Overall Residual Risk, Risk Tolerance and Simple Optimization

R.20-07-013 PHASE 4 STAFF PROPOSAL FOR WORKSHOP 2

November 6 2024

California Public Utilities Commission

Thanks to:

Insert contributor names

Contents

Figures

Tables

Executive Summary

One of the key pieces in need of refinement within the Commission's Risk-based Decision-making Framework (RDF) is the issue of risk tolerance. In this proposal, Safety Policy Division (SPD) begins by arguing that California's large investor-owned electric and natural gas utilities must present their respective overall residual risk for decision-makers to evaluate a utility's progress towards a risk tolerance standard. SPD recommends that the Commission should require the utilities in each Risk Assessment and Mitigation Phase (RAMP) filing to include a diagram and supporting workpapers demonstrating the change of overall residual risk since the utility's first RAMP filing.

SPD argues that tolerance for utility risk should be set at the State of California level, representing the residents of California. SPD recommends that a forum of key stakeholders be established whose consensus on risk tolerance would represent the residents of California. This will be called the California Utility Risk Tolerance Stakeholder (CURTS) Working Group. The utilities should engage with the CURTS Working Group while preparing its RAMP Application, include any recommendations from the CURTS Working Group with its RAMP Application along with a justification explaining why the utility did or did not integrate the CURTS Forum recommendations.

For a utility to determine the proper tolerance level for its enterprise risks it must first ensure that its risk models can address the following points:

1. Risk cannot be represented as a single number and must be represented by a probability distribution, which enables the arithmetically correct aggregation of risk and allows the utility to properly calculate average risk and tail average risk.
2. Risk tolerance is the stochastic representation of a subjective risk attitude. If the utility finds the residents of California are risk-averse towards a specific enterprise risk, then risk tolerance could be modified by an appropriate risk scaling function. Risk tolerance can be visualized probabilistically using exceedance curves. Similar to the first point, risk tolerance also cannot be represented as a single number in risk calculations and can be assessed according to both average risk and tail average risk.
3. 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.

Risk tolerance can be established for each of the three consequence attributes (re: safety, reliability, and financial consequences of a risk event). Additionally, risk tolerance can be set for the utility as a whole and/or for each distinct risk event. Establishing risk tolerances for utility risk will be more difficult as more tolerances are required. At a minimum, SPD recommends the Commission require the utilities to start by setting risk tolerance at the attribute level for wildfire risk and a bucket for all other risk events, which would require 6 tolerances (3 attributes x 2 risk event categories). As the utilities, evaluators and stakeholders gain experience working with risk tolerance and optimization, more tolerances should be added.

Establishing risk tolerances lends itself to simple optimization of portfolios of mitigations. SPD argues that the utilities must first begin grouping one or more risk mitigations into portfolios for reducing the risk of a given enterprise risk that can be compared to one another. Optimization can then be multi-dimensional and include assessment of affordability, benefit-cost ratios, and other trade-offs such as safety vs. reliability. Stochastic optimization, for instance, returns an efficient frontier of portfolios and enables the utility to optimize its portfolio according to average risk and tail average risk. This proposal considers how the utility might use linear programming according to the following two scenarios:

- Scenario 1: Minimize average overall residual risk for various mitigation cost levels.
- Scenario 2: Minimize tail average overall residual risk for various mitigation cost levels.

SPD recognizes that various approaches to optimizing portfolios of mitigations for each risk event could be available as long as the goal is minimizing overall residual risk towards Californian's risk tolerance according to various affordability constraints.

By addressing the issues of overall residual risk, risk tolerance and simple optimization, the Commission will guide utilities towards prioritizing the reduction of extreme safety and reliability consequences of risk events to Californians as well as establish quantifiable and practical affordability constraints on portfolios of risk mitigations. SPD argues that integrating these three aspects into the RDF will not only save ratepayers money but should also stimulate the utilities to innovate cost-efficient ways to reduce risk.

1. Background

In August 2016, as part of the Safety Model Assessment Proceeding (S-MAP, Application [A.] 15-05-002 et al.) the Commission published Decision (D.)16-08-018, which explores the topic of 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."[1] The decision also finds that at the time there were problems with the utilities' risk models that precluded them from implementing risk reduction and risk mitigation strategies consistent with the Commission's first risk focused Decision, D.14-12-025. Some major barriers included not having an explicit risk tolerance and no optimization of the portfolio of risk mitigation activities.[2]

In contemplating potential future steps, D.16-08-018 discusses the As Low as Reasonably Practicable (ALARP) framework.[3] Commission staff originally published a whitepaper on ALARP in the S-MAP Proceeding.[4] ALARP combines risk tolerance with a three-tiered optimization process and is focused on safety risk. The ALARP approach 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.[5] In the U.S., the Army Corps of Engineers has used ALARP,[6] as has the nuclear radiation industry.[7]

It was determined that ALARP may be a desirable end state but would not be possible to implement during that current phase of S-MAP. As will be discussed in this proposal, there are other ways to incorporate risk tolerance and optimization short of full implementation of ALARP. Even so, the decision determined that

---

[1] D.16-08-018 at 25.

[2] D.16-08-018 at 164.

[3] D.16-08-018 at 62.

[4] *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. (December 24 2015). https://docs.cpuc.ca.gov/PublishedDocs/Efile/G000/M157/K359/157359431.PDF

[5] 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

[6] 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

[7] 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

"the Commission should adopt explicit risk tolerance standards over time, but not before laying the groundwork in the development of probabilistic risk analysis."[8] The Commission included developing a risk tolerance framework and increasing the application of optimization among eight suggested long-term goals.[9]

D.16-08-018 noted that when the utilities do not provide an explicit risk tolerance in a RAMP or GRC filing, it handicaps the ability of other stakeholders to make an informed decision as to whether the utilities' rate case proposals are an appropriate strategy for reducing risk down to a level that is acceptable by Californians. [10] Additionally, the Commission noted that risk tolerance is necessary to ensure that a portfolio of risk mitigation activities are optimized to reduce risk within the context of limited budget resources and other constraints. The Commission clearly stated that 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 (rate shock, availability of trained personnel, and limitation of resources). And, optimization is always tied to risk tolerance. These concepts are all tied together."[11]

Sinc 2016, risk tolerance and optimization remained a topic of importance for the Commission's approach to risk. D.18-12-014 affirmed the importance of ALARP[12] and stated that the "settlement agreement does not preclude other long-term goals of the Commission, such as 'optimization' and 'explicit risk tolerance standards.'"[13] D.22-12-027 authorized technical working groups (TWGs) to propose recommendations regarding the application of risk tolerance.[14] In the RDF Proceeding Phase 4 Scoping Memo, the Commission noted that the methods for risk tolerance thresholds "can be established by regulators and these methods can include placing spending caps on risk Mitigations. Without such regulations, utilities implicitly set their own risk tolerance standards. The Commission is concerned that a risk tolerance goal that is too high or too low will yield suboptimal outcomes for ratepayer safety or ratepayer costs, respectively."[15]

The idea that an excessively low risk tolerance would have a significant impact on ratepayer costs is directly related to the concept of affordability. The Commission recognizes its legislative mandate to ensure Californians have affordable access to energy and defines affordability as "the degree to which a representative household is able to pay for an essential utility service charge, given its socioeconomic

---

[8] D.16-08-018 at 192.

[9] D.16-08-018 at 175.

[10] D.16-08-018 at 68.

[11] D.16-08-018 at 98.

[12] D.18-12-014 at 55.

[13] D.18-12-014 at 41.

[14] D.22-12-027 at 29.

[15] R.20-07-013, Phase 4 Scoping Memo at 4.

status."[16] There are no shortage of media assessments of what is now described as an "affordability crisis" for access to energy.[17] The recent legislative session discussed numerous bills with explicit reference to scrutinizing the cost of energy, two of which were signed by Governor Newsom.[18]

Academic studies have clearly documented the pathways through which unaffordable energy rates can have adverse health consequences.[19] In the California context, the Mussey Grade Road Alliance (MGRA) recently connected affordability and risk assessment by analyzing the relationship between income, life expectancy and estimated permanent annual rate increases thereby showing that rapid rate increases can result in "a general loss of human health and life".[20] MGRA is concerned that "rate increases required to implement massive wildfire mitigation programs may be approaching the level where they are impacting public health, especially for the poorest and most vulnerable."[21] In other words, there is a point at which reducing risk becomes excessively expensive, the costs of which can have life-threatening impacts to disadvantaged and vulnerable communities of California. Addressing risk on the utility's infrastructure should be balanced with the Commission's ESJ goal of focusing "resources on communities that have been underserved."[22] Requiring the utilities to establish a reasonable risk tolerance would ensure the Commission can achieve its mandated goal of reducing utility risk to an acceptable level while providing reliable and affordable energy to Californians.[23]

As noted by MGRA, a large concern for affordability is the increasing impact of wildfire mitigation costs on ratepayers. The Commission's affordability report notes that wildfire-related capital expenditures, such as installing covered conductor or undergrounding portions of a distribution system, have continued to gradually increase over the 2021 – 2024 period but are not yet a significant portion of the total revenue

---

[16] D.20-07-032 at CoL 1 and CoL 6.

[17] Habegger, Becca. *Unaffordable energy prices: What are California lawmakers doing about it?*, ABC10, August 30 2024. https://www.abc10.com/article/money/unaffordable-energy-prices/103-9310da70-edda-4c4c-9d58-6f28ec03c0cd, Madrigal, Alexis. *Historic PG&E Rate Increases Will Hit Hard in 2024*, KQED, January 4 2024. https://www.kqed.org/forum/2010101895461/historic-pge-rate-increases-will-hit-hard-in-2024

[18] AB-3265 requires the Commission to develop a framework for assessing, tracking, and analyzing total annual energy costs paid by residential households in California. https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB3264. AB-2847 requires utilities to provide their best estimation of the impact of proposed expenditures on the utility's authorized revenue for each year of the life of a capital asset, if the Commission determines it is necessary to provide such an estimate https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240AB2847.

[19] Hernández, Diana. "Understanding 'energy insecurity' and why it matters to health." *Social science & medicine* 167 (2016): 1-10.

[20] A.23-05-010, Exhibit: MGRA-01-E1 at 102

[21] A.23-05-010, Exhibit: MGRA-01-E1 at 99

[22] CPUC Environmental and Social Justice Action Plan 2.0 at 9

[23] See PUC 750, PUC 451, and PUC 739(d)(2).

requirement in rates.[24] Table 1-1 shows the wildfire-related portion of a bundled residential Non-California Alternate Rates for Energy (CARE) customer average monthly bill resulting from the year-end 2023 wildfire-related revenue requirement.

	Total Bill	Wildfire-Related Portion ($)	Wildfire-Related Portion (%)
PG&E	$190.94	$24.42	12.8%
SCE	$174.79	$17.73	10.1%
SDG&E	$182.82	$12.97	7.1%

Table 1-1 PG&E, SCE, and SDG&E Wildfire-Related Portion of Year-End 2023 Average Monthly Bill, Bundled Residential Non-CARE Customers[25]

The Commission's affordability report goes on to explain that only 0.1% of this wildfire-related portion of customers' bill includes undergrounding costs.[26] This percentage will continue to increase as the utilities request cost recovery for undergrounding as a part of their grid hardening strategy. It is not clear to SPD what role affordability plays as utilities design mitigation strategies for their RAMP and GRC filings.

