INCORPORATING RISK TOLERANCE AND SIMPLE OPTIMIZATION INTO THE RDF

FOR CALIFORNIA PUBLIC UTILITIES COMMISSION

10/07/2024

Title	Incorporating Risk tolerance and simple optimization into the RDF to improve the efficiency and effectiveness of risk mitigation selection		
Document #:			
Version:	Final Report	**Date:**	10/07/2024
Author(s):	Level 4 RSE Team		
Customer:	California Public Utilities Commission (CPUC)		
Contract:	CPUC 23NC0189		
Project:	Refining the Risk-Based Decision-Making Framework to Support an Expedited Undergrounding Program		
Deliverable ID:	0004B		

Version History

Version	Author	Date	Changes
DED	MSR	7/12/2024	Original DED
Working Paper	MSR	08/05/2024	Pre-final draft
Final Report	MSR	10/07/2024	Final Report

Approvals

Role	Name	Signature	Date
Author	Matthew Raphaelson		10/07/2024
CEO	Sam Featherstone		10/07/2024
COO	Erik Fike	Erik Fike	10/07/2024

Table of Contents

Table of Figures

1 Project Overview

1.1 Project Description

The Safety Policy Department (SPD) of the California Public Utilities Commission (CPUC) has engaged Level 4 as an expert contractor to provide complex and technical expertise on risk tolerance and simple optimization. This report will present Level 4's views on the importance of risk tolerance for making risk-based decisions, and how to use simple optimization techniques to improve the quality and transparency of risk mitigation selection. Level 4 will provide SPD with a set of recommendations on how the Risk-Based Decision-Making Framework (RDF) can be modified for risk tolerance and simple optimization.

1.2 Purpose

This contract addresses the requirements of California Senate Bill (SB) 884 to develop, administer, and enforce new standards for an expedited electric utility distribution infrastructure program. Critical efforts include refining the RDF to allow for an improved decision support process.

1.3 Approach

Level 4 will discuss the concepts of risk tolerance and simple optimization and how they relate to risk-based decision-making. Level 4 will

- Review the scientific literature and discuss risk tolerance and simple optimization using numerical examples and visualizations.
- Build a case for incorporating risk tolerance and simple optimization techniques in selecting mitigations within the RDF, with potential application for SB 884 and wildfire risk management in general.
- Develop recommendations for how to modify the RDF for risk tolerance and simple optimization, and how to transition to the modified RDF.

As much as possible, the numerical examples and visualizations will be based on wildfire risk and other risks faced by CPUC-regulated utilities, such as (but not necessarily limited to) cyber risk and hydro-power risk.

2 Executive Summary

2.1 Overview

This paper will explore the special relationship between a probabilistic view of risk, risk tolerance, and optimization. Our goal is to provide more technical insight into each piece and to suggest practical guidance on incorporating them more fully into the RDF

A probabilistic understanding of risk, risk tolerance, and optimizing risk-based decisions are three legs of a finely balanced stool. Building a case for risk reduction that focuses on any two while ignoring the third is at best a rickety proposition.

Risk is the chance of something bad happening. Embedded in those deceptively simple words are the laws of probability behind the word "chance," and the subjectivity behind the perception of "bad." The over-arching theme of this paper is that risk cannot be represented as a single number and must instead be represented by probability distributions. A companion theme is that risk "is in the eye of the beholder," and that subjectivity plays a crucial role in risk-based decisions.

Risk tolerance is a touchstone for deciding whether to mitigate a risk or to accept it. It is probabilistic and individualistic—risk tolerance is "shaped" by how an individual views small chances of very bad outcomes. Without an explicitly declared risk tolerance, evaluators can't assess whether risk mitigation objectives are being achieved.

The formal study of risk and risk tolerance emerged in the 16th century with the mathematician Gerolamo Cardano, who wrote the "Book on Chance and Games." The study was advanced in the 17th century through the correspondence between the great mathematicians Blaise Pascal and Pierre de Fermat.[1]

The study of risk and risk tolerance flourished in the 18th century. In 1738, Swiss mathematician and physicist Daniel Bernoulli presented his "Exposition of a new theory on the measurement of risk" to the St. Petersburg Academy. In what is now known as the St. Petersburg Paradox, Bernoulli showed how rational, individuals would not always be willing to pay the fair (or "expected value" (EV)) of a bet, especially if there was the possibility of large losses.[2] The St. Petersburg paradox led to the development of utility theory in economics, and the understanding that people will behave differently when facing risk depending on many factors, including wealth, income, age, gender, and past experience.[3]

---

[1] Peter L. Bernstein, "Against the Gods: The Remarkable Story of Risk," chaps. 4 and 6. (John Wiley & Sons, 1996).

[2] "Risk and uncertainty I: St. Petersburg paradox," Policonomics, 2024, https://policonomics.com/lp-risk-and-uncertainty1-saint-petersburg-paradox/

[3] Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jurgen Schupp, Gert G. Wagner, "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences," *Journal of the European Economic Association* 22, no. 3 (June 2011): 522-

A decision that makes sense for a person (or an organization) with a high-risk tolerance,[4] due to significant financial reserves, might be catastrophically irresponsible for another. Assessing whether that decision is sensible or catastrophic is impossible without knowing the underlying risk tolerance. Risk and risk tolerance will be covered in Chapter 4.

Optimization is a quantitative technique for determining the best choice, subject to conditions and constraints. It can be probabilistic (called stochastic optimization). For mitigations selection, optimization is superior to straight ranking when there are interrelationships between them, such as mutual exclusivity, synergies, or diminishing returns.

A simple optimization approach is based on portfolios of mitigations and draws on the experience of the finance and insurance industries.[5] Harry Markowitz, the Nobel laureate in economics and founder of Modern Portfolio Theory (MPT), made the fundamental insight that the performance of an individual asset in a portfolio is not as important as the performance of the entire portfolio. While Markowitz was focused on portfolios of financial stocks, the theory also holds for portfolios of other assets or projects, such as risk mitigations.

Different portfolios of mitigations can be plotted based on mitigation impact versus cost, which will reveal an "efficient frontier" of optimal portfolios. A portfolio on the efficient frontier achieves the highest amount of risk reduction for its level of expense—other portfolios may achieve higher risk reduction, but at higher cost. Since an efficient frontier normally contains multiple portfolios, the final selection of a single portfolio depends on budget and other considerations, such as available resources, safety versus reliability impact, the level of risk at different probabilities—and ultimately risk tolerance. We have come full circle. Portfolios of mitigations and simple optimization will be covered in more detail in Chapter 5.

2.2 Key Findings and Recommendations

Level 4 recognizes that the RDF is a journey, and noteworthy progress has been made. Many of the topics covered in this report have been raised and discussed in various CPUC decisions, but realistically deferred until future phases. Level 4 believes that the capabilities of the utilities have matured to the point where implementation can be accelerated.

Level 4 recommends that CPUC incorporate risk tolerance into the RDF and move toward a more quantitative optimization approach that reflects probabilistic concepts such as tail risk for evaluating and selecting risk mitigations.

---

550. https://www.swarthmore.edu/sites/default/files/assets/documents/user_profiles/dhuffma1/Individual_risk_attitudes_JEEA.pdf

4 We will define risk tolerance more formally later; for now, think of it as willingness to accept the possibility of large losses.

5 Ann Behan, "Harry Markowitz: Creator of the Modern Portfolio Theory." 2024. https://www.investopedia.com/terms/h/harrymarkowitz.asp

A summary of Level 4's recommendations is presented in Figure 2-1.

Recommendation	Description
(R1): *Use of probability distributions.*	Probability distributions describe the range and chance that a set of outcomes occurs within datasets and model results. Risk models must use probability distributions as inputs and return probability distributions as outputs.
(R2): *Include and define tail risk as a risk measure.*	In addition to using average risk, defined as the average of the probability distribution of risk, tail risk should be formally added for risk evaluation. The measure of tail risk should be tail average above a percentile (the percentile to be determined by the Commission in consultation with stakeholders). Tail average is preferred over other measures because it captures the entire tail of the distribution, is stable, and can be optimized using linear programming or other methods.
(R3): *Evaluation based on portfolios of mitigations.*	Risk reduction evaluation should be based on portfolios of risk mitigations to account for interrelationships between mitigations. Portfolio selection is well-suited to optimization (see R4).
(R4): *Portfolio selection based on simple optimization instead of ranking.*	Optimization ensures choosing the best portfolio of mitigations given the objective and constraints. It can, however, be a complex, computationally intensive, and time-consuming process. There are ways to simplify the optimization process such as limiting the number of optimization scenarios and choosing objectives that can be optimized using linear programming, which is computationally efficient and speedy compared to non-linear methods.
(R5): *Calculation of risk tolerance.*	Risk tolerance should be modeled as an exceedance curve and calculated by applying the risk neutral or risk averse scaling function to a constant risk exceedance curve.
(R6): *Establish risk tolerance representing the residents of California.*	Risk tolerance is the benchmark that determines whether utility risk levels are acceptable or not. Developing a set of acceptable risk levels that represents the risk tolerance of the residents of California requires an inclusive process, which should begin as soon as possible.

Figure 2-1. Summary of recommendations.

Detailed recommendations including proposed language changes to the RDF are presented in Chapter 7.

3 Background and Current State

3.1 Risk Tolerance and Simple Optimization in the RDF

In December 2014, the CPUC published Decision (D.)14-12-025, the "Decision Incorporating a Risk-Based Decision-Making Framework into the Rate Case Plan and Modifying Appendix A of Decision 07-07-004." The decision was in response to the 2010 natural gas explosion and fire in San Bruno and followed the Commission Rulemaking (R.)11-02-019 and several gas safety bills. The legislation required the Commission to "develop formal procedures to consider safety in a rate case application by an electrical corporation or gas corporation."[6]

D.14-12-025 launched two new procedures, the filing of a Safety Model Assessment Proceeding (S-MAP), and a Risk Assessment Mitigation Phase (RAMP). The decision acknowledges that the Commission "need(s) to require testimony in the General Rate Cases (GRCs)…an assessment of its risk tolerance, identifying areas of low risk and high risk…"[7]

In August 2016, the Commission published D.16-08-018, the "Interim Decision Adopting the Multi-Attribute Approach (or Utility Equivalent Features) and Directing Utilities to Take Steps Toward a More Uniform Risk Management Framework." This decision explores risk tolerance in detail and provides a definition: "maximum amount of residual risk that an entity or its stakeholders are willing to accept after application of risk control or mitigation."[8]

The decision finds that there are problems with the utilities' models that preclude them from implementing risk reduction and risk mitigation strategies consistent with D.14-12-025, including not having an explicit risk tolerance and no optimization of the portfolio of risk mitigation activities.[9]

There was broad agreement between the Commission, the utilities, and the intervenors that some form of risk tolerance is required. The Joint Utilities acknowledge the need for a risk tolerance framework.[10] The intervenors were also in general agreement about the importance of risk tolerance.[11]

In contemplating potential future steps, the decision discusses the As Low as Reasonably Practicable (ALARP) framework.[12] Commission staff originally published a whitepaper on ALARP as part of Workshop

---

6 CPUC decision D.14-12-025, page 4, footnote 3.

7 CPUC decision D.14-12-025, page 5.

8 CPUC decision D.16-08-018, page 25.

9 CPUC decision D.16-08-018, page 164.

10 CPUC decision D.16-08-018, page 47.

11 CPUC decision D.16-08-018, pages 70-80.

12 CPUC decision D.16-08-018, page 62.

#4 of Phase 1 of the S-MAP Proceeding in December 2015.[13] ALARP combines risk tolerance with a three-tiered optimization process and is focused on safety risk. It has been enshrined in the United Kingdom case law for the regulation of health and safety since 1949 and is also applied in other countries including Australia, Norway, and the Netherlands.[14] In the U.S., the Army Corps of Engineers has used it,[15] as has the nuclear radiation industry.[16]

It was determined that ALARP may be a desirable end state but would not be possible to implement at the time it was discussed. There are other ways to incorporate risk tolerance and optimization short of full implementation of ALARP; even so, the decision determined that "the Commission should adopt explicit risk tolerance standards over time, but not before laying the groundwork in the development of probabilistic risk analysis."[17] During Phase 2 of the S-MAP Proceeding, developing a risk tolerance framework and increasing the application of optimization were included among eight suggested long-term goals.[18]

Over the next several years, risk tolerance and optimization remained in the conversation. CPUC D.18-12-014, the "Phase Two Decision Adopting Safety Model Assessment Proceeding (S-MAP) Settlement Agreement with Modifications," issued in December 2018, stated that the "settlement agreement does not preclude other long-term goals of the Commission, such as 'optimization' and 'explicit risk tolerance standards.'"[19] It affirms that ALARP remains a priority topic.[20]

In the Risk-based Decision-making Framework Proceeding (R.20-07-013), the CPUC decision D.22-12-027, "Phase II Decision Adopting Modifications to the Risk-Based Decision-Making Framework Adopted in Decision 18-12-014 and Directing Environmental and Social Justice Pilots," issued in December 2022,

---

[13] *Safety and Enforcement Division Staff White Paper on As Low as Reasonably Practicable (ALARP) Risk-informed Decision Framework Applied to Public Utility Safety*. California Public Utilities Commission. (2015, December 24). https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M157/K359/157359431.PDF

[14] Noor Quddus, Denis Su-Feher, Christopher Gordon, Jyoti Sharma, and Troy O'Brien. "Risk Acceptance Criteria: Overview of ALARP and Similar Methodologies as Practiced Worldwide." Mary Kay O'Connor Process Safety Center, Texas A&M Engineering Experiment Station, 2020. https://psc.tamu.edu/wp-content/uploads/sites/2/2020/08/ALARP-Final-Paper-Publishing.pdf

[15] Isabella Dam Safety Modification Study, U.S. Army Corps of Engineers Response to Independent External Peer Review, (October 2012): 5. https://www.spk.usace.army.mil/Portals/12/documents/civil_works/Isabella/Final%20Agency%20Response%20to%20IEPR%20-%20Isabella%20Dam%20%282%29.pdf

[16] Regulatory Guide 8.10, U.S. Nuclear Regulatory Commission Office of Nuclear Regulatory Research, August 2016 (called ALARA, which is similar to ALARP). https://www.nrc.gov/docs/ML1610/ML16105A136.pdf

[17] CPUC decision D.16-08-018, page 192.

[18] CPUC decision D.16-08-018, page 175.

[19] CPUC decision D.18-12-014, page 41.

[20] CPUC decision D.18-12-014, page 55.

authorized technical working groups (TWGs) to propose recommendations regarding the application of risk tolerance.[21]

In September 2024, the assigned commissioner issued the Phase 4 Scoping Memo and Ruling as part of the RDF Proceeding.[22] Phase 4 priorities include risk tolerance standards and methodology, addressing overall residual risk, and mitigation selection optimization. The memo states that "a standard method is needed to integrate risk tolerance into the RDF and inform future RAMP and GRC filings" and "we are concerned that a risk tolerance goal that is too high or too low will yield suboptimal outcomes for ratepayer safety or ratepayer costs, respectively."[23] The memo further states that "the Commission should explore basic risk mitigation optimization techniques by requiring the utilities to identify and quantify the key constraints affecting their selection of mitigation options for implementation."[24]

3.2 Current State of Risk Tolerance and Simple Optimization

Based on its reviews, Level 4 has observed that risk tolerance and optimization are missing in utilities' RAMPs and Wildfire Mitigation Plans (WMPs). Level 4 further observes that these omissions detract from the RAMPs and WMPs consistent with the Commission decisions cited above. Specifically,

- In the RAMPs filed between 2020 and 2022, the utilities were in the initial stages of developing their risk mitigation modeling programs, and there is little mention of risk tolerance or optimization. San Diego Gas & Electric (SDG&E) is an exception to some extent, frequently referring to portfolios of mitigations and relating portfolio selection to Risk Spend Efficiency (RSE), though this falls short of optimization.
- In the 2022 and 2023 WMPs, none of the utilities adequately explained *why* the level of wildfire mitigation had been chosen. Why was the level of post-mitigation wildfire risk deemed satisfactory? This most basic question cannot be answered without acknowledging risk tolerance.
- In the 2022 and 2023 WMPs, none of the utilities adequately described *how* wildfire mitigations were chosen. While analytics were performed and RSE curves were constructed, it was unclear how these were used in mitigation selection; in fact, the selections appeared to be made mostly based on subject matter expert opinion.

Level 4 agrees with Commission staff that by not providing an "explicit specification of risk tolerance, the utilities are handicapping the ability of other stakeholders to make an informed decision as to whether the utilities' rate case proposals would have the desired risk reduction effect in relation to the desired level of

[21] CPUC decision D.22-12-027, page 29.

[22] See https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M539/K999/539999025.PDF

[23] Ibid., pages 3-4.

[24] Ibid., page 5.

risk tolerance…The utilities would in effect be asking the stakeholders to accept in blind faith that the proposed programs and projects are necessary and sufficient…to mitigate the risk down to a level that the utilities can tolerate, whatever that level is."[25]

Level 4 also agrees with Commission staff that "there is no optimization of the portfolio of risk mitigation activities. None of the utilities have a way to optimize their portfolio in a mathematically rigorous sense…Programs and projects are prioritized but not optimized…Inherent in risk management is the unavoidable fact of limited resources and other constraints. Without resource constraints, an operator could simply apply an infinite amount of an infinite number of risk mitigation activities and the risk would be driven to zero. Clearly, this is a reduction of the argument to an absurdity. Therefore, risk management always assumes recognition of some constraints…***optimization is always tied to risk tolerance. These concepts are all tied together*** (emphasis added)."[26]

We note that the very next sentence in the decision states that most parties agreed with the above statement and considered it a long-term priority. That was over seven years ago, and little if any progress has been made in incorporating risk tolerance or implementing optimization. In the following chapters, Level 4 will develop the case for moving ahead with risk tolerance and simple optimization, present tools and techniques for implementation, and suggest timing, pacing, and sequencing considerations.

[25] CPUC decision D.16-08-018, page 68.

[26] CPUC decision D.16-08-018, page 98.

4 Risk and Risk Tolerance

4.1 Risk is in the Eye of the Beholder

A Google search of "risk definition" results in hundreds, probably thousands of alternatives. We like the simple and direct "the chance of something bad happening" from the Cambridge Online Dictionary.[27] The CPUC definition of risk is similar, with additional detail "The potential for the occurrence of an event that would be desirable to avoid, often expressed in terms of a combination of various Outcomes of an adverse event and their associated Probabilities."[28]

In our simple definition, "chance" means a probability from 0% to 100%, and "bad" is a subjective interpretation of an outcome. The following parable will illustrate the probabilistic and subjective elements of risk and the interplay between them. We will use the Factor Analysis of Information Risk (FAIR) ontology for expressing risk, where risk equals the likelihood of risk event (LoRE) multiplied by the consequence of risk event (CoRE).[29]

Three venturers in a large metropolitan area approach a major bridge. Each venturer needs to reach their destination across the bridge on time or will lose something of value. Unfortunately, there was an accident on the bridge and traffic was backed up for miles. The three venturers happen to be near a heliport and have the option to take a helicopter into the city, for $150.

Venturer 1 (V1) spent $100 on tickets to a ballgame in the city. V1 sees the world in black-and-white terms and is certain that the traffic jam will result in missing the entire game. The risk of loss, as perceived by V1, is 100% likelihood x $100 consequence = $100.

Venturer 2 (V2) has several client appointments in the city and would lose fee income if late. V2 assesses there is a good chance that the traffic will clear up quickly based on years of commuting experience, confirmed by projections from a GPS app. V2 estimates there is a 50% chance of making all appointments on time, and a 50% chance of missing one appointment and $200 in income. The risk of loss, as perceived by V2, is 50% likelihood x $0 consequence + 50% likelihood x $200 consequence = $100.

Venturer 3 (V3) has become increasingly alarmed about the possibility of a cyberattack on local banks and depositors being robbed of their savings. V3 is concerned that because of the traffic, there is a 10% chance based on gut feel of not making it to the bank before the attack and will lose $1,000 in savings. The risk of loss as perceived by V3 is 90% likelihood x $0 consequence + 10% likelihood x $1,000 consequence = $100.

The three venturers represent three archetypes for perceiving risk. V1 perceives risk deterministically, with no sense of probabilities. V2 perceives risk probabilistically and assesses it based on data and experience. V3

---

[27] https://dictionary.cambridge.org/us/dictionary/english/risk

[28] CPUC R20-07-013, Appendix A, page A-5.

[29] See "The FAIR Standard," Risk Lens, https://www.risklens.com/cyber-risk-quantification/the-fair-standard

also perceives risk probabilistically but assesses it based on belief and instinct. The three perceptions of risk can be visualized in Figure 4-1.

Figure 4-1. Three perspectives - the subjective nature of risk

Though the three venturers assess the chances of something bad happening differently, for each the expected loss is $100, despite the very different "shape" of each risk assessment, as illustrated in Figure 4-1.

Does this mean that all risk assessments are correct, or conversely, all are wrong? While there is no such thing as a perfect risk assessment, some are better than others. Cognitive biases and inexperience in estimating likelihoods and consequences can lead to systematic errors in risk assessment. Techniques such as calibration training have been shown to improve the quality of risk assessments.[30]

More complete information and better models can also improve risk assessments, and if individuals are using the same information and the same models, we can expect the risk assessments to converge. This is not necessarily a good thing, since it can lead to group think.[31] The key point is that risk assessments are subjective, and even the best of them may differ.

Back to the three venturers. EV theory suggests that they should reject the price of the helicopter to mitigate the risk for $150, but their decision is more complicated. How will the venturers decide? To answer these the next two sections will introduce risk tolerance and related concepts.

4.2 Risk Tolerance (and Risk Attitude and Risk Scaling)

Before we delve into risk tolerance, risk attitude, and risk scaling, it is helpful to have a background on the history of risk and efforts to quantify it.

[30] For a quick overview, https://medium.com/@wadedeji/the-failure-of-risk-management-62aac1f5dd6d; for a summary of cognitive bias impact on assessing probabilities, https://www.researchgate.net/profile/Michael-Lindell-2/publication/278671139_Chapter_18_Judgment_and_Decision_Making/links/56845f6308ae1e63f1f1cdb4/Chapter-18-Judgment-and-Decision-Making.pdf; for improving risk assessment via calibration, Douglas W. Hubbard, *How to Measure Anything*, 3rd ed. (Wiley, 2014). Chapter 5.

[31] Based on criteria for the Wisdom of Crowds, https://www.crowdwisdomproject.org/the-wisdom-of-crowds/

4.2.1 A Brief Review of the History of Quantifying and Modeling Risk

The Italian mathematician Gerolamo Cardano's (1501-1576) "Book on Games of Chance," written in 1564 but not published until 1663, is considered the first systematic treatment of probability. Cardano used dice to illustrate basic concepts of probability and understood that the odds could be defined as the number of positive outcomes divided by the number of negative outcomes.[32]

The great French mathematicians Blaise Pascal (1623-1662) and Pierre de Fermat (1601-1665) never met in person and only corresponded for a short time in 1654, but that correspondence created the foundation for probability theory. The two discussed a puzzle called "the problem of points," proposed by an Italian monk named Pacioli in 1494. They worked out probabilistic calculations (Pascal using his famous "Pascal's Triangle"), performed combinatorial analyses, and established the principles of EV analysis and risk assessment.[33]

The mathematical study of probability flourished in the 18th century. In 1738, Swiss mathematician and physicist Daniel Bernoulli (1700-1782) presented a paradox that was first published by his cousin Nicolaus in 1713 to the St. Petersburg Academy. The paradox, which became known as the "St. Petersburg Paradox" imagines a game where the house flips a coin and pays out $1 if the first trial is heads and doubles the amount for each trial of tails thereafter, with the game stopping at the first heads. In theory, the EV is infinite—and yet most people are only willing to pay $1 to $2 to play.

Bernoulli's work led to the development of the expected utility theory where decisions are based on utility as opposed to monetary value, the diminishing marginal utility of wealth, and ultimately the concept of risk attitude and risk aversion.[34] Expected utility theory (not to be confused with public utilities regulated by the CPUC) refers to the idea that an individual's assessment of worth may not correspond with monetary value. Think of utility as meaning "usefulness." Closely related is the idea of diminishing marginal utility, which states that an extra $100 is worth less to a millionaire than to a college student. Another form of diminishing marginal utility is that for many people the pain of losing $100 is worse than the joy of winning $100, which is known as risk aversion.[35]

In the 20th century, fields such as behavioral economics have furthered the study of attitudes toward risk. Studies have shown that females are more risk averse, people become more risk averse as they age, and risk aversion increases for parents with more children. On the other hand, tall people are more risk-seeking, as

---

[32] Victor. J. Katz, *A History of Mathematics: An Introduction*, 3rd ed. (Boston: Pearson Education, 2009).

[33] Adrienne E. Lazes, "Pascal and Fermat: Religion, Probability, and Other Mathematical Discoveries," (Skidmore College, 2016). https://creativematter.skidmore.edu/cgi/viewcontent.cgi?article=1119&context=mals_stu_schol

[34] Benjamin Y. Hayden and Michael L. Platt, "The Mean, the Median, and the St. Petersburg Paradox," *NIH* (June 1, 2009). https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3811154/pdf/nihms422969.pdf

[35] The terms risk aversion, risk-seeking, and risk-neutral will be defined and discussed in Section 4.2.2.

are people with highly educated parents, and those reporting high life satisfaction. Risk-seeking increases with wealth.[36]

Individuals have different perceptions and attitudes toward risk as in the three wanderer's parable above, and so do businesses, public utilities, and government agencies. Society is the collection of all these groups, whose views on risk often conflict. As a result, organizations attempting to evaluate risk on society's behalf may have a different perspective on risk and risk tolerance.[37]

The study of risk combines mathematics, statistics, decision science, sociology, and psychology (among other disciplines). We will draw from these disciplines in our discussion of quantifying risk under uncertainty, risk tolerance, and risk-based decision-making.

