

NORTHROP GRUMMAN

DEFINING THE FUTURE

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

Presented at SCEA, June 2007






The Society of Cost Estimating and Analysis

**Eric R. Druker, Richard L. Coleman, Peter J.
Braxton, Christopher J. Leonetti**

Northrop Grumman Corporation

Motivation

- All risk analysis methodologies have their origins in mathematics
 - In many situations however, the practitioners of the analysis come from non-mathematical backgrounds
 - This can lead to methodologies that may have sound basis being applied incorrectly (albeit innocently) due to a lack of understanding of their underpinnings
- The purpose of this paper is to shed light on some of the common mistakes in the execution of risk analysis
 - It will also try to explain the math behind these mistakes and the mischief they can cause
- This paper is not intended to be, nor could it ever be, all-inclusive, but will discuss what seems to be the right mix of common and serious errors in the experience of the writers
- We have chosen to classify these mistakes in to three categories
 -  1. **Green Light – Small errors that will only have an effect on the analysis and will generally not give management a false impression of risk**
 -  2. **Yellow light – Larger errors that in certain situations could have a major effect on the analysis and have the potential to give management a false impression of risk**
 -  3. **Red Light – Errors that will always have a major effect on the analysis and/or will give management a false impression of risk**

Topics

- **Risk identification and quantification**
 - Continuous vs. discrete risks
 - Evaluating “below-the-line” (usually “cost on cost”) risks
 - Combining Triangular Risks
 - Understanding “confidence” in estimates
- **Risk Modeling**
 - Monte Carlo vs. Method of Moments
 - Modeling mutually exclusive events
 - Assessing Cost Estimating Variability
 - Breaking risks into categories
- **Somewhat related thought experiment**
 - The assumption of an underlying log-normal distribution
- **Conclusions**

Risk Identification and Quantification: Continuous vs. Discrete Risks



- **Although many risk methodologies account for both discrete and continuous risks, some analysts try to squeeze all of their risks into one of the two categories**
- **Pros:**
 - It's easier to model risks from the same family of distributions
 - It's easier to present risks to management when they all come from the same family
- **Cons:**
 - Unfortunately, rare is the case that risks can be properly categorized using one family of distributions
 - Improper categorizations cause distortions in risks, usually in their variation, less often in their mean
 - Unfortunately, variation is key to what is desired from risk analysis; it conveys a sense of the worst and best cases
 - Using only one family of distributions can thus lead to misguided management decisions brought on by a poor characterization of risk

Risk Identification and Quantification: Continuous vs. Discrete Risks



Continuous Distributions

- Continuous risks account for events where there is a range of possibilities for the cost impacts
- Example risks that tend to be continuous:
 - Below-the-line risks with estimates made using factors or regression
- Can be characterized by any number of distributions
 - Triangular, normal, and log-normal are three of the most common
- Characterizing continuous risks as discrete events causes these problems:
 - Gives management the false idea that we can totally eliminate a risk
 - Leaves out information that can show the opportunity side of the risk (if one exists)

Discrete Distributions

- Discrete distributions account for specific events with point estimates for their cost impacts
- Example risks that tend to be discrete:
 - Technical/Schedule Risks due to specific events
- Universally characterized as a Bernoulli or multi-valued discrete event, described by probability(ies) and cost impact(s)
- Characterizing a discrete event risk as continuous causes these problems:
 - Gives management the impression that they cannot avoid the risk and
 - Can show an opportunity where one does not exist

Choose the characterization of risks carefully, it makes a big difference!

Risk Identification and Quantification: Evaluating Below-the-Line Risks



- One of the most common mistakes we see is in the handling of below-the-line risks such as factors and rates
- Generally, one of two errors occurs
 - Applying the rate or factor risk to the non-risk-adjusted estimate
 - Using a discrete distribution to categorize this continuous risk
- To perform the analysis correctly, the distribution around the rate or factor must be found
- The next step is to apply this distribution to the risk-adjusted estimate
 - Application of a risk to a risk-adjusted estimate is called ‘functional correlation¹’
- The next page will show how these two errors can affect the results of the analysis

