

An Enterprise Portfolio Probabilistic Decision-Analytic Methodology that Unifies Planning and Operational Execution

Allan S. Benjamin, PhD

Distinguished Member of the Technical Staff, Sandia National Laboratories (Retired)

Consultant, NASA Office of Safety and Mission Assurance (Retired)

Contact Information:

✉ Email: asbenja2@gmail.com 📠 Telephone: +1 (505) 401-3773 🔗 Prof. Profile: [linkedin.com](https://www.linkedin.com/in/asbenja2)

Manuscript Status & Notice

- **Publication Track:** The original version of this manuscript is currently under active peer review at the **European Journal of Operational Research (EJOR)**.
 - **Version Note:** This preprint is an updated version of the original submission package (Abstract ID 7432739) with additions to highlight its applicability to resource allocation and its potential for extension beyond enterprise level to international issues including AI and climate change.
-

Practitioner & Consultancy Notice

- **Commercial Applications:** The enterprise portfolio framework and probabilistic decision-analytic models presented in this paper are engineered for immediate deployment across complex operational environments, high-consequence industries, and strategic risk-management sectors.
 - **Consulting Inquiries:** The author is **actively open to selective advisory, advisory-board, and corporate consultancy roles**. Industry practitioners looking to operationalize, adapt, or scale this decision-analytic methodology to optimize enterprise planning and risk mitigation are welcome to initiate contact directly via the email or telephone number listed above.
-

Abstract

The purpose of this paper is to introduce a fully integrated, risk-informed, probabilistic framework that carries over seamlessly from enterprise portfolio decision making to operational risk and opportunity management. It accomplishes this function by adhering to the following six principles and premises: (1) that quantitative, probabilistic risk analysis and management methods need to be used not only at the level of projects and programs, but also for enterprise portfolio planning and execution; (2) that success at enterprise level requires continuously considering not only the financial, safety, security, and technical gains that emanate from mission accomplishment, but also potential aspirational gains that are not directly quantifiable at the time planning occurs; (3) that the consideration of aspirational gains requires tapping into the insights and tacit knowledge of stakeholders and other affected parties by introducing and applying affective performance measures (e.g., their degrees of enthusiasm or concern); (4) that a key element of the analysis should be the initial derivation of an integrated resource allocation scheme and its subsequent updating to continually optimize the balance between competing objectives; (5) that the first four premises motivate the use of Monte Carlo methods to account for all important sources of uncertainty; and (6) that the overall methodology should be extendable to decision-making arenas that require national and international cooperation such as AI and climate change. The paper presents a formalism that accommodates these and other related principles and premises and demonstrates how the results of the analyses can be effectively communicated to management.

Keywords: enterprise portfolio planning; operational risk management; tacit knowledge; risk-informed resource allocation; Monte Carlo analysis.

1. Introduction

1.1 Purpose of the Paper

Most businesses, educational institutions, public service organizations, and government agencies lack a unified, fully integrated, robust quantitative basis for informing enterprise portfolio planning decisions, optimizing cross-program resource allocation, and performing operational risk and opportunity management in an ongoing manner. To address that need, this paper develops and explains a decision-analytic framework that can be applied seamlessly to these purposes.

More specifically, the intent of this paper is to proffer an integrated collection of quantitative risk analysis and management methods, augmented by quantitative/qualitative hybrid models where needed to address aspirational opportunities in addition to programmatic goals, and to demonstrate how these methods can be applied to enterprise-wide portfolios rather than just individual projects and programs so as to promote decision making and operational success for enterprises.

1.2 Principles and Premises

The following principles and premises form the basis of this formulation:

1. Quantitative, probabilistic risk analysis and management methods need to be used not only at the level of projects and programs, but also for enterprise portfolio planning and execution;
2. Success at enterprise level requires continuously considering not only the financial, safety, security, and technical gains that emanate from mission accomplishment, but also potential aspirational gains that are not directly quantifiable at the time planning occurs;
3. The consideration of aspirational gains requires tapping into the insights and tacit knowledge of stakeholders and other affected parties by introducing and applying affective performance measures (e.g., their degrees of enthusiasm or concern);
4. A key element of the analysis should be the initial derivation of an integrated resource allocation scheme and its subsequent updating to continually optimize the balance between competing objectives at the enterprise level; and
5. The first four premises motivate the use of Monte Carlo methods to account for all important sources of uncertainty and iterative resource allocation methods to account for the inter-program competition for limited resources.
6. The overall methodology should be extendable to risk-informed decision-making arenas that affect society at large and that require national and international cooperation, i.e., areas such as AI and climate change.

In addition:

6. Determination of the risk of slippages of critical milestones and the degrees of concern they present relative to the satisfaction of enterprise-level objectives should be another key element of the analysis;
7. The formulation should incorporate quantitative-qualitative hybrid models where needed to address the aspirational objectives rather than resorting to only quantitative models or only qualitative models;
8. The hybrid models should be informed by eliciting the judgments of both entities that are integral to the organization's mission (e.g., executive management, staff, customers, and partners) and entities

that are not directly part of the mission but could benefit from or otherwise be affected by its results (e.g., scientists, educators, and the public); and

9. The computational burden of searching for optimal resource allocations should be kept within manageable limits by selectively utilizing graded analysis approaches where needed in a manner that does not sacrifice overall rigor at the end.

1.3 Background

An earlier survey paper published in the European Journal of Operational Research (Liesiö et al., 2020) suggests that although many of the foundational items enumerated above have been considered by other authors in various papers, the need for and application of a unified quantitative probabilistic methodology that applies to enterprise-level portfolio decision analysis, resource allocation, and operational risk and opportunity management has been addressed rarely if at all. This judgment is validated by the present author's search of the risk management and operational research literature.

At present, many risk management leaders consider the ISO 31000 risk management guidelines report (ISO, 2018) and the COSO enterprise risk management report (COSO, 2017) to be the standards for achieving a risk-informed framework for enterprise-level portfolios, as they formally connect enterprise risk management to strategy-setting rather than treating it as a compliance activity. Many European and US Government agencies have adopted one or the other of these approach: e.g., the UK Orange Book (HM Treasury, 2023) and the US GAO Green Book (U.S. Government Accountability Office, 2025). However, the ISO and COSO frameworks do not recommend a defined quantitative methodology, leaving that decision up to the analysts. They furthermore encourage approaches that treat the expression of risk as a single outcome rather than a probability distribution, in order to promote both ease of implementation and management-friendly communication.

The same is basically true for other approaches, such as the Hoshin Kanri policy deployment methodology (Lean Enterprise Institute, 2020) and the balanced scorecard approach (Kaplan & Norton, 1992), both of which are noted for addressing aspirational objectives (e.g., building a world-class capability) more explicitly than the ISO 31000 and COSO methodologies.

Although probabilistic risk assessment and Monte Carlo analysis have become standard practice for supporting risk-informed decision making and continuous risk management for individual projects and programs, the same degree of rigor is not regularly applied to the analysis and execution of enterprise-level portfolios consisting of a set of mission-related programs, infrastructure development activities, partnership agreements, public service initiatives, and other efforts that serve the fundamental objectives of the enterprise. Yet the larger problem is at least as important as the smaller one, and arguably much more so because it concerns all the objectives of the enterprise rather than just programmatic ones.

