Alternative Risk Measures for Determining Program Reserves

Louis Fussell

Abstract

In 2005, NASA began requiring projects to statistically sum the cost of project components and the duration of their activities to determine joint confidence levels for total project cost and duration. This sum is typically accomplished by means of a Monte Carlo simulation. The simulation is executed with a cost-loaded schedule model augmented with probabilistic distributions assigned to costs and durations based on project risks and uncertainties. NASA policy requires that project managers reserve budget equal to a 50% joint cost and schedule confidence level (JCL) and that the managing directorate hold reserve to a 70% joint confidence level, with some exception. The 50% and 70% joint confidence levels are quantile risk measures. This paper will discuss the limitations of quantile risk measures and propose the use of alternative risk measures, namely superquantiles, for determining project reserve levels. Monte Carlo simulations for several NASA projects were executed and a comparison of quantile and superquantile risk measures is presented.

Keywords

Joint Confidence Level, cost, schedule, risk, Monte Carlo simulation

1 Introduction

According to the National Space Council (2010), the National Aeronautics and Space Administration (NASA) required projects to sum their costs using Monte Carlo simulation or other statistical techniques beginning in 2005. NASA Procedural Requirement (NPR) 7120.5C, *NASA Program and Project Management Processes and Requirements* (NASA, 2005), required that "project estimates shall include reserves, along with the level of confidence provided by the reserves." This confidence level requirement was strengthened and expanded over the subsequent years. The current requirement is that certain projects are required to complete a joint cost and schedule confidence level (JCL) analysis prior to completing specific lifecycle reviews. NASA Mission Directorates are required to budget these projects at no less than "70 percent JCL" and ensure funding for the projects at no less than a "50 percent JCL," unless otherwise approved by the decision authority (NASA, 2021). Perrino (2015) provides a comprehensive history of NASA's JCL policy from onset to 2015.

JCL analysis is process that integrates a project's cost, schedule, and risk into a single model. Alternative methodologies for conducting a JCL analysis have been suggested (Butts & Linton, 2009; Garvey et al., 2016; Smart, 2009), however NASA intends for the JCL analyst to perform a Monte Carlo analysis using a probabilistic cost-loaded schedule (PCLS) model (NASA, 2015). The JCL analysis process requires the analyst to identify the goals of the analysis, build a summary schedule of the project, load the schedule activities with costs, incorporate risk and uncertainty, execute the Monte Carlo simulation, and view the results. Figure 1 is an example of a portion of a PCLS model, built using Oracle's Primavera Risk Analysis. On the left side of the figure is a table of activities, along with their duration and cost. Activities in red are risks. The percent values in the yellow lozenges represent the likelihood the risk will occur, and the blue bars indicate the duration uncertainty.