When TURN recently asked SCE what affordability constraints it considered when developing its grid hardening scope for the 2025 GRC filing, SCE replied:

> There are no affordability thresholds or constraints specific to the planned grid hardening scope in the referenced volume of testimony… In other words, we strive to carry out our grid hardening projects in the most affordable manner that is reasonable, but the specific projects themselves are not undertaken to directly address affordability concerns. Instead, they are undertaken to directly address wildfire risk and public safety concerns.[27]

From SPD's perspective, SCE's response to TURN is troubling. The RDF requires quantifiable risk models and analysis to inform decision-making about how much money should be spent to reduce risk. If there are no affordability constraints, that implies any amount of money can justify reducing risk. As stated above, the Commission in D.16-08-018 that this would be an absurd position to take. In this Staff Proposal, SPD will argue that the concept of risk tolerance can address some of these affordability constraints. By integrating an approach to risk tolerance into the RDF, the Commission would not only ensure that the utilities prioritize

[24] CPUC 2024 SB 695 Report at 52.

[25] CPUC 2024 SB 695 Report at 53.

[26] CPUC 2024 SB 695 Report at 54. See also footnote 157 which states: SDG&E did not provide undergrounding costs or equivalent revenue requirement data, stating that with respect to its current GRC proceeding, it does not separately calculate the revenue requirements of specific/individual capital programs.

[27] SCE response to TURN data request, A.23-05-010 TURN-SCE-039 Question 2.a, January 9 2024 at 1.

reducing extreme safety and reliability outcomes of risk events to Californians, but they would also provide the foundational step for establishing quantifiable and practical affordability constraints. SPD argues integrating risk tolerance into the RDF will not only save ratepayers money but should also stimulate the utilities to innovate cost-efficient ways to reduce risk.

2. Overall Residual Risk

The RDF as stated in D.24-05-064 Appendix A does not require the utilities to report on the overall residual risk associated with each enterprise risk submitted in a RAMP or GRC filing. At present, utilities typically present their GRC Test Year Baseline Risk in a RAMP and GRC filing. The utilities then develop mitigation programs based on the estimated GRC Test Year Baseline Risk. The amount of risk remaining after implementing the mitigation programs authorized by a GRC decision that reduces the GRC Test Year Baseline Risk would be the residual risk only within the scope of that GRC application. Overall residual risk would include all the risk on the utility's assets or systems after taking account of the historical progress of risk reduction for every GRC cycle to date.

For the topic of risk tolerance to be properly situated within the RDF, SPD argues it will be necessary for the utilities to report on their progress of reducing overall residual risk for each enterprise risk addressed in a RAMP or GRC filing. This way the Commission can properly determine within the context of a given RAMP or GRC filing whether the utility has properly designed its mitigation programs to address any residual risk to a level Californians can tolerate and at a speed that recognizes the need for prioritizing safety, but also appropriately accounts for the limitation of affordability constraints. The Commission may need to require utilities to report on their progress of reducing overall residual risk in RAMP and GRC filings.

2.1 Example of Presenting Overall Residual Risk

When a utility prepares a RAMP application, it typically only submits risk assessment data for the assets and systems where it intends to implement mitigations. This risk assessment data is only a snapshot of the overall residual risk that exists for that particular enterprise risk. In some instances, the complete risk assessment data may be available through a data request.[28]

Overall residual risk can be presented in easily understandable graphs and Excel tables.[29] Figure 2-1 was generated by TURN and clearly denotes SCE's Overall Residual Wildfire Risk from 2018 out to the end of SCE's next GRC cycle (2028) after accounting for investments in grid hardening and Fast Curve.

[28] For an example see SCE response to MGRA data request, A.23-05-010 MGRA-SCE-006 Q.6-1, 6-1.a-b_Public_TURN-SCE-039_Q6.xlsx, February 7 2024

[29] For an example see SCE response to MGRA data request, A.23-05-010 MGRA-SCE-002 Q.2-2, MGRA-SCE-002_Q2.xlsx, January 22 2024.

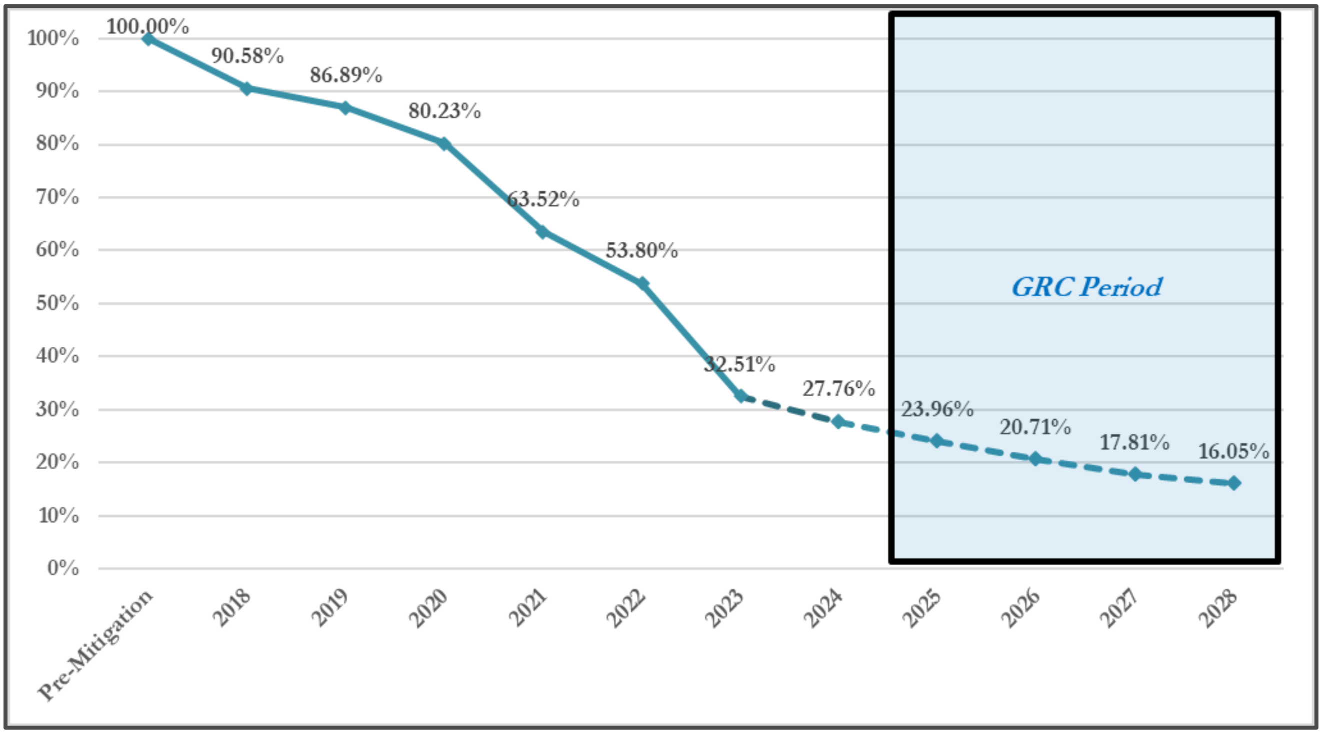

Figure 2-1 SCE Wildfire Risk Remaining After Grid Hardening and Fast Curve Settings[30]

SPD recommends that the Commission require the utilities to submit a similar graph and supporting Excel table for every risk included with each RAMP and GRC filing. Additionally, once a utility determines its risk tolerance, it must plot its exceedance curve on this graph. Doing so will allow parties to know how much closer the utility is to achieving an acceptable amount of risk on its assets and systems. The following sections discuss the procedures the utility should use to determine what that acceptable amount of risk is.

[30] A.23-05-010, TURN-12-E Clean at 21

3. Risk Tolerance

The Cambridge Online Dictionary defines risk as "the chance of something bad happening". In this simple definition, "chance" means a probability from 0% to 100%, and "bad" is a subjective interpretation of an outcome. The Commission's definition of risk is similar but includes specific detail: "the potential for the occurrence of an event that would be desirable to avoid, expressed in terms of a combination of various Outcomes of an adverse event and their associated Probabilities."[31] The Commission's definition corresponds with 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).[32] The following parable will illustrate the probabilistic and subjective elements of risk and the interplay between them.

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. This result is represented by Figure 3-1.

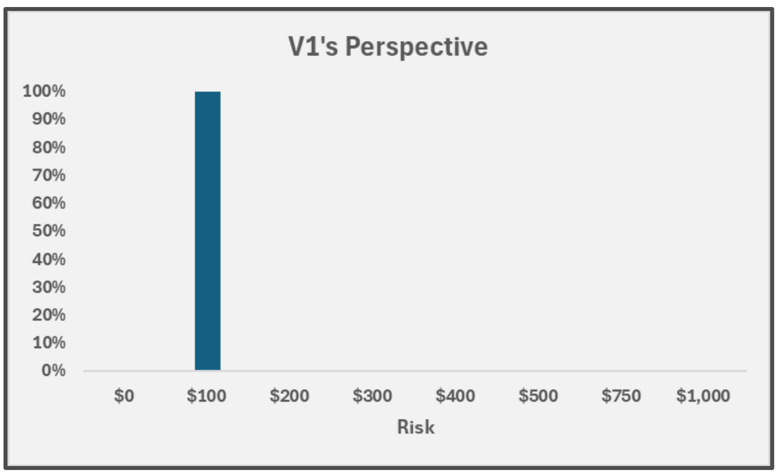

Figure 3-1Venturer 1's perceived risk of loss

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,

[31] CPUC R.20-07-013, D.24-05-064, Appendix A at A-5. Please note: As stated in the Phase 4 Scoping Memo, the Commission is currently considering clarifying the definition of risk stated in the RDF. In all the Phase 4 Staff Proposals, SPD will be preemptively making this change to ensure clarity going forward.

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

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. This result is represented by Figure 3-2.

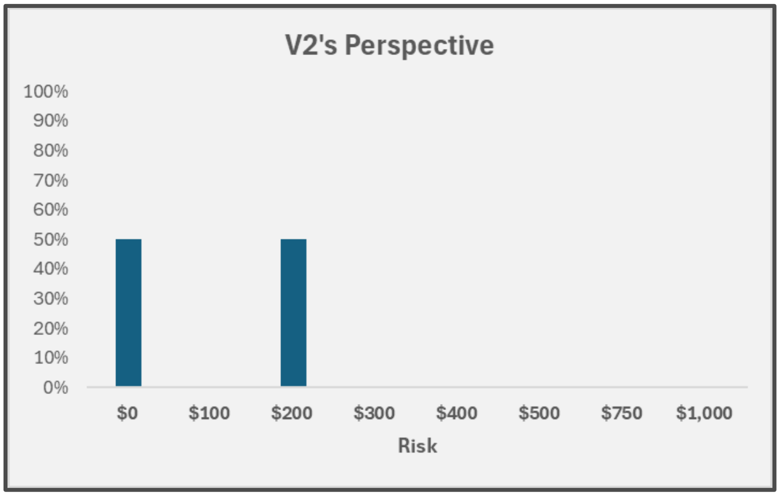

Figure 3-2 Venturer 2's perceived risk of loss

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. This result is represented by Figure 3-3.