Below are six anecdotal instances, extracted from a much larger database of examples, that highlight the need for more rigorous risk modeling at enterprise level:

- The European Space Agency has remarked on its difficulties in attaining an Agency-wide balancing of science, Earth observation, launcher, and exploration programs within its constrained industrial and financial resources (European Space Agency, 2021).
- The U.S. Government Accountability Office has reported on NASA's history with schedule slippages, most recently within its set of interdependent Artemis projects (U.S. Government Accountability Office, 2025b). NASA itself has cited its recurring problem with meeting milestones in the area of

technology maturation, which is the only source of concern within its list of performance goals that is currently rated as red (NASA, 2026), (NASA, 2024).

- The GAO has reported on the U.S. Department of Defense having experienced major schedule delays because of competition for resources among its multiple acquisition programs (U.S. Government Accountability Office, 2016). GAO has also stated that DoD's acquisition outcomes could be improved through an integrated portfolio management approach rather than treating weapon systems individually (U.S. Government Accountability Office, 2007).
- Airbus has remarked on its concern with meeting schedules because of the A350, A320neo, A220, and other aircraft development programs competing for aerodynamic engineers, software developers, certification specialists, manufacturing facilities, and suppliers (Airbus, 2025).
- Boeing has cited similar concerns regarding the interactions among concurrent activities such as its 787, 737 MAX, KC-46, and 777X programs (Boeing, 2025). Boeing also commissioned an independent study to determine how dependency matrices can be used to help evaluate and optimize a technology portfolio when the projects are highly coupled and are initiated in different funding cycles (Dickinson et al., 2021).
- Pharmaceutical companies such as Pfizer, Roche, Novartis, and Merck have noted the importance of portfolio management for clinical trials that compete for patients, physicians, manufacturing capacity, regulatory personnel, and specialized laboratories (DiMasi et al., 2016). At least one author has reported that risk quantification in pharmaceutical research and development often does not capture the full extent of a company's exposure to financial risk (Oreskovich & Gittins, 2016).

All of these examples testify to a tendency toward under-evaluating the risk caused by competition for necessary resources among the different activities that occur within the enterprise-wide portfolio. Improvement in that area of concern is one of the focal points of the methodology presented in this paper.

1.4 Accounting for Tacit Insights

When considering the benefits vs. costs of new technology developments or the potential for new knowledge to be gained from an enterprise-level portfolio, wherein the nature of new programs or activities that may be spawned in the future is not known, the expression of the tacit insights of those involved may take the form of intuitive degrees of enthusiasm, excitement, optimism, or concern informed by analysis and experience. It is natural for these to be adopted as additional performance measures to be considered in the overall evaluation, although there is not yet a widely accepted formalism for doing it.

There are real-world instances where placing emphasis on the values and intuitive knowledge of those most directly involved during decision making has led to far better results in the long term than relying only on a prediction of expected revenues or the achievement of near-term mission goals. A famous example emanates from the work of Ikujiro Nonaka (Nonaka, 2007), whose studies showed how an orientation toward knowledge creation can lead to much greater mid-term to long-term success than an orientation toward revenue. His work helped explain why companies like Mazda and other Japanese car companies in the 1970s and 1980s had much greater success than American and European car companies during that period. In the words of Valerie Lynch, chairman and technology consultant for AND Technology Research in Great Britain (Lynch, 2025), “the importance of capturing tacit knowledge has been highlighted through the study of Japanese companies where it was shown that tacit knowledge shared between team members is instrumental in achieving good outcomes from work in innovation.”

2. The Shared Importance of Programmatic and Aspirational Objectives

Enterprise-level executives operate with two fundamentally different objective types simultaneously: programmatic objectives (deliver Program X on time and within budget, maintain compliance with regulations, achieve operational key performance measures) and aspirational objectives (become the market leader in a new competitive segment, build resilience for an uncertain future, promote transitions to next-generation technology). Programmatic objectives have natural risk-quantification methods, whereas aspirational objectives mostly do not. They tend to be directional, long-horizon, and often unmeasurable in the same units as programmatic delivery risk.

Some example *programmatic* objectives that pertain to a mission oriented enterprise are as follows:

1. Mission Achievement Objective: Produce a stated set of mission-related results
2. Mission Reliability Objective: Achieve high mission reliability by exploiting designs and operational processes that are robust and by emphasizing integration of technologies that have become reasonably mature
3. Mission Sustainability Objective: Cultivate a broad base of support for ongoing future funding, have at most a manageable impact on pre-existing contracts, and enable a smooth transition from current operations
4. Mission Schedule Objective: Deliver the results on a schedule that is agreeable to the customer or funding agency and that accommodates sufficient margins to support unplanned delays
5. Mission Cost Objective: Deliver the results within the funding that is provided by the customer or funding agency and that accommodates sufficient margins to support unknown and underappreciated costs
6. Mission Safety Objective: Do not produce mission scenarios that have inordinately high risk and complexity with high potential impact on human safety and property protection
7. Cybersecurity Objective: Be amenable to safeguards and provisions that provide a specified level of protection against cyberattacks perpetrated by malevolent individuals, terrorist groups, or adversarial nations
8. Workforce and Infrastructure Objective: Be consistent with an ability to retain or develop the critical skills and research facilities needed to achieve the first seven objectives and do not seriously impact the status of critical skills or research facilities needed to support other national or enterprise priorities

Likewise, some corresponding example *aspirational* objectives that pertain to a mission oriented enterprise are as follows:

1. Extensibility Objective: Provide the opportunity to demonstrate technology, systems, and operations that will be important in future programs and ventures
2. Technology Innovation and Transfer Objective: Enable technology maturation, foster the development of new modes of operation, create new technologies and new engineering knowledge, and disseminate this technology and knowledge to multiple stakeholders
3. Public Engagement Objective: Cultivate the potential to inspire current and future generations, educate the public about the opportunities and societal benefits gained from the mission, and

motivate young people to pursue an education in science, technology, engineering, and mathematics

4. Global Partnership Objective: Provide the opportunity to strengthen and expand international partnerships

The best enterprise-wide methods for portfolio building and management are ones that build a communication network between these two objective types so that a single portfolio governance process can consider both and achieve a proper balance.

3. Methodological Framework for Addressing Risks to Enterprise-Level Programmatic and Aspirational Objectives during Portfolio Planning

3.1 Overview

The risk analysis methodology that is proposed in this paper is synopsisized in Figure 1. It depicts an integrated probabilistic analysis approach for evaluating programmatic and aspirational objectives founded on Monte Carlo sampling of uncertainty distributions and iterations on resource allocations.

As shown in the figure, traditional Monte-Carlo-based approaches that are applied in project-specific risk analyses have to be modified to simultaneously consider an entire portfolio of programs that interact with one another, as well as an entire set of enterprise-level programmatic and aspirational objectives that also interact with one another. This interactive formalism enables risks and opportunities to be aggregated from the lowest levels of the organization to the very top levels.

Given a proposed enterprise-level portfolio and a proposed allocation of resources across the programs and activities that constitute the portfolio, the first part of the risk aggregation formalism occurs in Charts A, B, and C across the top row within the Monte Carlo loop in the figure. The models generally pertain to the evaluation of performance measures that apply to technical performance, safety performance, physical security, cybersecurity, schedule, and cost, all evaluated at project and program level. The nuances of the evaluations in Charts A, B, and C will be explained at some level of detail in Sections 3.3 and 3.4.

