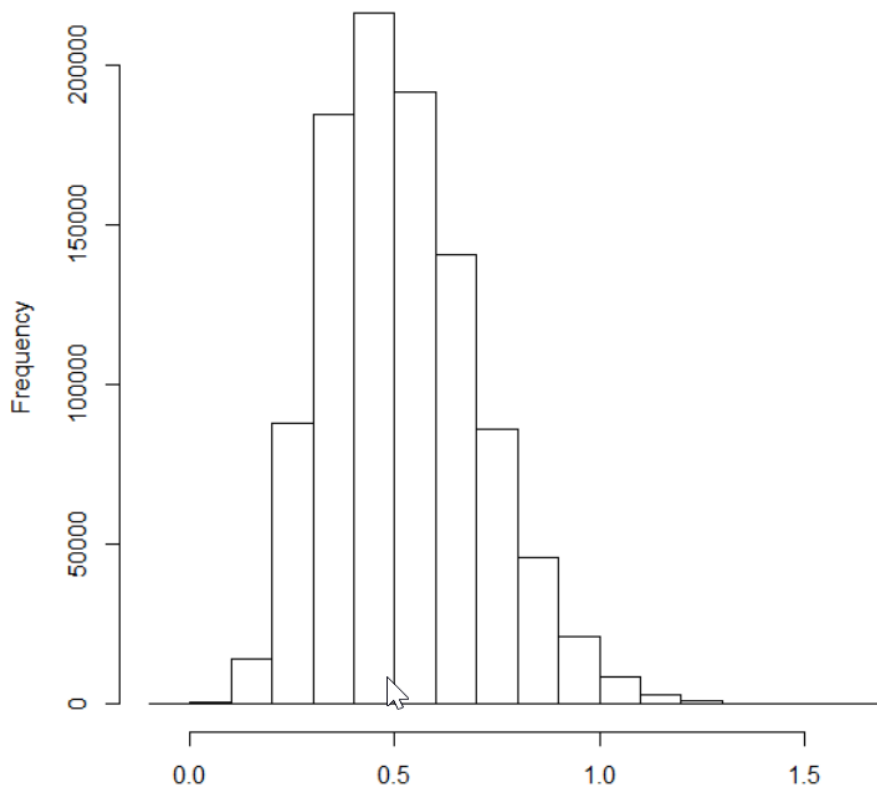

Exhibit 6. Histogram of Sampled Values.

In 1,000,000 trials, 37,678 had an average cost as small as the 13 recent missions, or less, which is approximately 3.8% of the trials.

Thus, there is some evidence that joint confidence level analysis is curtailing cost growth. However, using a standard statistical test and using a standard threshold for p-values that is commonly used in science, we cannot reject the hypothesis that there is no difference between the two samples. So, while these early results are promising, the sample size is too small to make definitive conclusions.

15

A Candidate Predictive Model

Most cost analysts are familiar with regression analysis, which when used to estimate cost, is a predictive technique that results in a continuous output variable, either dollars or hours. We have begun development of a model to predict whether or not extreme cost growth occurs, which we defined as cost growth in excess of 100%. This is a discrete outcome. This type of modeling is called classification. In order to model this we need to look at a technique different from classical regression analysis. The tool that we have used is called "logistic regression," which predicts log odds of success which ranges from 0 to 1. In order to classify the prediction into success or failure, we choose a cutoff point. Anything below that point we predict that it will not have extreme cost growth, and anything above that value we predict that it will experience extreme cost growth.

We estimate the outcome as $p(x) = \dfrac{e^{\beta_0 + \beta_1 x}}{1 + e^{\beta_0 + \beta_1 x}}$. The graph of p is the well-known S curve, and is bounded between 0 and 1. A graph of p is shown in Exhibit 7.
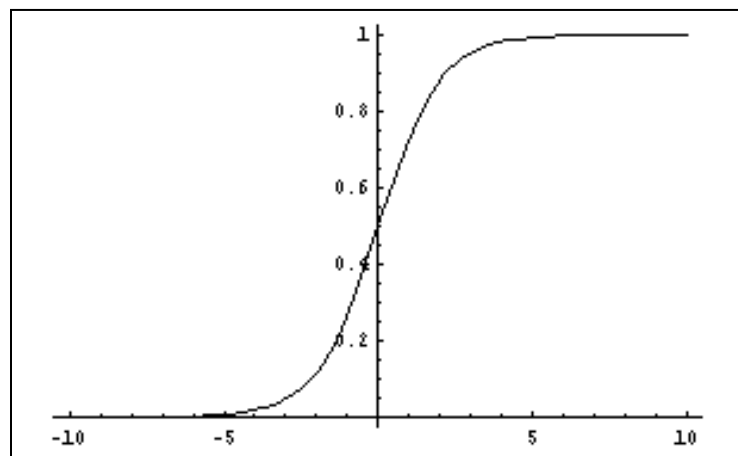