4.2.2 Defining and Applying Risk Tolerance (and Risk Attitude and Risk Scaling) to Risk Quantification

Risk tolerance, risk attitude, and risk scaling are often used in different ways. For this report, we will use the following meanings:

1. Risk attitude is a subjective expression of the willingness to accept risk.

- *Risk aversion* is the willingness to pay more than the EV of risk to avoid it (e.g., a person or organization is willing to pay $10 to avoid losing an EV of $5).
- *Risk seeking* is the willingness to accept risk instead of paying the EV to avoid it (e.g., a person or organization is willing to pay no more than $5 to avoid an expected loss of $10).
- *Risk neutral* is neither risk averse nor risk seeking, the willingness to pay exactly the EV of risk to avoid it (e.g., a person or organization is willing to pay $10 to avoid losing an EV of $10). Risk-neutral individuals are indifferent to extreme risk as long as the EV is zero or greater. Only EV matters. The implications of risk neutrality and indifference to extreme risk are discussed further in Section 4.4.

2. Risk scaling is the quantification of risk attitudes. This is defined by the CPUC as "a function or formula that specifies an attitude towards different magnitudes of Outcomes including capturing aversion to extreme Outcomes or indifference over a range of Outcomes."[38] In other words, risk scaling is how much one is willing to pay to avoid a risk—and more importantly how much an organization is willing to pay to avoid increasing amounts of risk. Risk scaling can be visualized as follows, in Figure 4-2.

---

[36] Tomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, et al. "Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences." *Journal of the European Economic Association* 9, no. 3(1 June 2011): 526-531. https://academic.oup.com/jeea/article/9/3/522/2298422

[37] John Adams and Michael Thompson, "Taking Account of Societal Concerns About Risk: Framing the Problem." Research Report 035. UK Health and Safety Executive. https://www.hse.gov.uk/research/rrpdf/rr035.pdf

[38] CPUC R20-07-013, Appendix A, page A-5.

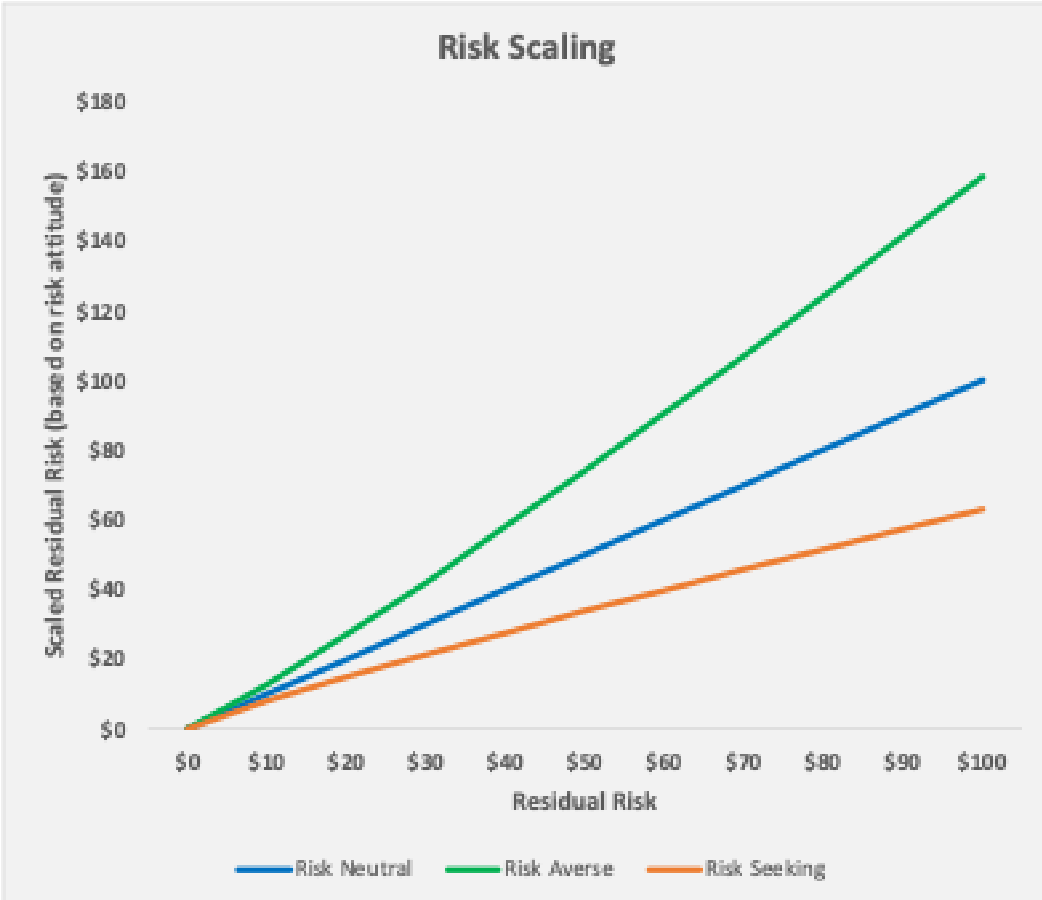

Figure 4-2. Risk scaling curves

All the curves have a positive slope. Risk neutrality is represented by a straight line with a slope equal to 1. This means an organization is willing to pay $1 to avoid losing $1 of risk, or willing to pay $100 to avoid losing $100 of risk.

Risk aversion is represented by a line or curve with a slope greater than 1. Scaled risk is perceived as higher than actual risk, consistent with risk aversion. The details vary depending on the chosen risk aversion curve, but one example of a risk aversion curve would mean that an organization is willing to pay $1 to avoid $0.90 of risk or $100 to avoid $60 of risk.

Risk seeking is represented by a line or curve with a slope less than 1—scaled risk is perceived as lower than actual risk. Similarly, the detail could vary depending on the chosen risk-seeking curve, but one example would be an organization willing to pay no more than $1 to avoid $1.10 of risk, or no more than $60 to avoid $100 of risk.

For utility risk, we are primarily interested in risk aversion and risk neutrality. We will leave further discussion of risk-seeking in fields populated by gamblers and excitement junkies.

3. Risk tolerance is the probabilistic expression of risk attitude. The CPUC definition is the "Maximum amount of Residual Risk that an entity or its stakeholders are willing to accept after application of risk Control or Mitigation."[39] Risk tolerance can be visualized with exceedance curves, as in Figure 4-3 below. The exceedance curve has a negative slope, and each point on the curve depicts the *maximum* level of acceptable risk for the associated probability. Since each point on the curve represents the same risk of $0.01, it is called the constant risk exceedance curve.[40]

The constant risk tolerance curve is useful for translating a risk scaling function to risk tolerance.

- For risk scaling, a risk-averse or risk-seeking function is multiplied against the risk-neutral curve.
- For risk tolerance, the constant risk curve is divided by the risk-averse or risk-seeking function. This transformation changes the shape of the curve, but not its interpretation. For example, the risk-averse scaling curve is convex, but the risk-averse tolerance curve is concave as in Figure 4-4.

The constant risk exceedance curve is not the same thing as risk neutral, which we will explain further in Section 4.4. The relationship between risk scaling and risk tolerance curves is discussed in more detail in Appendix E

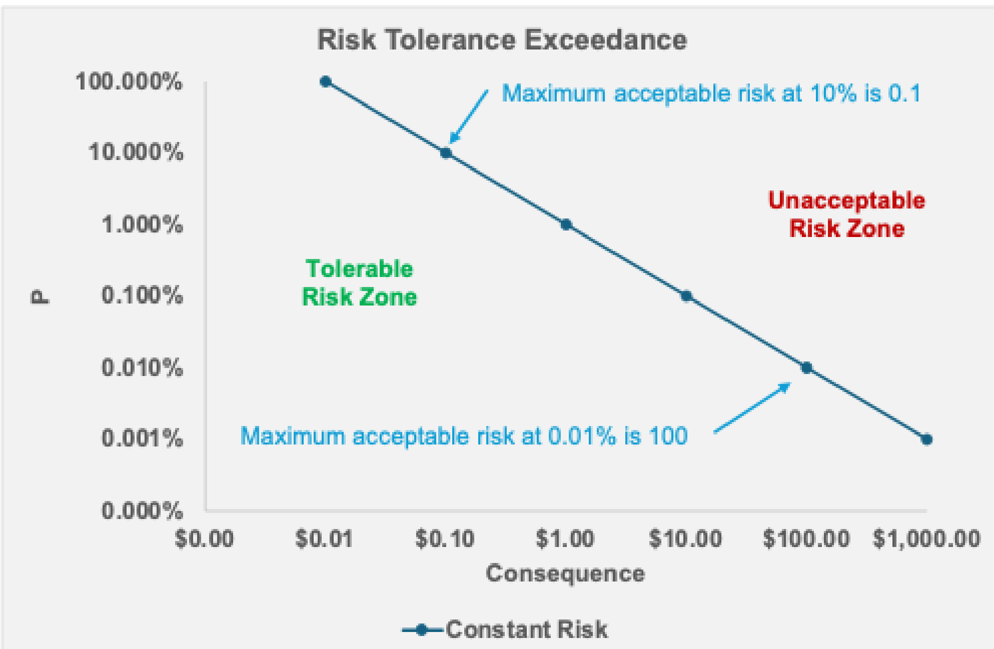

Figure 4-3. Exceedance curve example (in log-log space for readability).[41]

---

[39] CPUC R20-07-013, Appendix A, page A-5.

[40] Also known as iso-risk curve. See Rick Gorvett and Jeff Kinsey, "A Two-Dimensional Risk Measure" (Call Paper Program for 2006 Enterprise Risk Management Symposium). 7-8. https://citeseerx.ist.psu.edu/document?doi=bef8e5125d5dcede72b599c97c6644e520ed6520&repid=rep1&type=pdf

[41] An exceedance curve is the probabilistic representation of a single level of risk, in this case, risk = $0.10. Each point of risk on a risk scaling curve could have its own exceedance curve.

In the example above, the tolerable risk at 10% probability is no more than $0.10, which results in an EV of residual risk of $0.01. At 0.01% probability, the tolerable risk is no more than $100, again resulting in an expected residual risk of $0.01. More generally, any risk level above the curve is unacceptable (for example an expected residual risk of $0.02), while risk levels below the curve are within tolerance.

In Figure 4-4 below, we will add a risk-averse curve by applying a risk-averse scaling function. The risk-averse exceedance curve (green) is below the risk constant line, signifying a lower maximum acceptable level of residual risk for relatively infrequent but more extreme events.

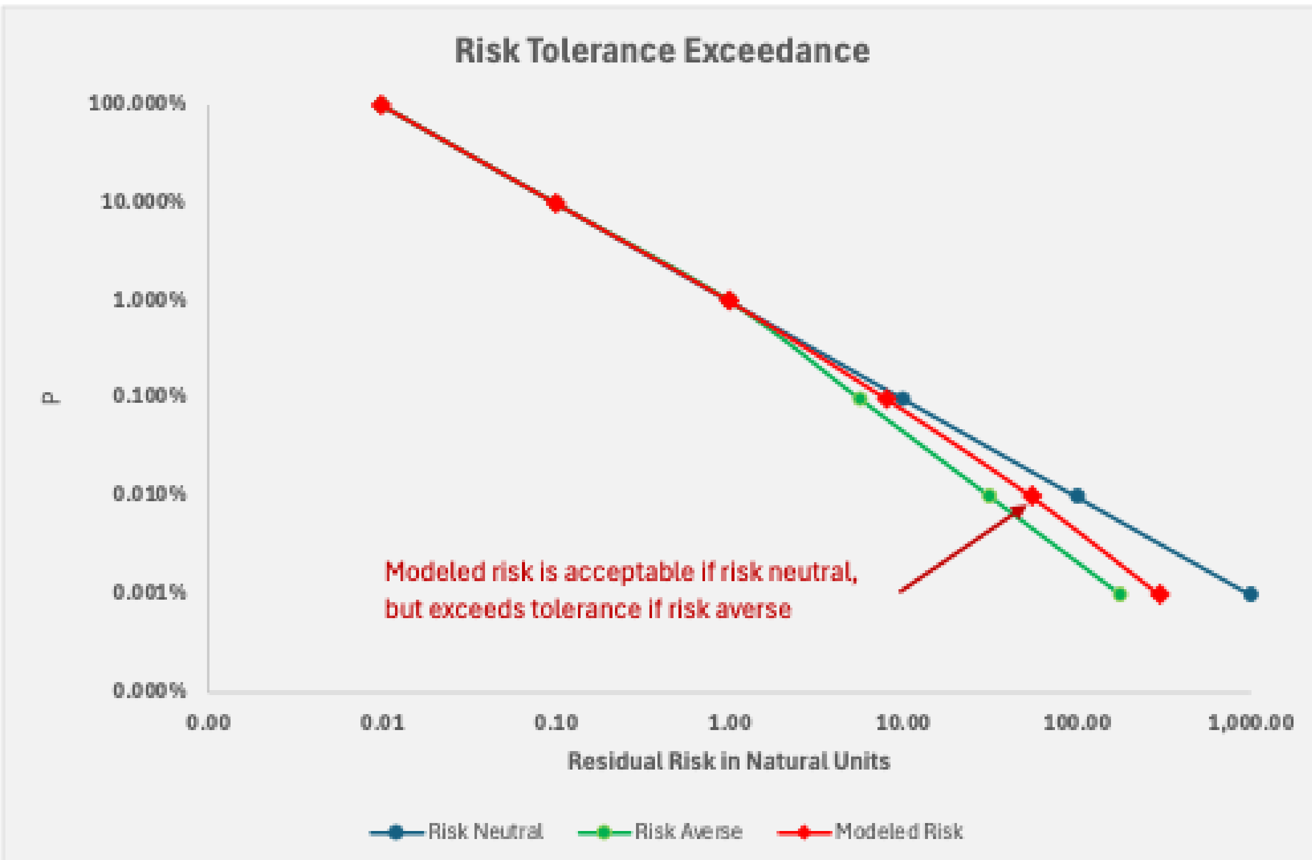

Figure 4-4. Impact of different risk tolerances in evaluating a modeled risk.

We can overlay a modeled risk, represented by the red line, which could represent the new level of risk after mitigations have been applied, and determine whether it meets risk tolerance standards. At below 1% exceedance, the red line lies below the blue curve, which would be acceptable for a constant risk tolerance. For the risk-averse curve, however, the red line would exceed tolerance.

Imagine if there were no risk tolerance curves in Figure 4-4. How would anyone accept that the level of risk reduction represented by the red line was sufficient? Different stakeholders may differ in their risk attitude, some may be risk-neutral, others risk-averse. Even among the risk-averse, some will be more averse than others. Without a risk tolerance standard, determining whether the red line marks an acceptable level of risk will require deliberations between the stakeholders *for every point along the line*, and if they are unaware of each other's risk attitude, many of their voices will argue past each other.

Making the risk tolerance standard explicit and transparent shifts the deliberative process to the standard rather than to each individual project. After the initial effort to determine risk tolerance standards, it is straightforward to assess whether residual risk is within tolerance.

4.3 Risk Cannot be Represented as a Single Number

The key concept underlying the discussion so far is that risk cannot be represented by a single number, such as an average or a percentile. Each of the three wanderers faces an average risk of $100, but the histograms in Figure 4-1 illustrate very different assessments of the consequences and probabilities. These are not the same risks.

Risk scaling, as discussed in CPUC decision D.24-05-064, does not solve the problem of single-number risk scores. It merely replaces one single number with another. A risk-scaled single-number is akin to selecting a percentile from the underlying probability distribution, which simplifies the calculations and decision-making. However, a risk-scaled single-number may oversimplify and cause non-optimal decision-making, especially when comparing different risk types. More specifically, when used in calculations, single-number risk scores and risk scaling run afoul of the laws of the arithmetic of uncertainty in three critical ways:

1. *The Flaw of Averages*. This is a systematic set of errors that occurs when using single numbers such as averages as inputs in complex models.[42] The Flaw of Averages is accentuated by non-linear functions, and especially for power law distributions used in modeling many types of risk including wildfire consequences. Appendix B will provide a fuller discussion of the Flaw of Averages.

2. *The Flaw of Extremes*. These are mathematical errors that occur when extreme results such as 90th percentiles are added as single numbers, related to the Flaw of Averages.[43] Adding single risk scores taken as percentiles from a distribution will likely result in a total risk that vastly overstates actual risk, which can lead to over-investing in risk mitigation. Appendix C will provide a fuller discussion of the Flaw of Extremes.

3. *Likelihood of Simultaneous Failures (LoSF)*. This is the probability that two risk events occur at the same time, which can greatly increase if the risks are interrelated. Often a factor in catastrophic events, e.g., "perfect storms."[44] It is impossible to capture simultaneous failures with single-number risk scores. Appendix D will provide a fuller discussion of LoSF.

The alternative to using single numbers is to use probability distributions. Probability distributions can be added, subtracted, and multiplied (including by scaling functions) using the arithmetic of uncertainty. There is compelling evidence that the large utilities have the underlying probabilities, which means the "raw

---

[42] See https://johnmjennings.com/beware-the-flaw-of-averages/ . Also, Sam L. Savage, *The Flaw of Averages*. (John Wiley & Sons, 2009).

[43] Sam L. Savage. *The Flaw of Averages*. Op.cit. Chapter 17.

[44] We attribute the term likelihood of simultaneous failure and LoSF to Dr. Sam Savage.

materials" for proper risk modeling are available.[45] Appendix A will provide a fuller discussion of the arithmetic of uncertainty.

The Flaw of Averages, Flaw of Extremes, LoSF, and the arithmetic of uncertainty will also be significant topics in our Guidance on Interrelationships Report.

4.3.1 Using the Whole Probability Distribution

Why rely on single numbers to represent risk when there is an entire probability distribution to work with? One can still calculate the average, the median, or any percentile for use in summary reports (i.e., the single number representation of risk). When adding different risks together or aggregating risk through an organization's hierarchy, using probability distributions ensures proper results.

In addition, the probability distribution includes the extreme risks present in the tail of the distribution, known as tail risk. Using the whole probability distribution allows us to use all these representations of risk—and use them simultaneously, such as the average risk and the tail risk. Once a probability distribution is reduced to a single number, it is no longer possible to model the effect of the most extreme catastrophic risks.

The next section will discuss tail risk and diverse ways to represent it in more detail.

4.3.2 Tail Risk Concepts

"*The problem with the standard way of thinking about risk is that it is focused on the average and ignores the impact of rare, extreme events. We are often too focused on the average or the median, and miss the importance of the outliers, the Black Swans, which are responsible for the majority of the impact.*" – Nassim Nicholas Taleb, The Black Swan

Extreme events are where ruin is found – Benoit Mandelbrot

The Commission, the utilities, and the intervenors all understand the importance of risk in the tails of the distribution. The question is how to incorporate tail risk into the RDF.

Level 4 believes that a single-number scaled risk score is not the best approach, preferring alternative methods for expressing and evaluating tail risk. Figure 4-5 below represents a hypothetical wildfire risk power law distribution.[46] Note that this is a visualization of pre-mitigated risk, not risk tolerance. The tail risk is the flat part of the curve extending to the right, which represents low probability—and potentially catastrophic—risks.

[45] Examples from WMP and RAMPs covered in more detail in Chapter 6.

[46] The power law is typically applied to the consequence attribute, but the resultant risk calculation will retain the power law shape.

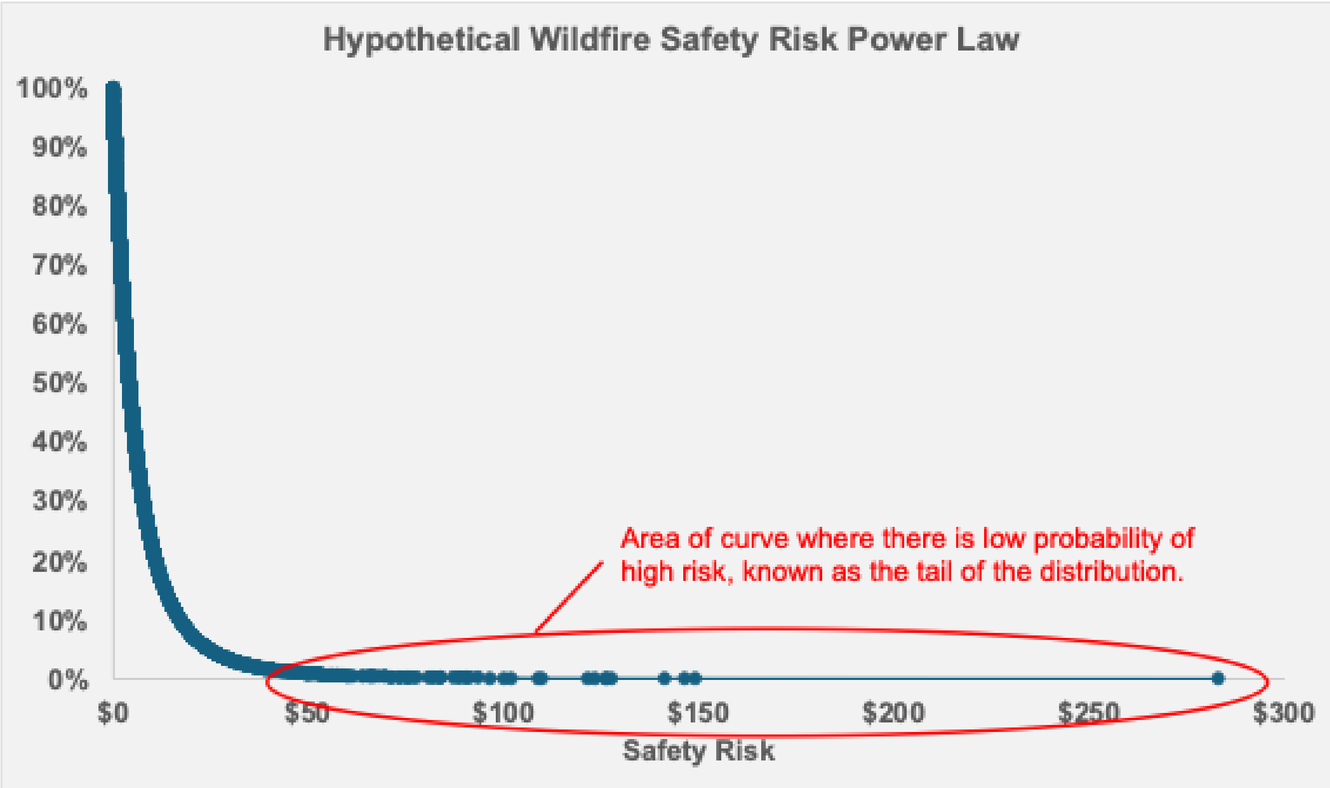

Figure 4-5. Hypothetical wildfire safety risk power law.

The average of the curve in Figure 4-5 is 7, which clearly is not a good representation of the risk, at least not in isolation. We are more interested in the small but non-zero chance of extreme risks located in the tail. Figure 4-6 below illustrates four ways of expressing tail risk:

1. *Scaling function.* A convex (risk-averse) function shifts the curve to the right, which increases the perceived risk.[47]

2. *Percentile.* A single value at a point on the curve. In Figure 4-5 above, the chosen percentile is 99% (the risk, which occurs 1% of the time), which equates to the point on the curve at $50.

3. *Tail average.* The average of the tail above a chosen percentile. In the example above, the tail average is defined as the risk above the 99th percentile (the risk occurs less than 1% of the time), which corresponds to the point on the curve at $50. All values at $50 and above are averaged, capturing the tail.

4. *The power law curve itself.* This is the same as using an infinite number of percentiles and is the same as the exceedance curves discussed earlier.

[47] Risk scaling processes the original risk into a new distribution. This can cause confusion between the scaled risk and the actual risk, and how to interpret. As will be discussed later, risk scaling is not Level 4's preferred approach.