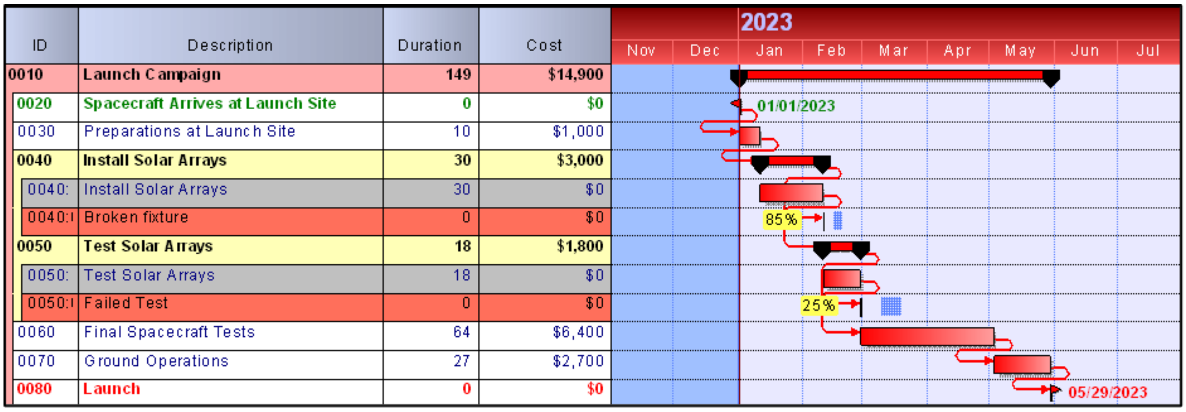

Figure 1. Example PCLS Model

Activity duration or cost uncertainty is modeled as a probabilistic distribution in the PCLS model. A risk is modeled as a mixed distribution: a binomial distribution describes the likelihood of the risk occurring and a second distribution describes the risk impact to duration and cost, conditional on the risk occurring. With each iteration of the Monte Carlo simulation, random numbers are generated which fit the distributions modeled and a total duration and cost is calculated for the project. If the Monte Carlo simulation executes *n* iterations, then it will generate *n* independent, identically distributed random samples from the joint distribution of the total project duration and cost. From these samples, the probability of possible project outcomes can be calculated.

When executed, the Monte Carlo simulation of the simple PCLS model (Figure 1) generates the total project duration and cost outcomes graphed in Figure 2. From the graph, one can see that duration and cost are positively correlated since cost increases as duration increases. This is due to a large portion of the project cost being time-dependent, primarily labor. Since duration and cost are correlated, NASA chooses to analyze their *joint* distribution rather than their separate duration and cost *marginal* distributions.

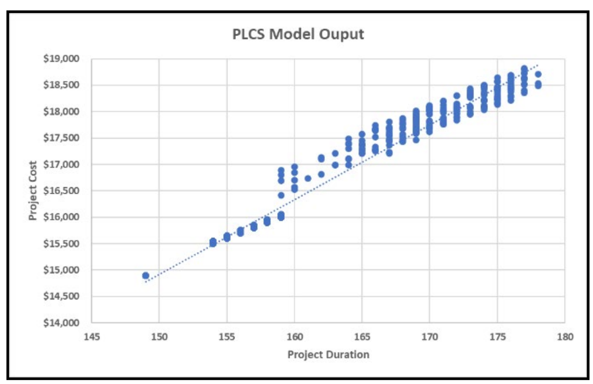

Figure 2. PCLS Output

2 Quantiles and Superquantiles Defined

In mathematical terms, a risk measure is a functional mapping of a distribution to the real numbers, $\mathbb{R}$. If the distribution is represented by a vector of random variables, $\mathbf{X}$, and ρ represents the risk measure functional, then $\rho(\mathbf{X}) \in \mathbb{R}$. The risk measure, ρ, is interpreted as the extra time or money that the project must *add* to the project estimate for the project to proceed "prudently" (Artzner et al, 1999). In other words, the risk measure indicates the amount of schedule or budget contingency required by the project to manage risk and uncertainty. Quantiles, such as NASA's JCL, and superquantiles are examples of risk measures.

Quantile and superquantiles are always specified with a given confidence level, α. As stated in the introduction, NASA specifies that JCL risk measures with α = 0.5 and α = 0.7 are to be calculated. Let *F(**X**)* represent the cumulative distribution function for a vector of random variables, **X**. In a univariate setting, i.e., if **X** is one-dimensional, the quantile, Q_α, is defined by

$$Q_\alpha(\boldsymbol{X}) \ := min[x \in \mathbb{R} \mid F(x) \geq \ \alpha].$$

In the univariate setting, the superquantile, $\bar{Q}_\alpha$, is defined by

$$\bar{Q}_\alpha(\boldsymbol{X}) \ := E[x \in \mathbb{R} \mid F(x) \geq \ \alpha].$$