Exhibit 7. The Logistic Curve.

The observed values of p, which we denote by y, are either 0 (failure) or 1 (success). Thus the conditional distribution of y does not follow a Normal distribution, but rather a binomial distribution. Therefore, linear regression cannot be used to estimate the parameters β_0 and β_1. However, the method of maximum likelihood can be used. The

likelihood function is $l(\boldsymbol{\beta}) = \prod_{i=1}^{n} p(x_i)^{y_i}[1 - p(x_i)]^{1-y_i}$. The log-likelihood of this function is

maximized by solving the equations $\sum_{i=1}^{n}[y_i - p(x_i)] = 0$ and $\sum_{i=1}^{n} x_i[y_i - p(x_i)] = 0$ for β_o and β_1.

These are nonlinear equations that do not have closed-formed solutions. Iterative numerical

methods must be used to solve these equations. We used the R statistical language for model calculation.

This is the case for one independent variable, or when one input parameter is used. In our model, several independent variables are used, and a natural extension of p is calculated as

$$p(x) = \frac{e^{g(\vec{x})}}{1 + e^{g(\vec{x})}}$$, where $g(\vec{x}) = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_n x_n$. The maximum likelihood method

also extends to multiple independent variables.

The use of logistic regression arose in epidemiological research, and is now commonly employed in business and finance, ecology, engineering, health policy, and linguistics.

We have some experience with using the logistic regression model for classification. Christian developed a logistic regression model in the early 2000s to predict technical mission success for robotic earth-orbiting and planetary spacecraft. (Smart 2002)

The hypothesis is that extreme cost growth is due in large part to disconnects between cost, schedule, and technical performance plans. For example, if the technical performance is high relative to the initial estimate, the cost will grow in order to meet that performance. Likewise, if the schedule is too short given the scale of the project, the schedule will slip, leading to cost growth. We used the cost and schedule data from the recent study by Andy Prince (2017) for the initial cost and schedule estimates, and Dr. Joe Hamaker's QuickCost version 5.0 for a variety of technical and programmatic parameters for use in the model.

The parameters we considered for the model are: initial cost; initial schedule; spacecraft complexity; instrument complexity; whether or not the mission is earth-orbiting or planetary; whether or not the budget was capped; and whether or not the mission was mostly designed and built in-house by the government.

We had 69 data points. Of those 69, 17 experienced cost growth in excess of 100%. Our model predicted 7 of those, slightly less than 60%. The model also predicted 11 instances of extreme cost growth when there was none. These results are displayed in Exhibit 8.

		Predicted	
		Extreme Growth	Not Extreme Growth
Actual	Extreme Growth	7	10
	Not Extreme Growth	11	41

Exhibit 8. Model Predictions – Actuals vs. Predicted.

There is clearly room for improvement but our initial model shows some promise that extreme cost growth can be predicted. One additional factor we plan to investigate that we did not

17

include in this version is interaction effects, which we believe will improve the accuracy of the model.

A Few General Observations

During our years in the field of cost estimating and analysis, we have made a few observations that appear, dare we say, obvious to us, yet sometimes seem to elude the most intelligent project managers and systems designers. Therefore, we offer these *heuristics* not as hard and fast rules, but rather as things to be on the lookout for when considering the possibility of extreme cost growth.

Our first observation is that if a system is more complex than previous similar systems, it should cost more than these predecessor systems. As a fellow cost engineer, you are probably thinking that this one is too obvious to even be mentioned. In our defense we offer the example in Exhibit 5.

The second observation is that the greater the number of cost-saving assumptions, the greater the probability of cost growth. This observation is really a summation of the problem of over-specification discussed above in the section titled "History as an Illusion," so we will not belabor the point.