1. An Overview of Correlation and Functional Dependencies in Cost Risk and Uncertainty Analysis, R. L. Coleman, S. S. Gupta, DoDCAS 1994


NORTHROP GRUMMAN

Risk Identification and Quantification: Evaluating Below-The-Line Risks



Assumptions			
Labor Point Estimate	\$	1,000,000	
Overhead Rate		8%	
Overhead Estimate	\$	80,000	
		Mean	St Dev
Historic Overhead Rate		10%	2%
		Mean	St Dev
Risk Adjusted Labor Estimate	\$	1,250,000	\$ 250,000

Outcome			
		Mean	St Dev
Bernoulli*	\$	15,000	\$ 6,495
		*Assumed pf of .75	
		Mean	St Dev
Normal (applied to non risk-adjusted estimate)	\$	20,000	\$ 20,000
		Mean	St Dev
Normal (applied to risk-adjusted estimate)	\$	44,974	\$ 35,771
*Approximated using Monte Carlo Simulation			





	Bernoulli	Normal 1	Normal 2
10%	\$ -	\$ (5,631)	\$ (868)
20%	\$ -	\$ 3,168	\$ 14,868
30%	\$ -	\$ 9,512	\$ 26,216
40%	\$ -	\$ 14,933	\$ 35,912
50%	\$ -	\$ 20,000	\$ 44,974
60%	\$ -	\$ 25,067	\$ 54,037
70%	\$ -	\$ 30,488	\$ 63,733
80%	\$ 20,000	\$ 36,832	\$ 75,080
90%	\$ 20,000	\$ 45,631	\$ 90,817

Risk Identification and Quantification: Combining Triangular Risks



- When developing a risk distribution for a portion of an estimate, analysts sometimes collect information on distributions at a lower level, and roll them up to obtain the risk distribution for the level where they are performing their analysis
- One of the mistakes we have seen is with triangular distributions for the lower levels of an estimate
 - Some analysts add the min/mode/max together to get the top level distribution
 - This incorrectly adds weight to the tails of the top level distribution
 - Percentiles and extrema do not add, only means add
- If possible, the lower level distributions should be run through a simulation to obtain the upper level distribution

Risk Identification and Quantification: Combining Triangular Risks



	Assumed Distribution		
	Min	Mode	Max
Distribution 1	-100	0	100
Distribution 2	-100	0	100

	Percentiles								
	10%	20%	30%	40%	50%	60%	70%	80%	90%
Summing Triangular Points	\$ (111)	\$ (74)	\$ (45)	\$ (21)	\$ 0	\$ 21	\$ 45	\$ 74	\$ 111
Actual Distribution	\$ (75)	\$ (50)	\$ (32)	\$ (16)	\$ 0	\$ 15	\$ 31	\$ 50	\$ 75
Difference	\$ (35)	\$ (24)	\$ (13)	\$ (5)	\$ 0	\$ 6	\$ 14	\$ 23	\$ 35



Summing Triangular Points
Actual Distribution
Difference

Risk Identification and Quantification: Understanding “Confidence”