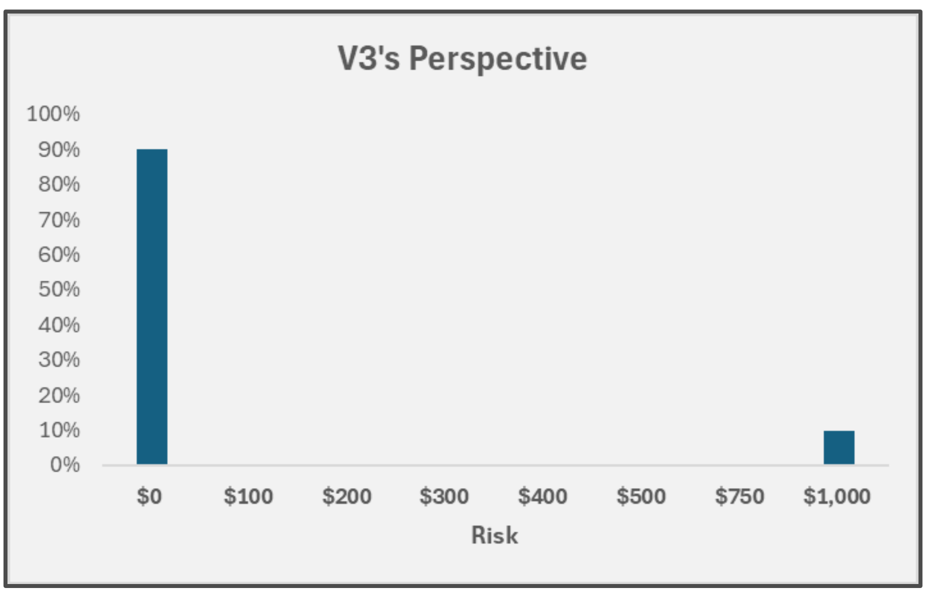

Figure 3-3 Venturer 3's perceived risk of loss

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 also perceives risk probabilistically but assesses it based on belief and instinct. The three perceptions of risk can be visualized in Figures 3-1, 3-2 and 3-3.

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 the three figures.

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.[33]

More complete information and better models can also improve risk assessments, and if individuals are using the same information and same models, we can expect the risk assessments to converge. This is not necessarily a good thing, since it can lead to group think.[34] The key point is that risk assessments are subjective, and even the best of them may differ.

Back to the three venturers. Expected Value (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.

3.1 Defining and Applying Risk Tolerance to Risk Quantification

Risk tolerance, risk attitude, and risk scaling are often used in different ways. For this report, SPD 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

[33] 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.

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

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 3.3.

2. Risk scaling is the quantification of risk attitudes. This is defined by the Commission 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."[35] In other words, risk scaling is how much one is willing to pay to avoid a risk—and more importantly how much a utility is willing to pay to avoid increasing amounts of risk. Risk scaling can be visualized as follows, in Figure 3-4.

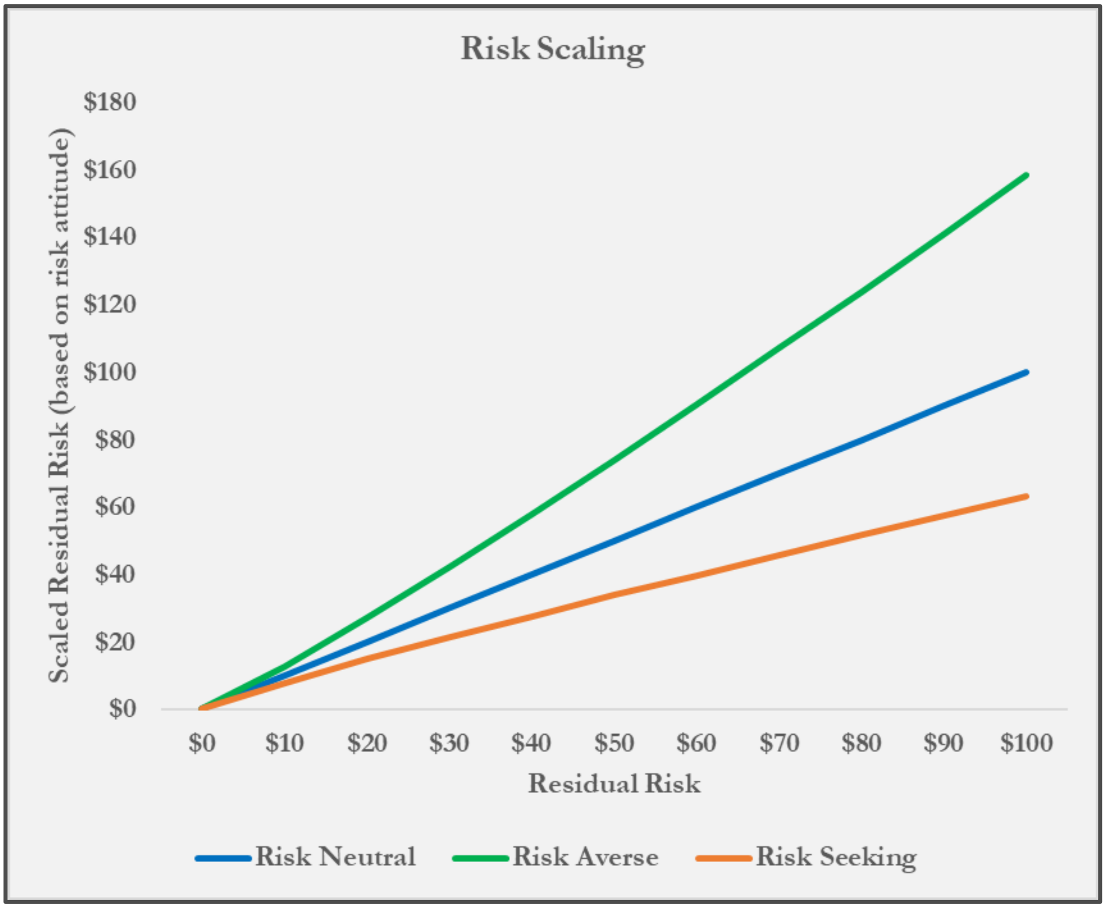

Figure 3-4 Risk scaling curves

All the curves have a positive slope. Risk neutrality is represented by a straight line with a slope equal to 1. This means a utility 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,

[35] D.24-05-064, Appendix A at A-5.

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. SPD 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 Commission definition of risk tolerance is the "Maximum amount of Residual Risk that an entity or its stakeholders are willing to accept after application of risk Control or Mitigation."[36] Risk tolerance can be visualized with exceedance curves, as in Figure 3-5. 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.[37]

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 3-6.

The constant risk exceedance curve is not the same thing as risk neutral, which the Staff Proposal will explain further in Section 3.3.[38]

---

[36] D.24-05-064, Appendix A at A-5.

[37] 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

[38] For more details about the relationship between risk scaling and risk tolerance curves, see Level 4 Incorporating Risk Tolerance and Simple Optimization into the RDF, Appendix E, November 1 2024, https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/safety-policy-division/meeting-documents/risk-tolerance-simple-optimization-level4-final_branded4_110124.pdf

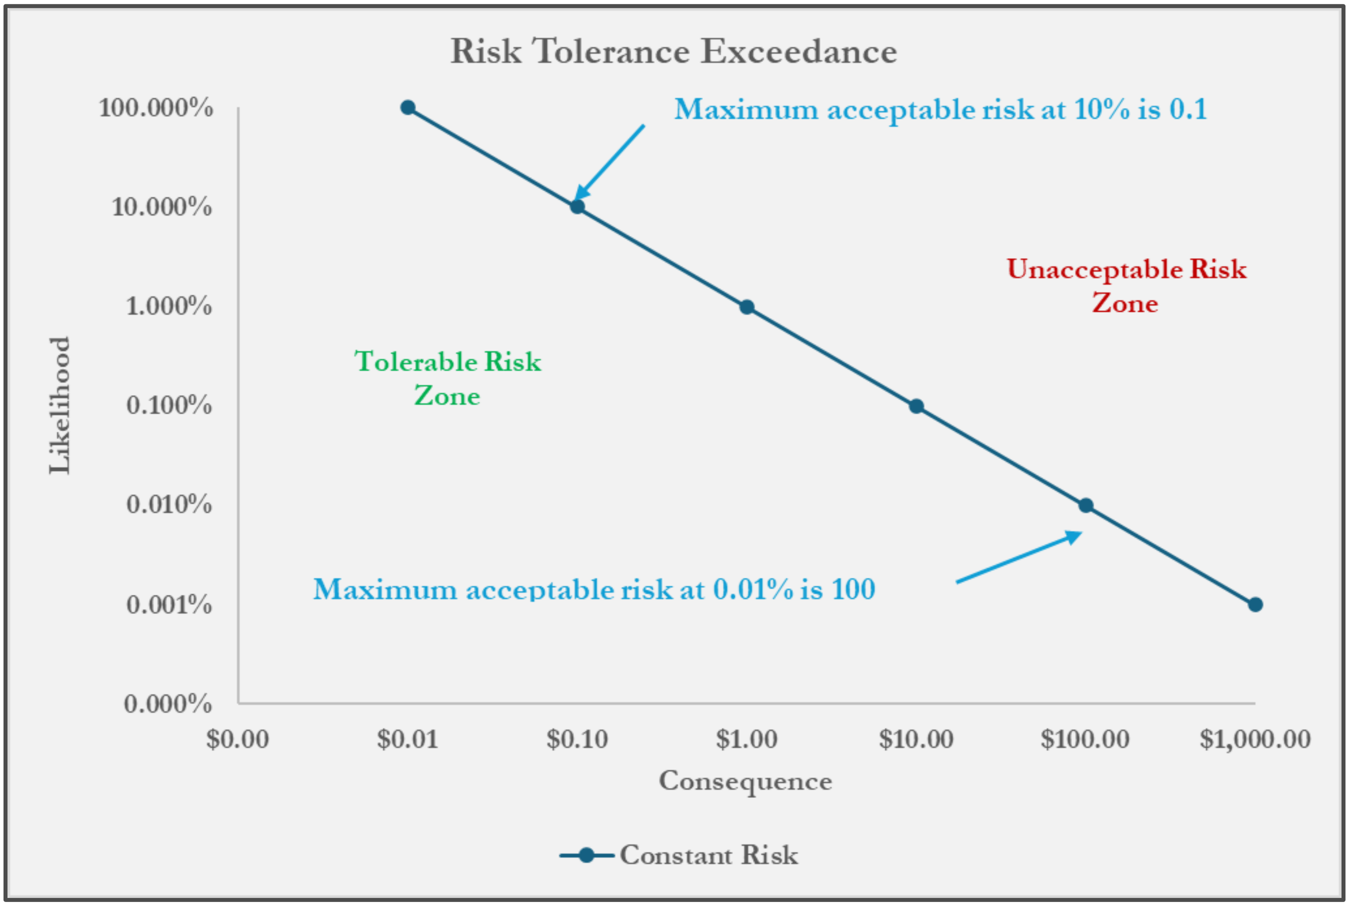

Figure 3-5 Exceedance curve example (in log-log space for readability).[39]

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 3-6, a risk-averse exceedance curve is established by applying a risk-averse scaling function to the constant risk exceedance curve. 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.