The third observation is that in general, technology advances will not reduce cost. We have observed over the years that development projects which rely on promising new technologies often realize greater cost overruns due to those technologies not being sufficiently mature at the start of development. Companies that product large, complex commercial products, such as aircraft and automobiles, typically only include two or three new technologies in a new model, so as to minimize the risk of development cost overruns and operational reliability issues. Early versions of the SpaceX Falcon 9 rocket relied heavily on existing technologies. Sometimes new technologies will have a cost and/or schedule benefit, so the cost analyst needs to be careful in applying this heuristic too liberally.

The fourth observation is that the more important the project is to the organization, the more it will cost. Important projects tend to be highly visible to outside stakeholders. Projects that have more visibility attract more resources, as organizations typically invest to insure success. Often these high-visibility projects attempt to achieve lofty goals, such as hitting a nuclear warhead with a projectile or designing a fighter jet that can satisfy the requirements of the Air Force, Navy, and Marines. Important and highly visible projects that expand capabilities and drive technology development are a recipe for extreme cost growth.

The final observation we want to share is this: if you are told up front what a project is going to cost, before any cost estimates have been done, it is a really, really bad sign. For an example, re-read the JWST story above. When you are told up front what the cost will be, needs to be, or even the budget wedge that the project needs to fit in, your estimate is being anchored.

Numerous studies show that anchoring influences cost estimates. It is very difficult to overcome an anchor. You cannot forget what you have learned. Even worse, your customer is now anchored. Once anchored it is very difficult to provide an objective cost estimate. We have even heard project managers say "It can't cost more than X, because that is the budget," as if the budget determines the cost estimate.

Conclusions

Last year we began this study with the hypothesis that it is the environment surrounding a project that creates the conditions for extreme cost growth. After studying the histories of NASA projects that had extreme cost growth, we have not found any evidence to disprove that hypothesis. In fact, what we found confirms that extreme cost growth is a failure of organizational leadership. The specific technical or programmatic factors leading to extreme cost growth (such as overestimating technology readiness or underestimating cost) have, at their root, a management culture that does not want or value independent cost and technical assessments as part of their decision-making process.

Alternative assessments by independent analysts is key to counteracting optimism and other biases inherent in successful project management. Ideally, these assessments are done by an organization that is independent of the advocacy organization. Application of the Joint Confidence Level (JCL) analysis to NASA projects, by independent review teams, appears to be a forcing function that is leading to good management practices. While it is too early to definitively say that JCL is a success, results so far are promising.

Randomness hinders our ability to identify *a priori* which projects will and will not experience extreme cost growth. Randomness also limits the ability of managers to assume their way to lower cost. Probability trumps determinism. Early research into a predictive model for identifying which projects will experience extreme cost growth produced mixed results. Better data, or more likely, a better way to model the conditions surrounding the formulation of a project, may ultimately bear fruit.

Finally, do not expect perfection, either in yourself or in your senior leadership. One of the great truths of life is that most of us are doing the best we can with what we got. Encourage good behavior when you see it. Stand up and be counted when you see harm being done. It takes great courage to point out that the emperor has no clothes. May we all be so courageous.

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

Hubbard, Douglas W., *The Failure of Risk Management,* New Jersey: John Wiley & Sons, 2009

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, "Being Certain about Uncertainty, Part 1," *Proceedings of the 2017 International Cost Estimating and Analysis Association Professional Development and Training Workshop,* Portland, OR, June, 2017

Prince, Frank, "The Dangers of Parametrics," *Proceedings of the 2016 International Cost Estimating and Analysis Association Professional Development and Training Workshop,* Atlanta, GA, 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, CA, 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, "Exploring the Limits of 'Faster, Better, Cheaper' with Mission Cost Risk Assessment," presented at the NASA Cost Symposium, Kennedy Space Center, November 2002.

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, *Antifragile, Things that Gain from Disorder,* New York: Random House, 2012

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

Appendix A
How heavy-tailed are cost risk distributions?