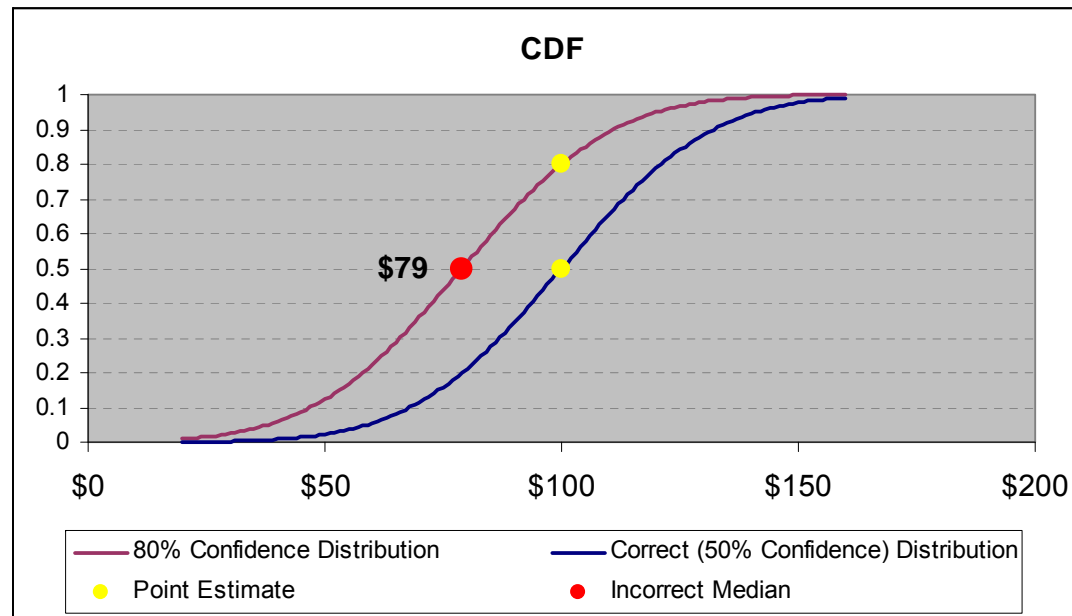
- **Some of the methodologies we see rely on an input of “confidence” in order to ultimately produce a distribution around the point estimate**
- **The problem lies in a simple breakdown of understanding somewhere in the chain between methodology developer and cost estimator**
- **What these models are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

Risk Identification and Quantification: Understanding “Confidence”



- There are three problems in using the approach of specifying confidence as an input that make it inherently dangerous
 1. It requires both the risk analyst and the estimator being interviewed to have a considerable level of statistical sophistication
 2. In the case where the risk analysis is being performed by an independent observer, it requires them to look deeper than the BOEs to obtain true confidence
 - Example: When BOEs are written to a target, the desired confidence should come from the method used to develop the target cost, not the justification used to support it
 3. In cases where actual risks do not constitute a large percentage of the total estimate, these “confidences in the estimate” can drive the entire analysis
 - The impact of this misunderstanding on the results of this analysis can be **substantial**

Risk Identification and Quantification: Understanding "Confidence"



Assuming 80% Confidence
Actual Distribution
Difference

		Percentiles								
		10%	20%	30%	40%	50%	60%	70%	80%	90%
	\$	47	58	66	73	79	85	92	100	111
	\$	68	79	87	94	100	106	113	121	132
	\$	(21)	(21)	(21)	(21)	(21)	(21)	(21)	(21)	(21)

- This methodology assumes a normal curve used to model the distribution around the point estimate
- The above analysis shows the effect of an analyst using 80% confidence where a 50% confidence is appropriate
- Management would receive two very wrong messages
 1. That the estimate has been created at an 80% confidence level
 2. That the 50th percentile for the actual costs will be much lower than the point estimate

Risk Modeling

- Now that we've discussed how to properly develop risks, it's time to look at how they are compiled into results for presentation to management
- There are two main ways of calculating the combined effects of a large number of risks
 - A Method of Moments Model
 - A Monte Carlo Simulation
- Both methods work equally well when applied correctly
- What follows is a quick summary of how each method works as well as the pros and cons of each



Risk Modeling: **Monte Carlo** vs. Method of Moments

- Monte Carlo arrives at the distribution of the combined effects of risks by simulating multiple, independent “runs of the contract” and portraying the range of outcomes
- Pros:
 - Most common approach
 - Will be understood by the largest audience
 - More intuitive than method of moments
 - Makes fewer assumptions than method of moments
- Cons:
 - Very difficult to apply correlation correctly
 - The output correlation matrix will rarely match the input correlation when multiple families of distributions are used
 - Can be time consuming/require somewhat heavy computing power
 - Thousands of runs are needed to converge to the actual distribution
 - Fewer runs are needed for the mean and 50th %-ile (a few hundred should do), progressively more runs for %-iles further out in the tails