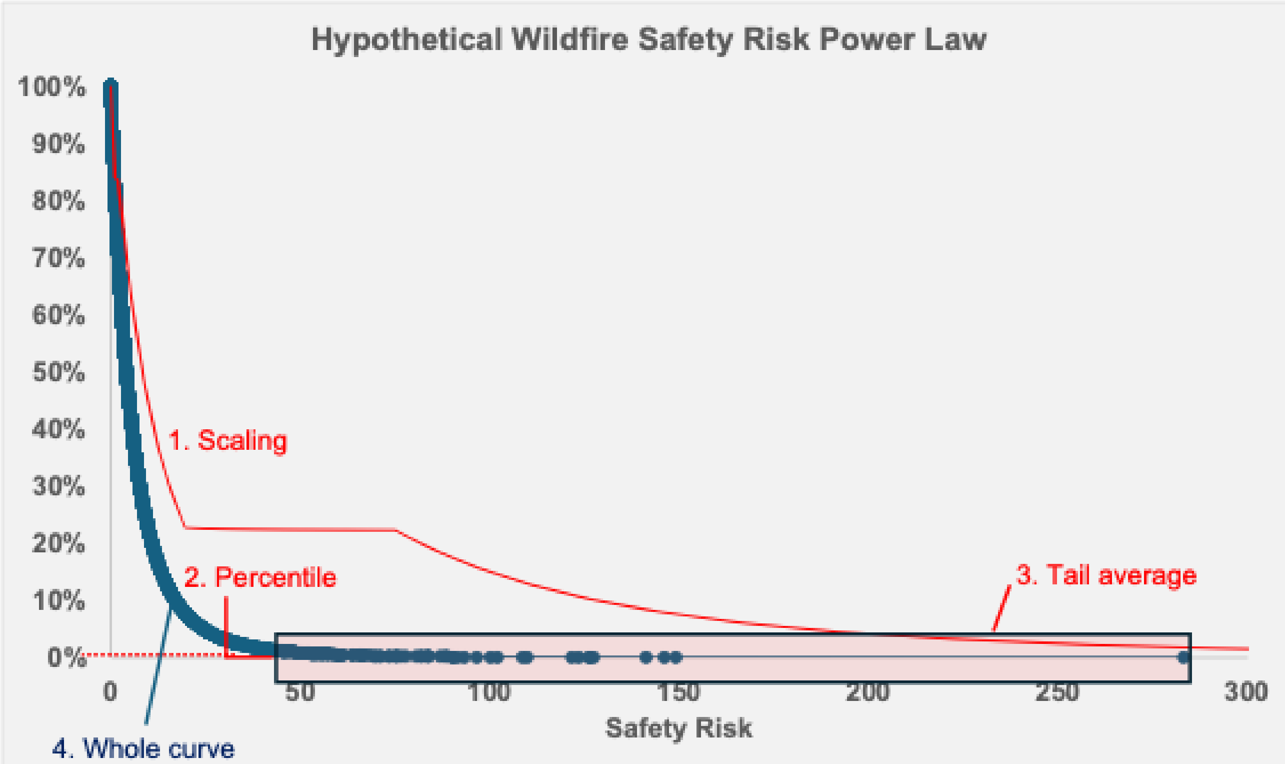

Figure 4-6. Four ways to visualize measures of tail risk.

Figure 4-7 below calculates the corresponding tail risk measures from Figure 4-6:

Safety Risk	
Average*	**$7**
#1. Scaled Average	$41
#2. P99th%	$50
#3. Tail average above P99th%	$70

**Average of entire risk curve (#4)*

Figure 4-7. Average risk and tail risk calculations. The row numbers correspond with the labels in Figure 4-6.

Each of the tail risk calculations is many multiples of the average risk, which will be of interest (and concern!) to risk-averse evaluators. They differ in important respects:

- *Average*. The average, or EV, of the entire risk curve. The average includes the tail but does not adequately represent it.
- *1. Scaled average*. The average of the scaled function. Though it places extra weight on the tail, like the average it blends the tail in with the rest of the curve and thus dilutes the tail. Depending on the formulation of the scaling function, it can be difficult or impossible to optimize. For tail risk evaluation, the scaled average functions like the percentile approach (discussed next), with the

disadvantage of not specifying the exact percentile. For example, the scaled average in Figure 4-7 above is 41, which is implicitly the same as using the value at the 98.5th percentile.

- *2. Percentile*. Also known as Value at Risk (Var), it has the benefit of being easy to calculate and may be the most stable measure of the tail in situations where there is concern about the validity of the most extreme events. Percentile values can be difficult to optimize when evaluating large portfolios of assets, which are in different areas and have diverse levels of inherent risk. A key disadvantage is that Var ignores risks above the chosen percentile, which could include catastrophic risks. For example, the Var at the 99th percentile in Figure 4-7 is $50, which excludes significant risks exceeding $100.
- *3. Tail average*. Also known as Conditional Value at Risk (Cvar), it captures the entire tail above the selected percentile. For that reason, it is more stable if the number of data points or simulation trials changes (as long as there is no concern about the validity of the most extreme events). It is also possible to optimize using linear programming, which greatly increases computational efficiency during optimization. For Figure 4-7 we calculated the tail average for the 99th percentile at $70, which includes all risks above $50. The tail average above a certain percentile will *always* be higher than the risk at that percentile.
- *4. The entire risk curve*. While attractive in theory, potentially having to assess risk along an infinite number of points is impossible. The alternative would be to choose several points along the curve, which is the same thing as choosing multiple percentiles. Conceptually, this is the approach taken for the Transparency Pilots that are part of CPUC decision D.24-05-064 and could be useful for sensitivity analysis.

  While using the entire risk curve to assess risk may be impractical, it is paramount to preserve the entire risk curve for aggregating risks in obeyance of the laws of the arithmetic of uncertainty.

Level 4's preferred approach for incorporating tail risk is to use the tail average, given its stability under many conditions and its beneficial optimization properties. Mitigations could be evaluated based on reducing average risk and tail average risk, which is discussed in detail in Section 5.7. For example, mitigations addressing the safety fire risk in Figure 4-7 would be evaluated based on how much and how cost-efficiently they reduce the average risk of $7 and the tail average risk of $70.

4.3.3 Risk Tolerance and Cost-Benefit Analysis

CPUC decision D.22-12-027 modified the RDF by replacing the multi-attribute value framework (MAVF) with the cost-benefit approach (CBA). In a CBA, the decision to invest in a project is based on the benefit-cost ratio (BCR). A BCR of 1.0 means that the benefits of a project exactly equal its costs, so a typical decision rule for investing in projects is a BCR greater than 1.0.

The BCR can be tied directly to risk attitude. Recall that the definitions for risk-neutrality, risk-aversion, and risk-seeking are based on willingness to spend to avoid risk. Risk-neutrality is the willingness to spend exactly the EV of risk to avoid it, risk-aversion is the willingness to spend more than the EV, while risk-seeking will only pay less than the EV. The risk attitudes can be visualized by the BCR curves in Figure 4-8.

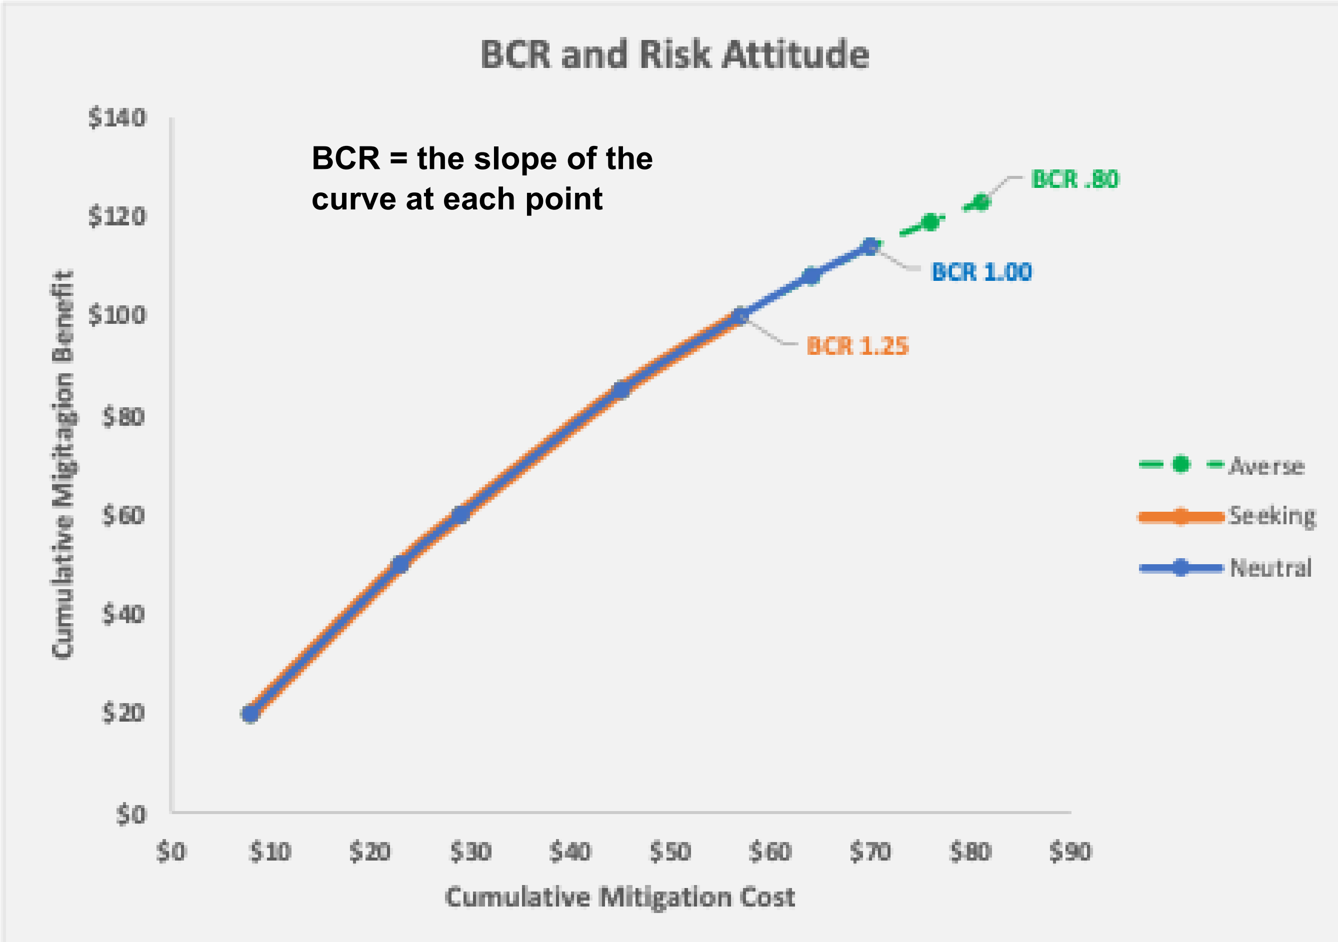

Figure 4-8. BCR and risk attitude.[48]

BCRs only make sense if the benefits (numerator) are based on average risk reduction. Using risk scaling or any measure of tail risk as the numerator to calculate BCR will result in significant over-investing. The example in Figure 4-9 illustrates our concern.

Portfolio-001	Average	Scaled Average	Tail Average
Risk Reduction	$80	$150	$320
Cost	$100	$100	$100
BCR	**0.80**	**1.50**	**3.20**

Figure 4-9 Hypothetical risk reduction BCR calculated 3 ways.

[48] To avoid risk, the risk-averse are willing to invest at BCR below 1.0, the risk-neutral are willing to equal a BCR of 1.0, and the risk-seeking will set a BCR threshold greater than 1.0.

Suppose the BCR threshold for selecting a portfolio of risk mitigations[49] is 1.0. Mitigation-001 has a BCR of 0.80 based on the ratio of average risk reduction and cost, which could result in the portfolio of mitigations being reassessed since costs exceed benefits. Perhaps costs can be further reduced, or a different portfolio with a slightly different set of mitigations and higher BCR might be considered.

Calculating BCR based on tail risk measures such as scaled average or tail average, however, will almost always result in BCRs above 1.0 creating the illusion of high cost-efficiency. Such calculations could be used to justify almost all mitigations, resulting in over-investment. In theory, a higher BCR threshold could be used for evaluating tail BCRs, but how would those be set? It is a slippery slope best avoided.

If the goal is to reflect risk aversion, it is better to use average risk reduction in the numerator and set the threshold for BCR to be less than 1.0.

Before we close out the discussion on tail risk, we will revisit the notion of risk neutrality and its special relationship with tail risk.

4.4 Risk Neutrality and Tail Risk

The tail risk discussion so far has been implicitly based on risk-averse tolerance. That is because risk neutrality means indifference to tail risk! A risk-neutral evaluator cares only about the EV of risk and ignores any potential downsides.[50] While a long-tailed risk curve such as a power law might impact the EV of the risk, the tail itself is of no interest. It therefore doesn't even need to be calculated, much less evaluated. This surprising implication of risk neutrality is demonstrated in Figure 4-10.

	Likelihood	Consequence A	Likelihood	Consequence B	Risk
Risk A	**100%**	**$1,000**	**0%**	**$0**	**$1,000**
Risk B	10%	$10,000	90%	$0	$1,000
Risk C	1%	$100,000	99%	$0	$1,000
Risk D	0.10%	$1,000,000	99.90%	$0	$1,000
Risk E	0.01%	$10,000,000	99.99%	$0	$1,000
Risk F	0.001%	$100,000,000	99.999%	$0	$1,000
Risk G	0.00001%	$10,000,000,000	99.99999%	$0	$1,000

Figure 4-10 A risk-neutral evaluator is indifferent to risks A to G since they all have the same EV of $1,000

Anyone who would tolerate an average risk of $1,000—but not say a 1% chance of losing $100,000 (risk C) or a 1 in 10,000 chance of losing $10 million (risk E)—is risk averse.

49 We will discuss evaluating portfolios of mitigations as opposed to individual mitigations in section 5.2.

50 Gordon Scott, "What is Risk Neutral? Definitions, Reasons, and Vs. Risk Adverse," *Investopedia* (2022). https://www.investopedia.com/terms/r/riskneutral.asp

It is now time to return to the three wanderers to tie the concepts of Chapter 4.

4.5 Epilogue: Risk Assessment, Risk Tolerance, Risk-Based Decisions, and the Three Venturers

Earlier, we left unanswered how the three venturers would approach deciding whether to accept the risk as they saw it or to mitigate the risk by paying the cost of the helicopter. The decision approach for each venturer can be visualized in Figure 4-11.

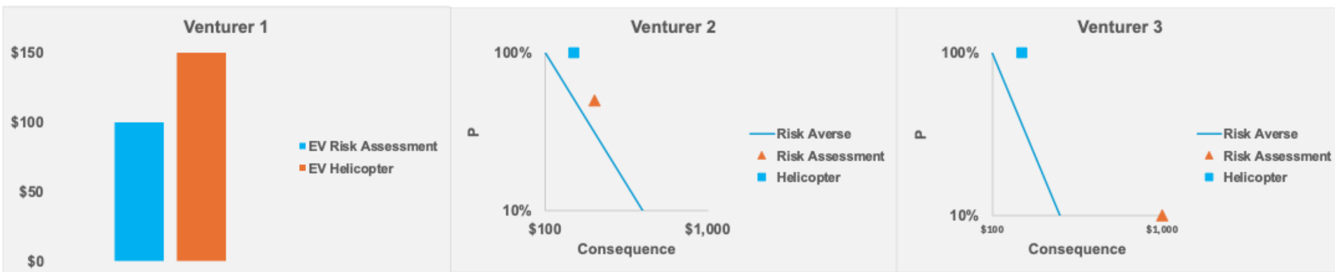

Figure 4-11. The three venturer decision approach.

- *Venturer 1*. V1 is risk neutral, and V1 will only pay the EV of the risk assessment to mitigate the risk. Since the EV of the helicopter is 33% higher than the risk assessment, without hesitation V1 turns off at the next exit and the $100 tickets go unused.
  - The visualization of V1's decision process in Figure 4-11 is a straight comparison of EVs; probabilities are not considered. That is why V1's visualization is different than the others. Even if V1's risk assessment was the same as the probabilistic ones made by V2 or V3, V1's decision wouldn't change since the EVs are the same.
- *Venturer 2*. V2 is moderately risk-averse, denoted by the risk-averse line whose slope is -1.7.[51] V2's assessment of a 50% chance of losing $200 is represented by the red triangle that is just slightly above the risk tolerance line. While the risk assessment exceeds tolerance, V2 has a decision to make. Is the cost of the helicopter worth mitigating the risk? V2 will likely choose to accept the risk in this instance since the risk only slightly exceeds risk tolerance compared to the cost of the helicopter.
- *Venturer 3*. Meanwhile, V3 is more risk-averse than V2. V3's risk tolerance curve has a slope of -2.5 and V3 has a much easier time deciding. The potential of a $1,000 loss—the tail risk—is so far beyond V3's risk tolerance, the distance between the red triangle and the risk-averse curve, that V3 has already made the decision to take the helicopter.

According to a straight interpretation of the expected value theory, none of the venturers should accept the cost of the helicopter. But that is before consideration of risk tolerance and tail risk. Once those are considered, one and possibly two of the venturers will choose the helicopter.

---

[51] Any line with a slope less than -1.0 is risk averse; greater than -1.0 would be risk seeking.

The parable demonstrates the interplay between probabilistic risk assessment, risk tolerance, and risk-based decisions. In Chapter 5, we will explore how simple optimization based on this interplay can work for utility risk management.

4.6 Chapter 4 Summary

- Risk cannot be represented as a single number when performing risk calculations. It must be represented by a probability distribution such as a power law, which enables the arithmetically correct aggregation of risk and allows working with average risk along with tail risk.
- Risk tolerance is the stochastic representation of a subjective risk attitude combined with an appropriate risk scaling function. Risk tolerance also cannot be represented as a single number in risk calculations and can be visualized probabilistically by exceedance curves.
- Risk-based decisions should be based on the relationship between the probability distributions of risk and risk tolerance. As a first step, this can be done by comparing the average and tail average of the risk probability distribution versus the average and tail average of the risk tolerance probability distribution.

5 Optimizing Risk-Based Decisions

5.1 The difference between ranking and optimizing

In Chapter 4 we made the case that a probabilistic view of risk and risk tolerance are inseparable for risk-based decisions. We now turn to the actual risk-based decision. If after assessing risk and applying risk tolerance there is only one course of action, the decision is easy. Most of the time, however, multiple options will remain in play.

CPUC decision D.16-08-018 established the MAVF framework and "explicitly asks for calculations of risk reduction and a ranking of mitigations based on risk reduction per dollar spent."[52] This decision rule is also known as risk spend efficiency (RSE). Ranking based on RSE does not qualify as optimization, but as The Utility Reform Network (TURN) pointed out in D.18-12-014 it can be viewed as an "optimization heuristic."[53] This holds true for ranking by BCR in a cost-benefit analysis.

Ranking based on RSE or BCR can lead to optimal decisions if the mitigations are independent, that is, the choice of one mitigation does not affect the choice or effectiveness of another. Independence between mitigations is rarely the case. It is common for mitigations to be interrelated, where mitigations may be mutually exclusive, synergistic, or exhibit diminishing returns.

- *Mutually exclusive*. Mitigations that cannot work together to reduce risk. Undergrounding and covered conductors on the same circuit segment are examples of mutual exclusivity. It would not make sense (or possible) to do both, even if they had the highest BCR ranks.
- *Synergistic*. Mitigations that work together to decrease the amount of risk. In cyber risk prevention, multi-factor authentication and security awareness training can create a more robust defense against cyber-attacks than either alone.[54]
- *Diminishing returns*. Mitigations that reduce risk together, but as investment in one increases, the need for the other mitigation is reduced. This is because each mitigation reduces the amount of risk that the other mitigation would be expected to eliminate. Reducing the risk of dam failure by increasing spillway capacity and raising the height of the dam is likely to have diminishing returns since the success of one reduces the risk that needs to be addressed by the other.

A simple numerical example drives home the point—suppose two mitigations each reduce risk by 60%. We would not expect that by employing both mitigations we would reduce risk by 120%. Furthermore, budget limitations can reduce independence between mitigations. For example, when approaching budget limits, smaller mitigations may replace larger ones even if the larger ones are ranked higher.

---

[52] CPUC decision D.16-08-018, page 3.

[53] CPUC decision D.18-12-014, page 57.

[54] "Initiative: Multi-Factor Authentication (MFA) and Security Awareness Training Expansion," The University of Memphis, https://www.memphis.edu/its/security/duo_training_expansion.php

A solution to evaluating interrelated mitigations is to construct portfolios of mitigations, which can be compared and ranked. Creating portfolios of mitigations is the topic of the next section.

5.2 Managing Risk Based on Portfolios of Mitigations

Borrowing from finance theory, a portfolio approach is one way to handle real-world uncertainty and the existence of interrelationships between projects.

A portfolio of mitigations could include any combination of feasible (i.e., non-mutually exclusive) mitigations. Suppose we are considering 3 mitigations, M1, M2, and M3. There could be a total of 7 different portfolios, as laid out in Figure 5-1.

Port_1	**M1**
Port_2	**M2**
Port_3	**M3**
Port_4	**M1, M2**
Port_5	**M1, M3**
Port_6	**M2, M3**
Port_7	**M1, M2, M3**

Figure 5-1. Possible portfolio combinations for three mitigations.

Seven portfolios for 3 mitigations assume none of the mitigations is mutually exclusive. If M1 and M2 were in fact mutually exclusive, such as undergrounding and covered conductor on the same section, the set of possible portfolios would be reduced, as shown in Figure 5-2.

Port_1	**M1**
Port_2	**M2**
Port_3	**M3**
Port_4	**M1, M3**
Port_5	**M2, M3**

Figure 5-2. Possible portfolios for three mitigations excluding mutually exclusive ones.

Within each portfolio, synergies and diminishing returns would be accounted for. Cost and benefit are calculated at the portfolio level. Portfolios can be evaluated against each other and the best one is chosen.

But first, a potential issue. The number of possible portfolios that can be constructed from N number of mitigations is 2^N-1. If there are 1,000 mitigations under consideration, that could mean as many as a 1 followed by 300 zeros (or 1x10^300) number of portfolios. Fortunately, the vast majority of portfolios are clear losers and don't need to be constructed, much less evaluated. Computational techniques such as linear optimization rapidly whittle the number of portfolios down to a manageable set.

There is also the question of whether portfolios should include mitigations across all risk events. In theory, this is the correct approach but in practice could prove cumbersome, especially when evaluating risk events separately. Portfolios may be created for each risk event, which might require additional optimization steps. There might be a slight reduction in optimality, but this loss would be minor compared to improved flexibility for addressing risk events individually.

5.3 Portfolio Optimization: Efficient Frontiers

The first and most important question in optimization is "What are we optimizing for?" Lowest cost? Highest efficiency? Lowest residual safety risk? Lowest residual total risk?

Based on decisions in the RDF Proceeding, the goal is to minimize residual risk within affordability constraints, which is different from maximizing. To understand the difference, consider two portfolios:

- Port_1 has a mitigation value (benefit) of $1,000 and costs $200, for a BCR of 1,000/200 or 5.0
- Port_2 has a mitigation value (benefit) of $1 billion and costs $0.5 billion, for a BCR of $1B/$0.5B or 2.0

If the goal is to minimize residual risk, $1 billion of risk reduction is better than $1,000, and Port_2 still has a BCR greater than 1, which is sufficient for neutral or averse risk tolerance. In optimization lingo, we are maximizing risk reduction, subject to a minimum threshold for BCR.[55]

Further constraints can be added. If there was a maximum spend constraint of $0.25 billion for affordability, then Port_2 would be reduced to a benefit of $0.5 billion at a cost of $0.25 billion, even if it meant forgoing an additional $0.5 billion of mitigation benefit.[56] Or, Port_2 could be replaced by another portfolio if there is one with a higher BCR at a $0.25 billion expense.

The above example shows why we cannot rank portfolios based purely on BCR. We also cannot rank based on benefits (the amount of risk mitigated) alone, what if the portfolio with the highest benefit had costs that were twice as high?

In 1952, Harry Markowitz solved this problem in his article "Portfolio Selection," published in The Journal of Finance.[57] In the article, Markowitz developed the concept of an efficient frontier of optimal portfolios, established the principle of evaluating the risk and return characteristics of the portfolio, not the individual

---

[55] Technically, we are minimizing level of residual risk, which is not always the same as maximizing mitigation impact. This nuance is discussed further in appendix F.

[56] This example assumes a constant BCR for Port_2 instead of diminishing returns for simplicity.

[57] Harry Markowitz, "Portfolio Selection," *The Journal of Finance* 7, no. 1 (March 1952): 77-91.

assets within the portfolio, and laid the groundwork for evaluating trade-offs between portfolios. Markowitz's principles can be applied to portfolios of real assets, not just financial assets.[58]

5.3.1 The Efficient Frontier

Markowitz's solution is elegant—if you plot each risk mitigation portfolio on an X, Y scatter plot with mitigation benefit on the Y axis and mitigation expense on the X axis, you will get a chart that looks something like Figure 5-3. Note that the units of the X and Y axes make up the components of the BCR.

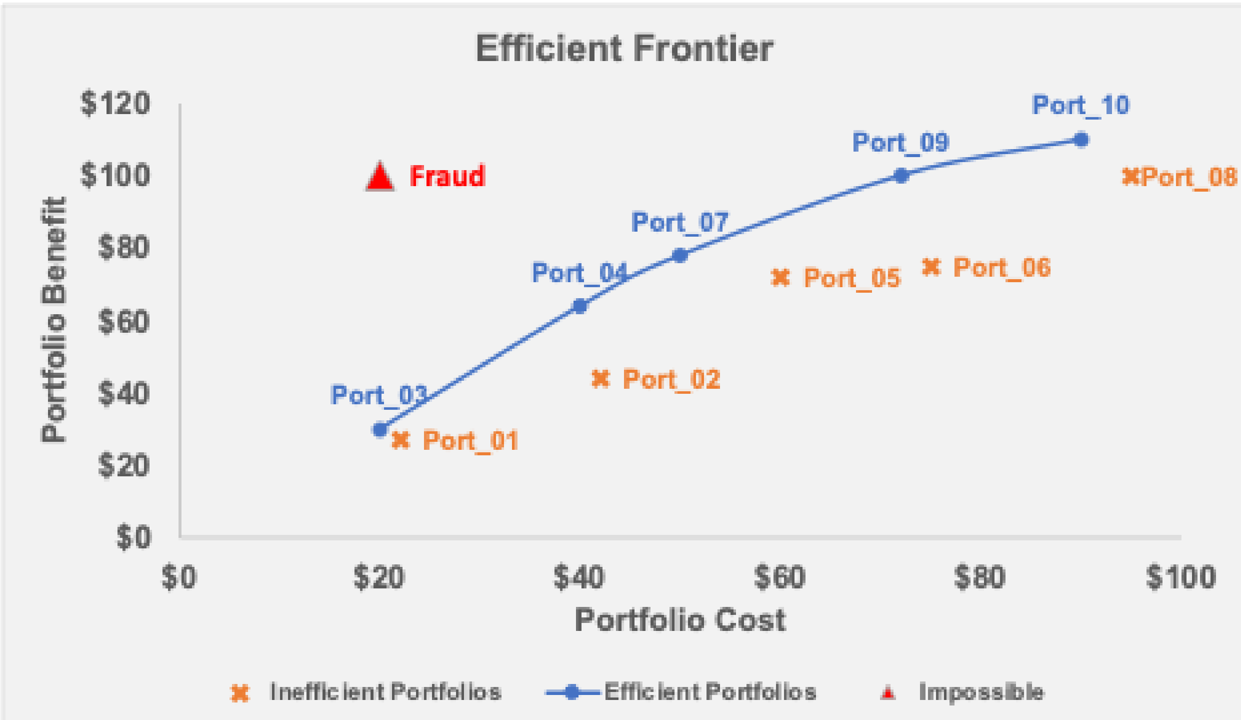

Figure 5-3. Efficient frontier, with optimal portfolios along the blue line, and suboptimal portfolios below the line. The impossible portfolio is an example of fraud.