The financial analyst and writer Nassim Taleb calls phenomena that follow a normal distribution as belonging to "Mediocristan," which means that these phenomena don't vary much around the mean or median. An alternative to Mediocristan is what Taleb terms "Extremistan," where there is extreme variation around the mean and median. A large number of phenomena fit in this category. Fluctuations is the stock market is a prime example, as is economic loss due to hurricanes (just think of the damage done by Hurricanes Harvey, Irma, and Maria in 2017). Income and net worth also belong to this category, just think of Bill Gate's net worth compared to your own. Book sales per author, populations of cities, and word usage in a vocabulary also belong in this category. (Taleb 2007)

If costs for defense and NASA programs belong to Mediocristan, it has significant policy and risk management implications. In a world in which costs have limited variation, they are easy to predict. There is also a pronounced portfolio effect, so it is easy to reduce the limited risk of a single program even further by combining these programs with other programs to achieve a portfolio effect. In such a world, we can fund projects to percentiles slightly above the mean and achieve high confidence levels for an entire organization.

However, as shown by one of the authors in previous papers (Smart 2009, 2012 (a) and (b)), there is more variation in cost than we would expect if costs follow a normal distribution, and there is no portfolio effect. Thus it is clear that historically cost risk does not belong to the low risk realm of Mediocristan.

There are two primary possibilities for cost risk in this case. Either cost risk follows a power law distribution, or it follows a lognormal distribution. The power law is a case where risk is the extreme fluctuations of Extremistan, while the lognormal represents a middle ground where the risks are more extreme than those of Mediocristan but they are not quite as wild as those of Extremistan.

A prominent example of a heavy-tailed distribution is the Pareto distribution, which is defined as

$$f(x) = \frac{\alpha \theta^\alpha}{(x + \theta)^{\alpha+1}}$$

for $x \geq 0$. The parameter α determines the degree of heaviness of the tail. When $\alpha \geq 2$ the first two moments (mean and standard deviation) are both finite. But when $1 \leq \alpha \leq 2$ only the mean is finite but the standard deviation is not, and when $0 \leq \alpha \leq 1$, neither the mean nor the standard deviation is finite. The cumulative distribution function is defined as

$$F(x) = 1 - \left(\frac{\theta}{x + \theta}\right)^\alpha$$

The lognormal distribution has finite mean and variance but has a heavier right tail than a normal distribution. Like a lognormal, cost cannot be less than zero. And cost uncertainty is skewed – there are more risks for growth than there are opportunities for cost savings, just like with a lognormal. The normal on the other hand is symmetric.

The probability density function of a lognormally distributed random variable **X** is defined as

$$f_X(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{(lnx-\mu)^2}{2\sigma^2}}$$

where μ and σ are mean and standard deviation, respectively, of the log transform of **X**, i.e., **ln(X)**.

The relationships between the log-space mean and standard deviation and their unit space counterparts, i.e., **E[X]** and **Var[X]**, are given by

$$E[X] = e^{\mu+\frac{\sigma^2}{2}}$$

$$Var[X] = \left(e^{\sigma^2} - 1\right)e^{2\mu+\sigma^2}$$

and

$$\mu = ln\left(\frac{E[X]^2}{\sqrt{E[X]^2 + Var[X]}}\right)$$

$$\sigma = \sqrt{ln\left(1 + \frac{Var[X]}{E[X]^2}\right)}$$

The p[th] percentile of a lognormal distribution, denoted by α_p, is defined as

$$\alpha_p = e^{\mu+\phi^{-1}(p)\sigma}$$

Where $\phi^{-1}(p)$ denotes the standard normal distribution evaluated at **p**.

While normally distributed random variable has no lower or upper bound, a lognormally distributed random variable has no upper bound and is bounded below by zero.

However, cost will not decrease from its initial starting point to a level that is close to zero, but that is possible and even probably with a two-parameter lognormal. Once cost is established there is a limited amount of variation to the downside. To account for this we can model cost risk as a three-parameter lognormal. The third parameter is a location parameter that fixes the minimum of the distribution to a fixed positive value.

As discussed in Smart (2018), cost risk better fits a three-parameter lognormal than a two-parameter lognormal.

We first examine the case for the lognormal. The theoretical argument for the lognormal is that changes in costs over time are proportional to prior costs. This makes sense. Cost is more likely to increase than decrease over time, as evidenced by numerous studies on cost growth that show that over 80% of government projects experience cost growth, and on average increase by over 50% (Smart 2015). So when we talk about cost changes, we almost always mean cost increases. Cost increases often do not result in funding increases in the short term due to funding constraints. Thus cost increases will result in longer schedules. Longer schedules imply a longer period in which the personnel devoted to a project will charge to that particular project. Larger projects have more personnel assigned to a project, meaning that increases in cost will result in a proportional increase in cost.