[39] 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.

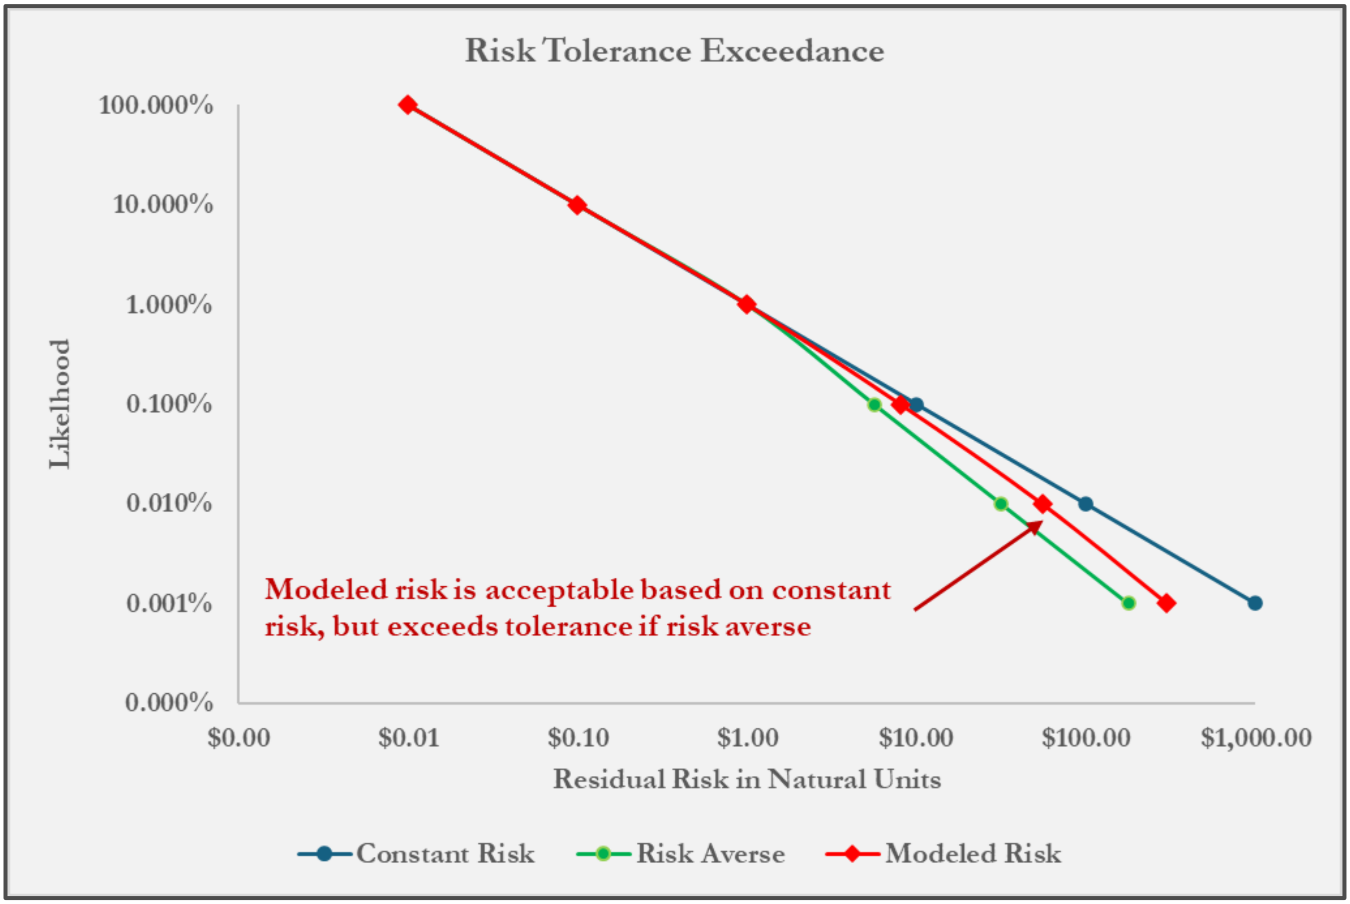

Figure 3-6 Impact of different risk tolerances in evaluating a modeled risk.

We can overlay 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. 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 3-6. 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 disaggregated unit within a mitigation program. After the initial effort to determine risk tolerance standards, it is straightforward to assess whether overall residual risk is within tolerance.

3.2 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 Figures 3-1, 3-2 and 3-3 illustrate very different assessments of the consequences and probabilities. These are not the same risks.

Risk scaling, as discussed in 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.[40] 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.

2. *The Flaw of Extremes*. These are mathematical errors that occur when extreme results such as 90th percentiles are added as single numbers and are related to the Flaw of Averages.[41] 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.

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. LoSF is often a factor in catastrophic events, e.g., "perfect storms."[42] It is impossible to capture simultaneous failures with single-number risk scores.

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 materials" for proper risk modeling are available.[43]

---

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

41 Sam L. Savage. *The Flaw of Averages*. Op.cit. Chapter 17.

42 The term likelihood of simultaneous failure and LoSF is attributed to Dr. Sam Savage.

43 Examples from WMP and RAMPs covered in more detail in Section 5.2. For more details on the arithmetic of uncertainty, the Flaw of Averages, the Flaw of Extremes, and Likelihood of Simultaneous Failure, see Appendices A-D in Level 4 Incorporating Risk Tolerance and Simple Optimization into the RDF, November 1 2024, https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/safety-policy-division/meeting-documents/risk-tolerance-simple-optimization-level4-final_branded4_110124.pdf

3.2.1 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 a utility'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.

3.2.2 Tail Risk Concepts

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?

SPD argues that a single-number scaled risk score is not the best approach, preferring alternative methods for expressing and evaluating tail risk. Figure 3-7 represents a hypothetical wildfire risk power law distribution.[44] 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.

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

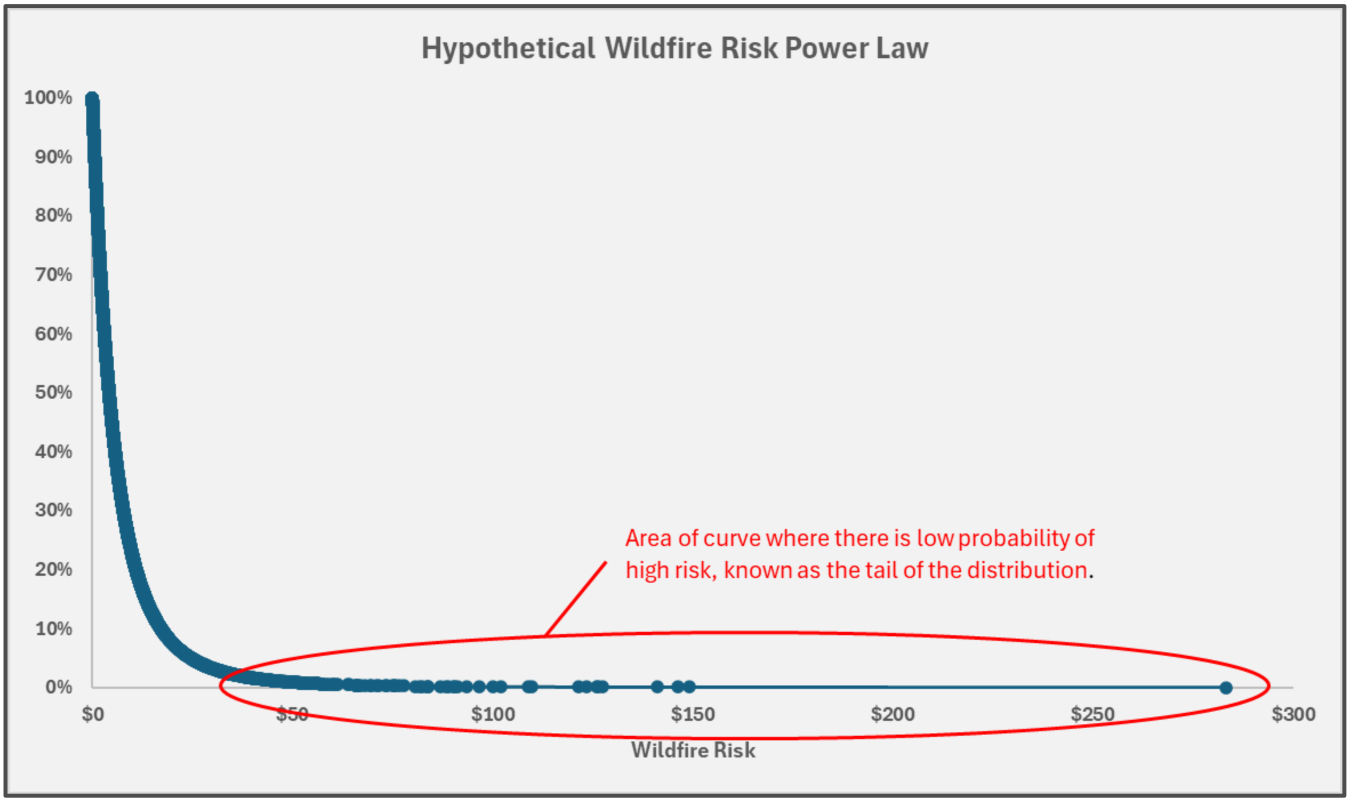

Figure 3-7 Hypothetical wildfire safety risk power law.

The average of the curve in Figure 3-7 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 3-7- 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.[45]

2. *Percentile*. A single value at a point on the curve. In the Figure 3-7 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.

---

[45] 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 it.