In Figure 5-3, the set optimal portfolios lie on the blue line, which is the efficient frontier. For each optimal portfolio, it is impossible to obtain higher benefits without paying more. There can be no portfolios above the efficient frontier. In other words, each of the portfolios on the blue line represents the maximum possible BCR at that level of risk reduction. Suboptimal portfolios (red x's) lie below the efficient frontier.

---

[58] For an example of water companies, see Manuel Mocholi-Arce, Ramon Sala-Garrido, Maria Molinos-Senante, and Alexandros Maziotis, "Performance assessment of water companies: A meta-frontier approach accounting for quality of service and group heterogeneities," *Socio-Economic Planning Sciences*, 74 (April 2021). https://www.sciencedirect.com/science/article/abs/pii/S0038012120307850

These are suboptimal because, for each of them, there is at least one portfolio that provides greater benefit for the same or lower cost. It would be irrational to choose a suboptimal portfolio.

Figure 5-3 does include a portfolio above the efficient frontier, which by definition is impossible and therefore fraudulent. In fact, applying the concept of efficient frontiers is how the authorities eventually caught Bernie Madoff.[59]

A feature of efficient frontiers is that the slope of the line decreases as you move up the curve. This means that BCRs are decreasing as portfolios increase in cost and benefit (diminishing marginal returns).[60]

A key observation is that the efficient frontier usually contains multiple portfolios—see Figure 5-4 for a more realistic visual of an efficient frontier. Many of the optimal portfolios will vary only slightly. The ultimate selection will depend on risk tolerance—the risk-averse may prefer one of the more expensive portfolios that generate higher benefits, albeit with lower BCRs—and also on safety vs. reliability impacts, or different Environmental and Social Justice (ESJ) impacts or other goals. Budget and available resources always play a critical role.

An example of a portfolio would be a series of circuit segments that are being mitigated for a risk. Each circuit segment would have its own targeted mitigations (like undergrounding or covered conductors). Together all of the costs associated and benefits (risk mitigated) for each circuit would be aggregated together at the level of a portfolio. A different portfolio may be the same, except one of the circuit segments may be mitigated with a different mitigation causing a slight difference in benefit (benefit=risk reduced) and cost. As one iterates through the potential combinations, one finds optimal portfolios, each with its own benefit and cost.

---

[59] Harry Markopolos (not to be confused with Harry Markowitz), *No One Would Listen: A True Financial Thriller* (John Wiley & Sons, 2010).

[60] This is not quite the same as diminishing marginal utility, which leads to risk aversion. Diminishing marginal returns for portfolios reflects that the number of high-return investments is limited, and at some point, adding more investments dilutes returns.

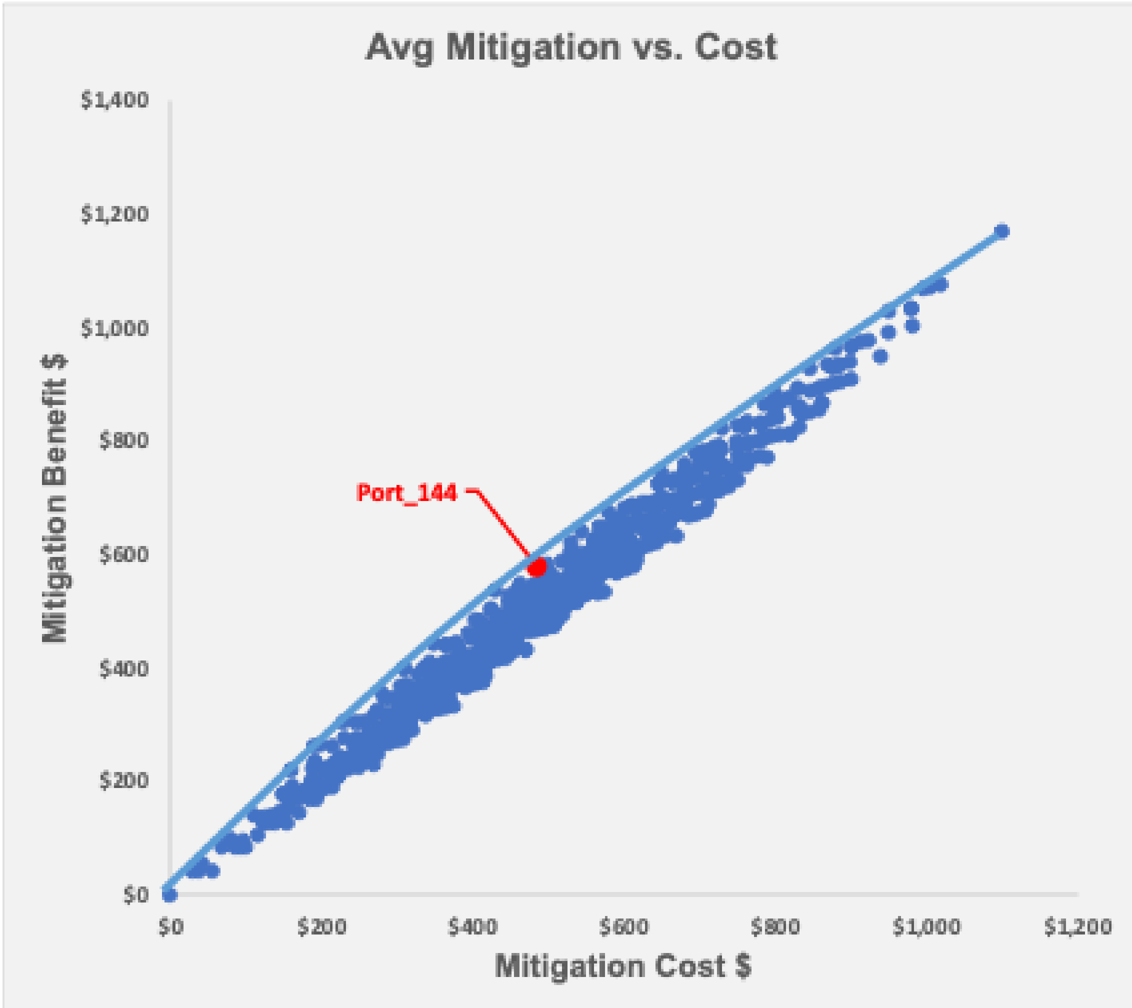

Figure 5-4. Efficient frontier calculated from over 700 portfolios. Multiple portfolios on or very close to efficient frontier provide an opportunity to make trade-offs. Portfolio Port_144 lies on the efficient frontier.

The selection of which optimal portfolio is a subjective one that weighs risk tolerance, other trade-offs, and impact on affordability.

5.4 Stochastic Optimization

The efficient frontier discussion in Section 5.3 didn't specify how portfolio benefit was defined, but the implication is that it represents an average benefit, which is appropriate when calculating BCRs. What about tail risk? Can you create efficient frontiers that take tail risk into account? Can you use tail risk in optimization? The answer to both questions is yes.

Figure 5-5 calculates an efficient frontier for tail average risk as our measure of tail risk and is presented along with the original efficient frontier based on average risk. The portfolios are the same on both frontiers—except for the third one (from the bottom of the curve). At around $50 portfolio cost, the optimal portfolio for average risk is Port_07, but for tail average the optimal portfolio is Port_05.

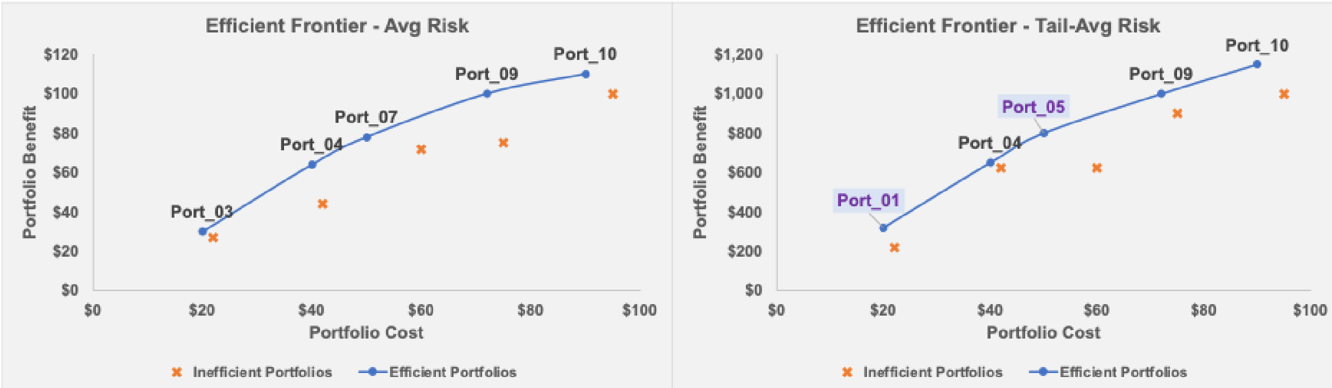

Figure 5-5. Efficient frontiers for average risk and tail average risk.

What does this mean? At the level of cost where Port_07 is optimal for average risk, a different set of mitigations is more optimal for reducing tail risk, and these make up Port_05. Which portfolio is selected will depend on risk tolerance—the risk-averse may prefer the greater reduction of tail risk in Port_05 because it is more likely to mitigate catastrophic events whereas Port_07 may be more likely to mitigate common events that are less risky. There can be other considerations as well. Fortunately, it is possible to optimize across multiple considerations, which mathematically are called "dimensions."

5.5 Optimizing for Multiple Considerations (Dimensions)

The decisions in the RDF Proceeding recognize the multi-dimensional nature of mitigating risk and do not require basing mitigation decisions solely on a single measure such as BCR.[61] There are other trade-offs that must be considered, for example, safety vs. reliability, affordability, ESJ impact, time exposure, or execution risk (and others). These trade-offs can be optimized quantitatively, for example by setting minimum requirements for safety improvement and reliability impact or a maximum rate increase during a single GRC cycle. For optimal portfolios, trade-offs between similar portfolios can be further evaluated subjectively, as visualized in the "herringbone" diagram[62] as seen in Figure 5-6.

---

[61] CPUC decision D.24-05-064, Appendix A, Row 26.

[62] We acknowledge Dr. Sam Savage for coming up with the idea of herringbone diagrams (named for the distinctive visual it creates).

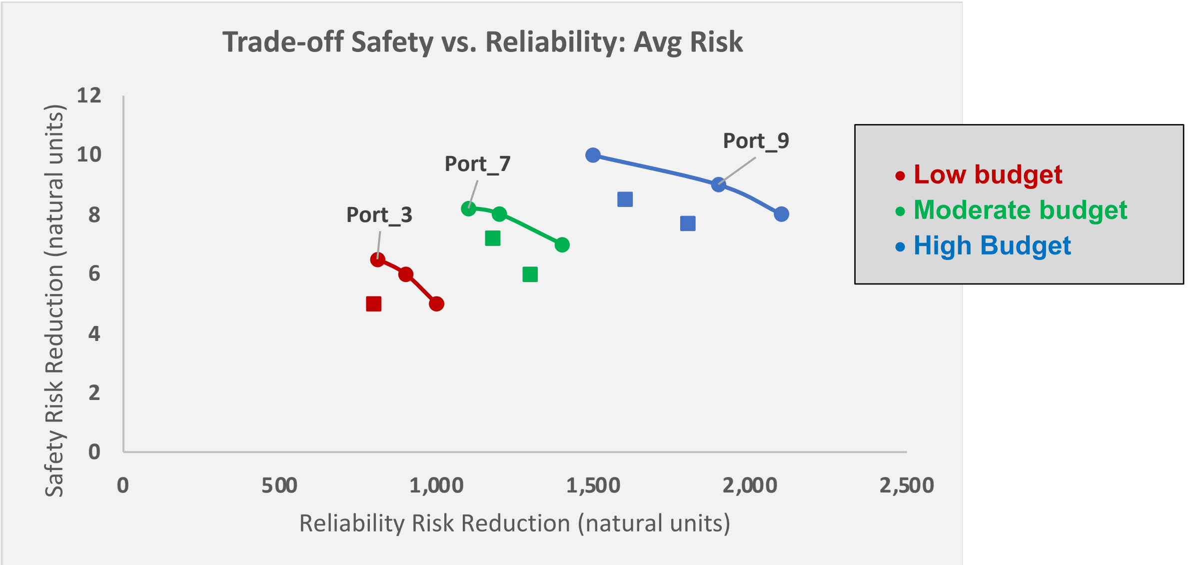

Figure 5-6. Herringbone diagram depicting safety and reliability trade-offs in natural units[63] at different budget levels.

The herringbone diagram visualizes trade-offs in three dimensions: safety, reliability, and budget. The portfolios on the connected curves are optimal (they are taken from the efficient frontier), while the portfolios marked by squares are sub-optimal, shown for context. Each color represents a budget range. The only way to achieve a higher level of safety than Port_3 (red) would be to jump to the next budget range of Port_7 (green).

Keeping with Port_3, it shares the efficient frontier with two other portfolios at the low (red) budget range. The other two portfolios trade off lower safety for higher reliability. All three portfolios are optimal, but the final selection would depend on how the evaluator weighs safety vs. reliability.

Herringbone diagrams can be used for any number of trade-off dimensions. It would be possible to create a dashboard of multiple herringbone diagrams to visualize all the trade-offs together.

Tail risk can also be represented in a herringbone diagram, for comparison with average risk as in Figure 5-7. In this example, the optimal portfolio for the moderate (green) budget is Port_5, which is different from the optimal portfolio for average risk (refer back to Figure 5-6 and Figure 5-5 above). Port_5 emphasizes reliability more than safety for tail risk reduction. Whether to select Port_7 or Port_5 will depend on the evaluator's risk tolerance and preferences for safety and reliability.

---

[63] It is possible to present the herringbone in monetized units—it looks exactly the same. The choice depends on whether evaluators would like to weigh the natural unit's impact between attributes or prefer comparing monetized values.

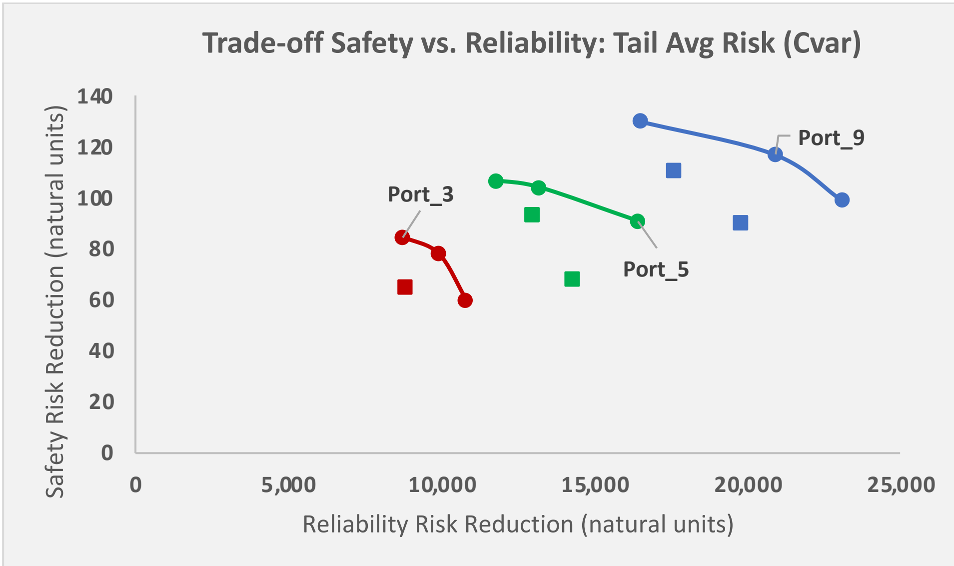

Figure 5-7. Herringbone diagram for safety vs. reliability trade-off based on tail average risk.

These examples highlight the subjective nature of risk and risk assessment, but it does not mean there is a license to make any decision according to any agenda. The efficient frontier greatly reduces the possible decisions to a manageable set of optimal (or near optimal) portfolios of mitigations. An explicit risk tolerance creates additional boundaries for justifiable decisions. Finally, portfolio optimization and risk tolerance improve the transparency of final decisions: it should be clear for all to see how close the decision aligns with the efficient frontier, how it stacks up against alternatives, and whether it is within risk tolerance.

5.6 Optimizing Frameworks

This section will provide an overview of the CBA decision framework and its two ranking or optimization methodologies (see Figure 5-8.)

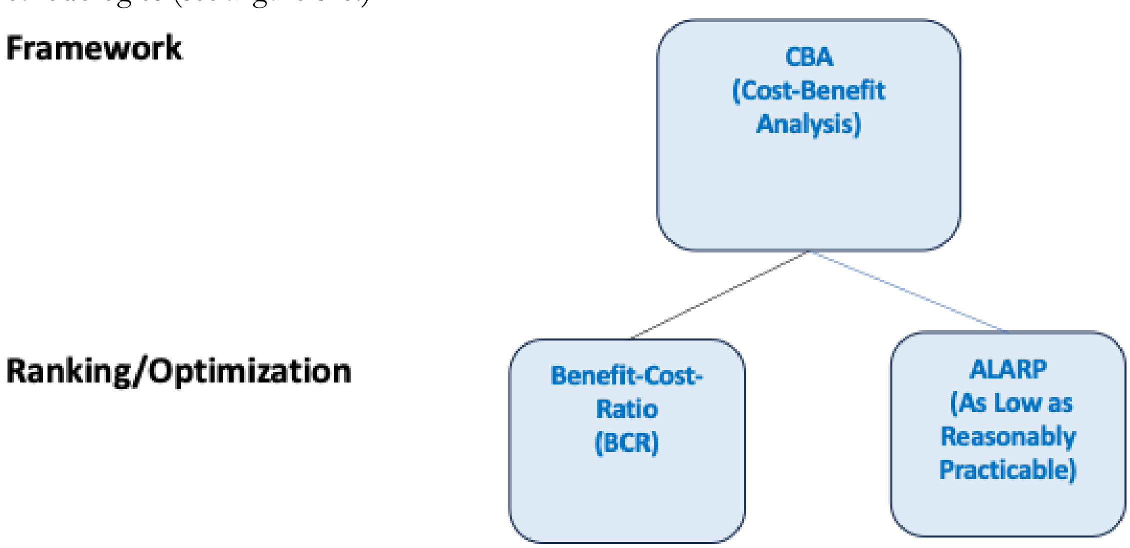

Figure 5-8. Risk decision framework.

CPUC D.22-12-027 modified the RDF to require CBA instead of MAVF. In CBA, benefits and costs of mitigations are calculated in dollars, using a monetization factor to translate the natural units into dollars. For decision-making, the benefit value is divided by the cost value, creating the benefit-cos ratio (BCR). A BCR greater than 1.0 means that the benefits exceed the costs and a BCR of 1.0 is often, though not always,[64] used for selecting mitigations. A BCR can be calculated for each mitigation option or used in optimization.

5.6.1 ALARP (As Low as Reasonably Practicable)

Another optimization process is called ALARP, which stands for As Low as Reasonably Practicable. What sets ALARP apart is its three-tiered approach to optimization (which includes BCR).

- Tier I. Risk exceeds maximum risk tolerance, mitigate immediately regardless of cost.
- Tier II. Risk level is within maximum risk tolerance, continue to mitigate if BCR is above a set threshold.
- Tier III. Risk level is at or below the accepted level of risk, no further action is taken (residual risk is accepted).

Figure 5-9 illustrates the ALARP methodology. The upper and lower bounds can be considered exceedance curves for maximum tolerable risk and acceptable risk. The white region in between is tolerable (which is not to say acceptable). Risk above the upper bound, the red zone, is considered intolerable and must be mitigated to at least tolerable levels *without regard to cost* (at least in theory). Once risk is within the tolerable range, it should continue to be mitigated as long as a BCR threshold is met. If the risk is within the accepted range, the green zone, no more mitigation is required, even if it is possible to do so above the BCR threshold.

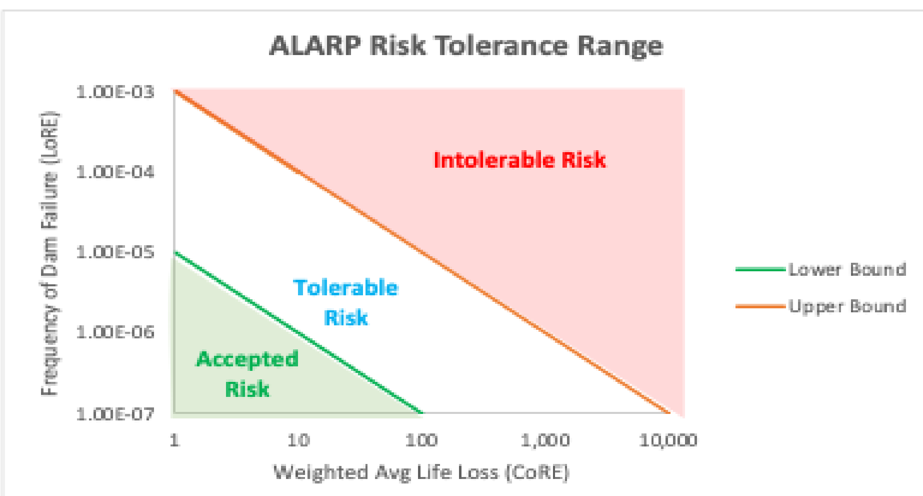

Figure 5-9. The three tiers of ALARP.

[64] As discussed in Section 4.3.3, risk aversion may lead to setting the BCR threshold below 1.0.

Figure 5-10 shows how ALARP works in practice. It shows the pre-mitigated risk exceedance curve for a potential cause of dam failure (i.e., the curve plots the likelihood of a failure, and the consequence at that level of failure). This risk is a low probability, high consequence risk—at higher LoREs, the risk is within the accepted range and does not need to be further mitigated, but at the other end of the curve, at a risk less than 1 in 1 million, the risk is deemed intolerable.

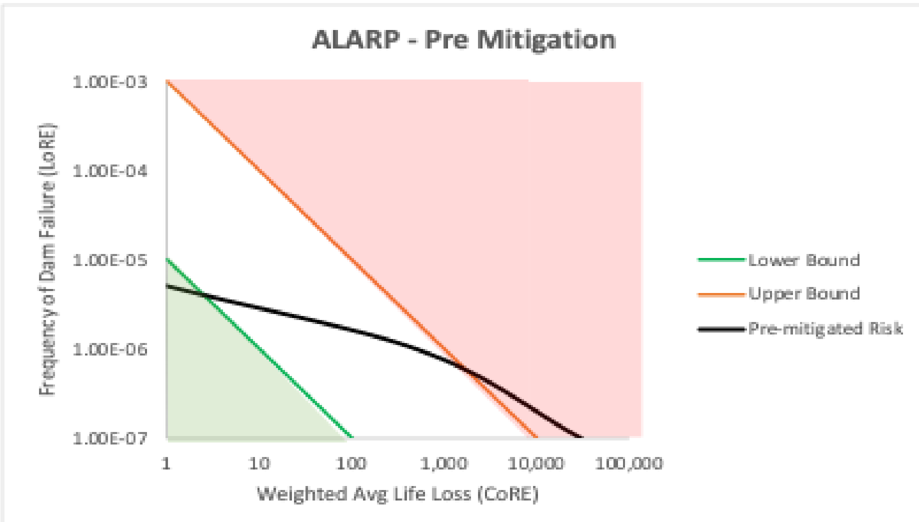

Figure 5-10. Pre-mitigated risk.

According to the principles of ALARP, the risk in the intolerable range must be mitigated below the upper bound into the tolerable range. Risk within the tolerable range should be mitigated according to the BCR threshold, based on risk tolerance. No more investment should be made to reduce the risk in the accepted range.