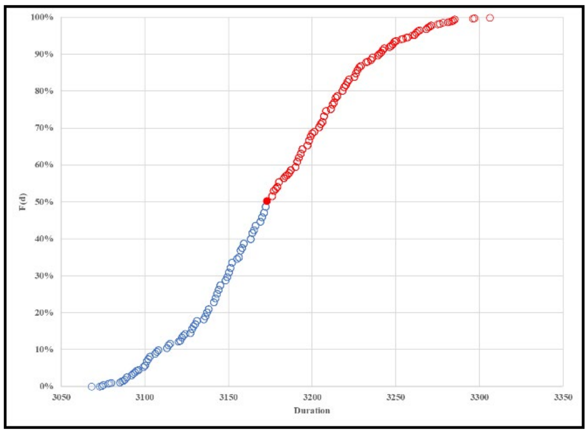

Figure 3. Plot of univariate quantile and superquantile data

Figure 3 depicts the relationship between the univariate quantile and superquantile graphically. For α = 0.5, the red circles are the outcomes wherein $F(d) \ \geq \ \alpha$, d is duration. The quantile, $Q_{0.5}(d)$, is the minimum of these outcomes and is shown in Figure 3 as a solid red circle. The superquantile, $\bar{Q}_{.5}(d)$, is the expected value, or average, of the outcomes represented by the red circles.

Defining quantiles and superquantiles in a multivariate setting is more difficult. In one-dimension, the outcomes are real numbers, which are totally ordered. This makes it simple to find the minimum value within a set of outcomes. In two-dimensions, the outcomes are vectors for duration and cost, (d, c), and are only partially ordered. For instance, given outcomes (d_1, c_1) and (d_2, c_2), (d_1, c_1) ≤ (d_2, c_2) if $d_1 \leq d_2$ and $c_1 \leq c_2$. However, if $d_1 \leq d_2$ and $c_1 > c_2$, the relationship between is (d_1, c_1) and (d_2, c_2) is not easily understood. The definition of quantile requires the identification of a minimum valued outcome. In two or more dimensions, the minimum valued outcome may not be unique. This has led to multiple attempts to define quantiles and superquantiles for multivariate distributions using various approaches (Serfling, 2002).

NASA uses a definition of multivariate quantile developed by Embrechts and Pucetti (2006) in the application of JCL. Since the focus here is on a bivariate quantile of duration and cost, we will refer to the bivariate quantile and superquantile as BQ_α and $\overline{BQ}_\alpha$, respectively. Recognizing that the quantile may not be unique, the multivariate quantile, BQ_α, is defined by

$$BQ_\alpha(\boldsymbol{X}) \ := \partial[X \in \mathbb{R}^k \mid F(X) \geq \ \alpha].$$

where ∂ symbolizes the boundary of the α-level sets for the random vector and $k > 1$ is the number of dimensions. Cousin and Di Bernadino (2014) go on to define the multivariate superquantile, $\bar{Q}_\alpha$, by

$$\overline{BQ}_\alpha(\boldsymbol{X}) \ := E[X \in \mathbb{R}^k \mid F(X) \geq \ \alpha].$$

Note: The definitions of multivariate quantiles and superquantiles offered by the authors mentioned are more complicated than stated to include cases where the distribution of the random vector is not exchangeable. Since the PCLS model produces independent, identically distributed outcomes, the outcomes are exchangeable. The simplified definitions are satisfactory in this context.

Figures 4 and 5 illustrate a quantile and superquantile in a bivariate setting germane to JCL analysis. In Figures 4 and 5, the outcomes satisfying the definition of quantile with $\alpha = 0.5$ are shown in red. Notice they are not unique. Figure 5 is this same distribution projected on to the (d, c)-plane. Where $F(d, c) \geq \alpha$ is shown by the red circles. The dashed curve represents the boundary of the α-level sets, and the solid red circles are the outcomes which lie on that boundary. The superquantile, $\overline{BQ}_{0.5}(d, c)$, is the expected value of the outcomes represented by the red circles