Risk Modeling: Monte Carlo vs. **Method of Moments**



- **Method of Moments arrives at the distribution of the combined effects of risks by relying on the central limit theorem (C. L. T.)**
 - The C. L. T. proves that a sufficiently large number of risks will eventually combine to a parent distribution (generally normal) whose moments match the combined moments of the child distributions
- **Pros:**
 - Very easy to use correlation
 - Assuming all distributions are normal allows random number draws from a normal random variable
 - Less computing power required
 - No simulation is needed since the mean, standard deviation and %-iles of the overall distribution are deterministic
- **Cons:**
 - Non-Intuitive
 - Understanding the moments of random variables requires considerable statistical sophistication
 - “Why is a Bernoulli risk being converted to a normal distribution?”
 - Makes several potentially dangerous assumptions
 - Assuming normality = assuming no skew in overall distribution
 - Assumes risks converge to C.L.T.
 - C. L. T. assumes there are many distributions all of which are independent and identically distributed
 - This is often not the case with risk registers

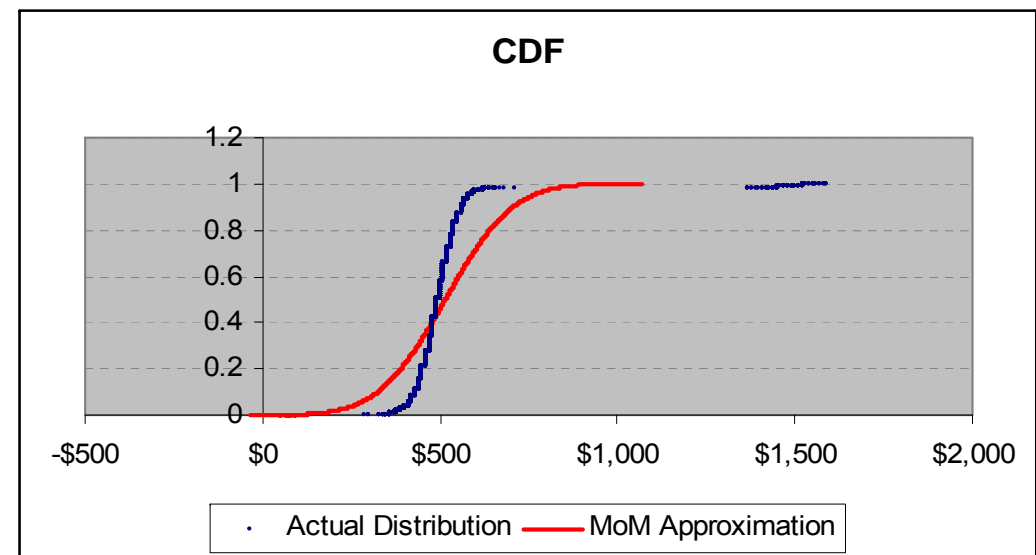


Risk Modeling : Monte Carlo vs. Method of Moments

- One very dangerous situation when using a Method of Moments technique occurs when there is a risk (or series of risks) that skew the distribution
 - This occurs when the risks in the risk register do not satisfy the Lyapunov condition
 - In cases like this, the Method of Moments will give management inaccurate total %-iles of risk
- This calls the viability of Method of Moments into question as a risk tool because:
 - This mistake cannot be caught without running a Monte Carlo simulation on the risk register and comparing the outputs to Method of Moments
 - At which point why use Method of Moments in the first place?
 - Without a math background, risk practitioners will be unaware that this mistake has occurred
- Below is an example of a risk register (exaggerated for clarity) that causes a skewed result
 - 99 risks with Pf of .5 and Cf of 10
 - 1 risk with Pf of .02 and Cf of 1000

	Actual	MoM
Mean	515	515
Standard Deviation	148.6	148.6

	MoM	Actual	Diff
10%	324.6	430.0	-105.4
20%	390.0	450.0	-60.0
30%	437.1	470.0	-32.9
40%	477.4	480.0	-2.6
50%	515.0	490.0	25.0
60%	552.6	510.0	42.6
70%	592.9	520.0	72.9
80%	640.0	540.0	100.0
90%	705.4	560.0	145.4