The second part occurs in Charts D through K, where the models include a combination of quantitative and qualitative analysis methodology at enterprise level. Part of the quantitative portion concerns an evaluation of the total monetary value of the entire portfolio of programs, covering all sources of monetary benefit that are quantifiable (Charts H and I). However, the quantifiable monetary value does not tell the whole story, because the ultimate value of the portfolio includes aspects that are resistant to monetary quantification. An example is the future value of new technologies that may have applications that cannot be known at present; another is the knowledge gained during operation from exploring new ground for which the potential for useful data is unknowable in the present (Charts F and G). Also, the portfolio planning evaluation incorporates an assessment of how schedule risks due to critical milestone slippages within the programs that comprise the portfolio impact the successful accomplishment of each enterprise-level objective (Charts D and E). Other measures of performance that influence portfolio planning are not specifically detailed in Figure 1 but are alluded to in Chart K of that figure and will be addressed later in Section 3.8. These include the overall cost risk, the overall safety risk, the overall cybersecurity risk, and the strategic benefit of collaborations with corporate partners, educational institutions, and international partners.




Figure 1. Synopsis of the analysis methodology for enterprise-level portfolio planning.

The more qualitative portion of the evaluation generally applies to the aspirational objectives of the enterprise. It relies on ascertaining the fundamental concerns, values, and tacit insights of the enterprise's stakeholders, customers, and partners. These are integrated into models that determine how the performance measures obtained from the results of the programmatic evaluation can affect the potential for success of the aspirational objectives. The nuances of the aggregation of risks and benefits from program level to enterprise level will be treated in more detail in Sections 3.5 through 3.8.

For attribution purposes, Figure 1 and the discussions in the following subsections have been informed by analogous figures and discussions published by the present author in (Benjamin, 2025); however, the contributions in this paper expand significantly beyond that original source. Both sources, in turn, are an outgrowth of guidance provided in the 2024 NASA Risk Management Handbook for which the present author was a co-author (Dezfuli et al., 2024a, 2024b).

3.2 The Overall Purpose and Composition of Resource Allocation

The purpose of allocating resources across an enterprise's portfolio is to ensure that the stakeholders' degrees of concern about not adequately achieving each programmatic objective are as low as reasonably achievable and their degree of enthusiasm about achieving each aspirational objective are as high as reasonably achievable. Inherent in this intent is a sense of balance: i.e., that none of the degrees of concern are inordinately high and none of the degrees of enthusiasm are inordinately low. The allocation typically comprises funding resources, human resources, facility and equipment resources, IT resources, testing resources, training resources, technology investment resources, supply chain resources, and operational resources. The rationale for balancing allocations across programs to best achieve all of the enterprise's stated objectives will be explored in more detail in Sections 3.8 and 5.1.

3.3 Program Model and Input Uncertainty Distributions

Regarding Chart A in Figure 1, the program-level uncertainties that may be sampled in the Monte Carlo process emanate from a variety of sources, which may be observed more clearly by referring to an enlarged and more detailed version of that chart (see Figure 2). First, there may be uncertainties about which of several alternative models best describes a system response or a physical process. Uncertainties in modeling are an important contribution to the overall uncertainty in the calculation of performance measure risks that is unfortunately left out of many risk analyses. The easiest way to account for them, when modeling uncertainty is considered to be significant (as it would be when new technology is being developed and implemented), is to propose several models that have the backing of subject matter experts, and let the experts weight them based on their judgment of their relative likelihood of being correct within the context of the program. The weights form a probability scale from which different models can be selected for different Monte Carlo iterations. The sampling of alternative models is shown schematically in the upper left graphic in Chart A.

Other input uncertainty distributions depicted in Chart A include system design parameters, environmental parameters, material properties, event occurrence probabilities or frequencies, and leading indicator values for the determination of the contribution from unknown and/or underappreciated risks. Uncertainty distributions for each of these are established by the risk analysts in conference with subject matter experts. The sampling of model input distributions is illustrated in the top right graphic within Chart A.

Chart A




Figure 2. More detailed rendering of Chart A in Figure 1.

Among the inputs that are sampled in the top-right graphic are the frequencies of occurrence of events or conditions that taken together, in various combinations, can lead to risk scenarios. The bottom-left graphic in Figure 2 suggests that in general, when dependencies between programs are taken into account, a separate sampling of the risk scenarios has to be made for each Monte Carlo iteration in order to determine the particular scenario(s) that apply for that iteration. This provision is often not necessary when the dependencies between programs are not of issue, as would be the case when the analysis is program-specific rather than portfolio-inclusive. In that case, the models can be framed to work with scenario frequencies as inputs, and the evaluation may be based on an aggregation of frequencies using Boolean algebra rather than an aggregation of occurrences. The practical necessity of working with occurrences rather than frequencies when inter-program dependencies are important leads to an increase of the number of Monte Carlo iterations that are needed but does not otherwise change the overall methodology.

Regarding the bottom-right chart in Figure 2, once leading indicator uncertainty distributions have been sampled, they may be applied through the use of analogous but different models to determine the incremental effect from unknown and underappreciated (U/U) scenarios upon the performance measure values previously calculated considering only known scenarios. Each model applies to a particular measure of performance and takes the form of a correlation between the U/U adjustment to be made for that performance measure and the relevant leading indicators. For example, there could be a separate U/U adjustment model for program cost, delivery date, technical maturity level, number or probability of fatalities, and time needed to recognize and resist a cybersecurity attack.

Two of the references, (Benjamin et al., 2016) and (Benjamin, 2025), provide details about the data and methods that can be used to develop correlations for U/U risk adjustments and present actual results that have been obtained for selected performance measures within the domains of cost, schedule, and safety performance. In general, it has been found that the potential magnitude of U/U risk adjustments compared to risks from known scenarios depends upon a composite of the following factors (Bearden, 2008), (Perrow,

1999), (Mars Climate Orbiter Mishap Investigation Board, 1999), (U.S. Government Accountability Office, 2015):

- Complexity in the design of the system or in the supply chain, system development, or operational procedures
- Use of tightly coupled subsystems that make it difficult for operators to comprehend the status of a system during real-time operation
- Scaling of existing technologies beyond the domain of knowledge
- Use of fundamentally new technology or fundamentally new application of an existing technology
- History of management oversight problems, or a hierarchical management style with insufficient attention paid to concerns at lower organizational levels
- Management priorities focused toward meeting schedules and costs over the other programmatic domains, causing high stress levels
- Frequent changes of requirements or funding profiles
- Supplier or contractor financial difficulties

The two Benjamin references cited in the preceding paragraph also provide a canvassing of military and commercial system reliability experience, human reliability experience, commercial nuclear reactor severe accident experience, and Government-sponsored space system experience to provide correlation suggestions. The results make it clear that the relative magnitude of the U/U adjustments can be large compared to the baseline calculation if these leading indicators combine synergistically.

Once they are derived, the U/U adjustment correlations themselves have uncertainties, as depicted schematically by the distribution shown in the lower right graphic of Figure 2. These correlation uncertainties are reflected in the fact that there is often more than one postulated equation form that may be applied to obtain a reasonable correlation and that the parameters within each equation have inherent uncertainties due to data sparsity. During the Monte Carlo analysis process, the correlation uncertainties can be accounted for by sampling alternative correlation equations per the top left graphic in Figure 2, and for each equation form, by sampling input parameter uncertainties per the top right graphic.