After mitigation efforts, the post-mitigation exceedance curve might look like Figure 5-11:

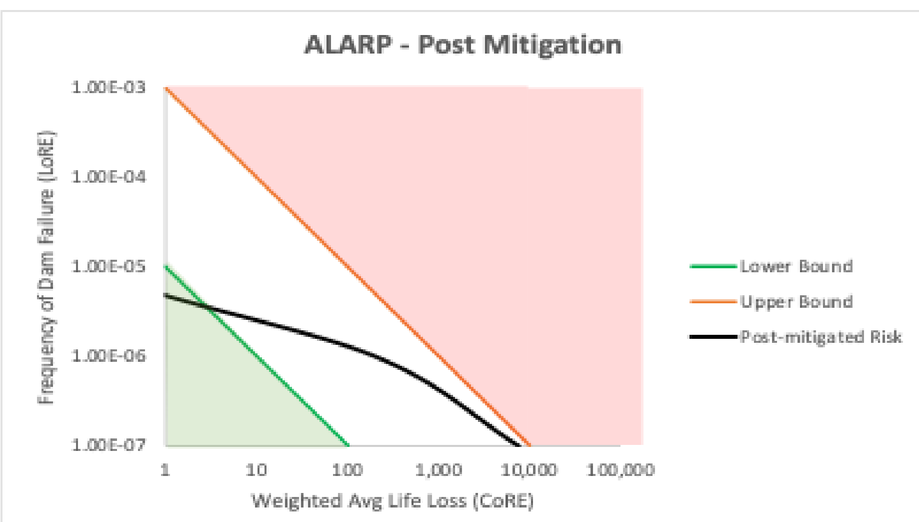

Figure 5-11. Post mitigation risk.

Figure 5-11 shows that ALARP has been satisfied. Most important, the low probability, high consequence risk has been brought down to tolerable levels. Risk in the tolerable range has been reduced as well, due to attractive BCRs for the mitigation projects. Risk in the acceptable range was left untouched.

5.6.2 ALARP History and CPUC Background[65]

The principle of ALARP has its origins in British case law dating back to 1949, dealing with safety in coal mines. It became law in the UK with the Health and Safety at Work etc. Act 1974. It applies to all industries in the UK.

Outside the UK, ALARP has been adopted to a lesser degree in Australia, Abu Dhabi, Belgium, Brazil, Denmark, Hong Kong, Ireland, Netherlands, Saudi Arabia, and the U.S. In the U.S., a similar principle has been used by the U.S. Nuclear Regulatory Commission since the 1950s and is included in two U.S. Federal regulations at the Department of Energy. The U.S. Army Corps of Engineers also uses it.

As discussed in Section 3, CPUC staff presented a white paper advocating for ALARP[66]. The paper was much discussed in CPUC decision D.16-08-018 and subsequent decisions. It was decided that the state of probabilistic modeling in the utilities is not advanced enough for an ALARP implementation, but it remains a priority topic.[67]

One of the most attractive elements of ALARP, its combination of risk tolerance and probabilistic risk modeling, also presents the greatest obstacle to adoption. ALARP requires establishing two risk tolerances for each attribute, safety, reliability, and financial.

The utilities have made substantial progress in the probabilistic modeling of risk, but it is unclear how close they are to being able to fulfill the requirements of ALARP. Expressing risk as probability distributions instead of single-number risk scores, capturing cross-cutting risks and other interrelationships, and correctly aggregating risks are all prerequisites for ALARP, and will be discussed by Level 4 in its Guidance on Interrelationships Report.

Finally, there are practical concerns with ALARP principles such as mitigating risk above the upper bound regardless of cost and the impact on competing concerns such as affordability.

Nonetheless, ALARP's holistic approach that combines the principles of probabilistic risk assessment and risk tolerance is a worthy aspiration. There are ALARP-type approaches that adopt key elements of ALARP that can be implemented sooner rather than later. We propose one such approach in the next section.

---

[65] This history draws on Steven Haine, *Safety and Enforcement Division Staff White Paper on As Low as Reasonably Practicable (ALARP) Risk-informed Decision Framework Applied to Public Utility Safety*. California Public Utilities Commission. (2015, December 24). https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M157/K359/157359431.PDF

[66] Steven Haine, *Safety and Enforcement Division Staff White Paper on As Low as Reasonably Practicable (ALARP) Risk-informed Decision Framework Applied to Public Utility Safety*. California Public Utilities Commission. (2015, December 24). https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M157/K359/157359431.PDF

[67] CPUC Decision D.18-10-014, page 55.

5.7 Simple Optimization: One Approach

The goal is to optimize post-mitigated risk by reducing it to an acceptable level, given affordability and other constraints. We have argued that post-mitigated risk must be thought of as a probability distribution. How do you optimize an entire probability distribution, especially if the goal is simple optimization, at least initially?

One approach is to perform a two-step linear optimization, one for average risk and one for tail average risk.

- Average risk is an important representation of the probability distribution since it is required for calculating BCRs
- Tail average risk is a good measure of the tail of the distribution. As discussed in Section 4.3.2, it is stable unless there are invalid data points in the tail, and it can be optimized using linear programming.[68]

A key point is that optimization by reducing residual risk is not exactly the same as optimization by maximizing mitigation impact. The distinction is subtle but could lead to suboptimal mitigation selection. An example of how maximizing mitigation impact can lead to different results than minimizing residual risk is presented in Appendix F.

At the end of the day, we are still interested in mitigation impact, so the correct formulation is:

Optimized mitigation impact = pre-mitigated risk – optimized post-mitigation risk.

The two efficient frontiers can be evaluated together, and an optimal mitigation that satisfies average risk and tail risk reduction goals, along with any other trade-off considerations can be selected. The next section will discuss the risk-based decision-making process further.

5.8 Making Optimal Risk Reduction Decisions

Suppose we have performed an optimization that results in the following efficient frontiers for average risk mitigation and tail risk mitigation. The organization is interested in Port_161 because it sits on the efficient frontier for average mitigation and tail average mitigation (see Figure 5-12).

---

[68] Sergey Sarykalin, Gaia Serraino, and Stan Uryasev, "Value-at-Risk vs. Conditional Value-at-Risk in Risk Management and Optimization," Tutorials in Operations Research, INFORMS, 2008. https://www.ise.ufl.edu/uryasev/files/2011/11/VaR_vs_CVaR_INFORMS.pdf

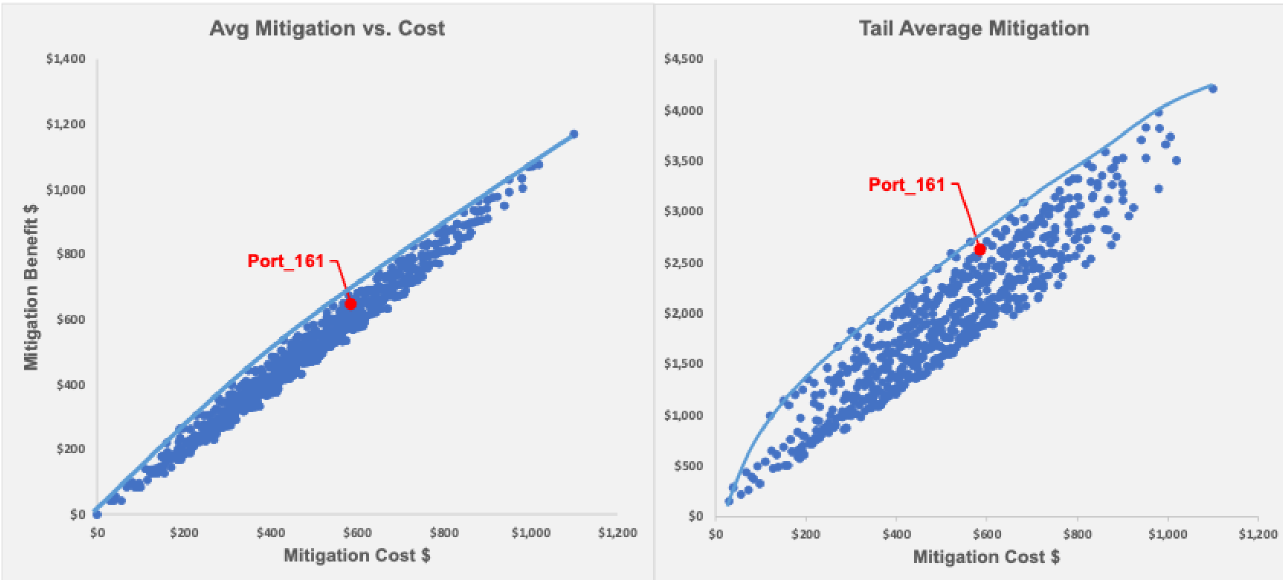

Figure 5-12. Hypothetical optimal mitigation selection example.

The evaluators are also interested in the safety versus reliability risk reduction of the chosen portfolio, which is represented in the herringbone diagrams in Figure 5-13.

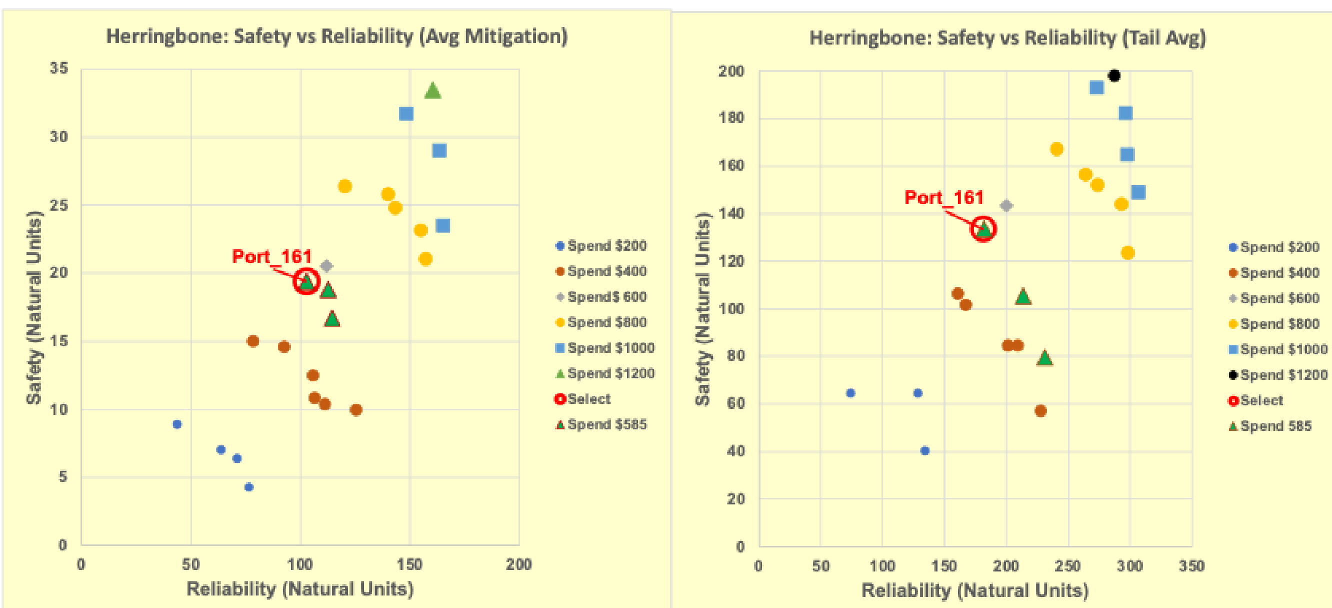

Figure 5-13. Herringbone representation of safety vs. reliability trade-off.

Port_161 has the highest safety impact of the alternative optimal portfolios at the given budget level (green triangles).

Figure 5-14 calculates the key statistics from the optimizations, including applying the relevant quantification of risk tolerance. Post-mitigated risk (row C) is $718 for average risk and $2,095 for tail-average risk. Risk tolerance is shown on rows D (neutral) and E (averse).

		Average	Tail Avg @95%
A	**Pre-mitigated Risk**	$1,364	$4,716
B	**Mitigation Benefit**	$646	$2,621
C=A-B	**Post-mitigated Risk**	**$718**	**$2,095**
D	**Risk Tolerance - Neutral**	$800	
E	**Risk Tolerance - Averse**	$700	$1,800
	BCR of Mitigation Benefit	**1.11**	

Figure 5-14. Simple optimization outcomes.

The key takeaways are that whether this portfolio is acceptable or not depends on whether an organization is risk-neutral or risk-averse. This portfolio meets the threshold for a risk-neutral organization — the average post-mitigated risk of $718 (row C) is below the neutral risk tolerance level of $800 (row D).

For a risk-neutral organization, there is no comparison with tail average risk; by definition, risk neutrality does not distinguish between average and tail risk. For a risk-averse organization, however, post-mitigated risk is above risk tolerance—$718 (row C) versus $700 for the average (row E) and $2,095 (row C) vs. $1,800 for tail risk (row E), requiring further mitigation.

The selected portfolio BCR is 1.11, which exceeds 1.0 by a healthy margin. It may warrant further conversation on whether the budget can be increased to reduce post-mitigated risk closer to the risk-averse threshold at BCR greater than 1.0 while assessing the impact on affordability.

A question sometimes arises if it is possible for risk-averse tolerance to be the same as risk neutral for average risk. The answer is yes, as shown in Figure 5-15. Tolerance for average risk is $800 for risk-averse and risk-neutral, and post-mitigated average risk is within tolerance for risk-neutral and risk-averse. Tail average risk remains out of tolerance for risk averse.

		Average	Tail Avg @95%
A	**Pre-mitigated Risk**	$1,364	$4,716
B	**Mitigation Benefit**	$646	$2,621
C=A-B	**Post-mitigated Risk**	**$718**	**$2,095**
D	**Risk Tolerance - Neutral**	$800	
E	**Risk Tolerance - Averse**	$800	$1,800
	BCR of Mitigation Benefit	**1.11**	

Figure 5-15. Simple optimization outcomes - risk-averse tolerance equals risk-neutral tolerance for average risk.

5.9 Chapter 5 Summary

Chapter 5 has related important concepts in risk assessment, risk tolerance, and optimal decision-making:

- Risk cannot be represented as a single number for use in risk calculations, probability distributions such as power laws should be used instead.
- Risk reduction cannot be assessed without understanding risk tolerance.
- Sophisticated frameworks such as ALARP are attractive because they combine risk tolerance and decision rules for mitigating risk.
- It is possible to develop a simpler framework that requires only one risk tolerance threshold instead of ALARP's two and optimizes based on average risk and tail average risk from the post-mitigation risk probability distribution.
- Optimization can be multi-dimensional and includes affordability, BCR, and other trade-offs such as safety vs. reliability.

6 Modifying the RDF, and Ensuring a Manageable Transition

6.1 Risk Tolerance: Gaining Consensus

All the parties to the RDF proceedings agree that risk tolerance is important, but two key questions remain:

1. Whose risk tolerance?
2. How should the risk tolerance be set and used?

Level 4 believes that tolerance for utility risk should be set at the State of California level, representing the residents of California. It would not be equitable for one utility to have a higher tolerance than another utility for safety risk, which would imply that safety depends on where someone lives in California.

6.1.1 Risk Tolerance Considerations

There are a number of things that need to be decided even before the challenging work of quantifying tolerance begins.

- Should tolerance be set at the aggregated post-mitigated risk level in dollars, as was done in the example in Section 5.8?
- Alternatively, should tolerance be set at the attribute level, in natural units? This would mean setting individual risk tolerances for safety, reliability, and financial risk.
- Should risk tolerance be set for each risk (e.g., wildfire, cyber-risk, hydropower, gas containment, etc.)?
- If risk tolerance is set for total risk, does it need to be apportioned out somehow to each risk category? For example, would it be considered okay if in a year total risk was within tolerance, but wildfire risk consumed 99% of the total that year?

Level 4 recommends that to start, risk tolerance should be set at the highest level that makes sense, either in aggregated dollars for total risk or possibly at the attribute level. Appendix G describes options for setting risk tolerance in more detail.

6.1.2 Developing Risk Tolerance Standards: A Process

Setting risk tolerance on behalf of the residents of California requires input from the many constituents of California. This would include at minimum regulatory agencies, intervenors, and the utilities. There could be several public workshops and technical sessions that would include:

- Training on probabilistic risk assessments, LoRE and CoRE, how to understand a probability distribution for post-mitigated risk, and the difference between average and tail risk.
- Deciding on which tolerances are needed, as discussed in Section 6.1.1.

- What levels to set risk tolerance for average risk and tail average risks? This would include debates on whether the state should be risk neutral or risk averse, and if averse, how averse?

Once the initial risk tolerance standards have been established by the utility, the regulatory agencies would need to determine if those standards need to be codified and how they would be enforced. This would include how quickly a utility would be required to remedy exceeding a risk tolerance, and at what cost.

The process is a lengthy one—we would expect it could take 12-18 months or longer. In the meantime, there is value in having the utilities declare, quantify, and justify the risk tolerance they are using to make mitigation decisions. The utilities would have to decide which tolerances are needed as laid out in Section 6.1.1, which may lead to learning on behalf of the state. How they set tolerance levels for risk neutrality and risk aversion will also be instructive. Finally, it will make evaluating RAMPs, WMPs, and other risk processes much more transparent.

6.2 Simple Optimization: Data and Model Requirements

In parallel with the risk tolerance process, there is work that needs to be done to ensure that the utilities have the technical capacity to perform simple optimization. Fortunately, the progress made over the past several years makes us feel confident that the raw materials are in place. Level 4 recommends a series of workshops on the technical requirements of simple optimization to ensure consistency and proper methodology across the utilities.

Workshop #1. Assessment of current use of probability distributions. It is clear from RAMPs and WMPs that the utilities are already working with probability distributions and are storing them. Figure 6-1 is an example from PG&E's 2024 RAMP, which presents Monte Carlo trials for safety risk outcomes, reliability risk outcomes, and financial risk outcomes.

(PG&E-2)
TABLE 2-17

SAMPLE BOW TIE: SIMULATED SEVERE OUTCOMES VALUES IN NATURAL UNITS AND ATTRIBUTE CORE CALCULATIONS[(a)]

	Safety				Reliability				Financial			
Trial	Sim Natural Unit (EF)	Monetization Factor	CoRE ($M risk adj.)	Implied Risk Adj. Factor	Sim Natural Unit (1k Cust)	Monetization Factor	CoRE ($M risk adj.)	Risk Adj. Factor	Sim Natural Unit ($M)	Monetization Factor	CoRE ($M risk adj.)	Risk Adj. Factor
1	8	15.23	228	1.88	108	1570	329	1.93	999	1	1,988	1.99
2	14	15.23	746	3.50	92	1570	278	1.92	831	1	1,651	1.99
3	8	15.23	228	1.88	111	1570	337	1.93	959	1	1,908	1.99
4	5	15.23	137	1.80	104	1570	316	1.93	969	1	1,928	1.99

5	11	15.23	404	2.41	93	1570	279	1.92	1088	1	2,651	2.44
6	11	15.23	404	2.41	99	1570	298	1.92	1004	1	2,018	2.01
7	12	15.23	518	2.83	99	1570	300	1.92	989	1	1,968	1.99
8	11	15.23	404	2.41	101	1570	307	1.93	818	1	1,627	1.99
9	9	15.23	259	1.89	102	1570	310	1.93	1192	1	3,431	2.88
10	12	15.23	518	2.83	100	1570	303	1.93	1116	1	2,860	2.56
	Safety CoRE 475				Reliability CoRE 302				Financial CoRE 2,208			
Sum of Attribute Values: 2,985												
(a) The Attribute CoRE is the average of the CoRE per trial for that Attribute.												

Figure 6-1. Sample PG&E simulation results.

PG&E's table in Figure 6-1 shows data for each trial in a Monte Carlo simulation, a useful technique for building risk models that uses probability distributions to run the simulation. In our Guidance on Interrelationships Report, we will cover in detail the proper use of these stored probability distributions, including not collapsing them into averages for input into other calculations as being done here.

For now, it is encouraging that these probability distributions have been developed and are being used. We still need to determine the extent to which all the necessary probability distributions exist for risk modeling, including those for LoRE.

Workshop #2. Data libraries. Depending on the granularity of probabilistic models, the data storage requirements could be immense, especially if a large number of Monte Carlo simulation trials for each distribution must be stored. For example, SCE reports that it simulates matchstick ignitions for 29 million ignition points.[69] Advanced tools for efficient storage of simulation data such as metalogs[70] and sparse Monte Carlo[71] may be explored.

Workshop #3. Maintaining probability distribution coherence. A critical feature of storing probability distributions is making sure they remain coherent, that is interrelationships between the distributions are preserved. For example, suppose several models are based on the relationship between temperature and the likelihood of a risk event. If the temperature data point for trial #9 is 102° Fahrenheit, then all the models that include temperature should all have 102° for trial #9. This allows us to have a full picture of risk across all models on 102° days.

[69] SCE 2023-2025 Wildfire Mitigation Plan, page 131.

[70] http://www.metalogdistributions.com/

[71] https://analytica.com/decision-technologies/monte-carlo-simulation-software/ Scroll down to "More efficient variants" section.

Workshop #4. Training in the arithmetic of uncertainty. A key feature of risk modeling is that risks need to be aggregated, and if doing this is based on single numbers, it will be aggregated incorrectly. Probability distributions can be added (or subtracted or multiplied) as long as done following the arithmetic of uncertainty, and proper "order of operations."

Topics 1-4 will be covered in greater detail in our Guidance on Interrelationships Report.

Workshop #5. Tail risk concepts and methodologies. Every probability distribution has a tail, some longer than others, such as power laws. There are multiple ways to calculate them, especially tail average (Cvar) with some to be preferred over others.

Workshop #6. Simple optimization techniques. The use of average risk and tail average makes possible linear optimization, which is consistent with "simple" and improves computational efficiency. Nonetheless, we will need to evaluate what type of optimization is feasible given the substantial number of mitigations under consideration, and the large number of data points in each probability distribution.

Workshop #7. Communicating the results of simple optimization and explaining risk-based decisions. Reporting conventions will need to be developed so it is clear to evaluators how the final mitigation selection relates to the model results. This will be covered in greater detail in our Risk Mitigation Accountability Reporting (RMAR) Report.

7 Recommendations

These recommendations are based on the preceding sections and assume the reader has read and understood those sections. Text in red-underline (deletions) and blue-underline (additions) represent proposed changes to the Risk-Decision Framework.

Recommendation 1 (R1): *Use of probability distributions*. Probability distributions describe the range and chance that a set of outcomes occurs within datasets and model results. Risk models must use probability distributions as inputs and return probability distributions as outputs.

- Likelihood is stated as a probability and can be represented in simulation models as a distribution of zeros and ones, (the ones representing risk event occurrences[72]).
- Consequence is represented as a probability distribution.
- Risk = LoRE x CoRE and represented as a probability distribution.

These definitions are consistent with D.24-05-064 Appendix A Rows 10, 11, and 13, with the clarification that Likelihood, Consequence, and Risk are based on probability distributions, not single numbers.

The utilities have made considerable progress in their use of probabilities and probabilistic modeling, but single-number representations of LoRE, CoRE, and Risk are still prevalent. An immediate first step should be ascertaining how each utility is capturing, storing, and using probability distributions for risk modeling, wherein the modeling process is the probability distributions collapsed into single numbers, and what utilities must do to replace the use of single numbers in their risk models with the underlying probability distributions.

Building on this we recommend the following updates to definitions:

Consequence (or Impact): the effect of the occurrence of a Risk Event. Consequences affect Attributes of a Cost-Benefit Approach and can be presented in the natural units of the attribute or monetized. Consequence is represented as a probability distribution.

Likelihood or Probability: the chance that an event will occur, quantified as a number between 0% and 100% (where 0% indicates impossibility and 100% indicates certainty). The higher the Probability of an event, the more certain we are that the event will occur. Likelihood of an event will be represented in simulation models as a distribution of zeros and ones whose average is the chance that the event will occur.

Probability Distribution: the range and chance that a set of outcomes occurs within datasets and model results.

Risk: The potential for the occurrence of an event that would be desirable to avoid, ~~often~~ expressed in terms of a combination of various Outcomes of an adverse event and their associated Probabilities. Risk is the product of LoRE and CoRE and is represented as a probability distribution.

[72] It is possible for LoRE to be expressed as zeros and integers greater than one if multiple risk events per trial are possible. This requires additional steps for the LoRE x CoRE calculation.

Additionally, Level 4 recommends the following changes to Rows 10, 11, and 13 of D.24-05-064, Appendix A:

10.	Identification of Potential Consequences of Risk Event	The identified potential Consequences of a Risk Event should reflect the unique characteristics of the utility and will be represented as a probability distribution. For each enterprise risk, the utility will use actual results, available and appropriate data (e.g., Pipeline and Hazardous Materials Safety Administration data), and/or Subject Matter Experts (SMEs) to identify potential Consequences of the Risk Event, consistent with the Cost-Benefit Approach developed in Step 1A. The utility should use utility-specific data, if available. If data that is specific to the utility is not available, the utility must supplement its analysis with subject matter expertise. Similarly, if data reflecting past results are used, that data must be supplemented by SME judgment that considers the Benefits of any Mitigations that are expected to be implemented prior to the GRC period under review in the RAMP submission. For each enterprise risk, the utility must explain how they derived the probability distribution for Consequence of a Risk Event.