Risk Modeling : Assessing Cost Estimating Variability



- Risks and opportunities shift the mean of the S-Curve and contribute to its spread, but no risk analysis is complete without an assessment of cost estimating variability
- In other words; ignoring risk, how much error exists in the cost estimate?
- As discussed previously, data-driven estimates often contain the information needed to assess this variability
- In cases where data is not available, such as estimates made using engineering judgment, it is not uncommon to see variability assessed through an interview with the estimator
 - This variability is generally evaluated at the estimate level using a normal or triangular distribution around the point estimate
- In the following slides, we will:
 - Give an example of assessing cost estimating variability for data driven estimates
 - Show the danger of assessing cost estimating variability at too low a level when data is not available

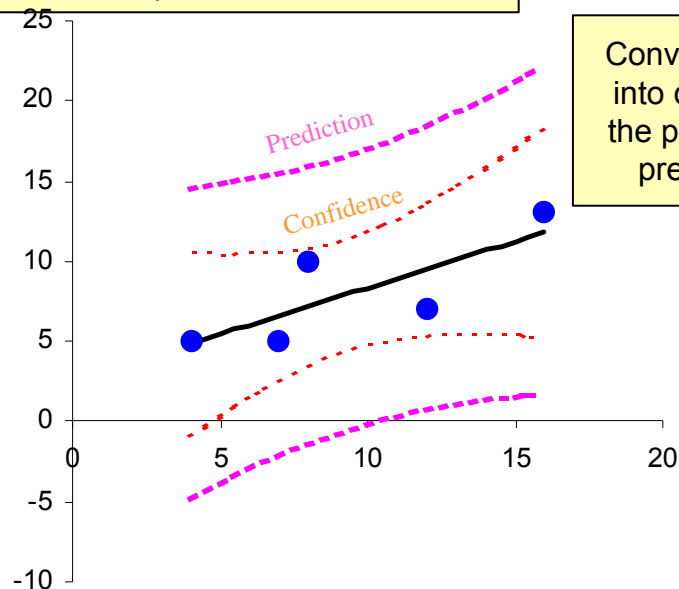
Risk Modeling : Assessing Cost Estimating Variability



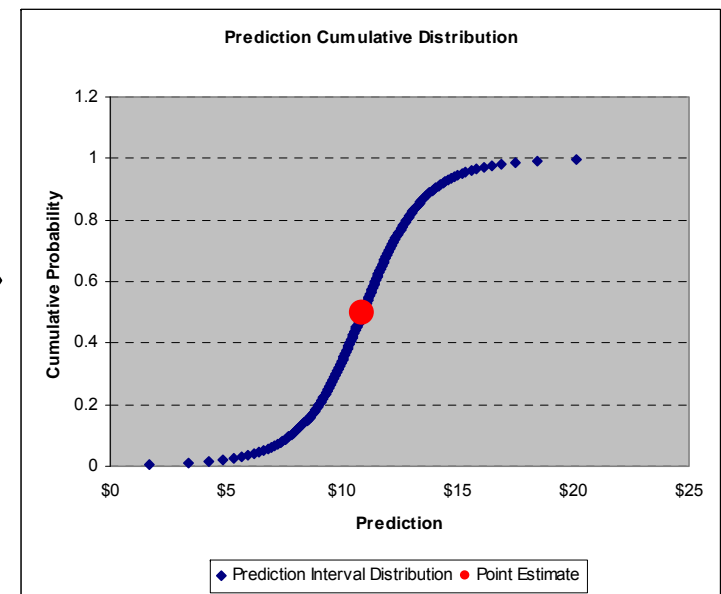
- For data-driven estimates, cost estimating variability is often a direct product of the analysis needed to produce the estimate
- When estimating using CERs, the prediction interval can be used to assess cost estimate variability
 - The distribution that is uncovered can then be placed into the Monte Carlo simulation

Regression Reminder:

Confidence Intervals give bands for the mean value of a prediction,
Prediction intervals for the value of the prediction itself



Convert Prediction Interval into distribution by finding the prediction band for the prediction at all %-iles



Risk Modeling : Assessing Cost Estimating Variability



- When data is unavailable, interviews are often conducted to assess estimating variability
 - The outcome of this is generally a triangular or normal distribution around the point estimate
- Assessing this variability at too low a level is one of the pitfalls that can occur using this method
 - While the analyst may believe they are achieving a greater level of granularity, their practice is artificially removing variability from the estimate
 - In general, for similar distributions, CV decreases by a factor of $1/\sqrt{\text{(number of distributions)}}$
 - Correlation can mitigate this, but only to a certain extent
 - As a separate issue, it is in doubt whether estimators can accurately assess cost estimating variability at low levels
 - It is likely that they are applying their perception of top level variation to the lower level estimates