3.4 Characterization of Inputs, Models, and Outputs Applicable to the Assessment of Program Risks

Chart B of Figure 1, which is enlarged in Figure 3 below, provides a schematic look at the performance measure inputs, models, and outputs that are typical in an evaluation of interfacing program-level risks. At the program level, these elements generally pertain to the risk domains of technical performance, schedule, cost, safety, physical security, and cybersecurity.

The schematic is of course simplified in that it only begins to depict the complexity involved in dealing with models that interact heavily with one another. The degree of complexity in this endeavor is suggested by Figure 4, which provides a partial list of commonly-used models that may comprise such an analysis, along with some of the corresponding sampled input parameters that need to be provided to the models and some of the output metrics whose probability distributions evolve from the Monte Carlo execution of the models.

Figure 5 illustrates more specifically the interactions between models, inputs, and outputs for a particular hypothetical example relevant to NASA, in which it is planned for a science package to be transported to and put into orbit around a distant planet or moon within our solar system. While still simplified, this example helps clarify the nature of the interactions that need to be accounted for in the calculation of program-level risks.

3.5 Characterization of Degrees of Concern for Enterprise-Level Portfolios Caused by Schedule Risks

One of the principal factors to consider during enterprise portfolio planning and operational risk and opportunity management is the risk to each enterprise objective that is promulgated by the potential for significant slippages in meeting critical milestones. Many enterprises consider how well the organization is doing by concentrating on lagging indicators (i.e., the amount of slippage that has occurred during the prior year) rather than looking at the risk of slippages in future years. This makes it difficult to be proactive in understanding the potential future effect of milestone risks on enterprise objectives.

The relationship between schedule risk and risk to objectives is usually clear at the level of individual programs, where the time to deployment and operation of a system is in itself a programmatic objective. However, the relationship is less clear for enterprise-level objectives, which are much broader in scope than program objectives. For example, slippages in critical milestones during the development phase of an exploratory program whose objectives include knowledge discovery do not easily translate through quantitative models into degrees of concern about meeting the long-term objective of stimulating the public and inspiring students to undertake an education in STEM, which is one of the enterprise-level objectives listed in Section 2.

Chart B




Figure 3. More detailed rendering of Chart B in Figure 1.




Figure 4. Synopsis of input distributions, performance models, and outputs that may be included in the programmatic performance measure evaluation portion of the Monte Carlo analysis.




Figure 5. Synopsis of the interactions between models, inputs, and outputs for a hypothetical example relevant to NASA.

While the risk of individual milestone slippages is determined quantitatively within the programmatic analysis depicted in Charts A, B, and C, understanding the corresponding risk to enterprise objectives requires a more qualitative analysis. The importance of individual schedule slippages on enterprise objectives depends on the relative importance of each program, individually and in combination, with respect to each enterprise objective, which introduces an element of subjectivity.

The treatment of this issue is most easily approached by eliciting information from those most directly affected, namely executive management, their committees of expert advisors, and the principal stakeholders. Implicit in this elicitation is the assumption that the most directly affected people have an intuitively sound basis for registering degrees of concern about how various combinations of slippages in critical milestones affect the viability of the enterprise's ability to meet each of its objectives.

The elicitation approach and building of functional relationships consistent with the elicitation results are schematically illustrated in Charts D and E of Figure 1, which are reproduced in a more readable size in Figure 6 below. The respondents are asked to provide their degrees of concern for a handful of combinations of milestone slippages, with separate results obtained for each of the enterprise-level objectives. Formal expert elicitation methods such as the Delphi method (Linstone & Turoff, 1975), the Sheffield Elicitation Framework (Gosling, 2018), the Analytical Hierarchy Process (Saaty, 1987), and Multi-Attribute Utility Theory (Weirich, 2012) may be used to help guarantee that the results obtained are consistent and representative of the respondents' true beliefs.




Figure 6. More detailed rendering of Charts D and E in Figure 1.

The equation at the top of Chart E reads as follows: “The schedule degree of concern ($SDOC_j$) for enterprise objective j is a function of the slippage amounts of the various critical milestones ($CMSS_i$) that relate to objective j .” Also, the dimensionality of the graphics in Figure 6 is notional. It could be higher or lower than the two dimensions shown in the figure.

The elicitation of degrees of concern from different stakeholders will almost certainly result in different results for Charts D and E if considered individually as opposed to being amalgamated into an average over all the respondents. As was discussed relative to Chart A in Figure 2, and in particular the upper-left graphic within that chart, different functional relationships resulting from different stakeholders can be treated as a source of modeling uncertainty and can be incorporated into the Monte Carlo analysis protocol as such.

To properly account during the sampling process for stakeholders' diversity of opinions regarding the functional relationships between schedule degrees of concern and enterprise-level programmatic performance measures, a credibility weighting factor would need to be determined for each respondent. The weight could be obtained by using accepted means for assigning credibility to respondents based on the respondent's experience, repute, and record of success (e.g., Cooke, 1991). These weights would govern the relative number of times each elicited functional relationship would be selected during the Monte Carlo process.

The results obtained from this enterprise-level schedule analysis can be plotted in the form of a set of heat maps, as shown schematically in Figure 7. The varying intensity within each bar reflects the density function for the schedule degree of concern obtained from the Monte Carlo analysis.




Figure 7. Heat maps illustrating hypothetical results from the Monte Carlo analysis pertaining to schedule degree of concern at enterprise level.

Additionally, to aid communication between executive management and risk management technical leads and practitioners, the mean values of SDOC for each objective can be plotted on a spider chart, as shown in Figure 8. The spider chart format introduces regions of risk tolerability denoted as acceptable (green), marginal (yellow), and intolerable (red). The term *risk*, in this context, is taken to be synonymous with the term *degree of concern*.

3.6 Characterization of Degrees of Enthusiasm or Excitement Associated with the Development of New Technology and the Potential for Discovery of New Knowledge

A risk-informed evaluation of the relative merits of alternative enterprise-level portfolios and the operational success of a selected portfolio depends upon a clear understanding of how the risks are balanced against the opportunities. While the preceding section presented a methodological formulation for assessing enterprise-level risks, the present section provides a comparable formulation for enterprise-level opportunities. Charts F, G, H, I, and J in Figure 1 summarize the methodology.

Charts F and G are reproduced below in Figure 9 and address the potential scientific and motivational value of new technology development and new knowledge discovery. Both of these attributes have monetary value, and it would be ideal if they could be evaluated in monetary terms. This is not realistically plausible, however, because many of the future uses of the new technology and the specifics of the new knowledge to be gained are not yet known and will only become known in the future. Therefore, the figures of merit are more accurately represented by attributes that emanate from tacit information possessed by the enterprise's stakeholders and partners, as well as other beneficiaries such as educational and scientific

institutions and the public at large. These tacit insights can best be mined by eliciting their degrees of enthusiasm or excitement for the potential future uses of the new technology to be developed or the new knowledge to be gained.




Figure 8. Spider chart illustrating the mean values of hypothetical results from the Monte Carlo analysis pertaining to schedule degree of concern at enterprise level.