11.	Identification of the ~~Frequency~~ Likelihood of the Risk Event	The identified ~~Frequency~~ Likelihood of a Risk Event should reflect the unique characteristics of the utility and will be represented in simulation models as a distribution of zeros and ones. Likelihood of a Risk Event is the average of the distribution of the ones and zeroes. Frequency is the number of risk events over a defined period based on likelihood and can be presented for readability. For each enterprise risk, the utility will use actual results and/or SME input to determine the annual Frequency of the Risk Event. The utility should use utility-specific data, if available. If data that is specific to the utility is not available, the utility must supplement its analysis with subject matter expertise. In addition, if data reflecting past results are used, that data must be supplemented by SME judgment that considers the Benefits of any Mitigations that are expected to be implemented prior to the GRC period under review in the RAMP submission. For each enterprise risk, the utility must explain how they derived the probability distribution for Likelihood of a Risk Event. The utility will consider all known relevant Drivers when specifying the ~~Frequency~~ Likelihood of a Risk Event. Drivers should reflect current and/or forecasted conditions and may include both external actions as well as characteristics inherent to the asset. For example, where applicable, Drivers may include the presence of corrosion, vegetation, dig-ins, earthquakes, windstorms, or the location of a pipe in an area with a higher likelihood of dig-ins.

13.	Calculation of Risk	For purposes of the Step 3 analysis for each enterprise risk assessed in the RAMP, pre- and post-mitigation risk will be calculated by multiplying the distribution representing Likelihood of a Risk Event (LoRE) by the probability distribution of Consequences of a Risk Event (CoRE) and represented as a probability distribution. The CoRE is the sum of each of the ~~Risk-Adjusted~~ Attribute Values probability distributions monetized using the utility's full Cost-Benefit Approach.

Recommendation 2 (R2): *Include and define tail risk as a risk measure.* In addition to using average risk, defined as the average of the probability distribution of risk, tail risk should be formally added for risk evaluation. The measure of tail risk should be tail average above a percentile (the percentile to be determined by the Commission in consultation with stakeholders). Tail average is preferred over other measures because it captures the entire tail of the distribution, is stable, and can be optimized using linear programming or other methods.

Level 4 recommends adding the following definitions to the RDF:

- Expected Value: the sum of all values in the probability distribution divided by the count of values in the probability distribution. Expected Value can be calculated for LoRE, Attributes of CoRE, and Risk.
- Tail Average: the sum of all the values in the probability distribution above a specified percentile divided by the count of values within that same specified percentile of the probability distribution. For example, Tail Average at the 95th percentile is the sum of all values above the 95th percentile in the probability distribution divided by the count of values above the 95th percentile in the probability distribution. Tail average can be calculated for Attributes of CoRE and Risk.
- Tail Risk: a measure of low probability, high consequence occurrences, which are represented in the extremities of the probability distribution, known as the tail. The tail is typically defined as the values above a specified percentile, such as the 95th percentile. Tail risk can be evaluated for Attributes of CoRE and Risk.

Based on R2, Level 4 recommends that D.24-05-064, Appendix A Row 5 be rewritten as

5.	Cost-Benefit Approach Principle 4 – Risk Assessment	~~When~~ Attribute Levels that result from the occurrence of a Risk Event are uncertain.~~, assess the uncertainty in the Attribute Levels by using expected value or percentiles, or by specifying well-defined probability distributions, from which expected values and tail values can be determined.~~ This uncertainty must be represented as a probability distribution and must be described by using the Expected Value of the probability distribution and the tail average above a specified percentile of the distribution. Monte Carlo simulations, other simulations (including calibrated subject expertise modeling), and output from machine learning models, among other tools, may be used to satisfy this principle.

Recommendation 3 (R3): *Evaluation based on portfolios of mitigations.* Risk reduction evaluation should be based on portfolios of risk mitigations to account for interrelationships between mitigations. Portfolio selection is well-suited to optimization (see R4 below).

Level 4 recommends adding the following definition related to R3:

Mitigation Portfolio: a collection of one or more risk mitigations for reducing the risk of a given enterprise risks. Costs, benefits, and benefit-cost ratios can be calculated for each portfolio, and portfolios can be compared to one another.

Mitigation Group: the combining of two or more mitigations that exhibit either synergy, meaning the mitigations result in mutually reinforcing risk reduction efficiency, or diminishing returns, meaning as one mitigation reduces risk it limits the efficiency of the other mitigation to reduce risk.

Based on R3, Level 4 recommends that D.24-05-064, Appendix A include a new row after Row 25 and before Row 26 on portfolio construction, as well as revisions to Row 26.

25.1	Portfolios of Risk Mitigations	Utilities must construct portfolios of risk mitigations for each Risk as identified in Row 8. Mitigations in each portfolio should account for interrelationships between them, such as mutual exclusivity, synergies, and diminishing returns. • Mutually exclusive mitigations must be avoided, only one or the other can exist in the same portfolio. • Synergies and diminishing returns can be captured by combining two or more mitigations, called a mitigation group. Synergies or diminishing returns can be calculated for the mitigation group. For example, a wildfire mitigation portfolio could include for a given circuit segment: covered conductor as mitigation, vegetation management as a mitigation, or covered conductor with vegetation management as a mitigation—but not covered conductor and vegetation management as separate mitigations since their benefits are not additive (may exhibit diminishing returns).
26	Mitigation Strategy Presentation in the RAMP and GRC	The utility's RAMP filing will provide a ranking of all RAMP Mitigations by ~~Cost-~~Benefit-Cost ~~r~~Ratios. Additionally, the utility must present a set of optimal portfolios for reducing each enterprise risk. Mitigation Groups defined in Row 25.1 can also be ranked within each portfolio. The utility must justify the portfolio selection, optimization, and structure of Mitigation Groups. In the GRC, the utility will provide a ranking of Mitigations by ~~Cost-~~Benefit-Cost Ratios, as follows: (1) For Mitigations addressed in the RAMP, the utility will use risk reduction estimates, including any updates, and updated costs to calculate ~~Cost-~~Benefit-Cost Ratios and explain any differences from its RAMP filing; (2) For Mitigations that require Step 3 analysis under and consistent with Row 28, the utility will include the ~~Cost-~~Benefit-Cost Ratios, calculated in accordance with Step 3, in the ranking of Mitigations by ~~Cost-~~Benefit-Cost Ratios. In the GRC, the utility will provide an updated presentation of a set of optimal portfolios for reducing each enterprise risk if an update is necessary. Any differences in the set of optimal

		portfolios from the RAMP filing must be clearly explained by the utility in its GRC filing. In the RAMP and GRC, the utility will clearly and transparently explain its rationale for selecting Mitigations for each enterprise risk and for its selection and optimization of its ~~overall~~ portfolio of Mitigations for each enterprise risk. The utility must explain how the Benefit-Cost Ratios constraint and other constraints factored into the utility's portfolio selection. ~~The utility is not bound to select its Mitigation strategy based solely on the Cost-Benefit Ratios produced by the Cost-Benefit Approach.~~ Mitigation selection and Mitigation Portfolio optimization can be influenced by Benefit-Cost Ratios and other factors including, but not limited to, funding, labor resources, technology, planning and construction lead time, compliance requirements, Risk Tolerance thresholds, operational and execution considerations, and modeling limitations and/or uncertainties affecting the analysis. In the RAMP and GRC, the utility will explain whether and how any such factors affected the utility's Mitigation selections. In the RAMP and GRC, the utility must also implement and justify a systematic way to integrate these other factors into the optimization of its Mitigation Portfolios. GRC Post-Test Year Reporting: All Controls and Mitigation programs must include Benefit-Cost Ratios in each of the GRC post-test years as well as aggregate Benefit-Cost Ratios for the entire post-test year period and the entire GRC period, by Tranche.

Recommendation 4 (R4): *Portfolio selection based on simple optimization instead of ranking*. Optimization ensures choosing the best portfolio of mitigations given the objective and constraints. It can, however, be a complex, computationally intensive, and time-consuming process. There are ways to simplify the optimization process such as limiting the number of optimization scenarios and choosing objectives that can be optimized using linear programming, which is computationally efficient and speedy compared to non-linear methods. There are three components to our simple optimization recommendation:

- *Stochastic optimization*: Stochastic optimization is optimizing using the entire probability distributions, not single numbers. It typically returns an efficient frontier and enables optimizing for average risk and tail risk (see next bullets). Linear programming is one method for performing stochastic optimization, but the utilities may use their preferred method.
- *Efficient frontier*: An efficient frontier is the set of optimal and near-optimal portfolios based on a two-dimensional trade-off, such as risk reduction versus mitigation cost. Efficient frontiers enable trade-off analysis and alternative analysis as defined in D.24-05-064, Appendix A Page A-3.

- *Two scenarios*: Two efficient frontiers can be created, one for each of two stochastic optimization scenarios:
    - *Scenario 1*. Minimize average post-mitigation risk[73] for various mitigation cost levels.
    - *Scenario 2*. Minimize tail average post-mitigation risk for various mitigation cost levels.

We acknowledge it may be necessary to start with optimizing average risk (scenario 1) and incorporating tail risk (scenario 2) in a later cycle, depending on the utility's expertise in stochastic optimization.

Recommendation 5 (R5): *Calculation of risk tolerance.* Risk tolerance should be modeled as an exceedance curve and calculated by applying the risk neutral or risk averse scaling function to a constant risk exceedance curve.

- *Risk tolerance* is the maximum amount of residual, or post-mitigated, risk that an entity or its stakeholders are willing to accept after the application of risk Control or Mitigation. Risk tolerance can be influenced by legal or regulatory requirements.
- *Exceedance curves* depict the maximum acceptable Consequence for a given probability of a risk event. Risk attitudes such as risk neutrality or risk aversion can be applied to exceedance curves by applying an appropriate scaling function. After the application of the scaling function, an exceedance curve is the probabilistic representation of risk tolerance.
- *The Constant Risk Exceedance Curve*[74] is the curve that results in the same Expected Value of Risk for every probability. For example, for an Expected Value of $10 risk, the Constant Risk Exceedance Curve would include the points 10% Likelihood of $100 Consequence; 1% Likelihood of $1,000 Consequence; and 0.1% Likelihood of $10,000 Consequence.

This recommendation significantly modifies D.24-05-064, Appendix A Row 7, which applies the scaling function to an attribute Consequence. R5 enables the comparison of the actual probability distribution of Consequence to risk tolerance in the form of a scaled exceedance curve. The scaling function is more intuitively applied to the constant risk exceedance curve for an attribute, not to the attribute Consequence itself.

Based on R5, Level 4 recommends adding the following definitions to the RDF:

Exceedance Curve: A function that depicts the maximum level of acceptable Consequence for an attribute for a given probability that the Risk Event will occur.

Constant Risk Exceedance Curve: the curve that results in the same Expected Value of Risk for every probability. For example, for an Expected Value of $10 risk, the Constant Risk Exceedance Curve would include the points 10% Likelihood of $100 Consequence; $1% Likelihood of $1,000 Consequence; and 0.1% Likelihood of $10,000 Consequence.

---

[73] Defined as residual risk in D.24-05-064, Appendix A, page A-5

[74] Also known as iso-risk curve. https://citeseerx.ist.psu.edu/document?doi=bef8e5125d5dcede72b599c97c6644e520ed6520&repid=rep1&type=pdf See page 7.

Based on R5, we recommend that D.24-05-064, Appendix A, a new Row be added after Row 6 and Row 7 be revised as follows:

6.1	Cost-Benefit Approach Principle 6: Attribute Exceedance Curves	Establish a Constant Risk Exceedance Curve for each attribute relevant to a given risk event. Each Attribute Level Constant Risk Exceedance Curve must depict the *maximum* level of acceptable Consequence for the associated probability that a given Consequence occurs. Each point on the curve represents the same Expected Value of risk. It will inform the establishment of the Constant Risk Exceedance Curves for Risk Events in Row 13.1.
7	Cost-Benefit Approach Principle 6 – Applying ~~Risk~~ Scaling Function to the Attribute Exceedance Curves	~~Apply a Risk Scaling Function to the Monetized Levels of an Attribute or Attributes (from Row 6) to obtain Risk-Adjusted Attribute Levels.~~ For each enterprise risk included in the RAMP, the utility may apply a Scaling Function reflecting Risk Attitude to the Attribute Level Constant Risk Exceedance Curve (from Row 6.1) to obtain a Scaled Attribute Exceedance Curve. The Scaled Attribute Exceedance Curve (which represents Risk Tolerance, see Row 13.1) is obtained by dividing the Attribute Level Constant Risk Exceedance Curve by the Scaling Function. The ~~Risk~~ Scaling Function is an adjustment made in the ~~risk~~ model due to different magnitudes of Outcomes, which can capture aversion or indifference towards those Outcomes. The ~~Risk~~ Scaling Function can be linear or convexly non-linear. For example, the ~~Risk~~ Scaling Function is linear to express indifference if avoiding a given change in the Monetized Attribute Level does not depend on the Attribute Level. Alternatively, the ~~Risk~~ Scaling Function is convexly non-linear to express aversion if a change in the Attribute level results in an increasing rate of change in the ~~Risk-~~Adjusted Monetized Attribute Level as the Level of the Attribute increases. When completing Rows 5 and 24 in the RDF, if a utility chooses to address tail risk using the power law or other statistical approach and chooses to present ~~Risk-~~Adjusted Attribute Levels by relying on a convex scaling function, then it must supplement its analysis by also presenting ~~Risk-~~Adjusted Attribute Levels by relying on a linear scaling function.

13.1	Risk Tolerance	Utilizing the Attribute Level Constant Risk Exceedance Curves from Row 6.1, establish a Constant Risk Exceedance Curve for each enterprise risk assessed in the RAMP. The Constant Risk Exceedance Curve must depict the *maximum* level of acceptable Risk for the associated probability that a given Risk Event occurs.

		Since each point on the curve represents the exact same level of risk, it is called the Constant Risk Exceedance Curve. The goal of the RDF is to reduce Attribute Consequence Levels below each Risk Tolerance, which is the Scaled Attribute Exceedance Curve. No later than one month after the utility's pre-RAMP workshop, the utility must present its preliminary Attribute Level Exceedance Curves and Constant Risk Exceedance Curve for each enterprise risk assessed in the RAMP to the California Utility Risk Tolerance Stakeholder (CURTS) Forum. Within 21 days of the CURTS Forum discussion, stakeholders of the CURTS Forum should make recommendations to the utility for ensuring that the Attribute Level Exceedance Curves and Constant Risk Exceedance Curve appropriately represent the risk tolerance of the residents of California. The utility must submit these recommendations with its RAMP Application along with a justification explaining why the utility did or did not integrate the CURTS Forum recommendations into its RAMP Application.

Recommendation 6 (R6): *Establish risk tolerance representing the residents of California*. Risk tolerance is the benchmark that determines whether utility risk levels are acceptable or not. Developing a set of acceptable risk levels that represents the risk tolerance of the residents of California requires an inclusive process that should begin as soon as possible. The process should include the following components:

1. *Participants*. Establish a forum of key stakeholders whose consensus on risk tolerance would represent the residents of California. This will be called the California Utility Risk Tolerance Stakeholder (CURTS) Forum. The forum should be established by July 2025 with the goal of informing the SCE 2026 RAMP (see # 2 below).
2. *Timing, pacing, and sequencing*. Develop a timeline for the implementation of risk tolerance standards. Initial implementation should be SCE 2026 RAMP, PG&E 2028 RAMP, and SEMPRA 2029 RAMP.
3. *Number of tolerances to be set*. Determine which tolerances are needed, for example, one for each attribute and for which risks.
4. *Interim tolerances determined by each utility*. While the process for determining State-wide tolerance levels is playing out, requires each utility to declare and justify a risk tolerance, and evaluate risk reduction based on this risk tolerance. See #5 for how many risk tolerances the utilities need to declare in the first cycle.
5. *Phased approach*. Consider initially setting risk tolerance for one risk event, such as wildfire, and for each attribute, and then adding tolerance if desirable over time.
6. *Long-term vision*. With some experience working with risk tolerance and simple optimization, discuss whether to move ahead with more sophisticated frameworks such as ALARP.

Based on R6, we recommend that D.24-05-064, Appendix A, a new Row between 13 and 14 be added:

13.2	Test Year Risk Tolerance	The utility must determine how much risk can be reduced in the next GRC cycle to approach the Constant Risk Exceedance Curve or Scaled Exceedance Curve for each enterprise risk assessed in the RAMP filing.

8 Appendices

8.1 More Information and Details for Selected Topics

8.1.1 Appendix A: The Arithmetic of Uncertainty

Foremost among the first principles of risk tolerance and simple optimization of utility risk reduction is "don't use single numbers as inputs to risk models." Single numbers are prevalent in risk modeling because they are easy to use. Everyone knows how to add, subtract, and multiply single numbers according to the rules of arithmetic.

Correct risk modeling, however, requires the use of probability distributions to avoid serious errors, which are described in subsequent appendices. Probability distributions can also be added, subtracted, and multiplied according to the arithmetic of uncertainty[75] (along with other mathematical operations, but these are the most common ones), which allows them to be used as inputs in risk models.

The arithmetic of uncertainty is defined as performing arithmetic on one or more probability distributions with the calculations resulting in a new probability distribution.

Technically, in mathematics, this is called "functions of random variables."[76] Historically, this meant analytical solutions that could be applied to only a few probability distributions and had limited practical use. In the 1980s when computers became sufficiently powerful, financial engineers in banking and insurance developed proprietary methods for applying the arithmetic of uncertainty. In the last ten years, personal computers and software such as MS Excel have become powerful enough to perform the arithmetic of uncertainty, and ProbabilityManagement.org has created an open standard to further enable it.[77]

Uncertainty can be represented as probability distributions displayed as a list of numbers. In Excel or any programming language, the list can be entered in a column. Two probability distributions can be entered as two columns, as in Figure 8-1. Let's assume these distributions are drawn from a risk model of cybersecurity events at two data centers.

---

[75] For more on arithmetic of uncertainty, see "The Arithmetic of Uncertainty," Chance Age Webinar Services: Probability Management, https://www.probabilitymanagement.org/arithmetic-of-uncertainty?rq=arithmetic

[76] See Gordon Zitkovic, "Lecture 4: Functions of Random Variables," lecture in Mathematical Statistics, last updated September 25, 2019. https://web.ma.utexas.edu/users/gordanz/notes/functions_color.pdf

[77] Sam L. Savage, *Chancification: How to Fix the Flaw of Averages*, (published by author 2021). Chapter 3.

	Distribution 1	Distribution 2
Trial 1	3	3
Trial 2	4	3
Trial 3	2	5
Trial 4	2	1
Trial 5	3	4
Trial 6	5	4
Trial 7	6	6
Trial 8	2	2
Trial 9	4	3
Trial 10	2	3
Trial 11	3	4
Trial 12	1	2

Average:

Distribution 1 = 3.1

Distribution 2 = 3.3

Figure 8-1. Two probability distributions are represented as a list or column of numbers.

Each row is a trial, perhaps representing a modeled year, taken from a large risk simulation. In trial 3, there were 2 risk events in Distribution 1 and 5 risk events in Distribution 2. What if we wanted to know the total risk between the two data centers? Those unfamiliar with the arithmetic of uncertainty might fall back on the more familiar grounds of single numbers and calculate, then sum, the averages of the two distributions, 3.1 + 3.3 = 6.4.

However, the arithmetic of uncertainty allows adding the two distributions as follows:

	Distribution 1	Distribution 2	Distribution 1 + Distribution 2
Trial 1	3	3	6
Trial 2	4	3	7
Trial 3	2	5	7
Trial 4	2	1	3
Trial 5	3	4	7
Trial 6	5	4	9
Trial 7	6	6	12
Trial 8	2	2	4
Trial 9	4	3	7
Trial 10	2	3	5
Trial 11	3	4	7
Trial 12	1	2	3

On Trial 3, Distribution 1 + Distribution 2 = 2+5=7

Figure 8-2. Probability distributions can be added, creating a new probability distribution.

Each trial is summed, creating the new distribution (i.e., Distribution 1 + Distribution 2). We can calculate the average of the new distribution, which happens to be 6.4—the same as the single number approach, though this is not always so as we shall soon see. But even in the case where the average of the new distribution does not differ from the single number approach, on trial 7 the total risk is 12, almost twice the average, and on trial 6 the total risk is 9, nearly 50% more than the average. This critical tail risk information is destroyed when the distributions are collapsed into single numbers.

Here is a visualization of the arithmetic of uncertainty in action:

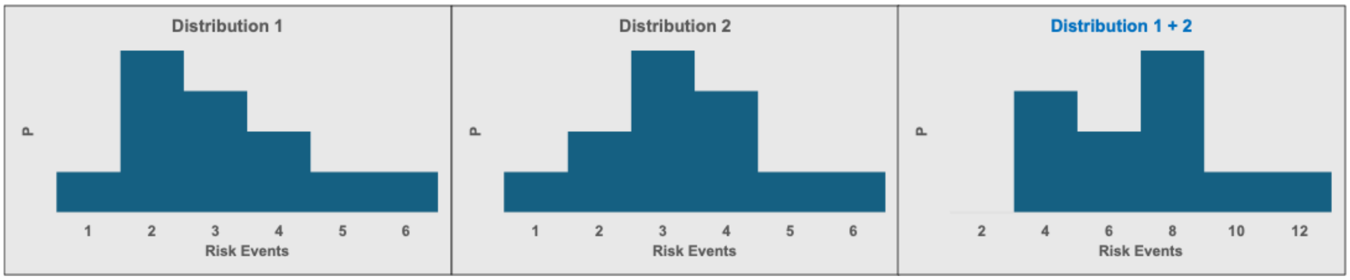

Figure 8-3. Distributions as histograms.

Distribution 1 is different from Distribution 2 and adding them together creates a new distribution with an entirely different shape! It is unnecessary to "name" the distributions based on classical statistics. Is Distribution 1 lognormal? It doesn't matter. Is Distribution 2 normal? It doesn't matter. Nobody would've known what the sum of the distributions, Distribution 1+2, would look like.

Another key point is that the new Distribution 1 + 2 can be used in further risk calculations, such as calculating the distribution of customers impacted by the risk events.

So far, we have demonstrated how the arithmetic of uncertainty works for adding probability distributions. However, risk modeling is often multiplicative; no problem, we can also take the product of probability distributions.

Let's make a slight change in our example—suppose Distribution 1 is a number of risk events as before but now Distribution 2 is a consequence (for simplicity assume that all events on a trial will have the same consequence). We are interested in total risk = events x consequence, which is depicted in Figure 8-4.

	Distribution 1	Distribution 2	Distribution 1 x Distribution 2
Trial 1	3	3	9
Trial 2	4	3	12
Trial 3	2	5	10
Trial 4	2	1	2
Trial 5	3	4	12
Trial 6	5	4	20
Trial 7	6	6	36
Trial 8	2	2	4
Trial 9	4	3	12
Trial 10	2	3	6
Trial 11	3	4	12
Trial 12	1	2	2

Figure 8-4. Distribution 1 x Distribution 2.

Multiplying two probability distributions also creates a new probability distribution and Distribution 1 x Distribution 2 looks like this:

Figure 8-5. Distribution 1 x Distribution 2 histogram.

We are certain that this distribution does not have a name in classical statistics, and few people would've guessed its shape before performing the arithmetic. Fortunately, modern computing doesn't care about names, and it calculates the shape for us. As before, this new distribution may be used in subsequent risk calculations, such as aggregating risk across all risk events.

A final comment on Distribution 1 x Distribution 2: Using the single number approach, average risk = average of Distribution 1 x average of Distribution 2 = 3.1 x 3.3 = 10.2. The average of Distribution 1 x

Distribution 2 = 11.4. This is a case where using single numbers doesn't even get the averages right. As before, using single numbers fails to alert anyone that there is a 1/12 chance of risk = 36, more than 3 times the average.

The arithmetic of uncertainty plays a critical role in risk modeling using the FAIR ontology, where Risk = Likelihood of Risk Event x Consequence of Risk Event, or LoRE x CoRE. Unfortunately, FAIR models are typically based on single numbers, for example, LoRE 40% x CoRE $200 = Risk $80. The good news is the FAIR approach is fully compatible with probability distributions and the arithmetic of uncertainty, as in Figure 8-6.

	LoRE	CoRE	Risk
Trial 1	0	$100	$0
Trial 2	**1**	**$200**	**$200**
Trial 3	0	$50	$0
Trial 4	**1**	**$350**	**$350**
Trial 5	0	$500	$0
Trial 6	0	$300	$0
Trial 7	0	$25	$0
Trial 8	**1**	**$150**	**$150**
Trial 9	**1**	**$100**	**$100**
Trial 10	0	$225	$0
Average	**40%**	**$200**	**$80**
P90th%		**$365**	**$215**

Figure 8-6. FAIR ontology and arithmetic of uncertainty.

LoRE, CoRE, and Risk can be visualized as probability distributions shown in Figure 8-7. In addition to average risk, tail risk can be computed from the distributions, which is impossible when using single numbers.

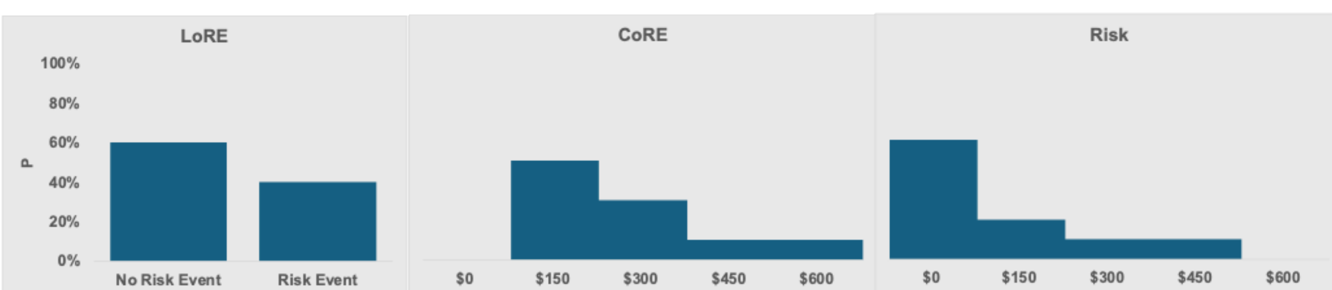

Figure 8-7. LoRE, CoRE, LoRE x CoRE histograms.