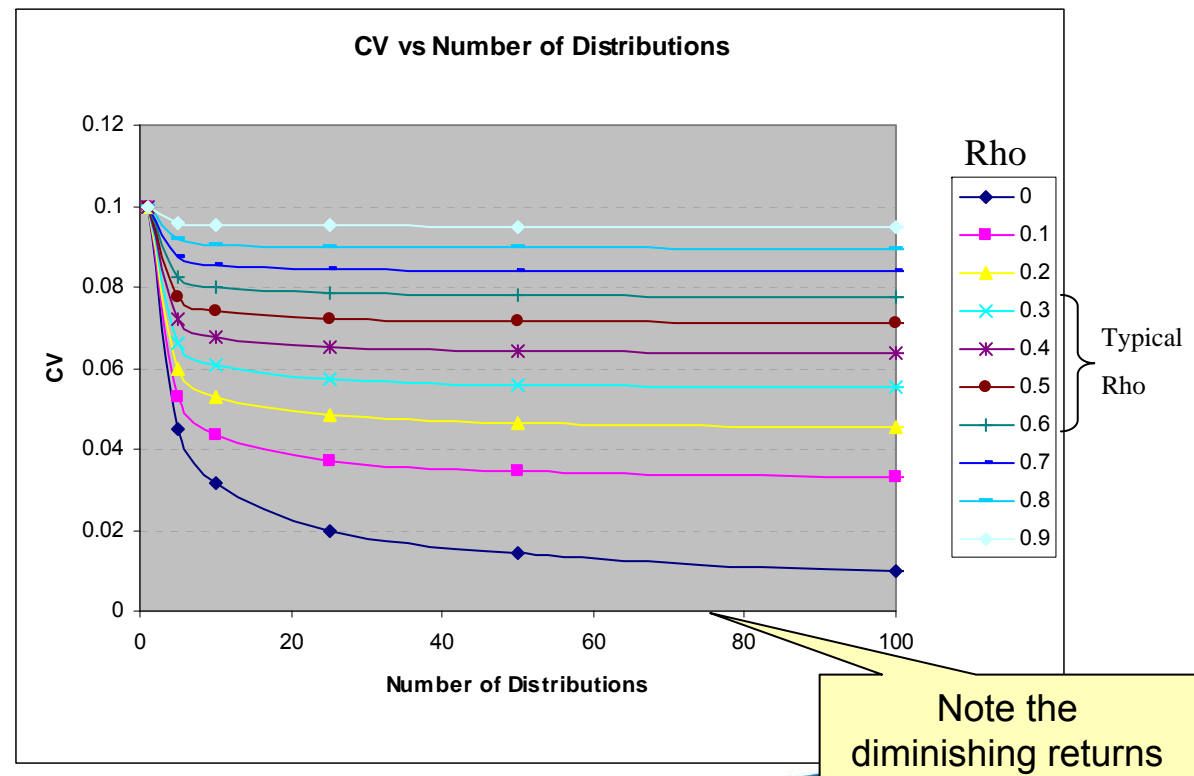
Assumptions

All Distributions are $N(10,100)$ (CV of 10%)

CV shown on graph is CV of the sum of the distributions

Note

With a ρ of 0.0 (no correlation), CV of the sum of the distributions = $10\% \times \sqrt{\text{(number of distributions being summed)}}$