The respective equations in Figure 9 denote that technology development degree of enthusiasm ($TDOE$) is a function of the vector of programmatic performance measures that pertain to the performance of the new technology ($\overrightarrow{NTPM}$), and the knowledge discovery degree of excitement ($KDOE$) is a function of the vector of programmatic performance measures that pertain to the success of data collection during operation ($\overrightarrow{DCPM}$). Those programmatic results emanate from Charts A, B, and C.

Chart F

New technology development potential: degree of stakeholder & partner enthusiasm ($TDOE$) as a function of new technology performance measures ($NTPM$)

$$TDOE = F(NTPM \rightarrow)$$

Scale of 0 to 1

Chart G

New knowledge discovery potential: degree of public & scientist excitement ($KDOE$) as a function of data collection performance measures ($DCPM$)

$$KDOE = G(DCPM \rightarrow)$$

Scale of 0 to 1

Figure 9. Enlargement of Charts F and G in Figure 1.

As an example, if the intent of a technology initiative is to develop and implement a new type of automobile engine that exceeds the performance capability of existing engines and can be implemented within a reasonable timeframe, the performance measures whose values could affect the degree of enthusiasm of the executives in the automotive company and their customers would likely include (a) the engine's energy conversion efficiency, (b) the achievable thrust and range, and (c) the probability that the new engine would not induce an automotive failure leading to a fatal accident. Arrangements would be made to include these in the set of programmatic performance measures evaluated in Charts A through C.

In addition, certain schedule and budgetary performance measures, such as (d) the time it would take to achieve operability and satisfy regulations, and (e) the availability of the required supply chain, would also likely affect executive and customer enthusiasm for the new technology, but these would likely be lower-order effects. Evidence for this is provided by NASA's experience, wherein enthusiasm for the Space Shuttle and James Webb Space Telescope initiatives remained high during program development and operation and after program completion, even though the schedules and costs of both programs greatly exceeded the initial plan. Thus, these measures would be included in the programmatic analysis but would be secondary considerations in the elicitation processes for determining customer degrees of enthusiasm or excitement.

The process for eliciting information from the stakeholders, customers, partners, educators, scientists, and public and converting the elicitation results into functional relationships is basically the same as the process described in the preceding section for evaluating enterprise-level schedule degrees of concern, and so the graphics in Figure 6 (Charts D and E) apply here in concept as well. However, in addition to the expert elicitation methods listed in Section 3.5 for characterizing degrees of concern about schedule impacts on enterprise-level objectives (i.e., Delphi, Shelf, AHP, and MAUT), there are other methods that are especially applicable for mining stakeholders' tacit knowledge in order to rate degrees of enthusiasm and degrees of excitement for aspirational objectives.

In particular, Adaptive Comparative Judgment (ACJ) processes take advantage of the fact that pairwise comparisons turn out to be much more accurate than absolute rankings for aspirational objectives when it is not possible to dissect the experts' tacit knowledge into specific measures of performance. Examples of the ACJ approach include the Thurstone Case V model (Thurstone, 1927) and the Bradley-Terry model (Bradley & Terry, 1952). As a further enhancement, some researchers have combined the ACJ approach with a value-function anchoring process in order to increase stability by anchoring the rankings to defined best-case and worst-case reference values. Others, however, have warned that this approach has the potential to introduce unwanted biases (e.g., Keeney & Raiffa, 1976). One way around this may be to have each stakeholder propose their own anchor and to use this information to randomize the anchor during the Monte Carlo process in order to minimize individual biases. This would be a good topic for additional research.

The results obtained from the analysis can be plotted in the form of a set of heat maps, as shown schematically in Figure 10. As before, the varying intensity within each bar reflects the density function for the enterprise-level performance measures (*TDOE* and *KDOE*) obtained from the Monte Carlo analysis.

3.7 Calculation of Quantifiable Monetary Value and Characterization of Associated Degrees of Enthusiasm

Charts H and I of Figure 1, which are reproduced in Figure 11 with additional details included, depict the process for evaluating the quantifiable portion of the monetary value of each portfolio. Contributions to the quantifiable value include the following, as examples:

- Near-term financial return from sales to customers or from additional funding from sponsoring agencies
- Longer-term financial return from sales or additional funding based on existing knowledge of how the programs in the portfolio will spawn new programs
- Projections based on current knowledge of how technology transfer to commercial enterprises will translate to financial gains for them
- Projections of how public enthusiasm about the programs in the portfolio will lead to STEM commitments among the younger generation with resultant quantifiable increases in human resources for the enterprise or its partners




Figure 10. Heat maps notionally illustrating new technology degrees of enthusiasm and knowledge discovery degrees of excitement obtained from results of the Monte Carlo analysis.

Chart H




Chart I




Figure 11. More detailed rendering of Charts H and I in Figure 1.

As discussed earlier, they do not include sources of revenue from future activities that cannot even be conceived of at the present time.

The monetary evaluation is based on verifiable and validated quantitative models. The inputs to the evaluation include both economic and financial inputs as well as the applicable programmatic performance measure results obtained from Charts A, B, and C. The latter provides information about the status of performance for each program in the domains of technical achievement, safety, physical security, cybersecurity, program schedule, and program cost, all of which affect the evaluation of monetary return. In equation form, the portfolio quantifiable monetary value (MV) is a function of the vector of economic and financial inputs (EFI) and the vector of applicable programmatic performance measures (PPM), as follows:

$$MV = H(\overrightarrow{EFI}, \overrightarrow{PPM}) \quad (1)$$

Similar to the results depicted earlier in Figures 7 and 10, the results obtained from the monetary analysis can be plotted as a set of heat maps, shown schematically in Figure 12. The varying intensity within each bar reflects the probability density function for MV obtained from the Monte Carlo analysis. It is useful to show the MV density function side-by-side with the density function of the cost of the portfolio, so that decision makers can readily infer the ratio of benefit to cost.




Figure 12. Heat maps notionally illustrating the enterprise-level portfolio cost and the quantifiable monetary value of the portfolio obtained from results of the Monte Carlo analysis.

Finally, by applying the specialized elicitation processes discussed in Section 3.6, each stakeholder's degree of enthusiasm about the quantifiable monetary value of the portfolio can be specified in terms of a constructed function. For example, the relationship $MDOE = H(MV - C)$, originally presented in Chart J of Figure 1, specifies this degree of enthusiasm in terms of the *quantifiable monetary profit* (value minus portfolio cost). Alternatively, the relationship $MDOE' = H(MV/C)$ expresses it in terms of the *cost-normalized monetary value* (value divided by cost). Either approach provides a basis for comparing $MDOE$ with the aspirational figures of merit in Figure 9, namely $TDOE$ and $KDOE$.

3.8 Enterprise-Level Comparison of the Balance Achieved between Programmatic and Aspirational Risks and Opportunities

For the example considered in this section, eight key attributes are assumed to affect the risk-informed decision making process during enterprise-level portfolio selection. The first four concern performance measures associated with the degrees of enthusiasm or excitement generated by (1) new technology developments, (2) the potential for new knowledge discovery, (3) the benefit of strategic partnerships, and (4) the quantifiable portfolio monetary profit. The second four involve performance measures associated with degrees of concern about the aggregated enterprise-level risks associated with (5) schedule slippages, (6) cost overruns, (7) safety shortcomings, and (8) cybersecurity vulnerabilities. The objective is to achieve a proper balance between these eight attributes.