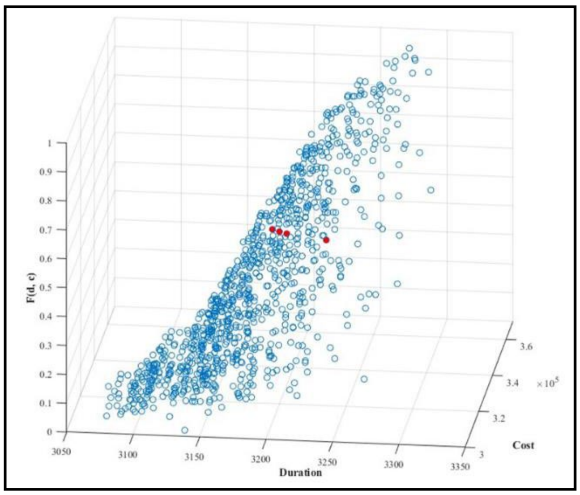

Figure 4. Three-dimensional plot of bivariate quantile and superquantile data.

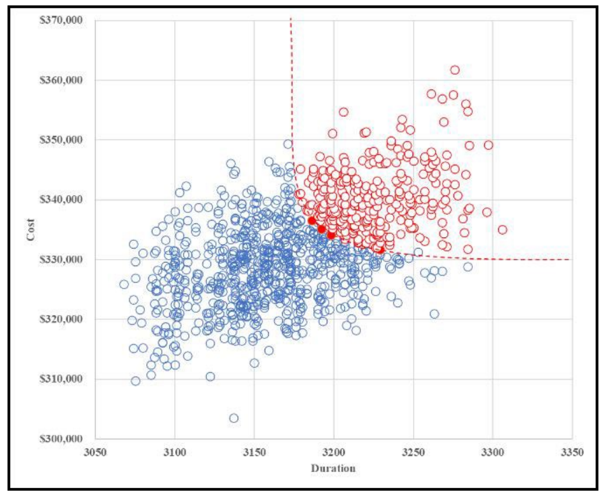

Figure 5. Plot of bivariate quantile and superquantile data projected onto the (d, c) plane.

3 Coherent Measures of Risk

3.1 Definition of Coherent Risk Measure

Artzner et al. (1999) devised a set of axioms that should apply to any risk measure used to manage risk. Risk measures that meet the axioms are deemed to be coherent risk measures. According to Artzner et al., a coherent risk measure, ρ, must satisfy the following four axioms:

1. Translation invariance: $\rho(X + c) = \rho(X) - c$. In the context of this research, translation invariance implies that adding a risk-free activity, c, to the project does not change the level of reserves required.
2. Monotonicity: If $X < Y$ for each scenario then $\rho(X) < \rho(Y)$. For example, when we execute the Monte Carlo simulation for the PCLS, if Activity X has a shorter duration than Activity Y in every iteration, then the risk of X must be less than the risk of Y.
3. Positive homogeneity: $\rho(cX) = c\rho(X)$. In the NASA context, this implies the cost risk of building two identical instruments ($c = 2$) is twice the cost risk of building one unit. Of course, this is only true if both instruments are built simultaneously without benefit of learning.

4. Sub-additivity: $\rho(X + Y) \leq \rho(X) + \rho(Y)$. This axiom is also referred to as the diversification principle. As an example, suppose an instrument is added to a spacecraft. The reserves needed for the project should not increase more than the additional reserves required to add that instrument.

3.2 Sub-additivity Example

Quantile risk measures satisfy the first three axions, however quantile risk measures are not coherent because they are not sub-additive (Artzner et al., 1999). JCL is a quantile risk measure and is not sub-additive. This is important because the reason why JCL is not sub-additive leads to one limitation and the impact of JCL not being sub-additive leads to another limitation. Due to this importance, an demonstration of the sub-additivity axiom is provided.

The PCLS model shown in Figure 1 models the following scenario:

- Due to previous technical issues, the project is required to install solar arrays on the spacecraft after it is delivered to the launch site.
- The solar arrays must be installed and then tested.
- There is a risk, Risk 1, that a fixture may be broken impacting installation. Risk 1 has a likelihood of occurrence of 85%. The impact to task duration is described with a uniform distribution of 5 to 10 days. The cost impact is described with a uniform distribution of \$100 to \$150.
- There is another risk, Risk 2, that the solar arrays may fail a test impacting testing. Risk 2 has a likelihood of occurrence of 25%. The impact to task duration is described with a uniform distribution of 10 to 20 days. The cost impact is described with a uniform distribution of \$500 to \$1000.
- The other activities in the launch campaign are risk-free.