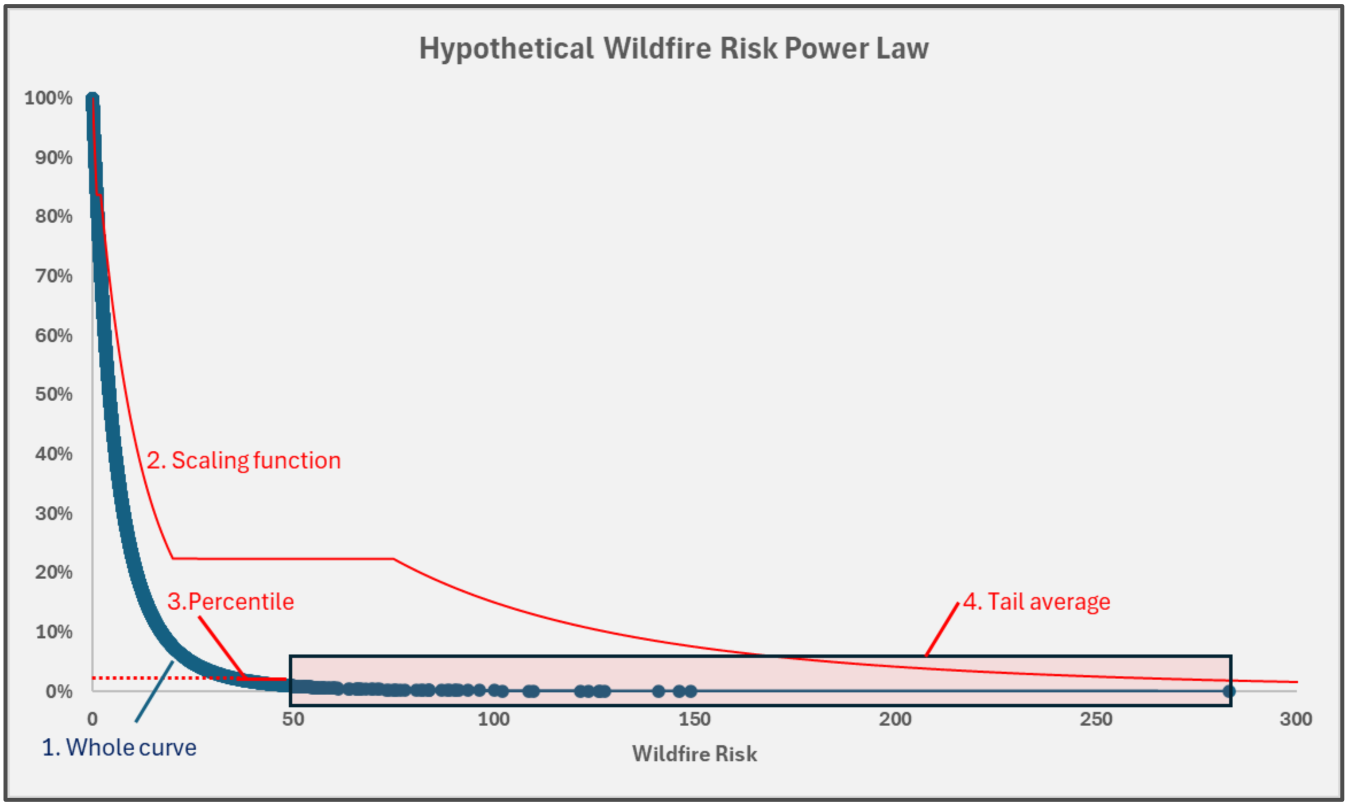

Figure 3-8 Four ways to visualize measures of tail risk.

Table 3-1 calculates the corresponding tail risk measures from Figure 3-8:

Wildfire Risk	
1. Average of entire risk curve	**$7**
2. Scaled Average	**$41**
3. 99th percentile	**$50**
4. Tail average above 99th percentile	**$70**

Table 3-1Average risk and tail risk calculations.

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.
- *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 according to Table 3-1 is $41, which is implicitly the same as using the value at the 98.5th percentile.

- *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 according to Table 3-1 is $50, which excludes significant risks exceeding $100.
- *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 3-8, the tail average for the 99th percentile is $70, which includes all the data points on the curve above $50. The tail average above a certain percentile will *always* be higher than the risk at that percentile.
- *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 adopted in 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.

SPD recommends using the tail average to incorporate tail risk, 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 4.7. For example, mitigations addressing the safety fire risk in Figure 3-8 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.

3.2.3 Risk Tolerance and Cost-Benefit Analysis

In the Cost-Benefit Approach (CBA) adopted by the Commission, the decision to include a proposed mitigation is informed by the benefit-cost ratio (BCR).[46] A BCR of 1.0 means that the benefits of a mitigation exactly equal its costs, so a typical decision rule for investing in a mitigation is a BCR greater than 1.0.

---

[46] D.22-12-027, FoF 11. Please note: As stated in the Phase 4 Scoping Memo, the Commission is currently considering clarifying the nomenclature used in the RDF from Cost-Benefit Ratio to Benefit-Cost Ratio (i.e. Benefit divided by Cost). In all the Phase 4 Staff Proposals, SPD will be preemptively making this change to ensure clarity going forward.

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 a risk to avoid it. Risk-aversion is the willingness to spend more than the EV, while risk-seeking individuals will only pay less than the EV. The risk attitudes can be visualized by the BCR curves in Figure 3-9.

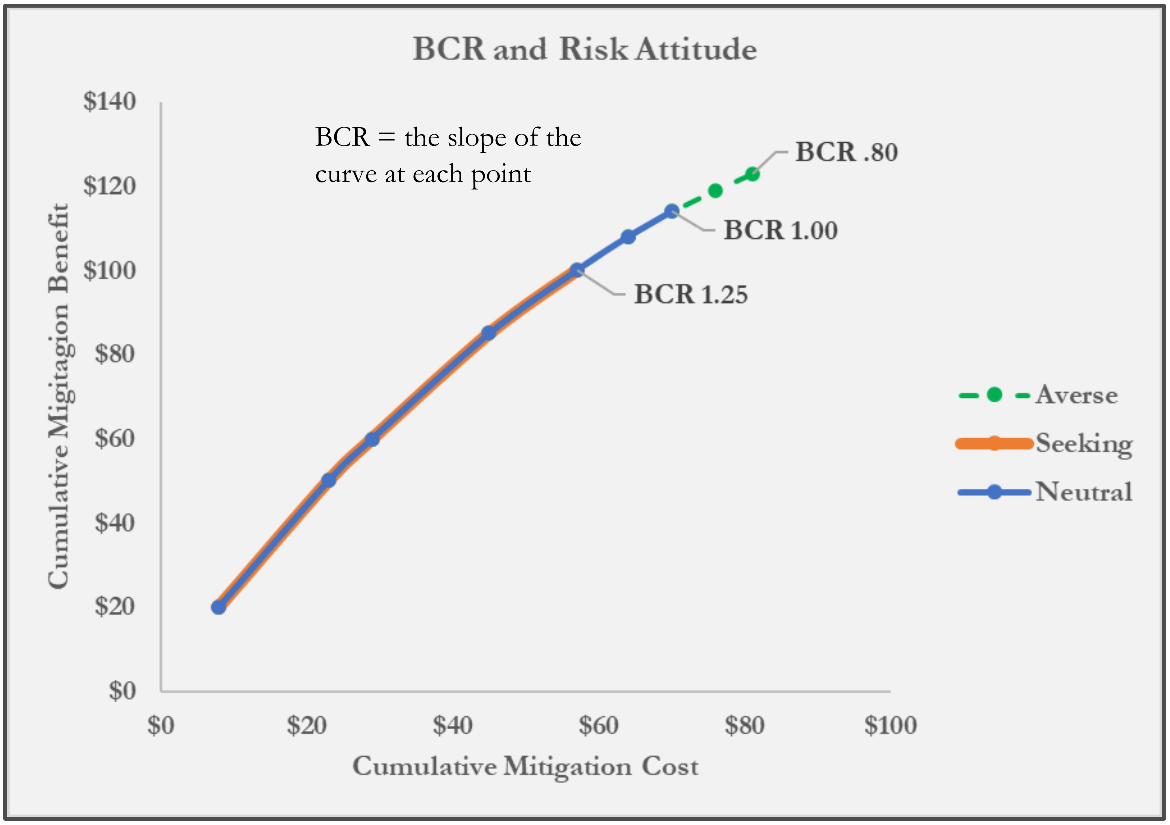

Figure 3-9 BCR and risk attitude.[47]

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 Table 3-2 illustrates our concern.

[47] 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.

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

Table 3-2 Hypothetical risk reduction and BCR calculated 3 ways.

Suppose the BCR threshold for selecting a portfolio of risk mitigations[48] is 1.0. The portfolio 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? This topic is a slippery slope that SPD argues is best to be 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 closing out the discussion on tail risk, SPD will revisit the notion of risk neutrality and its special relationship with tail risk.

3.3 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.[49] 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 does not even need to be calculated, much less evaluated. This surprising implication of risk neutrality is demonstrated in Table 3-3.

[48] The Staff Proposal will discuss evaluating portfolios of mitigations as opposed to individual mitigations in Section 4.2

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

	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	99.9999900%	$0	$1,000

Table 3-3 Risk neutrality in a long-tailed risk curve

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.

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

3.4 Risk Tolerance and the Three Venturers

Earlier, the Staff Proposal 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 ($150). The decision approach for each venturer can be visualized in the three figures below.

- *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 to the baseball game go unused.

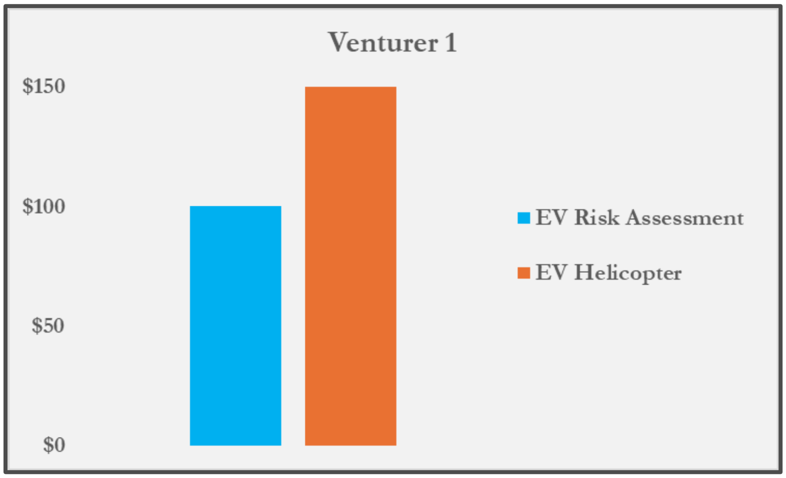

Figure 3-10 Venturer 1 decision-making process for not paying for the helicopter[50]

- *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 of fees from clients 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, which is analogous to an insurance premium, worth mitigating the risk? V2 will likely choose to accept the risk of losing fees from clients in this instance since the risk only slightly exceeds risk tolerance compared to the cost of the helicopter.

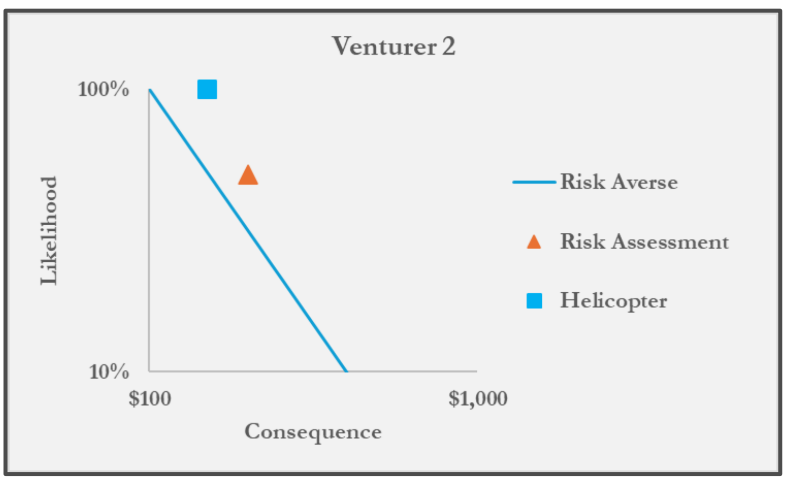

Figure 3-11 Venturer 2 decision-making process for not paying for the helicopter

---

[50] The visualization of V1's decision-making process in Figure 3-10 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 would not change since the EVs are the same.