Mathematically the change in cost from time **t-1** to time **t** can be represented as

$$X_t - X_{t-1} = \epsilon_t X_{t-1}$$

where the $\epsilon_t's$ are mutually independent and independent of X_{t-1}. Rearranging, we have that

$$\frac{X_t - X_{t-1}}{X_{t-1}} = \epsilon_t.$$

Summing over **t** we find that

$$\sum_{t=1}^{n} \frac{X_t - X_{t-1}}{X_{t-1}} = \sum_{t=1}^{n} \epsilon_t.$$

Proportional changes can be approximated as

$$\sum_{t=1}^{n} \frac{X_t - X_{t-1}}{X_{t-1}} \approx \int_{X_0}^{X_n} \frac{dX}{X} = ln(X_n) - ln(X_0)$$

Thus

$$ln(X_n) - ln(X_0) \approx \sum_{t=1}^{n} \epsilon_t$$

Rearranging terms we find that

$$ln(X_n) \approx ln(X_0) + \sum_{t=1}^{n} \epsilon_t$$

According to the Central Limit Theorem the sum of many random variables is normally distributed. Thus for large values of **n**, **ln(X_n)** is normally distributed. Thus by definition **X_n** is lognormally distributed.

However, only a slight change to this generative model is needed to turn the lognormal into a power law. An additional assumption that cost has a lower bound is all that is needed to change

the distribution to a power law. (Mitzenmacher 2003). This is a reasonable assumption since once a contract is signed, it is not likely for costs to decrease at all.

Empirically, there is evidence for both power laws and lognormal distributions. Based on a sample of 289 data points, Smart (2015) found that the lognormal tail was a better fit than a power-law type distribution. See Figure 1.
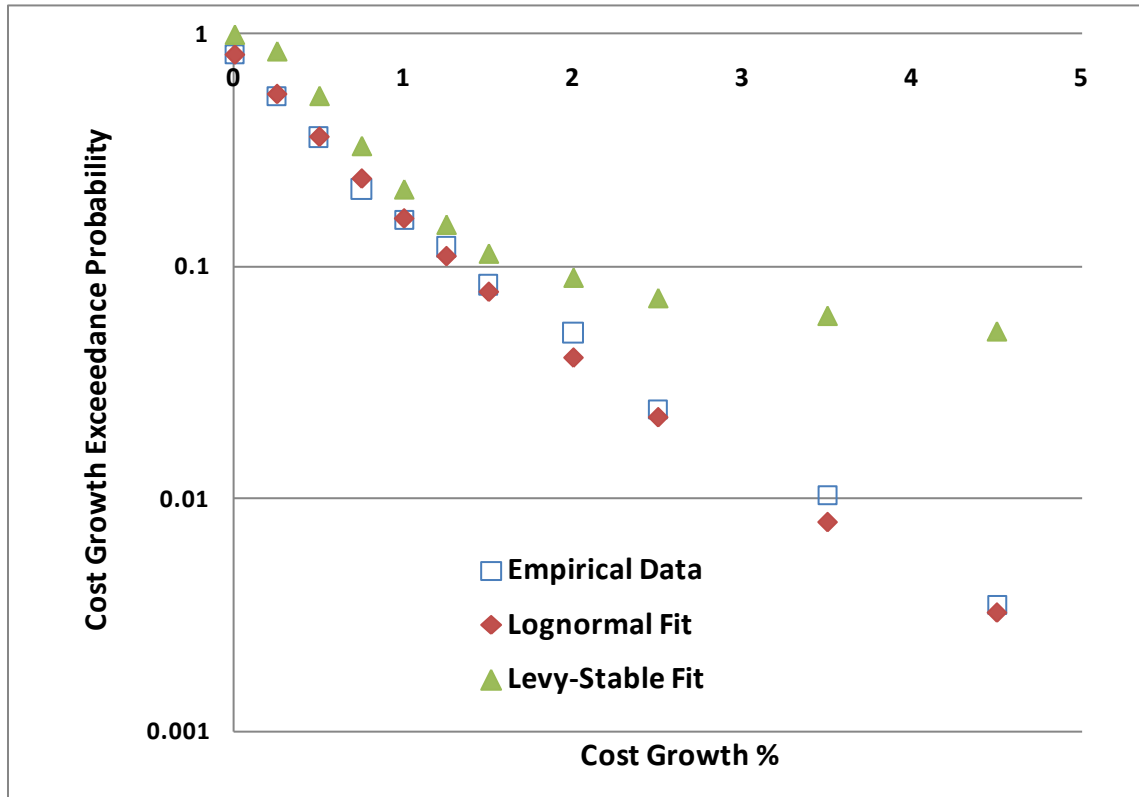
Figure 1. A lognormal is a better fit for the tail than a power-law distribution. Source: Smart 2015