To provide a consistent comparison between them, it is best to convert the conversation about enterprise-level programmatic risks from degrees of *concern* to degrees of *acceptance*. The latter is the complement of the former (e.g, $SDOA = 1 - SDOC$, where $SDOA$ is the schedule risk degree of acceptance and $SDOC$ is the schedule risk degree of concern). The degree of acceptance reflects the stakeholders' confidence that the enterprise's overall schedule risk, cost risk, safety risk, and cybersecurity risk will be satisfactorily contained. The reason for this inversion is that it allows the results for the programmatic risk attributes to be compared directly with the results for the other attributes, which are positive rather than negative.

In addition, for summary purposes only, it is often best to use a single figure of merit for $SDOA$ which includes combined consideration of the several values of $SDOC_j$ (one for each enterprise objective j). The purpose is to balance the presence of the various attributes in the portfolio selection process; otherwise schedule risk would be over-represented. One method for doing this is to define an *overall* schedule degree of acceptance as follows:

$$OSDOA = 1 - \sum_j w_j(SDOC)_j \quad (2)$$

where the weighting factors, w_j , are obtained from the respondents in the elicitation process and represent their assessment of the relative importance of each objective j to the enterprise. The linear weighting may not be the perfect representation but it is expected to be adequate for this purpose.

To facilitate communication about the results of the attribute evaluation to management, it is again possible to take the mean values obtained from the Monte Carlo analysis and present them on spider charts, as shown in Figure 13. The upper part of the figure portrays a spider-web contour for each of the different resource allocations made within the outer loop in Figure 1, with a separate chart being provided for each candidate portfolio. For Portfolio 1 in this example, Allocation 1 results in high ratings for the technology development, knowledge discovery, and strategic partnership attributes but low values for the monetary value, schedule risk, cost risk, safety risk, and cybersecurity risk attributes, whereas Allocation 2 results in the opposite. Allocation 3 provides a balancing of the attributes, which would constitute the optimum allocation according to Principle 4 in Section 1.2.

An important benefit of the displays of the results for different resource allocations is that it assists management in implementing a risk-informed process for allocating resources across its programs and other activities without necessarily committing them to the specific “near-optimal” resource allocation that results from the analysis after completion of the outer iteration loop. In other words, the intent is for the deliberation on resource allocation to be risk-informed rather than for the decision to be dictated by the analysis.




Figure 13. Example spider chart comparing the results from the Monte Carlo analysis pertaining to eight key attributes for three candidate enterprise-level portfolios.

The lower part of Figure 13 portrays the spider-web contour of the attribute results for each candidate portfolio after the resource allocation for the portfolio has been optimized. Thus the contour for Portfolio 1 in the lower part of the figure corresponds to the contour for Resource Allocation 3 in the upper part of the figure.

For this hypothetical set of results, the recommendation would be in favor of Portfolio 3 over Portfolio 1 or 2 according to Principle 4, because that portfolio best optimizes the balance between competing objectives at the enterprise level. This follows from the fact that the minimum score over the eight attributes for Portfolio 3 is the highest for all the portfolios, having a value of about 0.78 out of 1.00 for the overall cybersecurity risk degree of acceptance and no less than that for the other seven attributes. This value is better than the result for Portfolio 2 at about 0.63 (for the safety risk degree of acceptance) and for Portfolio 1 at about 0.43 (for the new technology degree of enthusiasm).

4. Applying the Methodological Framework to Portfolio Risk Management during Operational Execution

After an enterprise-level portfolio has been accepted and put into operation, the evaluations and results depicted notionally in Figures 7, 8, 10, 12, and 13 can be updated on a recurring basis to incorporate new data and changes in understanding that come to light during operation. This updating and the actions that result from it are a part of enterprise portfolio operational risk and opportunity management.

For this purpose, as stated earlier, one of the values of spider charts like Figure 8 that incorporate risk tolerance rankings is that they readily convey which risks are of most concern and need to be managed through controls or risk mitigation actions. In the example of Figure 8, the risks to the enterprise's mission and extensibility objectives for candidate portfolio 1 are shown to be intolerable, and so the critical milestone slippages that matter to these objectives would need to be addressed. How they should be addressed hinges on identifying the events or conditions that drive those milestone slippages, and this can be determined by examining the logic models that comprise Chart B in Figures 1 and 3.

For example, Figure 14 presents a fault-tree-based depiction of the events and conditions leading to slippage in a project milestone that has been determined to be of concern. Specific risk scenarios, common root causes of the scenarios, and leading indicators that cause the milestone to slip to an excessive level of discomfort are identified in this fault tree. Their relative importance in producing the slippage can be determined by performing sensitivity analyses in which the relevant risk scenarios, common root causes, and leading indicators are deleted, individually or in selected combinations, from the modeling and the effects of these deletions on the milestone slippage are determined via a reduced Monte Carlo analysis. The results provide the basis for suggesting controls and mitigation actions that should prove to be successful in reducing the enterprise-level risks that are of most concern.

This process may lead to more than one candidate control or mitigation alternative. Each alternative can be evaluated separately by repeating the full Monte Carlo analysis to obtain a chart comparable to Figure 8. An example is shown in Figure 15. Decision makers can then select from among the alternatives knowing how much each improves the level of risk tolerability over all objectives.

During operational risk and opportunity management, a useful refinement is to present the schedule risk to each enterprise objective not only "en masse" but also on a yearly basis to provide more insight into the timing of the risks. For that purpose, the equation at the top of Chart E in Figure 6 changes to the following:

$$SDOC_{j,k} = f_{j,k}(\{CMSS_i | i \in j, k\}) \quad (3)$$

which states that the overall schedule degree of concern for year k relative to enterprise objective j is a function of the slippage amounts of the various critical milestones that occur during year k and relate to objective j ." A spider chart can now be presented to illustrate the schedule degree-of-concern at enterprise level for each year following the present, as shown in Figure 16.




Figure 14. Example fault tree detailing the events and conditions leading to slippage in a critical project milestone.




Figure 15. Example spider charts comparing the schedule degree-of-concern at enterprise level for two candidate mitigation alternatives.




Figure 16. Example spider charts comparing the year-by-year schedule degree-of-concern at enterprise level for the first three evaluation years.

If an analysis actually led to the hypothetical results shown in Figure 16, a risk analyst would infer that the enterprise's cybersecurity and partnership objectives will be in future danger of becoming at least borderline

intolerable because of one or more critical milestone slippages occurring during Year 3. An analysis of the specific milestones, risk scenarios, leading indicators, and common-cause contributors leading to that undesirable condition would reveal what mitigation or control actions would be needed to rectify it and when they would need to be instituted. A subsequent reanalysis using the methods discussed in this paper would reveal the monetary cost of the mitigations and control actions, the levels of residual risk in the areas of concern, and the collateral effect on other enterprise-level objectives. This would provide the information needed by the authorities to make a risk-informed decision about the actions to be pursued and the timing of those actions.