The arithmetic of uncertainty requires preserving probability distributions at every step. Every output probability distribution is a potential input distribution for a future calculation. Once a summary statistic such as an average or a percentile is calculated that number may be used in reports or dashboards but can never be used again as an input to a model.

We have begun to touch on a major problem with using single numbers instead of the arithmetic of uncertainty when modeling risk, known as the Flaw of Averages, which is the topic of the next appendix.

8.1.2 Appendix B: The Flaw of Averages

The Flaw of Averages is the set of systematic errors that occurs when people use single numbers (usually averages) to describe uncertain future quantities.[78] When these single numbers are used as inputs to complex models, they can produce erroneous results. The Flaw of Averages is magnified when model inputs are multiplied, exponentiated, interrelated, or some combination of the three.

A wildfire risk model example will illustrate the Flaw of Averages. It combines all three Flaws of Average magnifiers—multiplication, exponentiation, and interrelationships. Suppose we are modeling the impact of wind speed on wildfires, where higher windspeeds are associated with an increased likelihood of ignition and increased wildfire consequences, as shown in Figure 8-8.

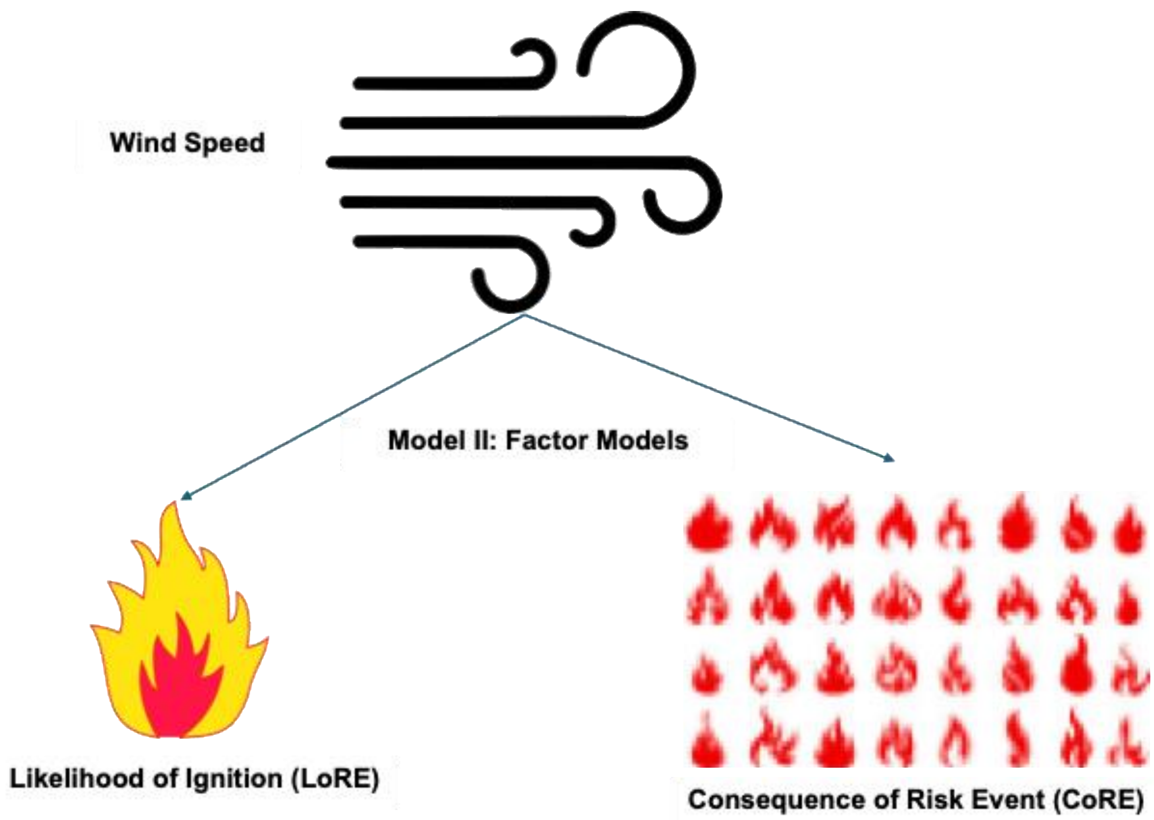

Figure 8-8. Simple wildfire factor model.

---

[78] See https://johnmjennings.com/beware-the-flaw-of-averages/ . Also, Sam L. Savage, *The Flaw of Averages*. (John Wiley & Sons, 2009).

- Wind Speed is typically modeled as a Weibull distribution, which is exponentiated and often results in a "fat-tailed" distribution.
- Wildfire consequence is often modeled as a power law, which is exponentiated and results in an extreme "fat-tailed" distribution.
- In the FAIR ontology, risk is multiplicative—the product of likelihood and consequence.
- Likelihood and consequence share the same cross-cutting risk factor—windspeed—and are therefore interrelated.

The distributions of our hypothetical model are visualized in Figure 8-9 below.

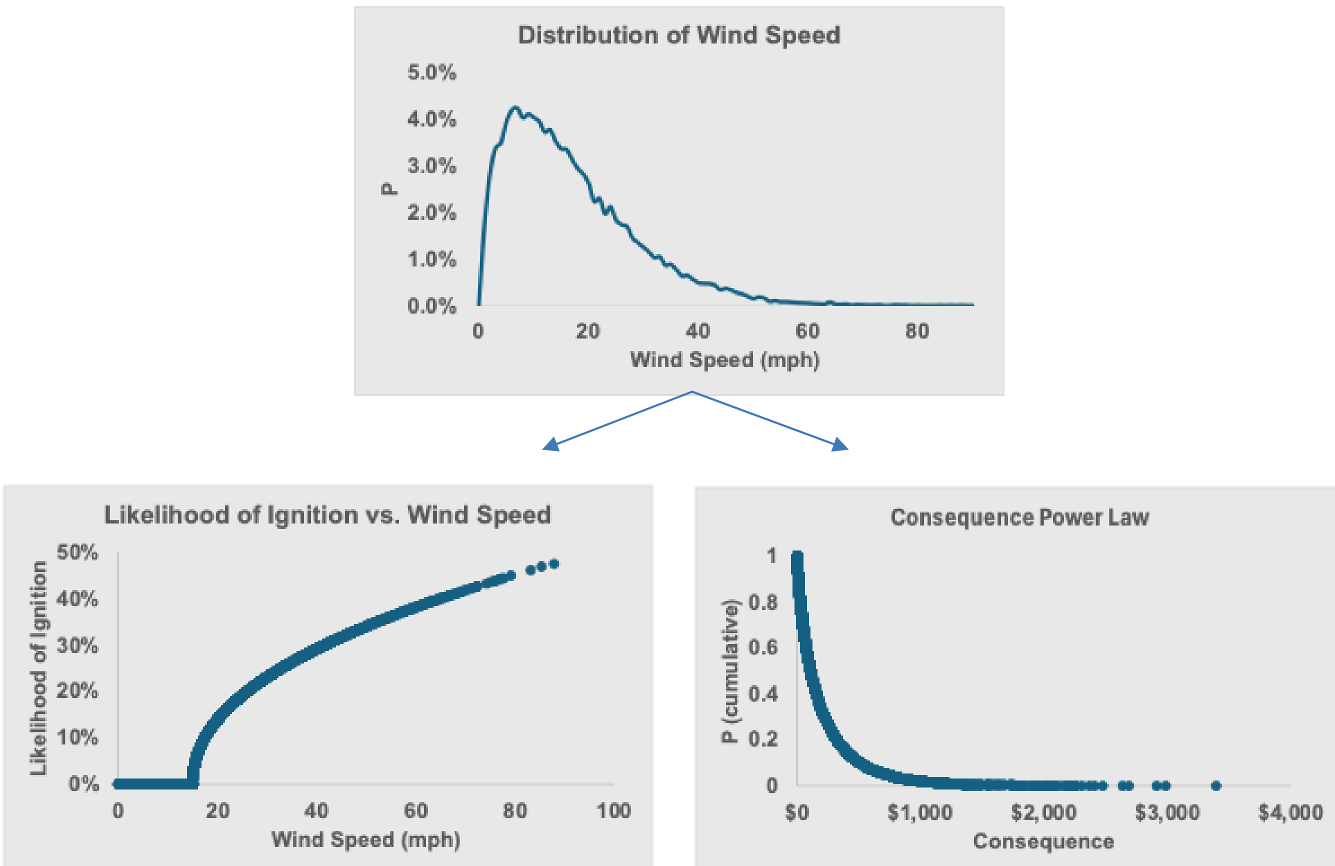

Figure 8-9. Factor model distributions.

The interrelationship between likelihood and consequence produces a Spearman rank correlation of roughly 0.5, as shown in Figure 8-10. A correlation of 0.5 is moderate to high; as wind speeds increase, the likelihood of ignition increases as do the consequences of a wildfire risk event.

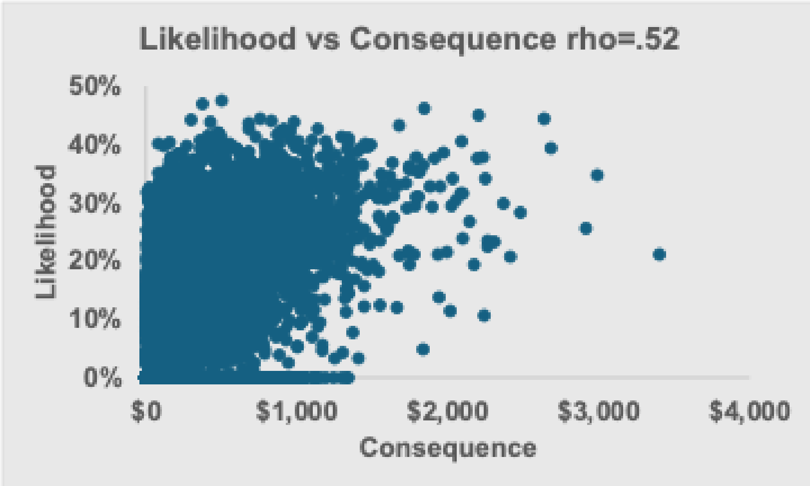

Figure 8-10. Interrelationship between wind speed and consequence.

As discussed in the arithmetic of uncertainty in Appendix A, these distributions serve as the inputs to our wildfire model, which produces a distribution of wildfire risk as shown in Figure 8-11 below. While the average is $30.56, up until around the 92 percentile, the Risk is zero and then increases rapidly. The 99th percentile risk is $750.57, about 25x the average.

	Risk
Average	$30.56
Median	$0.00
92%	$24.97
95%	$202.05
99%	$753.57
99.5%	$985.88
99.9%	$1,489.40

Figure 8-11. Distribution of wildfire risk, descriptive statistics.

Figure 8-11 results are calculated from the distribution of risk, based on a Monte Carlo Simulation of 20,000 trials. Figure 8-12 below examines alternative risk calculations based on LoRE and CoRE:

Risk Calculations	LoRE	CoRE	Risk
Single Number Input (a)	7.4%	$147.22	$10.86
Avg LoRE x Avg Core (b)	8.3%	$200.96	$16.58
Avg (LoRE*CoRE) (c)			$30.56

Figure 8-12. LoRE, CoRE, and Risk Calculations.

- The most common—and incorrect—alternative is to take a single number, typically the average, and plug it into a model. This approach (row (a) in Figure 8-12) takes the average windspeed and plugs it into the LoRE and CoRE models, and then multiplies LoRE and CoRE model output as single numbers to calculate risk.
- Less common, but also incorrect, the distributions of LoRE and CoRE are created and the single number averages from those distributions are then multiplied to calculate risk, as in row (b).
- The correct approach, in row (c), is to multiply the distributions for LoRE and CoRE, and take the average from the new distribution.

None of this would matter if these alternative methods produced the same or even related results. They often do not. In our hypothetical example, the single-number approach understates average risk by a factor of 2.8x. Multiplying the averages of LoRE and CoRE understates average risk by 84%.

Given the construct of risk models—the product of LoRE and CoRE, the exponentiated (fat-tailed) distributions, and interrelationships between the inputs—the incorrect approaches will *systematically* underestimate risk.

If these errors aren't bad enough, it gets worse. Once the single numbers are used in the calculations, it becomes impossible to evaluate other measures of risk such as tail risk. In the first two calculations, it is no longer possible to evaluate risk at the 95th percentile or the 99th percentile. The evaluator would never know that there was a 5% chance of risk more than 6x the average, or a 1% chance of risk 25x the average.

8.1.3 Appendix C: The Flaw of Extremes

A close relation to the Flaw of Averages is the Flaw of Extremes, which results from combining or aggregating abnormal results, such as 90th percentiles, minimums, or maximums, or other results from the tails of probability distributions.[79] For example, adding the 95th percentile from two different distributions will not produce the 95th percentile from the aggregated distribution unless they are perfectly correlated.

For example, suppose a utility is developing a plan to reduce risk for its portfolio of five hydropower dams. The utility has developed probabilistic risk models for each dam and wishes to address tail risk at the 98th percentile, equivalent to a 1-in-50-year event. The five dams are geographically dispersed from each other and share no common cross-cutting risk factors, and their pre-mitigated risks are considered independent. The pre-mitigated risk levels for each dam are represented in Figure 8-13.

[79] Sam L. Savage, *The Flaw of Averages*. (John Wiley & Sons, 2009). Chapter 17.

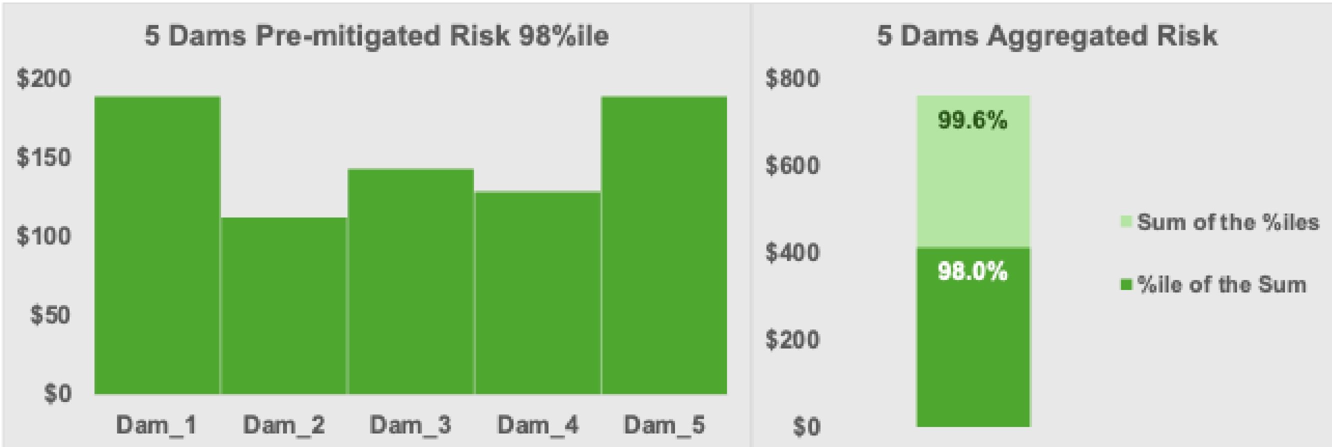

Figure 8-13. Pre-mitigated risk levels at 98th percentile for 5 dams.

The left-hand chart shows the 98th percentile pre-mitigated risk levels for each of the five dams. The right-hand chart shows the aggregated risk level. The light green (top) bar on the right-hand chart is the result produced by adding up the 98th percentile risk level for each of the five dams. The dark green bar on the right-hand chart adds up the probability distributions of pre-mitigated risk levels for the five dams and then takes the 98th percentile, which is the correct approach.

Why is adding the individual 98th percentiles incorrect? It assumes that the 98th percentiles all occur at the same time (or on the same Monte Carlo simulation trial), which if the risks are independent, is extremely unlikely. In this example, the risk calculated by adding the 98th percentile can be expected to occur only once every 250 years (99.6%). Figure 8-14 compares the Flaw of Extremes (the sum of the 98th percentile) with properly applying the arithmetic of uncertainty (the 98th percentile of the sum).

	Risk	Actual Pth%	Frequency
Sum of P98th%	$762.37	99.6%	1 in 250 years
P98th% of Sum	$413.02	98.0%	1 in 50 years

Figure 8-14. The Flaw of Extremes vs. the arithmetic of uncertainty.

The aggregated pre-mitigated risk level is overstated by 85%, which may lead to substantial over-investment in mitigations.

This example demonstrates the Flaw of Extremes by aggregating only five risks. The Flaw scales rapidly as the number of risks being aggregated increases. Once again, avoiding the Flaw of Extremes depends on properly applying the arithmetic of uncertainty to probability distributions first, and only then calculating summary statistics such as averages or percentiles.

8.1.4 Appendix D: Likelihood of Simultaneous Failure (LoSF)

Catastrophic risk events are seldom the result of a single failure. In most cases, there are multiple failures simultaneously, which can be thought of as a "perfect storm." We call it "likelihood of simultaneous failure" or LoSF.[80] Understanding if and when LoSF occurs can be the key to effective risk mitigation.

Determining LoSF is impossible if risks are reduced to single numbers. Once again, the arithmetic of uncertainty provides us with a way to uncover LoSF.

Suppose we are evaluating two independent risks, e.g., wildfire and cyber. They are independent because they do not share any cross-cutting risk factors. For simplicity, the wildfire risk and cyber risk have the same probability distribution. The sum of the two independent distributions is shown on the far-right hand chart in Figure 8-15.

Figure 8-15. Independent distribution of hypothetical wildfire and cyber risk.

Suppose we are also evaluating wildfire risk and hydropower risk, which are not independent. They are interrelated because wildfires can cause damage to reservoirs and dams.[81] We will assume that hydropower also has the same probability distribution as wildfire (and cyber), and the interrelated wildfire and hydropower risk is shown in Figure 8-16.

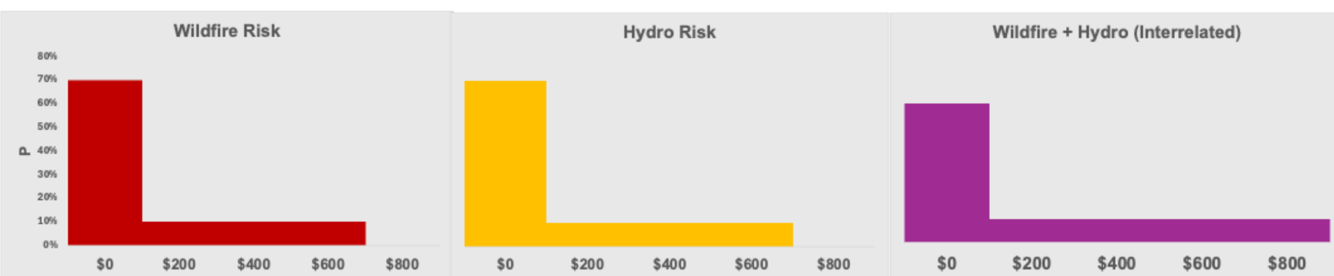

Figure 8-16. Interrelated wildfire and hydropower hypothetical risk.

The aggregated wildfire and hydro risk (far-right hand chart) has a significantly longer tail. What is the root cause? Let's look at the probability distributions depicted as columns of numbers in Figure 8-17 to find out.

---

[80] We attribute the term likelihood of simultaneous failure and LoSF to Dr. Sam Savage.

[81] For an example of wildfire consequence impact on hydropower and other clean energy capacity, see https://www.ncpa.com/wp-content/uploads/2018/01/Issue_Paper_Wildfires.pdf

Independent Risks

	WF Risk	Cyber Risk	Total
Trial 1	$0	$500	**$500**
Trial 2	$0	$0	**$0**
Trial 3	$100	$0	**$100**
Trial 4	$0	$300	**$300**
Trial 5	$0	$0	**$0**
Trial 6	$500	$0	**$500**
Trial 7	$0	$0	**$0**
Trial 8	$300	$0	**$300**
Trial 9	$0	$0	**$0**
Trial 10	$0	$100	**$100**
Average	**$90**	**$90**	**$180**
P90th%			**$500**

Interrelated Risks

	WF Risk	Hydro Risk	Total
Trial 1	$0	$500	**$500**
Trial 2	$0	$0	**$0**
Trial 3	$100	$0	**$100**
Trial 4	$0	$0	**$0**
Trial 5	$0	$0	**$0**
Trial 6	**$500**	**$300**	**$800**
Trial 7	$0	$0	**$0**
Trial 8	**$300**	**$100**	**$400**
Trial 9	$0	$0	**$0**
Trial 10	$0	$0	**$0**
Average	**$90**	**$90**	**$180**
P90th%			**$800**

Figure 8-17. Distribution detail of independent and interrelated risks.

The tables show the probability distributions for each risk as a set of Monte Carlo simulation trials. Each risk occurs on 3 of 10 trials for a likelihood of 30%, and each averages $90. The combined risk for the independent risks and interrelated risks are also the same, $180. As we saw on the histograms, however, the tail risk at the 90th percentile is much higher for the interrelated risks, $800 vs. $500. Examining the trials shows why: the interrelated risks occur simultaneously on trial 6 and trial 8. Perfect storm! For the independent risks, there are no trials in which both risks occur.

While these examples are highly illustrative, interrelated risks are more likely to occur simultaneously and are more likely to produce extreme events.

How might understanding LoSF inform risk mitigation? In Japan, it has been well understood for decades that earthquakes can damage critical infrastructure such as power for nuclear reactors, and regulators have required backup generators typically installed in the basements of nuclear plants. It was also well known that earthquakes can cause tsunamis. Regulators failed to evaluate what happens when an earthquake knocks out power and causes a tsunami at the same time. In March 2011, such an earthquake occurred. The backup generators in the basements survived the initial quake but were flooded by the tsunami and incapacitated, which contributed to the meltdown of the reactor.[82] Had LoSF been understood, might the backup generators have been placed in a higher location?[83]

[82] See "Fukushima Daiichi Accident," World Nuclear Association, last updated April 29, 2024, https://world-nuclear.org/information-library/safety-and-security/safety-of-plants/fukushima-daiichi-accident

[83] In fact, one generator in a higher location was undamaged by flooding (see above footnote).

8.1.5 Appendix E: Relationship Between Risk Scaling and Risk Tolerance

In section 4.2.2, we describe risk scaling as the quantification of risk attitude and risk tolerance as the probabilistic representation of risk attitude. Risk scaling and risk tolerance are related, that is the same risk aversion (or risk-seeking) function can be applied to both and produce the same results. There can be some confusion when moving between risk scaling and risk tolerance since risk aversion is represented as a convex curve for the former and concave for the latter.

Risk scaling. A risk scaling function is applied to the risk-neutral curve to make "perceived" risk higher than actual risk, as in Figure 8-18. By definition, for the risk-neutral curve (the blue line) perceived risk is the same as actual risk—actual risk of 1 is perceived as 1, actual risk of 100 is perceived as 100, and actual risk of 10,000 is perceived as 10,000.

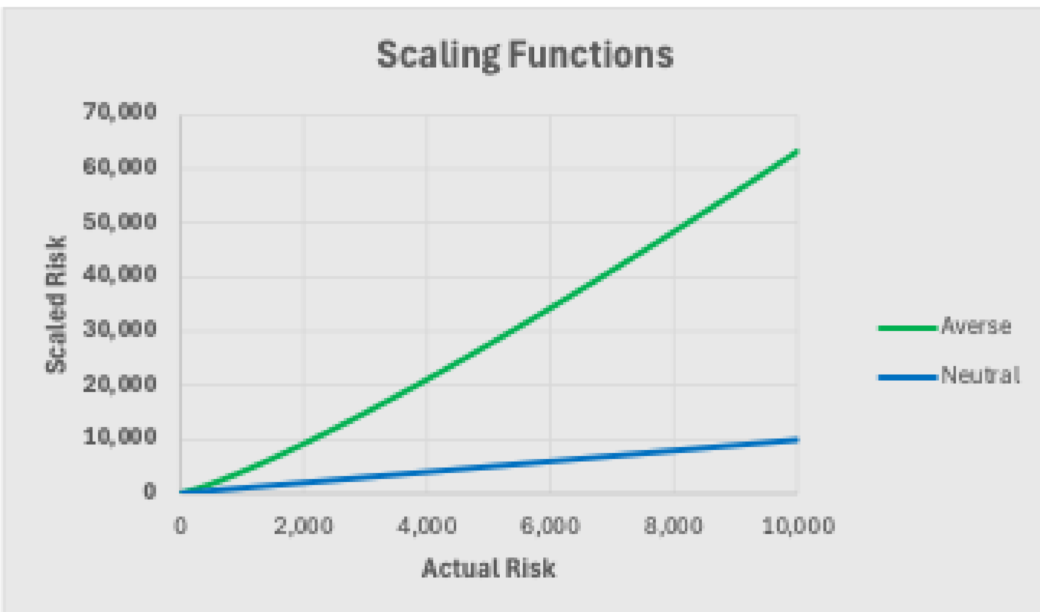

Figure 8-18. Risk-averse vs. risk-neutral scaling function.