Figure 1 displays the results of a lognormal fit to the right tail in comparison with a type of heavy-tailed distribution called the Levy-Stable. The Levy Stable distribution has a right tail governed by power laws.

How well does the lognormal fit the 133 NASA data points that is considered in this study? See Figure 2.
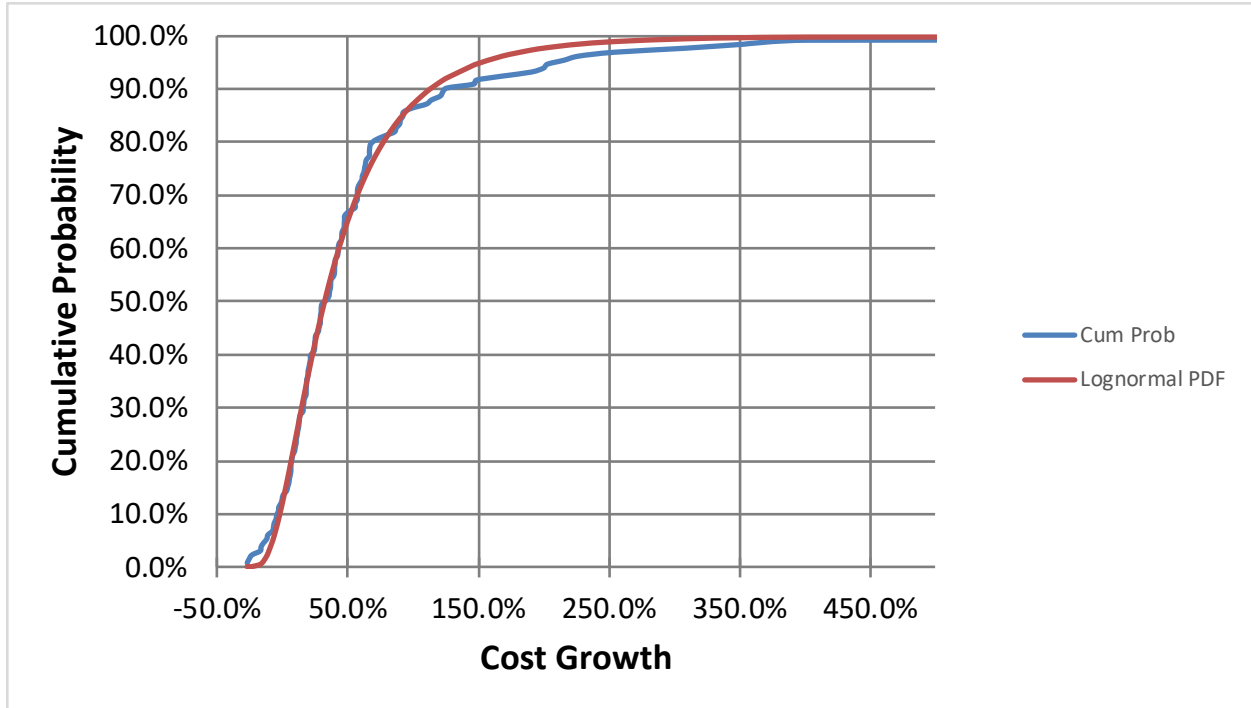
Figure 2. Lognormal fits the body of the distribution well, but underestimates the right tail.