The results of the simulation of the PCLS model are show in Figure 6. Duration and cost for activities without risks do not change. Risk 1 has an impact of 7 days and \$824. Risk 2 has no impact on total project outcome because its likelihood of occurrence, 25%, is less than $1 - \alpha$. At $\alpha = 0.7$, Activity 0050, Test Solar Arrays, shows no impacts from the risk of a failed test. All the impact of Risk 2 occurs in the tail of the distribution. The total impact to the project from the two risks is 10 days and \$1,127. $JCL_{0.7}$(Project) = (159, \$16,207) > (156, \$15,724) = $\Sigma JCL_{0.7}$(Activities), demonstrating that JCL is not sub-additive.

ID	Description	Deterministic Duration	Deterministic Cost	JCL Duration $\alpha = 0.7$	JCL Cost $\alpha = 0.7$
0010	Launch Campaign	149	\$14,900	159	\$16.027
0020	Spacecraft Arrives at Launch Site	0	\$0	0	\$0
0030	Preparations at Launch Site	10	\$1,000	10	\$1,000
0040	Install Solar Arrays	30	\$3,000	37	\$3,824
0050	Test Solar Arrays	18	\$1,800	18	\$1,800
0060	Final Spacecraft Tests	64	\$6,400	64	\$6.400
0070	Ground Operations	27	\$2,700	27	\$2.700
0080	Launch	0	\$0.00	0	\$0
	Sum of Activities			156	\$15,724

Figure 6. Sub-additivity example data.

4 Limitations of Joint Confidence Levels

4.1 JCL Lacks Tail Information

One reason why quantile risk measures are not sub-additive is that they do not consider the impacts of low likelihood risk events. Low likelihood events are those with a probability of occurrence less than $1 - \alpha$ They occur in the tail of the distribution. Numerous researchers have pointed out that a major weakness in quantile risk measures is that they are not reflective of large impacts beyond the α-level (Cossette et al., 2013; Gustafson, 2004; Kaye, 2005; Sarykalin et al., 2008; Smart, 2021; Sollis, 2009). In the example, “low likelihood” is a risk with a 25% likelihood of occurrence. Whether a project manager would accept that risk depends on the consequences of the risk and their risk tolerance.

4.2 JCL Misleads Decision Makers

The impact of JCL not being sub-additive is that the analyst may underestimate the impact of a risk and relay inadequate information to decision makers. In the example provided, the JCL values for Risk 1 and Risk 2 may lead one to assume that all the project risk impact is attributed to Risk 1. A program manager might then decide to allocate additional resources to solar array installation, say by ordering extra work shifts, and not reserve any additional resources for testing. So, sub-additivity facilitates prioritization of risks and allocation of resources.

In addition, the decision maker may assume there is a diversification benefit derived from managing the risks at a project level and not individually. In other words, the manager can hold fewer reserves since some of the risks in the project may not occur. However, the diversification benefit may not be communicated using quantile risk measures because they are not sub-additive. The total project risk may seem to be greater than the contribution of the individual risks.

4.3 JCL is Subject to Bias

Another weakness of JCL results from the property that it is not unique, except when $\alpha = 0$. In practice, the analyst must select a JCL point from the α-level boundary based on their expert opinion. As pointed out by Perrino (2015), an analyst may select a JCL value, or be directed by management to select a JCL value, that is perceived to be “better” than others based on political, administrative, or political reasons. It may not be politically viable to select a JCL value with a high cost or longer than the desired duration. For planetary missions, it is important that the project is ready to launch during a launch window determined by the physics of the mission. Missing the launch window could mean a delay in the launch by months or years. In these situations, it is typical for a JCL point to be selected that is within the launch window or soon thereafter. Selecting a JCL point to satisfy the analyst’s political, administrative, and scientific narratives is an example of confirmation bias. It ignores the fact that all JCL points on the boundary are possible and distorts the information provided to the decision maker.