For the risk-averse scaled risk (green curve), the perceived risk is higher than the actual risk at each point. The curve is upward-sloping, or convex. The same curve can be shown in log-log space in Figure 8-19, which will make comparing multiple curves easier. For the risk-averse curve, the interpretation is the same—perceived risk after scaling is higher than actual risk.

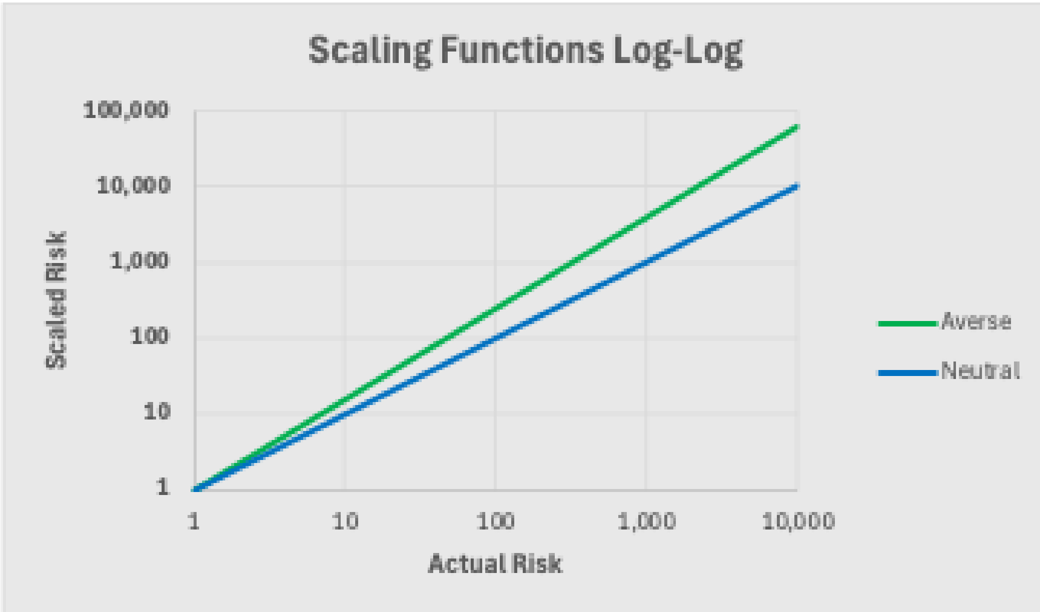

Figure 8-19. Scaling functions are presented in logarithmic space for readability.

Risk Tolerance. Rather than adjusting actual risk, risk tolerance sets the level of acceptable risk, to which actual risk can be compared. The risk-averse tolerance will be set below the risk-neutral line in Figure 8-20.

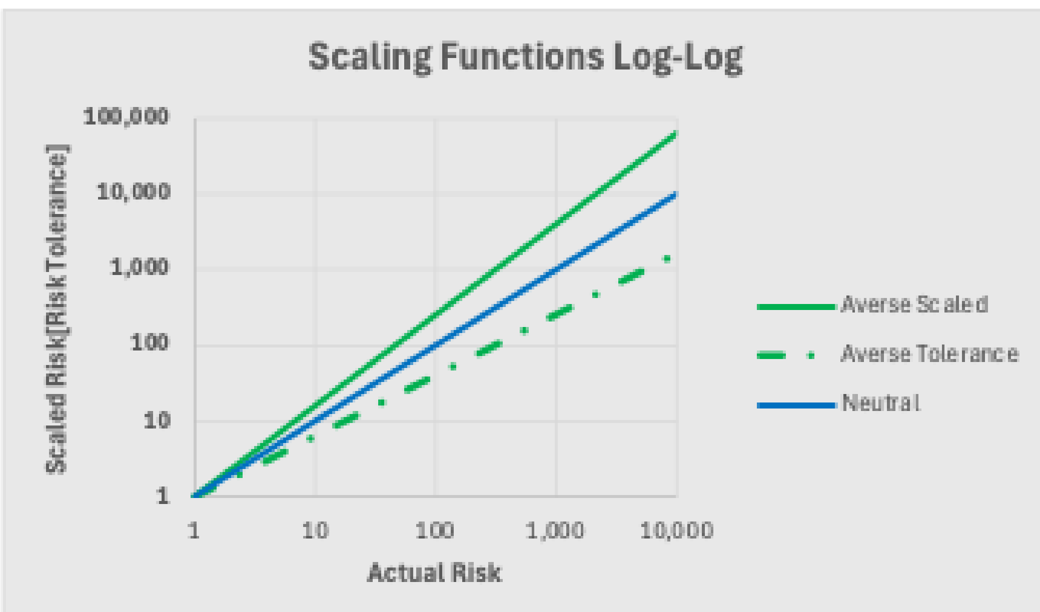

Figure 8-20. Applying risk scaling functions to risk tolerance.

For a given risk-averse function, the risk tolerance line (dotted green) is the mirror image of the averse scaled risk line around the risk-neutral line. This is why the risk-averse curve in risk scaling is convex, but for risk tolerance it is concave. Mathematically, the relationship is shown in Figure 8-21.

Relationship between risk-averse scaled risk function and risk tolerance:

$$\frac{Scaled\ Risk}{Actual\ (Neutral)\ Risk} = \frac{Actual\ (Neutral)\ Risk}{Risk\ Tolerance}$$

$$\frac{25}{10} = \frac{10}{4}$$

Figure 8-21. Mathematical relationship between scaled risk and scaled risk tolerance.

Risk Tolerance, probabilistic. Risk tolerance is often represented probabilistically as exceedance curves. This means that risk tolerance and scaled risk, which are plotted on the y-axis above, move to the x-axis, and the y-axis becomes the cumulative probability distribution as shown in Figure 8-22 below.

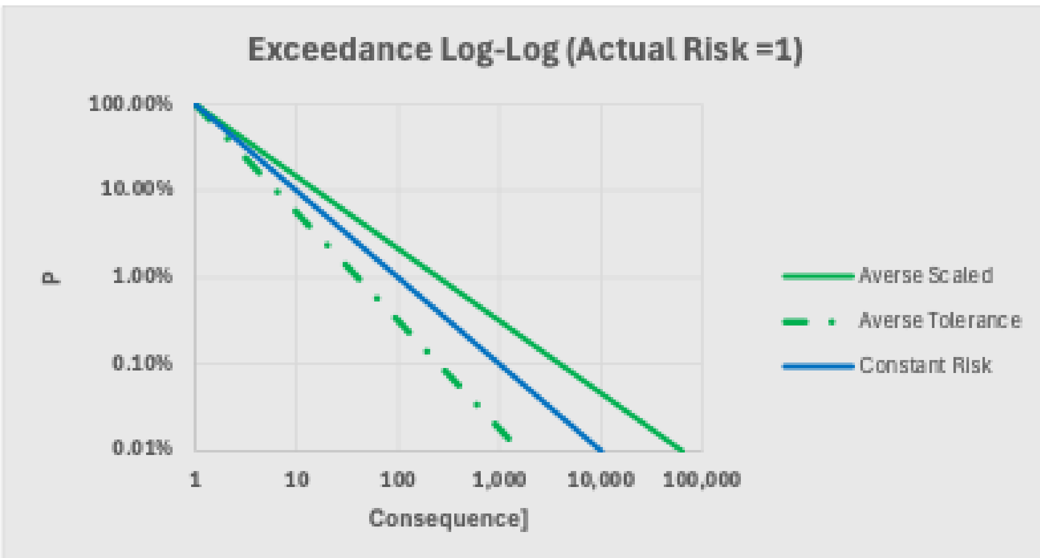

Figure 8-22. Transforming risk scaling function to exceedance curve.

The exceedance curve rotates the chart 90 degrees, preserving the "mirroring" of the tolerance and scaled curves. Two key differences:

- The exceedance curve presents likelihood and consequence values for a specific level of actual risk, in this case, risk =1. Each level of risk would have its own exceedance curve.
- The mirroring axis is not the risk-neutral line—there is no risk-neutral line because there is no such thing as a risk-neutral probability distribution (see Section 4.4). Instead, the blue line is called the constant risk line since at every point probability x consequence = 1.

The key takeaway is that risk scaling and risk tolerance can be related to each other by the same risk aversion (or risk-seeking) formula.

8.1.6 Appendix F: Maximizing Mitigation Value or Minimizing Post-Mitigated (Residual) Risk

The goal of risk management is to reduce risk to some acceptable level, for average risk and tail risk. This can be achieved by minimizing post-mitigated risk or by maximizing risk reduction, in both cases subject to constraints such as cost-benefit, affordability, etc. We might prefer maximizing risk reduction since mitigation actions are how we achieve risk reduction. If we minimize post-mitigation risk, we have to derive the mitigations associated with that level of risk, which is more cumbersome from a modeling standpoint.

For average risk, minimizing post-mitigation risk and maximizing risk reduction are the same thing. Not necessarily so for tail risk.

It turns out that minimizing post-mitigation tail risk may require a different set of mitigations than maximizing tail risk reduction, and that minimizing post-mitigation tail risk is the only way to guarantee optimal tail risk reduction. Said another way, it is possible that maximizing mitigation tail risk-benefit does not result in minimized post-mitigation tail risk.

This is a surprising and unintuitive result that requires a demonstration.

Consider a simple utility substation with two circuits in a high fire-threat area. The pre-mitigated risk statistics for the two circuits and the total for the substation are presented in Figure 8-23 below. The average risk for Circuit 1 is higher, $360 vs. $212 but the Cvar (or tail average risk) is higher for Circuit 2, $1,468 vs. $1,354.

Pre-mitigated Risk Level	Circuit 1	Circuit 2	Total
LoRE	50%	40%	
Average risk	$360	$212	$572
Cvar @ P95th%*	$1,354	$1,468	$2,113
**Tail risk is not additive*			

Figure 8-23. Summary risk statistics for two circuits.

The utility must choose between two mitigations, one reduces LoRE by 50% for Circuit 1 and the other reduces LoRE by 50% for Circuit 2. The mitigation benefit of the two mitigation choices is as follows, in Figure 8-24.

Mitigation Benefit	Circuit 1 Mitigation 1	Circuit 2 Mitigation 2
Higher is better		
LoRE reduction	50%	50%
Average mitigation impact	$182	$105
Cvar @ P95th% mitigation impact	$1,179	$1,121

Figure 8-24. Comparing mitigation benefits (higher is better).

Based on mitigation benefit, Mitigation 1 reduces more average risk and also has more Cvar risk reduction, defined as the tail average of the risks mitigated. This definition of mitigation Cvar is very important as we will soon see.

The post-mitigation risk, which is pre-mitigated risk minus mitigation impact, is presented in Figure 8-25. The total post-mitigation risk (the substation level) is shown for each of the two mitigation options.

	TOTAL RISK, given choice of:	
Post Mitigation Risk Level	**Mitigation 1**	**Mitigation 2**
Lower is better		
Average risk	$390	$466
Cvar @ P95th%	$1,894	$1,824

Figure 8-25. Comparing post-mitigated risk (lower is better).

For average risk, the lowest post-mitigated risk is for Mitigation 1, which agrees with the mitigation impact results above. Not so for Cvar: Mitigation 2 results in the lowest post-mitigation risk for tail risk.

How can this be? Recall the definition of Cvar for mitigated risk—it is the tail average of the risks mitigated. This does not include the risks that remain unmitigated! It is possible to have high mitigation benefits and also leave the highest risks unmitigated. The Cvar for post-mitigated risks captures the tail risk of remaining risks directly.

Rank Cvar of mitigation impact =/= Rank Cvar of post-mitigated risk

However, it is possible to impute the risk impact of mitigations by calculating the difference between pre-mitigation and post-mitigation tail risk for average risk and for tail average (Cvar) risk, which we will call "net mitigation benefit."[84] Net mitigation benefit is presented in Figure 8-26, and agrees with the ranking of mitigation impact in Figure 8-25 above.

	Net Benefit	
Net Mitigation Benefit	**Mitigation 1**	**Mitigation 2**
Higher is better		
Average net mitigation benefit	$182	$105
Cvar @ P95th% net mitigation benefit	$219	$289

Figure 8-26. Comparing the net mitigation benefit as the difference between pre-mitigated and post-mitigated risk (higher is better).

[84] One must be careful with performing any type of calculation with single numbers from tail risk as discussed in the Flaw of Extremes appendix, but it is justified in this specific case.

Rank (Cvar of pre-mitigated risk minus post-mitigated risk) = Rank Cvar of post-mitigated risk

It is therefore recommended that optimization is defined as minimizing post-mitigated risk, and mitigation benefit imputed from pre-mitigated risk minus post-mitigation risk. Minimizing post-mitigated tail risk will not always be different than maximizing mitigation tail risk impact, but as we have shown, it is possible.

8.1.7 Appendix G: Approaches for Setting Risk Tolerance

Optimizing risk mitigation based on risk tolerance requires risk tolerance to be set at some level in the organization. That level can be set for the organization itself, for every combination of risk events and attributes, or somewhere in between. Figure 8-27 below shows the range of options.

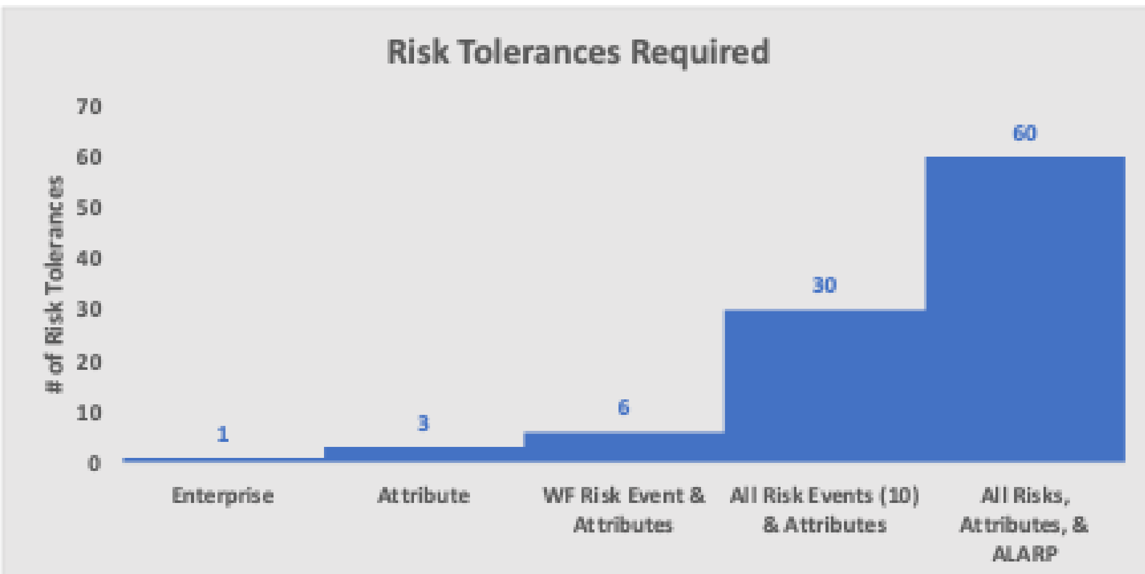

Figure 8-27. Range of risk tolerances required.

Setting a tolerance means establishing the entire exceedance curve, from which average risk tolerance and tail risk tolerance can be calculated.

Establishing risk tolerances for utility risk has not been attempted before and will require a process that includes education, debate, consensus, and a decision. Such a process will be more difficult as more tolerances are required. As many as 60 different tolerances could be required if set at 10 risk events x 3 attributes x two ALARP tolerances. While requiring a lower number of tolerances would be desirable—at least to begin with—this doesn't mean setting a single total risk tolerance at the enterprise level would be sufficient.

Suppose a single tolerance at the utility level has been set at $100 million. In one year, there is a single risk event, say a cyberattack, that results in a massive power outage and $99 million of risk, mostly due to the monetized value of reliability. Technically, total risk is within tolerance—and yet nobody would feel that this was an acceptable outcome.

At a minimum, we would suggest starting by setting risk tolerance at the attribute level and at least for wildfire risk, which would require 6 tolerances: 3 attributes x 2 risk events (wildfire, and a bucket for all other risk events). Each additional risk event added would increase the number of tolerances by 3. For example, setting tolerances for 3 attributes, and wildfire, gas events, hydropower, and all others, would require a total of 12 tolerances (24 for ALARP).

As the utilities and evaluators gained experience working with risk tolerance and optimization, more tolerances can be added as needed.

8.2 Acronyms

Term	Definition
ALARP	As Low as Reasonably Practicable
BCR	Benefit-Cost Ratio
CBA	Cost-Benefit Approach
CoRE	Consequence of Risk Event
CPUC	California Public Utilities Commission
CURTS	California Utility Risk Tolerance Stakeholder
Cvar	Conditional value at risk
ESJ	Environmental and Social Justice
EV	Expected Value
FAIR	Factor Analysis of Information Risk
GPS	Global Positioning System
GRC	General Rate Case
LoRE	Likelihood of Risk Event
LoSF	Likelihood of Simultaneous Failure
MAVF	Multi-Attribute Value Function
MPT	Modern Portfolio Theory
MS	Microsoft
PG&E	Pacific Gas & Electric
RAMP	Risk Assessment and Mitigation Phase

Term	Definition
RDF	Risk-Decision Framework
RMAR	Risk Mitigation Accountability Report
RSE	Risk Spend Efficiency
S-MAP	Safety Model Assessment Proceeding
SCE	Southern California Edison
SDG&E	San Diego Gas & Electric
SME	Subject matter expert
SPD	Safety Policy Division
TURN	The Utility Reform Network
TWG	Technical Working Group
Var	Value at risk
WMP	Wildfire Mitigation Plan

8.3 Definitions

Term	Definition
ALARP	Stands for "As Low as Practicably Reasonable" and is a three-tiered optimization method within a cost-benefit analysis.
Arithmetic of uncertainty	The rules of arithmetic for summing, multiplying, or subtracting probability distributions, are used for calculating risk and for risk aggregation.
Attribute	An observable aspect of a risky situation that has value or reflects a utility objective such as safety or reliability. Changes in the levels of attributes are used to determine the consequences of a risk event (CoRE). The attributes in a cost-benefit approach should cover the reasons that a utility would undertake risk mitigation activities.
Benefit	The reduction in risk, as measured by the changes in attribute levels, which would occur when a program or set of activities is implemented.

Term	Definition
Benefit-Cost Ratio (BCR)	The ratio of monetized benefits (numerator) and costs (denominator) calculated by dividing the dollar value of mitigation benefit by the mitigation cost estimate used in a cost-benefit analysis.
Conditional value at risk (Cvar)	A measure of tail risk, calculated as the average of risks above a given percentile. Also called tail average risk.
Consequence (or Impact)	The effect of the occurrence of a risk event. Consequence affects attributes of a cost-benefit.
CoRE	Estimated dollar value of the consequence of a risk event.
Cost-Benefit Approach	A decision-analysis tool for comparing the monetized benefits of a program, or set of activities, against the costs of the program, or set of activities, to create a measurement of value.
CPUC	California Public Utilities Commission
Cross-cutting risks	Risk drivers or factors that impact more than one risk area. An example would be a seismic event that affects the risk of wildfire and dam failure.
Deterministic	The use of single numbers such as averages in a model, without consideration of any randomness.
Diminishing returns	Occurs when the value of combining two or more mitigations is less than the sum of each of them individually.
Driver	A factor that could influence the likelihood of occurrence of a risk event, the consequence of a risk event, or both. A driver may include external events or characteristics inherent to the asset or system.
Efficient Frontier	The set of optimal portfolios that offer the highest mitigation value for a defined level of cost.
Exceedance curve	A graph that shows for specified probabilities, the level of CoRE that will be equaled or exceeded at each probability.
Expected Value (EV)	The average of all the values in a probability distribution.
Factor model	A type of statistical model that makes predictions based on factors or drivers and takes the form $y = f_1x_1 + f_2x_2 + f_nx_n + e$. Factor models work well with portfolios and influence diagrams.

Term	Definition
Flaw of Averages	A set of systematic errors when using single numbers such as averages as inputs into complex models.
Flaw of Extremes	Mathematical errors that occur when extreme results such as 90th percentiles are added as single numbers. Related to the Flaw of Averages.
Frequency	The number of events generally defined per unit of time. (Frequency is not synonymous with probability or likelihood.)
Gross disproportionality	A concept in some implementations of ALARP where risks must be averted unless there is a gross disproportion between the costs and benefits of doing so. In other words, the BCR is weighted to favor carrying out the safety improvement.
Herringbone	A multi-dimensional visualization of trade-offs.
Independent	In statistics the absence of any influence or causality between variables.
Inherent risk	See pre-mitigated risk.
Interrelated (interrelationship)	In statistics, where the occurrence of one event influences the occurrence of another. The two events can directly influence each other or be jointly influenced by another event. Does not necessarily imply causation.
Likelihood or Probability	The relative possibility that an event will occur, quantified as a number between 0% and 100% (where 0% indicates impossibility and 100% indicates certainty). The higher the probability of an event, the more certain we are that the event will occur.
Likelihood of Risk Event (LoRE)	Likelihood of Risk Event
Likelihood of simultaneous failure (LoSF)	The likelihood of two or more risk events occurring at the same time (as in a "perfect storm"). Impossible to determine when risk is represented by single numbers.
Metalog	A versatile probability distribution that can replicate most continuous distributions without having to know the specific parameters of any particular distribution.
Mitigation	Measured or activity proposed or in process designed to reduce the impact/consequences and/or likelihood/probability of a risk event.

Term	Definition
Mitigation value	The impact on the risk of a mitigation. The difference between pre-mitigated risk and post-mitigated risk.
Monetization	The process of converting a risk consequence into a monetary unit, such as dollars.
Monetized levels of an attribute	E.g., monetized levels of safety attribute. The representation, in dollars, of the potential outcomes that an attribute is exposed to, obtained by converting from the natural units of the attribute levels using an appropriate conversion factor or function.
Monte Carlo simulation	A modeling technique that uses uncertain inputs in the form of probabilities and returns outputs as probability distributions.
Natural unit of an attribute	The way the level of an attribute is measured or expressed. For example, the natural unit of a safety attribute may be fatalities. Natural units are chosen for convenience and case of communication and are distinct from monetized levels of attributes.
Outcome	The final resolution or end result of a risk event.
Optimization	See stochastic optimization.
Portfolios	In the context of risk optimization, a collection of mitigations that captures interrelationships such as synergies and diminishing returns. By definition, mutually exclusive mitigations are excluded.
Post-mitigated risk	The risk that remains after mitigations are applied. Can be thought of as an accepted risk.
Power law	A type of probability distribution that often has a “fat tail” representing very low likelihoods of very high consequence, often catastrophic, events.
Pre-mitigated risk	The current level of risk, before any new mitigations are applied.
Probability distribution	The chances of different outcomes for the event. The probabilities of a complete distribution must add up to 100%.
Residual risk	See post-mitigated risk.

Term	Definition
Risk	The chance of something bad happening, often expressed in terms of a combination of various Outcomes of an adverse event and their associated probabilities.
Risk attitude	Risk attitude is a subjective description of one's willingness to take on risk. Risk attitudes range from risk aversion to risk neutral, to risk seeking.
Risk-averse	A risk attitude where an individual is willing to pay more than the value of risk reduction in order to avoid the risk. It can be represented by a convex scaling function.
Risk event	The occurrence of risk. It can be thought of as when the possibility of a risk becomes a certainty, i.e., the risk occurs. In particular, the occurrence of a risk event changes the levels of some or all of the attributes of a risky situation.
Risk neutral	A risk attitude where an individual is willing to pay the exact amount of the value of risk reduction, no more, no less. It can be represented by a linear scaling function.
Risk scaling function	A function or formula that specifies an attitude towards different magnitudes of outcomes including capturing aversion to extreme outcomes or indifference over a range of outcomes.
Risk seeking	A risk attitude where an individual is only willing to pay less than the value of risk reduction to avoid a risk. It can be represented by a concave scaling function.
Risk spend efficiency	The ratio of a risk score for an initiative in a MAVF framework to the cost for that initiative.
Risk tolerance	The maximum amount of residual risk that an entity or its stakeholders are willing to accept after the application of risk control or mitigation. Risk tolerance can be influenced by legal or regulatory requirements. Risk tolerance is a stochastic function of risk attitude. It may be represented as an exceedance curve of acceptable consequence for each probability (e.g., 50%, 10%, 5%, 2%, 1%, 0.1%, etc.).
Sensitivity analysis	Analysis and statistical tests to determine how various sources of uncertainty in a mathematical model contribute to the model's overall uncertainty.
Scaling function	See Risk scaling function.

Term	Definition
Sparse Monte Carlo	A subset of Monte Carlo simulation that includes the trial information of outcome events only, which improves computational efficiency for extremely low likelihood risks.
Stochastic	Another word for probability.
Stochastic optimization	Optimizing based on a probability distribution, such as average risk and Cvar (tail average risk)—as opposed to optimizing on single numbers, which is deterministic optimization.
Synergy	Occurs when the value from combining two or more mitigations is higher than the sum of the mitigations individually.
Tail average risk	Another term for Cvar.
Tail risk	Risk that is reflected in the tails of a probability distribution. Tail risk focuses on the consequences of rare events.
Trade-offs	The act of giving up something of value to gain something else of value expending the same level of resources.
Tranche	A logical disaggregation of a group of assets (physical or human) or systems into subgroups with like characteristics for purposes of risk assessment.
Uncertainty	The state where it is impossible to exactly describe current conditions or future outcomes.
Value at risk (Var)	A measure of tail risk, represented as the minimum risk at a given percentile.