5. Additional Considerations

5.1 Using a Graded Analysis Approach for the Resource Allocation Loop

As shown originally in Figure 1, the formulation includes an iterative outer loop for re-allocating resources among programs so as to minimize risks and maximize benefits across enterprise-level objectives. Section 3.8 presented this concern in terms of achieving a balance between enterprise portfolio risks and opportunities, with the measures of such risks and opportunities being expressed in terms of stakeholder degrees of concern about programmatic risks such as potential milestone slippages or degrees of enthusiasm about programmatic or aspirational opportunities such as the potential for new technology applications.

A practical concern is the potential computational burden that may arise due to the many iterations implied by the use of Monte Carlo methods together with a resource allocation loop. If the performance measure models in Chart B and their interfaces are complex, that may become a problem. In such a case, it would be advisable to use a graded analysis approach, wherein faster-running models are utilized at intermediate stages without sacrificing rigor in the final analysis. For example, artificial neural networks and/or recursive equation systems could be used while conducting the resource allocation iterations until convergence is obtained in the search for an optimal allocation scheme. These faster-running models would be trained to simulate the results obtained from the more complex integrated models and would only be applied for testing resource allocations until one was selected. After that, the process would revert to the use of the more complex models.

Chart B' in Figure 17 illustrates a generic recursive equation system that could be used to simulate the more complex functional relationships in Chart B without ignoring important dependencies between variables. It applies in situations where the progression of performance measure evaluations follows a forward-directional path without feedback loops, similar to the example that was shown in Figure 5. The advantage of a recursive equation system over neural networks is that the relationships between inputs and outputs are transparent, whereas the advantage of neural networks over recursive equations is that complicated feedback loops are not neglected. For this reason, it may be best to use both of these faster-running models to conduct the resource allocation iterations rather than just one of them.

It may be noted that the recursive equation set could include both standard algebraic equations and first-order differential equations to provide maximum flexibility in accounting for the sudden onset of steep gradients such as tipping points (Strogatz, 2024). The use of differential equations goes a step beyond the relationships provided in Chart B'.

Chart B





Chart B'



Key to Symbols

E = Event	J = Risk Scen. Index	P = Occurrence Probability	$\vec{x}$ = Vector of input Parameters	Π $\equiv$ Product
f, g = Functions	k = Perf. Meas. Index	PMV = Perf. Meas. Value	$\vec{y}$ = Vector of Leading Indicators	$ $ $\equiv$ Conditional Upon
i = Event Index	n = Iteration Number	$\vec{J}$ = Vector of Risk Scenarios	$\in$ $\equiv$ Membership	

Figure 17. An equation-based form of the integrated performance analysis flowchart that is suitable for graded analyses in support of resource allocation.

5.2 Extending the Methodology to National and International Arenas such as AI and Climate Change

Although the formulation presented in this paper is framed as a means for risk-informing the planning and execution of portfolios of multiple programs at enterprise level, it should be readily extendable to societal arenas requiring cooperation at national and international levels. For example, the diversity of AI platforms created and operated by multiple private organizations and government agencies can be considered to be a “portfolio” of programs that need to be coordinated nationally and internationally so as to ensure a balancing between benefits and risks.

The main distinguishing feature between enterprise portfolios of mission-related programs and international portfolios of AI-related platforms lies in the diversity in the range of opinions of the stakeholders, especially regarding their degrees of concern about safety. At a cursory level, one might believe that expert elicitation regarding AI safety risks is of limited use because of the disparities between stakeholders. After further thinking, however, it becomes clearer that the wide diversity of opinions can be used as the input needed to determine how much time and money should be spent on safety research to narrow the subjective gaps. The contribution of the methodology, and in particular the decision-analytic protocols for resolving biases,

is that it helps determine how resources can be allocated to approach a near-optimal balance between monetary gain, societal benefit, and safety.

Similarly, there is a wide diversity of opinion about how much attention should be given at national and international levels to combatting climate change as opposed to more immediate societal concerns such as the robustness of the economy and the price of commodities. Some believe that the future advance of technology will alleviate the longer-term concern, whereas others do not want to take the risk that it may not. Once again, the methodology in this paper can be applied to elucidate how resources can be allocated to approach a near-optimal balance between competing factors, in this case between concerns about the near-term economy and price of goods and concerns about longer-term safety and security and the uncertainty surrounding future technology developments. The fact that peoples' values differ on this subject does not negate the value of rigorous probabilistic analysis in helping to resolve questions of resource allocation.

6. Discussion

As described in the preceding sections, this paper has emphasized the need for and application of an integrated quantitative probabilistic methodology that (1) focuses on the entire spectrum of enterprise-level objectives, including both programmatic and aspirational objectives, (2) accounts in full for the dependencies between projects, programs, infrastructure development activities, and other initiatives contained within the enterprise portfolio, (3) leads to the determination of near-optimal resource allocations, (4) applies equally well to enterprise-level portfolio decision analysis and operational risk and opportunity management, and (4) is extendable to arenas requiring national and international cooperation, such as AI and climate change.

The formulation is conducted within a Monte Carlo framework, but traditional Monte-Carlo-based approaches that are applied in project-specific risk analyses are modified to simultaneously consider an entire portfolio of programs that interact with one another, as well as an entire set of enterprise-level programmatic and aspirational objectives that also interact with one another. This interactive formalism enables risks and opportunities to be aggregated from the lowest levels of the organization to the top levels without loss of rigor.

Input distributions for the Monte Carlo analysis account for uncertainties emanating from a large variety of sources, including design parameters, environmental parameters, material properties, event occurrence probabilities or frequencies, logical and phenomenological modeling assumptions and approximations, operational reliability, human error probabilities and impacts, expert elicitation variability or bias, the potential for unknown and/or underappreciated (U/U) risk scenarios, leading indicator values, and derived correlations between U/U risks and leading indicators.

The integrated approach uses detailed quantitative models when available together with hybrid quantitative-qualitative models when needed, the latter being informed by the insights and tacit knowledge of the stakeholders. Stakeholder elicitation is relied upon in any area that is difficult to characterize quantitatively because of a lack of data at the time the analysis is being performed. One example is in gauging how critical milestone slippages translate to concern about meeting enterprise-level objectives, which are typically much broader in scope than project-level or program-level objectives. Another is in estimating how much the development of new technology or the discovery of new knowledge during the execution of the portfolio may lead to monetary or nonmonetary benefit in future applications that are not knowable at present. These lead to the inclusion into the analysis of affective performance measures, such as degree of stakeholder concern or stakeholder enthusiasm or stakeholder excitement. Depending on the context, the stakeholders

may include the enterprise's leadership and its advisory committees, its customers and partners, regulatory authorities, the scientific community outside of the enterprise, educational institutions, and the public at large.

The outputs of the analysis include performance measure distributions that can be used to compare how well one candidate portfolio satisfies the enterprise's objectives compared to others during the portfolio selection process, as well as to determine how well an already implemented portfolio is tracking against each enterprise objective, what kinds of alterations in the form of improved controls, new mitigations, or revised resource allocations need to be put forward and decided upon, and in what timeframe these controls, mitigations, or re-allocations need to be implemented. For communication purposes, example spider charts and heat maps that appeal to a wide variety of parties have been schematized and illustrated.

To validate that the methodology recommended in this paper produces verifiable results with proven practical usefulness, future efforts should be oriented toward pursuing large-scale case studies using this methodology and comparing the results to empirical data. For that purpose, the analysts charged with conducting these studies and analyzing the results will need access to all the portfolio details, design information, test protocols, operational procedures, organizational setups, resource allocations, and operational results that are pertinent to the study. For this reason, it would be easiest for each case study to be performed by or managed by the organization that owns the portfolio from which the case study emerges.