In order to provide a better fit in the tail, we fit a lognormal distribution up to the 80th percentile, and a Pareto above the 80th percentile. This provides a better fit for the tail while still modeling the body of the distribution well. See Figure 3 for a graphical comparison.
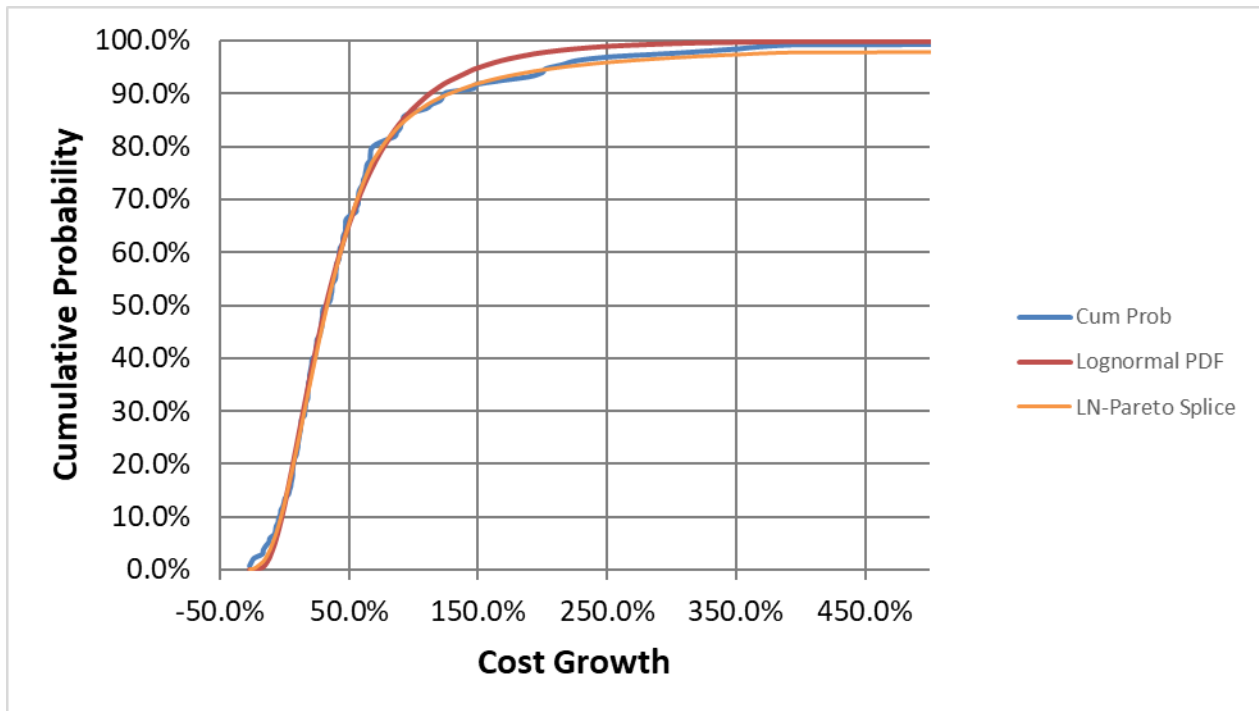
Figure 3. Lognormal-Pareto splice provides a better tail fit while still modeling the body of the distribution well.

Thus the lognormal was a better fit for the bulk of the distribution (up to the 80^h percentile), but a Pareto distribution better fit the right tail, indicating the extreme risks follow a power law distribution. In this case, the alpha parameter is between 1 and 3, indicating finite mean but an infinite variance. This model puts cost growth decidedly in the territory of Extremistan. This helps explains why we as cost estimators have such a challenge with providing accurate estimates – in the case of extreme risk, point estimates are not useful from a predictive standpoint.

See Table 1 for a comparison of the odds of cost growth greater than a variety of large percentages. This table compares the Pareto tail from Figure 3 with the lognormal tail from Figure 2.

Growth Greater Than:	Lognormal Tail	Pareto Tail
100%	1 in 8	1 in 7
250%	1 in 94	1 in 25
500%	1 in 1,800	1 in 63
1,000%	1 in 100,000	1 in 156
2,000%	1 in 15 million	1` in 393

Table 1. Comparison of the Lognormal and Pareto tails for modeling cost growth.