5 Joint Confidence Level Alternatives

5.1 Superquantiles Include Tail Information

Superquantiles include tail information since it is the average of those tail events beyond the α-level. Alternative names for a superquantile include Cumulative Tail Expectation and Tail Value-at-Risk, which emphasize the inclusion of tail information.

5.2 Superquantiles Are Sub-additive

Superquantiles are sub-additive if the risks are continuously distributed, and all the components are independent (Lee & Prékopa, 2013). The PCLS model outcomes are independent, as previously mentioned. Continuous distributions, especially triangle and lognormal distributions, are predominantly used to model risk. However, discrete distributions, such as empirical distributions, are necessary on occasion. If a programmatic risk analyst includes discrete distributions in the PCLS model, they should be mindful to check if sub-additivity still holds.

5.3 Superquantiles are Unique

Calculation of the superquantile does not depend on the selection of a particular point on the quantile boundary, as it did with JCL. So, for any given α-level, the superquantile will be unique. This eliminates potential bias when selecting the JCL value to report to management.

6 Practical Evaluation

To evaluate the utility of superquantiles, the $\bar{\bar{B}}\bar{Q}_{0.5}$ and $\bar{\bar{B}}\bar{Q}_{0.7}$ values are calculated for ten NASA projects and compared to the $JCL_{0.5}$ and $JCL_{0.7}$ values. PCLS models for the ten projects were collected from various NASA programmatic analysts. Monte Carlo simulations with 1000 iterations were executed to obtain 1000 ordered pairs of duration and cost representing the possible outcomes for each project. Given this data, six risk measures are calculated in the following manner:

- $JCL_{0.5}/JCL_{0.7}$ – Cossette et al. (2012) propose two methods for choosing a JCL value from the JCL curve. The orthogonal method chooses the point on the JCL curve that is closest to the marginal quantile points. The proportional method chooses a point that preserves the ratio of marginal quantile points. As shown in our sub-additivity example, it is possible for these marginals to be zero and the ratio to be undefined. So, the proportional method is not recommended. The orthogonal method is used here to choose JCL points form the JCL α-level boundary.
- $\bar{\bar{B}}\bar{Q}_{0.5}/\bar{\bar{B}}\bar{Q}_{0.7}$ – Having determined $JCL_{0.5}/JCL_{0.7}$, the expected shortfall from $JCL_{0.5}/JCL_{0.7}$ is the average of all iteration outcomes, (d, c), above the JCL α-level boundary. Percentage change from the JCL value is given.

The ten models were built to analyze the projects after preliminary design was completed and before the project was approved to begin detailed design. The results of the analysis are presented at the project's Preliminary Design Review (PDR). It is at this point in the project lifecycle that the MA and the ABC are determined.

Figure 7 shows the values for $JCL_{0.5}$ and $\overline{BQ}_{0.5}$. Figure 8 shows the values for $JCL_{0.7}$ and $\overline{BQ}_{0.7}$. Based on the ten projects analyzed, the risk measure alternatives do not result in a significantly different MA or ABC values. One possible explanation for this is the inclusion of uncertainty in the PCLS models examined. Risks have a likelihood component, along with a probabilistic distribution, leading to some risk events occurring in the tail of the risk distribution. Uncertainties are characterized solely by a probabilistic distribution without a likelihood component. Uncertainties typically do not exhibit extreme values in the tails of the distribution. Before the PDR, when the ten NASA models were built, there is a great amount of uncertainty due to the preliminary nature of the projects. So, the models have large duration and cost uncertainties and the effects of risks on the model results are muted.