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

- *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.

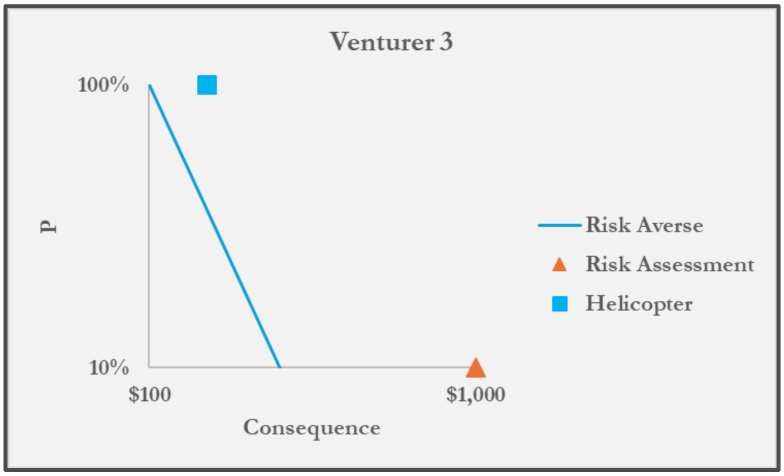

Figure 3-12 Venturer 3 decision-making process for paying for the helicopter

According to a straight interpretation of 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.

The parable demonstrates the interplay between probabilistic risk assessment, risk tolerance, and risk-based decisions. In Chapter 4, the Staff Proposal will explore how simple optimization based on this interplay can work for utility risk management.

4. Simple Optimization

In Section 3, the Staff Proposal made the case that a probabilistic view of risk and risk tolerance are inseparable for risk-based decisions. This section now turns 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.

4.1 The difference between ranking and optimizing

Ranking based on 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. Implementing the wildfire mitigations of undergrounding and covered conductor on the same electric grid asset 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.[52]
- *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%. No one would expect that by employing both mitigations they 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.

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.

4.2 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 mitigations.

[52] "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 portfolio of mitigations could include any combination of feasible (i.e., non-mutually exclusive) mitigations. Suppose a utility is considering 3 mitigations, M1, M2, and M3. There could be a total of 7 different portfolios, as laid out in Table 4-1.

Portfolio	Mitigations
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

Table 4-1 Possible portfolio combinations for three mitigations.

Seven portfolios for 3 mitigations assume none of the mitigations is mutually exclusive. If M1 and M2 are in fact mutually exclusive, such as implementing undergrounding and covered conductor on the same electric grid asset, the set of possible portfolios would be reduced, as shown in Table 4-2.

Portfolio	Mitigations
Port_1	M1
Port_2	M2
Port_3	M3
Port_4	M1, M3
Port_5	M2, M3

Table 4-2 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 chosen.

However, there is 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 do not need to be constructed, much less evaluated. Computational techniques such as linear optimization rapidly reduce 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, including mitigations across risk events within a portfolio is the correct approach, but in practice it 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.

SPD wants to make one thing clear: designing portfolios of mitigations are helpful for optimizing risk reduction as will be detailed below. When it comes to reporting the benefits and costs of mitigations, the utilities must continue to present both by program and activity in RAMP and GRC filings. This approach to presenting benefits and costs by program and activity should be maintained for cost recovery applications to the Commission where the utility requests determinations of reasonableness and incrementality.[53] The portfolio approach described in this Staff Proposal must not be used as a justification for reducing transparency of the benefits and costs of mitigations in any GRC or other cost recovery Applications to the Commission. Additionally, workpapers and data templates should still be available at the Risk Reporting Unit, as was recommended in SPD's Staff Proposal.[54]

4.3 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 overall residual risk within affordability constraints, which is different from maximizing cost-efficiency. 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 overall 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, this approach is 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

---

53 See D.24-03-008, CoL 28 and 29.

54 Definition of Scoped Work and the Risk Reporting Unit, SPD Staff Proposal, November 5 2024

55 Technically, this approach is minimizing the level of overall residual risk, which is not always the same as maximizing mitigation impact. This nuance is discussed further in Level 4 Incorporating Risk Tolerance and Simple Optimization into the RDF, Appendix F, November 1 2024, https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/safety-policy-division/meeting-documents/risk-tolerance-simple-optimization-level4-final_branded4_110124.pdf

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. After all, 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 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]

4.3.1 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 4-1. Note that the units of the X and Y axes make up the components of the BCR.

[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.

[58] For an example on water companies, see Manuel Mocholi-Arce, Ramon Sala-Garrodo, Maria Molinos-Senante, and Alxandros Maziotis, "Performance assessment of water companies: A metafrontier 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

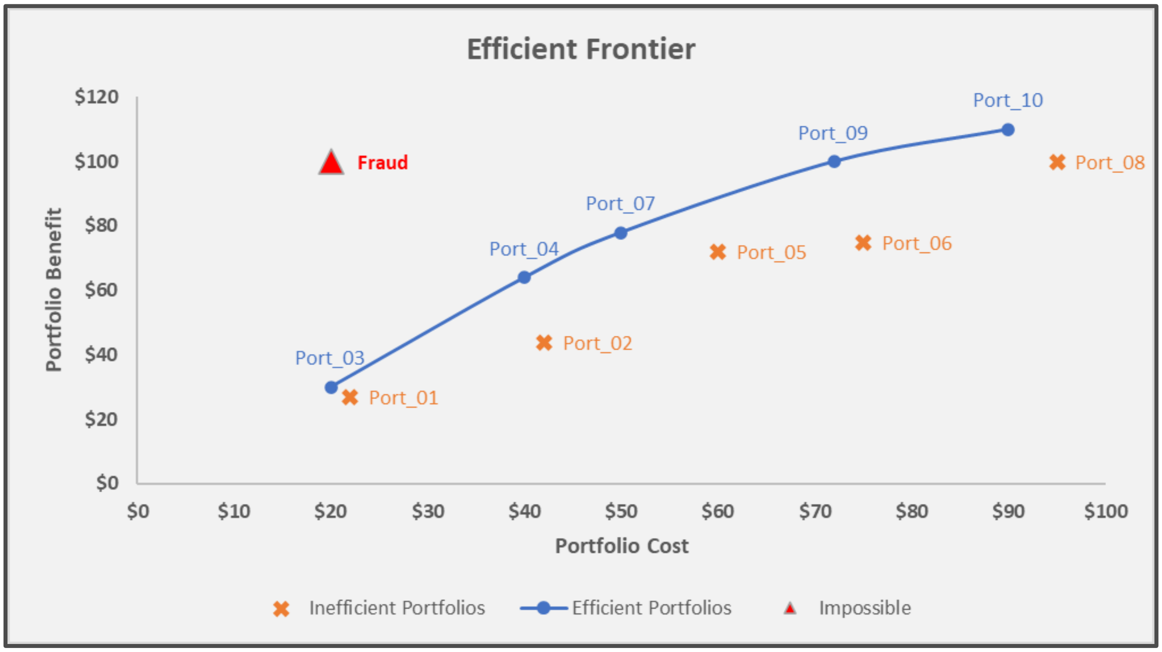

Figure 4-1 Efficient frontier[59]

In Figure 4-1, 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. 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 4-1 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.[60]

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).[61]

---

[59] Optimal portfolios along the blue line, and suboptimal portfolios below the line. The impossible portfolio (red triangle) is an example of fraud.

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

[61] 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.

A key observation is that the efficient frontier usually contains multiple portfolios—see Figure 4-2 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 the costs 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 cost and benefit (benefit=risk reduced). As one iterates through the potential combinations, one finds optimal portfolios, each with its own benefit and cost.

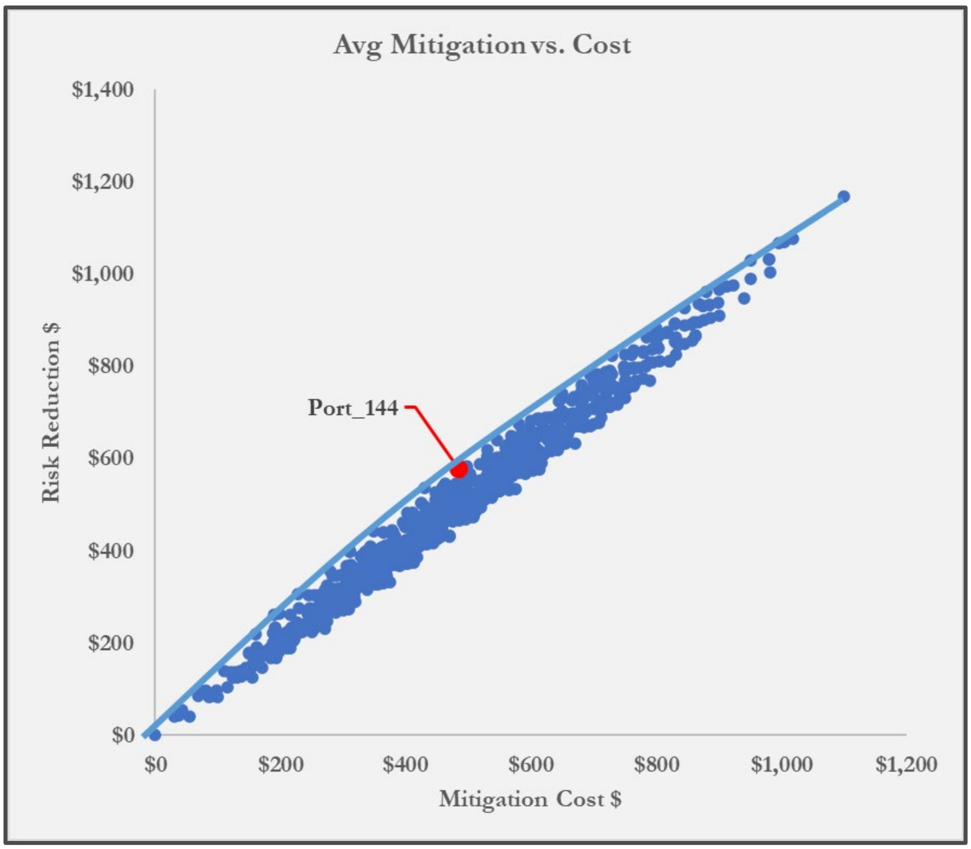

Figure 4-2 Efficient frontier calculated from over 700 portfolios.[62]

[62] 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 impacts on affordability.