Risk Modeling: Modeling Mutually Exclusive Events

- Sometimes, risk practitioners are faced with two outcomes for a risk
- Most of the times, these are meant to be mutually exclusive events
- Consider a risk with two possibilities:
 - A 20% chance of a \$20,000 risk
 - A 20% chance of a \$10,000 opportunity
- Modeled as two line items without taking into account exclusivity, the risk is actually categorized as such:
 - A 16% chance of a \$20,000 risk
 - 20% chance of \$20,000 risk x 80% chance of no opportunity
 - A 16% chance of a \$10,000 opportunity
 - 20% chance of \$10,000 opportunity x 80% chance of no risk
 - A 64% chance that nothing happens
 - 80% chance of no opportunity x 80% chance of no risk
 - A 4% chance of a \$10,000 risk
 - 20% chance of \$10,000 opportunity x 20% chance of \$20,000 risk
- **Although this does not change the expected value of the item, it does change the standard deviation**
 - Modeled as exclusive events the standard deviation is \$9,797
 - Modeled as above the standard deviation is \$8,944
- **Repeated enough times, this mistake will lead to incorrect percentiles of the overall risk distribution**





Risk Modeling : Breaking Risks into Categories

- **One of the biggest hurdles in presenting risk analysis results lies in the fact that subcategories of risk will never sum to the total**
- **Several methodologies contain processes for adjusting the results by category so that they sum to the total**
- **We believe that an understanding of why categories don't sum to the total can be given through a simple (and more importantly, quick) explanation**
 - We agree that in general, management does understand this fact; but, giving decision makers some of the basic tools needed to understand our analysis increases its usefulness to them
- **We will propose a simple way of presenting the information**



Risk Modeling : Breaking Risks into Categories

- **Example: The Dice Game:**
- **Suppose I have one die and roll once**
 - The probability of getting a 1 is 1/6 (There is an equal probability of landing on any side)
- **Now suppose that I have one die and roll twice**
 - What is the probability of having the total of two rolls equal 2?
 - The only way this can happen is if I roll a 1 twice
 - Probability of rolling 1 on first throw: 1/6 
 - Probability of rolling 1 on second throw: 1/6 
 - Because each roll is independent, the probability of the rolls summing to 2 is $(1/6) \times (1/6) = 1/36$
- **This is the same logic that needs to be applied to each category**
 - Assuming the categories are independent, the probability of having ALL worst case scenarios is close to zero!
 - This same logic applies to categories of risk
 - Percentiles will not add because the probability of having EVERYTHING (or most everything) go wrong (or right) is very small





Risk Modeling : Breaking Risks into Categories

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

- **Risk analysis is generally only a piece of the puzzle when decision makers receive a presentation on a program**
 - This generally leads to the risk assessment results being compressed onto a couple of slides
 - It is therefore critical that we present the information in a way that is both compressed and evocative
- **The above chart shows how categories can be presented along with the bottom line**
- **The point estimate is included for reference, along with the 20th/50th/80th percentiles**
- **Risk \$s and Risk %s (based on the 50th percentile) are shown off to the right**
 - This allows decision makers to see the risks from both important perspectives

A Thought Experiment: The Assumption of Log-Normality

- Many studies have asserted the CGF distribution across many DoD programs to be distributed log-normally, an example is Arena and Younossi¹
- A paper by Summerville and Coleman² presented a risk approach that recommended applying a normal distribution with a mean and standard deviation based on a weighted average risk score based on several objective measures
- Could it be possible that the log-normal distribution described in the Arena and Younossi paper is due to the risk scores from the Summerville and Coleman² paper being distributed log-normally?
- This would give the illusion of an underlying log-normal distribution when the actual distribution is normal with a mean and standard deviation dependent on the technical score
- We're not necessarily advocating dropping the umbrella log-normal assumption that is being used in many methods, especially when the technical score is unknown
- We present this as a thought experiment that could be expanded on at a later date

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

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

Conclusions

- One of the biggest problems with risk analysis is that it is impossible to catch all mistakes just by looking at %-iles or an S-Curve
- Catching mistakes requires looking at not just the models and their outputs, but the methods used to produce the inputs
 - We all know that **Garbage in = Garbage out**
 - We forget that **Good data into bad methods = Garbage out**
- Due to the mathematical knowledge required to catch many of these mistakes we advocate the vetting of all risk analysis performed within an organization with someone (or some group) who understands both the process and the math behind it
- Normally, a few days to a week is all that is needed to catch problems like the ones discussed in this paper
- Once problems have been caught, they can generally be quickly fixed in order to present the most accurate information available to management