As is evident from the table, the odds of extreme cost growth is very small for the lognormal. The odds that a program would experience 1,000% cost growth for example, is practically impossible for a lognormal, while for a Pareto the odds are only a little less than 1%. The Pareto may be a little conservative, but the lognormal does not provide realistic odds of extreme growth. For example, the James Webb Space Telescope is one program that is not included in the 133 data points. This next generation space telescope is a successor to Hubble and truly cutting edge technology, just as was Hubble when it was launched. Like Hubble, it has also experienced cost growth of approximately 5.5% to date, more than any program in the database we have analyzed in this paper. The lognormal predicts the odds of this amount of cost growth to be 1 in 2,940, while the odds with a Pareto are much higher, 1 in 70. If we were to add this data point to the database it would be 1 of 134 data points, so the empirical probability would be between that predicted by the Pareto and lognormal, but closer to the Pareto than the lognormal.

The case for the lognormal is that programs that experience extreme cost growth will likely be cancelled. There are many such instances in history. The counterargument is that in many instances, even when a program experiences extreme cost growth it is such a high priority that even though funding constraints may limit the cost spent in one year, the schedule will slip, and resource will be re-allocated from other lower priority programs in order to pay for the

27

program. An example of this is the James Webb Space Telescope. Other notable examples are the Apollo program, and National Missile Defense, which the George W. Bush administration poured funding into to meet a tight deadline. The case for the lognormal is also buttressed by the fact that the wild risks of the stock market and natural disasters cannot be easily mitigated while government projects are more easily controlled. Scope can be removed in many cases, and fixes can be made to improve project performance. The counterargument again is that in many cases, the technology is so cutting edge that little can be done to reduce costs – it will cost what it costs.

In the end I believe that the answer to the question "does cost risk follow a lognormal or a Pareto distribution, at least in the tails?" is, in so many cases, "it depends." My hypothesis is that if a program is not a high priority, then it will be cancelled if costs spiral out of control, and that costs can be reined in as long as the technology is not too cutting edge. Otherwise we are looking at a power law case. See Table 2.

	Technology is within State of Practice	Technology is Cutting Edge
Normal Priority	Lognormal	Lognormal
Extremely High Priority	Lognormal	Pareto

Table 2. Cases when extreme cost risk is a lognormal vs. a Pareto.

If we can either cancel a program or rein in the cost performance with mitigation, then cost risk follows a lognormal. Otherwise, we are looking at a truly wild risk, Taleb-type Extremistan risk for a program.

References:

Mitzenmacher, M., "A Brief History of Generative Models for Power Law and Lognormal Distributions," Internet Mathematics, 2003.

Smart, C.B., "The Portfolio Effect and the Free Lunch," presented at the International Society of Parametric Analysts and Society of Cost Estimating and Analysis Joint Annual Conference, St. Louis, June, 2009.

Smart, C.B., "The Fractal Nature of Cost Risk," Journal of Cost Analysis and Parametrics, 2012 (a).

Smart, C.B., "Here There Be Dragons: Incorporating the Right Tail in Risk Management," Journal of Cost Analysis and Parametrics, 2012 (b).

Smart, C.B., "Covered with Oil: Incorporating Realism in Cost Risk Analysis," Journal of Cost Analysis and Parametrics, 2015.

Smart, C.B., "Enhancing Risk Calibration Methods," presented at the International Cost Estimating and Analysis Association Conference, Phoenix, June, 2018.

Taleb, N.N., *The Black Swan:The Impact of the Highly Improbable*, Random House, New York, 2007.

Appendix B
Data Sources for NASA Cost Growth Study

The following table provides the source material for Exhibit 4, Summary of NASA Cost Growth Studies.

Sources:
GAO Report: Need for improved reporting & Cost Estimating on Major Unmanned satellite projects (NASA)
GAO Report: Financial Status of Major Federal Acquisitions
GAO Report to Congress March 1973 Cost Growth in Major Weapons Systems
Rand Report: Acquisition Policy Effectiveness October 1979
An Analysis of DOD/NASA Cost Growth Profiles for the Congressional Committee of Gov't operations January 1980
NASA Project Management Study January 1981
Office of Comptroller: New Project Estimates Study August 1985
Office of Comptroller: Lessons Learned on Cost/Schedule June 1990
NASA Program/Project Planning Study November 1992
NASA Cost Growth: A look at recent performance January 2004
GAO Work on DOD Space Acquisitions Dec 2006
GAO Report: NASA: Long Term Commitment to and Investment in Space Exploration July 2006
GAO Report: NASA: Lack of Disciplined Cost-Estimating Processes Hinders Effective Program Management May 2004