4.4 Stochastic Optimization

The efficient frontier discussion in Section 4.3 did not 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 4-4 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 (Figure 4-3), but for tail average risk the optimal portfolio is Port_05 (Figure 4-4).

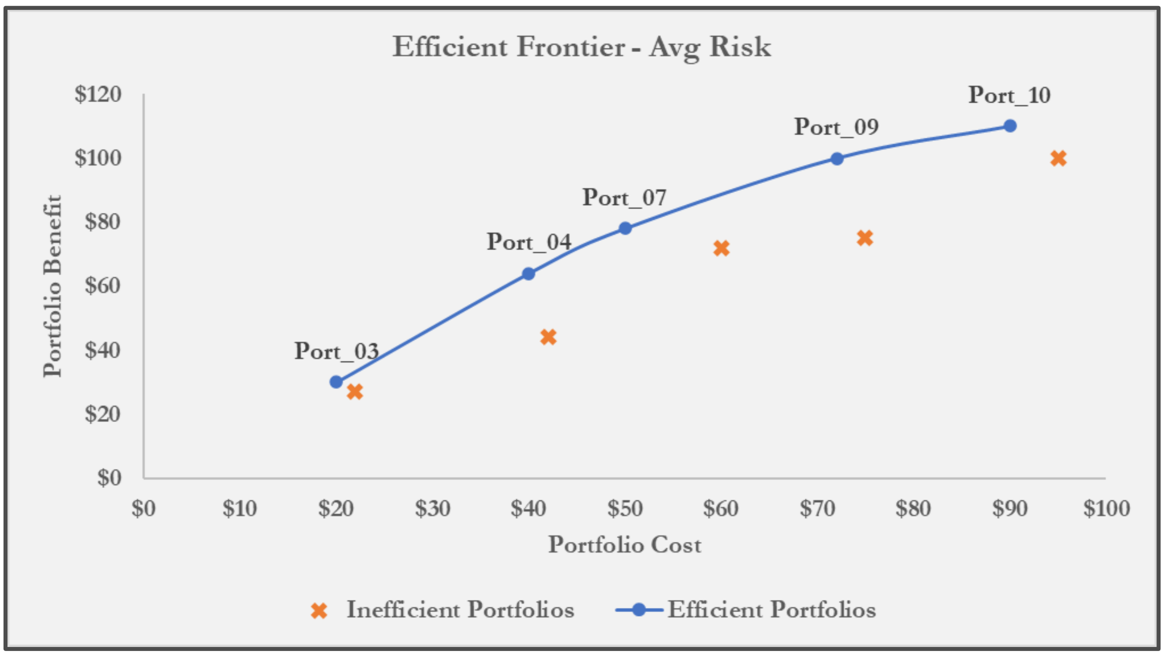

Figure 4-3 Efficient frontiers for average risk.

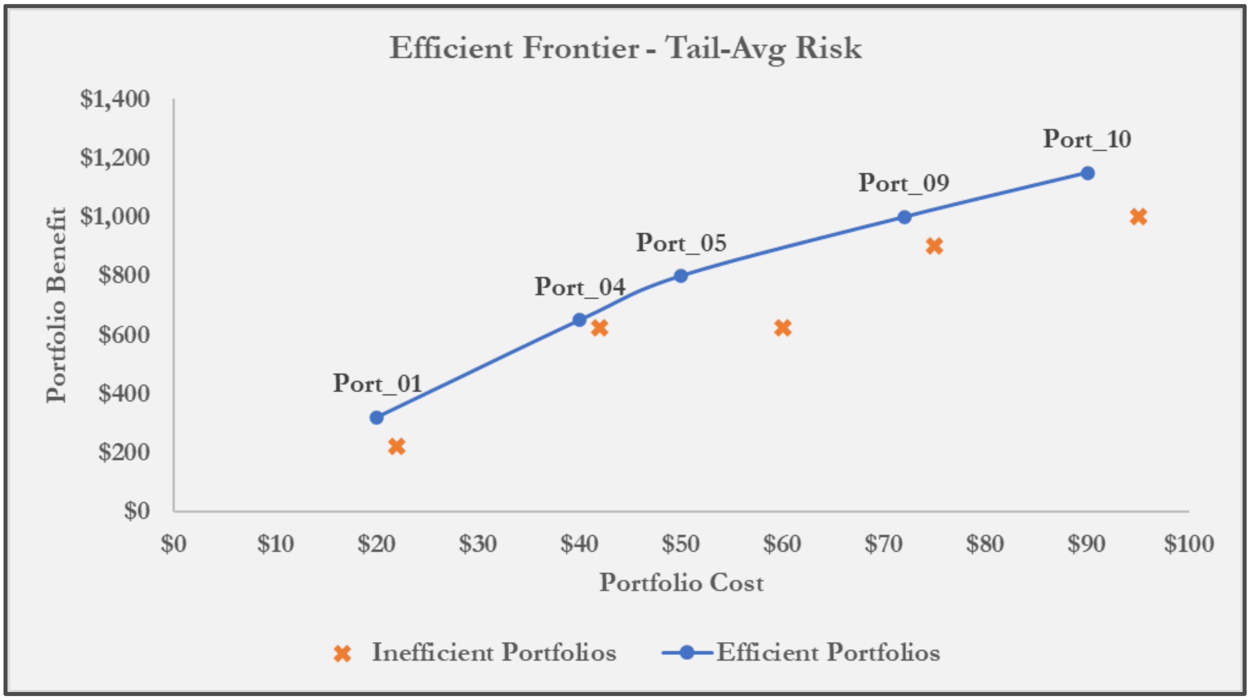

Figure 4-4 Efficient frontiers for tail average risk.

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 makeup 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."

4.5 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.[63] There are other trade-offs that must be considered, including, but not limited to, safety vs. reliability, affordability, ESJ impact, and time exposure. 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 as seen in Figure 4-5.

[63] D.24-05-064, Appendix A, Row 26.

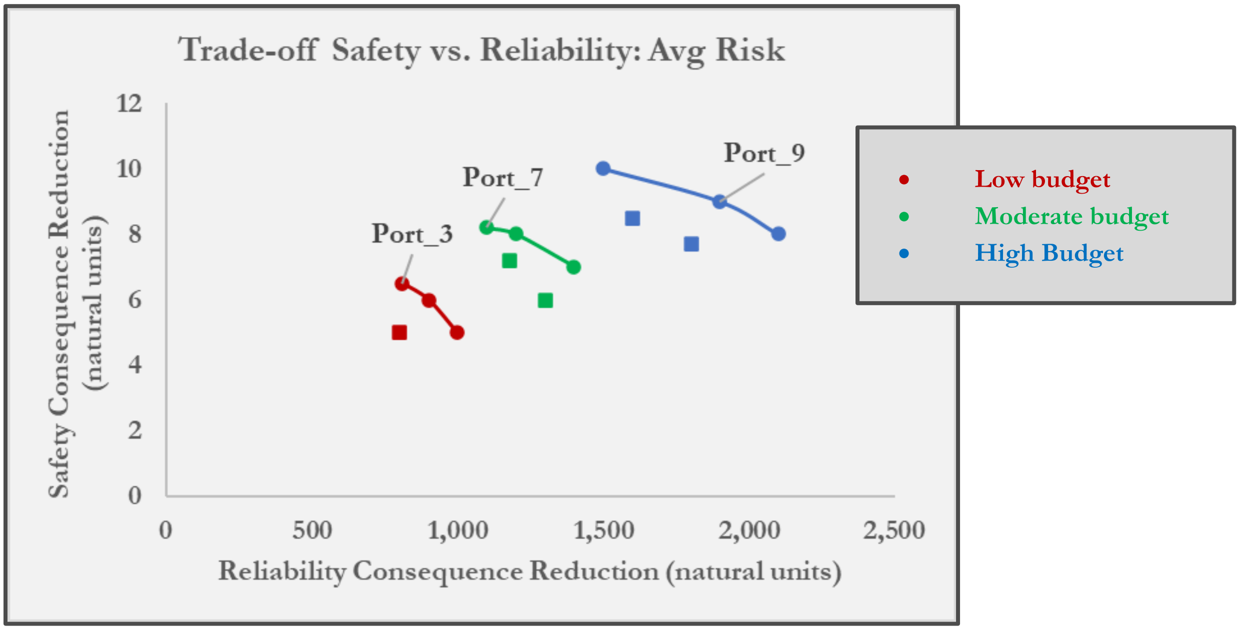

Figure 4-5 Herringbone diagram for safety vs. reliability trade-off based on average risk.[64]

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 4-6. In this example, the optimal portfolio for the moderate (green) budget is Port_5, which is different from the optimal portfolio for average risk seen in Port_7 (refer back to Figure 4-3 and Figure 4-4 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.

[64] 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.

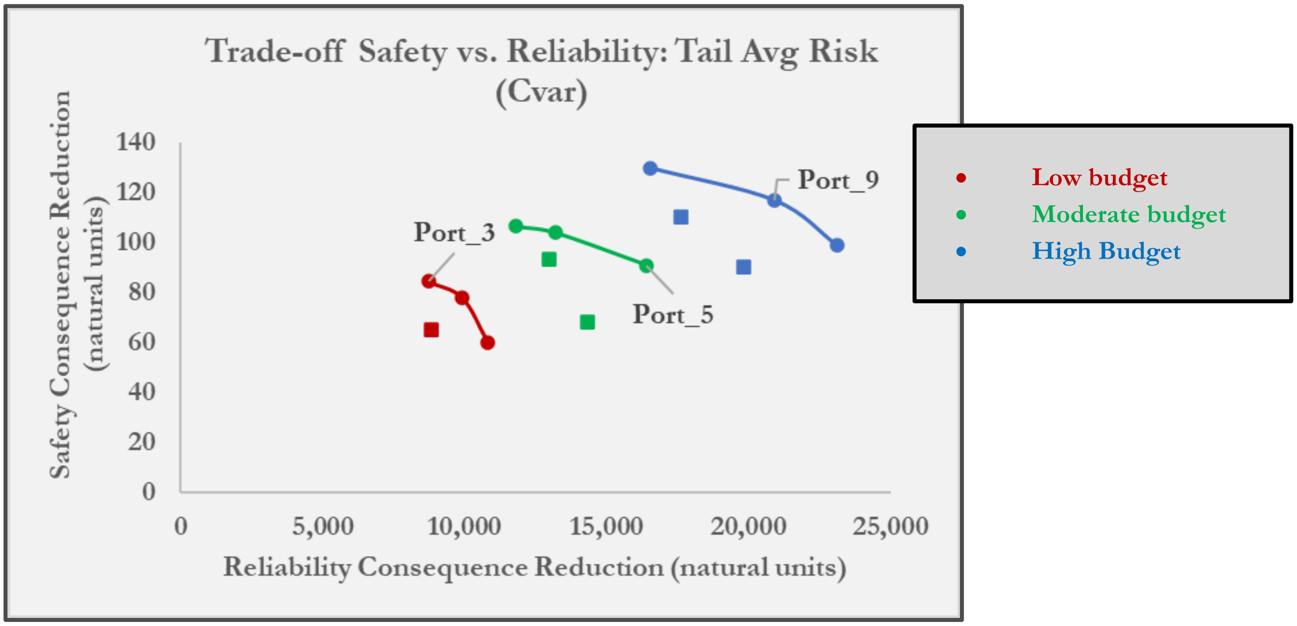

Figure 4-6 Herringbone diagram for safety vs. reliability trade-off based on tail average risk.