As an alternative, it would be possible for the case studies to be performed by an independent investigator, but this would require the willingness of people in positions of leadership within each subject enterprise to share the needed information. While that action in some cases might strain proprietary concerns, it would allow for a level of transparency that would be beneficial to the technical and scientific communities as well as benefitting society as a whole.

References

- Airbus. (2025). Airbus annual report 2024. <https://www.airbus.com/sites/g/files/jlcbta136/files/2025-04/Airbus%20Annual%20Report%202024.pdf>.
- Bearden, D. (2008, June 3). *Perspectives on NASA mission cost and schedule performance trends* [Conference presentation]. NASA Goddard Space Flight Center Systems Engineering Symposium.
- Benjamin, A., Dezfuli, H., & Everett, C. (2016). Developing probabilistic safety performance margins for unknown and underappreciated risks. *Reliability Engineering & System Safety*, 145, 329–340. <https://doi.org/10.1016/j.res.2015.07.021>
- Benjamin, A. S. (2025). *Integration of project, infrastructure, and enterprise risk management: How to evaluate, aggregate, and manage high-tech performance risks* (Risk, Governance and Society, Vol. 21). Springer. <https://doi.org/10.1007/978-3-031-85070-7>
- Boeing Company. (2025). 2024 annual report.
- Bradley, R.A., & Terry, M.E. (1952) Rank analysis of incomplete block designs: I. The method of paired comparisons". *Biometrika*, 39 (3/4): 324–345.
- Committee of Sponsoring Organizations of the Treadway Commission [COSO] (2017). *Enterprise risk management – Integrating with strategy and performance*. COSO.
- Cooke, R.M. (1991). Experts in uncertainty; Opinion and subjective probability in science. *Oxford University Press*.

- Dezfuli, H., Guarro, S., Everett, C., Benjamin, A., & Coan Skow, M. (2024a). *NASA risk management handbook: Version 2.0, Part 1* (NASA/SP-20240014019). National Aeronautics and Space Administration.
- Dezfuli, H., Guarro, S., Everett, C., Benjamin, A., & Coan Skow, M. (2024b). *NASA risk management handbook: Version 2.0, Part 2: Guidance for an integrated framework of risk leadership and management across the NASA organization: Application examples* (NASA/SP-20240014326). National Aeronautics and Space Administration.
- Dickinson, M.W., Thornton, A.C., & Graves, S.C. (2001) Technology portfolio management: Optimizing interdependent projects over multiple time periods.” *IEEE Transactions on Engineering Management*, 48(4): 518-527.
- DiMasi, J. A., Grabowski, H. G., & Hansen, R. W. (2016). Innovation in the pharmaceutical industry. *Journal of Health Economics*, 47, 20–33.
- European Space Agency (2021). *ESA agenda 2025* [.-unclassified for public release]. https://www.esa.int/About_Us/ESA_Publications/Agenda_2025.
- Gosling, J. P. (2018). SHELF: The Sheffield elicitation framework. In L. C. Dias, A. Morton, & J. Quigley (Eds.), *Elicitation: The science and art of structuring judgement* (pp. 61–93). Springer.
- HM Treasury. (2023). “The Orange Book: Management of risk—Principles and concepts,” HM Treasury. <https://www.gov.uk/government/publications/orange-book>
- International Organization for Standardization [ISO]. (2018). *Risk management — Guidelines* (ISO Standard No. 31000:2018), <https://www.iso.org/standard/65694.html>
- Kaplan, R. S., & Norton, D. P. (1992). The balanced scorecard—Measures that drive performance. *Harvard Business Review*, 70(1), 71–79.
- Keeney, R. & Raiffa, H. (1976). *Decisions with multiple objectives: Preferences and value tradeoffs*. Wiley.
- Lean Enterprise Institute. (2020). *Understanding Hoshin Kanri*. <https://info.lean.org/understanding-hoshin-kanri>
- Liesiö, J., Salo, A., Keisler, J. M., & Morton, A. (2020). Portfolio decision analysis: recent developments and future prospects. *European Journal of Operational Research*, 293, 811-825, <https://doi.org/10.1016/j.ejor.2020.12.015>
- Linstone, H. A., & Turoff, M. (Eds.). (1975). *The Delphi method: Techniques and applications*. Addison-Wesley
- Lynch, V. (2025, January 3). Emerging technologies – preserving and growing value over multiple innovation timelines. *R&D Today*. <https://www.rndtoday.co.uk/project-valuation-and-selection/how-to-value-innovation-in-emerging-technologies>
- Mars Climate Orbiter Mishap Investigation Board. (1999). *Mars Climate Orbiter Mishap Investigation Board phase I report*. National Aeronautics and Space Administration.
- National Aeronautics and Space Administration. (2026). *FY 2027 Volume of Integrated Performance (VIPer)* [Performance plan and report]. <https://www.nasa.gov/wp-content/uploads/2026/08/nasa-fy-27-volume-of-integrated-performance-508compliant-final8-3.pdf>.
- Nonaka, I. (2007). The knowledge-creating company. *Harvard Business Review*, 85(7/8), 162–171.

- Oreskovich, A.-M., & Gittins, J. (2016). The evaluation of risk in pharmaceutical research: A study of current models and techniques. *R&D Management*, 46(5), 900-913. doi.org.
- Perrow, C. (1999). *Normal accidents: Living with high-risk technologies*. Princeton
- Saaty, R. W. (1987). The analytic hierarchy process—What it is and how it is used. *Mathematical Modelling*, 9(3–5), 161–176. [https://doi.org/10.1016/0270-0255\(87\)90473-8](https://doi.org/10.1016/0270-0255(87)90473-8)
- Strogatz, S. H. (2024). *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering*. Third edition. CRC Press. <https://doi.org/10.1201/9780429398490>
- Thurstone, L.L. (1927). A law of comparative judgment. *Psychological Review*, 79, 281-299.
- U.S. Government Accountability Office. (2015). *Defense acquisition process: Military service chiefs' concerns reflect need to better define requirements before programs start* (GAO-15-469).
- U.S. Government Accountability Office. (2025). *Standards for internal control in the federal government* (GAO-25-107721). <https://www.gao.gov/products/gao-25-107721>
- Weirich, P. (2012). Multi-attribute approaches to risk. In S. Roeser, R. Hillerbrand, P. Sandin, & M. Peterson (Eds.), *Handbook of risk theory*, 517–543. Springer.

ACKNOWLEDGEMENTS

In researching and writing this paper, I relied on the insights and assistance of several people. Chief among them is Dr. Homayoon Dezfuli, the Technical Fellow for System Safety and Risk Management at NASA's Office of Safety and Mission Assurance (recently retired), who was the project lead for the preparation of NASA's procedural requirements and associated handbooks on these two subjects. In close proximity are Sergio Guarro and Chris Everett, who were co-authors with me on several versions of these handbooks. I also wish to acknowledge the many people who worked with me over the years when I was either leading or performing risk analyses and writing about it for various government agencies, national laboratories, and commercial contractors. Finally, I want to thank my family (Gail, Corinna, and Frank), who allowed me the latitude to work on this article at home more-or-less full time with relatively few interruptions.