Management Agreement Guidance						
	$JCL_{0.5,d}$	$JCL_{0.5,c}$	$\overline{BQ}_{0.5,d}$		$\overline{BQ}_{0.5,c}$	
Project 1	5646	$2,490M	5822	3%	$2,697M	8%
Project 2	3360	$228M	3383	1%	$229M	0%
Project 3	4258	$11,085M	4338	2%	$11,124M	0%
Project 4	1641	$766M	1705	4%	$794M	4%
Project 5	3141	$500M	3333	6%	$531M	6%
Project 6	2750	$1,096M	2756	0%	$1,107M	1%
Project 7	2952	$716M	3041	3%	$709M	-1%
Project 8	3368	$487M	3439	2%	$506M	4%
Project 9	1682	$274M	1713	2%	$289M	5%
Project 10	3192	$335M	3326	1%	$340M	2%

Figure 7. Alternative risk measures for determining the Management Agreement.

Agency Baseline Commitment Guidance						
	$JCL_{0.7,d}$	$JCL_{0.7,c}$	$\overline{BQ}_{0.7,d}$		$\overline{BQ}_{0.7,c}$	
Project 1	5694	$2,490M	5917	4%	$2,825M	13%
Project 2	3367	$238M	3397	1%	$232M	-2%
Project 3	4391	$11,103M	4378	0%	$11,141M	0%
Project 4	1663	$788M	1735	4%	$808M	3%
Project 5	3194	$547M	3441	8%	$551M	1%
Project 6	2772	$1,101M	2763	0%	$1,113M	1%
Project 7	3000	$714M	3079	3%	$721M	1%
Project 8	3427	$498M	3472	1%	$518M	4%
Project 9	1700	$283M	1747	3%	$295M	4%
Project 10	3268	$336M	3244	-1%	$344M	2%

Figure 8. Alternative risk measures for determining the Agency Baseline Commitment.

7 Conclusions

To summarize, JCL presents three limitations for NASA projects. First, JCL diminishes the impact of risk events that occur beyond the α-level threshold. These risk events are not necessarily "black swan" events (Taleb, 2012). A risk event with a 25% likelihood of occurrence is not improbable. Management may be misled since the impact of the risk is not communicated appropriately. The lack of tail information also contributes to JCL not being sub-additive. Second, since JCL is not sub-additive, it is possible that the decision makers may be given inadequate or misleading information. Third, requiring the analyst or decision maker to choose a JCL value from the JCL curve introduces bias into the decision-making process. Employing superquantiles instead of JCL solves these issues, with the one caveat that superquantiles may not be sub-additive if discrete distributions are used.

The overall benefit of superquantiles is the ability to better communicate the true impact of risks to management. The production of greater or lesser MA and ABC values was never a motivation for switching to alternative risk measures. In fact, drastic changes to MA and ABC would result in mixed reactions. A reduction in MA and ABC is seen as limiting management options in responding to occurring risks. An increase in MA and ABC may push a project past certain funding thresholds so that the project requires greater oversight. Too large of an increase could also threaten project approval.

The superquantile risk measures are satisfactory alternatives and are computationally simple. Given the duration and cost data from the Monte Carlo simulation, the risk measures can be calculated with a few Microsoft Excel formulas. Note that improving the risk measure for determining the MA and ABC does not improve the predictive ability of the PCLS model. The validity of the PCLS model depends on the quality of the schedule, cost, and risk inputs. Poor input will produce poor output no matter how the output is presented. It is also erroneous to believe that improving the risk measure will improve the cost and schedule performance of the project. A 2014 Government Accountability Office (GAO) study attributed an improvement in smaller project performance to the JCL modelling requirement to "quantify potential risks and calculate cost, schedule, and reserve estimates based on all available data." (GAO, 2014) The improvement was attributed to improved quality of the PCLS inputs, not the quality of the output.

Improving the management of project risks should improve project performance. Key to risk management is the ability to prioritize risks. Denault (2001) introduced the concept of coherent risk allocation and Kaye (2005) acknowledges that coherent risk allocation requires a coherent risk measure. Kaye identified several techniques for risk allocation that could be used to prioritize risk based on their project impact using the risk measures from the PCLS model. A potential future research topic is the use of the superquantile risk measures for this purpose. Differences in risk ranking due to the selection of risk measure is of interest also.

References

Adamko, P., Spuchľáková, E., & Valášková, K. (2015). The history and ideas behind VaR. *Procedia Economics and Finance, 24*, 18-24. https://doi.org/10.1016/S2212-5671(15)00607-3

Artzner, P., F. Delbaen, J. M. Eber, & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance, 9*, 203-228.

Butts, G. & Linton, K. (2009). The joint confidence level paradox – a history of denial. *2009 NASA Cost Symposium.*

Cossette, H., Mailhot, M., Marceau, É., & Mesfioui, M. (2013). Bivariate lower and upper orthant value-at-risk. *European Actuarial Journal, 3*(2), 322-357. https://doi.org/10.1007/s13385-013-0079-3

Denault, M. (2001). Coherent allocation of risk capital. *Journal of Risk, 4*(1), 1-34. https://doi.org/10.21314/JOR.2001.053

Embrechts, P. & Puccetti, G. (2006). Bounds for functions of multivariate risks. *Journal of Multivariate Analysis, 97*, 526-547. https://doi.org/10.1016/j.jmva.2005.04.001

Garvey, P. R., book, S. A., & Covert, R. P. (2016). *Probability Methods for Cost Uncertainty Analysis: A Systems Engineering Perspective*. Chapman and Hall/CRC. https://doi.org/10.1201/b19143

Government Accountability Office (GAO). (2014). *NASA: Assessments of Selected Large-Scale Projects: GAO-14-338SP.*

Gustafson, C. R. (2004). Limitations of value-at-risk (VaR) for budget analysis, *Agribusiness & Applied Economics: Report 23612*. North Dakota State University, Department of Agribusiness and Applied Economics.

Kaye, P. (2005). A guide to risk measurement: capital allocation and related decision support issues. *Risk Management in Insurance.* Casualty Actuarial Society.

NASA. (2005). *NASA Program and Project Management Processes and Requirements: NPR 7120.5C*.

NASA. (2015). *NASA Cost Estimating Handbook, Version 4.0: NASA CEH v4.0.*

NASA. (2021). *NASA Program and Project Management Processes and Requirements: NPR 7120.5F*.

National Research Council. (2010). *Controlling Cost Growth of NASA Earth and space Science Missions.* The National Academies Press. https://doi.org/10.17226/12946

Newey, W. & Powell, J. (1987). Asymmetric least squares estimation and testing. *Econometrica*. 55. 819-47. https://doi.org/10.2307/1911031

Perrino, R. (2015, August). Emergence, preliminary assessment, and cross-agency applicability of NASA joint confidence level (JCL) policy. *AIAA SPACE 2015 Conference and Exposition*. https://doi.org/10.2514/6.2015-4572

Rockafellar, R.T. & Royset, J.O. (2010). On buffered failure probability in design and optimization of structures. *Reliability Engineering & System Safety, (95)*5, 499-510. https://doi.org/10.1016/j.ress.2010.01.001

Sarykalin, S., Serraino, G., & Uryasev, S. (2008). Value-at-risk vs. conditional value-at-risk in risk management and optimization. *State-of-the-Art Decision-Making Tools in the Information-Intensive Age*, 270-294. https://doi.org/10.1287/educ.1080.0052

Serfling, R. (2002). Quantile functions for multivariate analysis: approaches and applications. *Statistica Neerlandica (56)*2, 214-232. https://doi.org/10.1111/1467-9574.00195

Smart, C. (2009). Comparison of parametrics with resource-loaded schedule for joint confidence analysis assessment: a case study. *2009 NASA Cost Symposium.*

Smart, C. (2021). Solving for project risk management: understanding the critical role of uncertainty in project management. *2021 NASA Cost and Schedule Symposium.*

Sollis, R. (2009). Value at risk: a critical overview. *Journal of Financial Regulation and Compliance, (17)*4, 398-414. https://doi.org/10.1108/13581980911004370

Taleb, N. (2012). *The black swan.* Random House.