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, the 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 brings the utility within risk tolerance.

4.6 Optimizing Frameworks

This section will provide an overview of the CBA decision framework and its two ranking or optimization methodologies (see Figure 4-7).

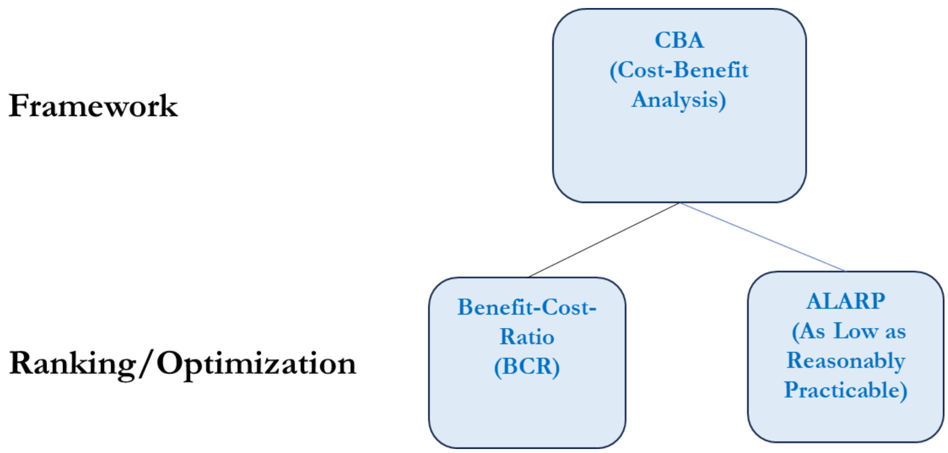

Figure 4-7 Risk decision framework.

D.22-12-027 modified the RDF to require benefits and costs of mitigations are calculated in dollars, the CBA, using a monetization factor to translate the natural units into dollars.[65] 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,[66] used as a threshold for selecting mitigations. A BCR can be calculated for each mitigation option or used in optimization.

4.6.1 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 4-8 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. 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.

[65] D.22-12-027, CoL 5

[66] As discussed in Section 3.2.3, risk aversion may lead to setting the BCR threshold below 1.0.

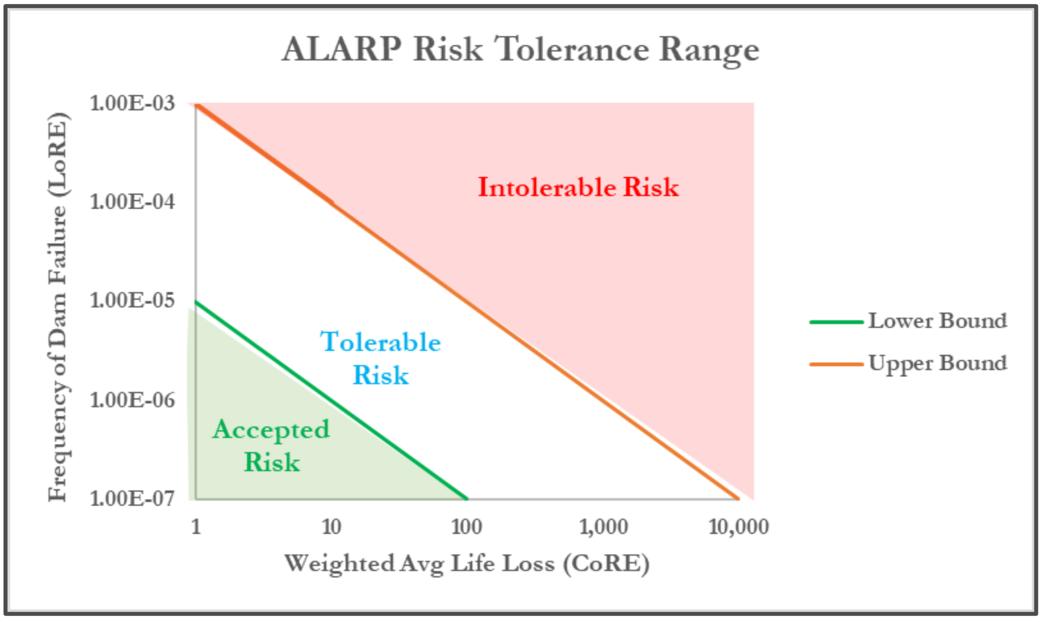

Figure 4-8 The three tiers of ALARP.

Figure 4-9 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 certain 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.

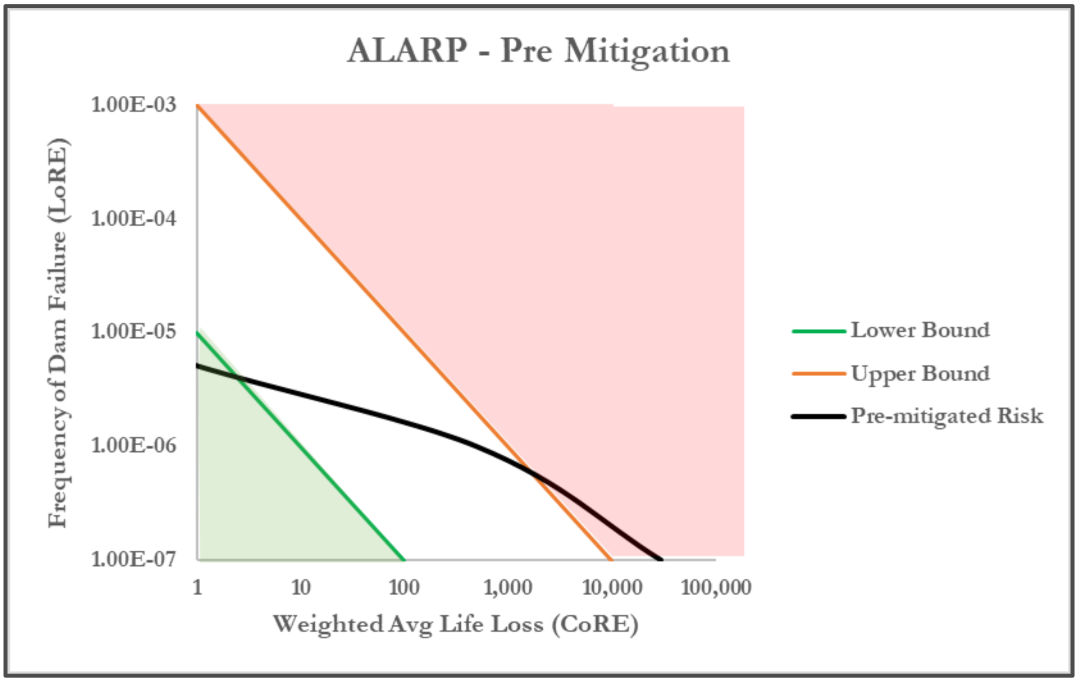

Figure 4-9 Pre-mitigated risk, ALARP Chart.

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 4-10:

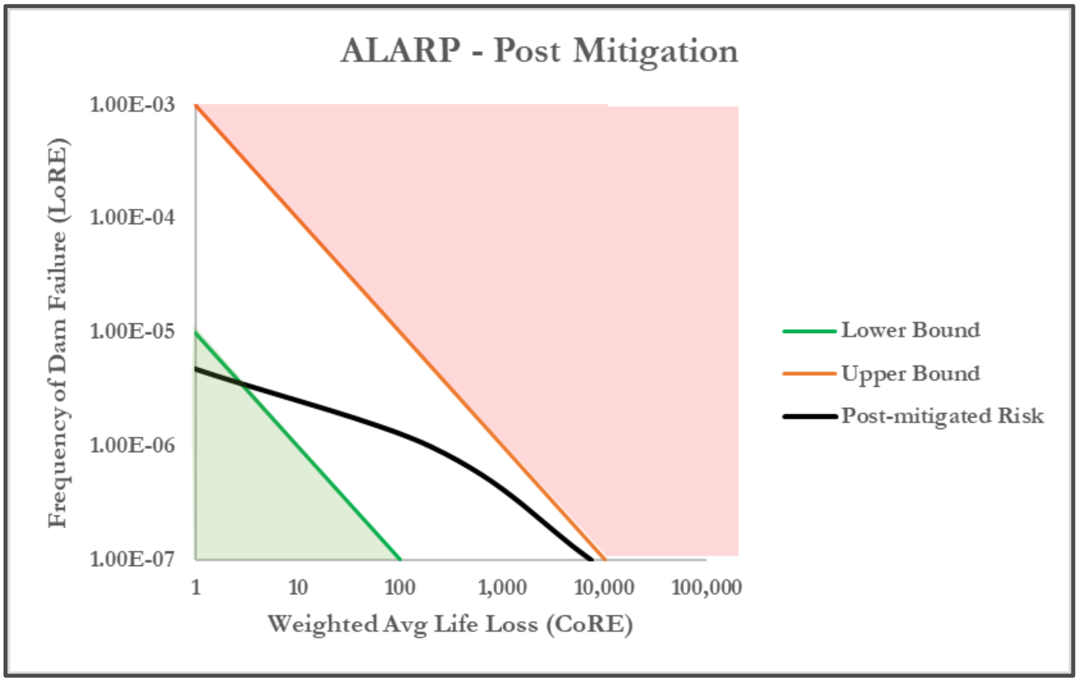

Figure 4-10 Post mitigation risk, ALARP Chart.

Figure 4-10 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. Risk in the acceptable range was left untouched.

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 of the three attributes (re: 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.[67]

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. SPD proposes one such approach in the next section.

[67] The topic of interrelationships of risks may be addressed in a future phase of this proceeding or its successor.

4.7 Simple Optimization: One Approach

The goal is to optimize overall residual risk by reducing it to an acceptable level, given affordability and other constraints. SPD has argued that overall residual 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 3.2.2, tail average risk 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 overall residual risk is not exactly the same as optimization by maximizing mitigation impact. The distinction is subtle but could lead to suboptimal mitigation selection.[69] Ultimately, we are still interested in mitigation impact, so the correct formulation is:

Optimize mitigation impact = pre-mitigated risk – optimized overall residual 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 further discuss the risk-based decision-making process of optimization.

4.8 Making Optimal Risk Reduction Decisions

Suppose a utility performed an optimization that results in the efficient frontiers for average risk mitigation and tail risk mitigation seen in Figure 4-11 below. The utility is interested in Port_161 because it sits on the efficient frontier for average mitigation and tail average mitigation (see Figure 4-11).

---

[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

[69] For a detailed example of how maximizing mitigation impact can lead to different results than minimizing residual risk, see Incorporating Risk Tolerance and Simple Optimization into the RDF, Appendix F, November 1 2024, https://www.cpuc.ca.gov/-/media/cpuc-website/divisions/safety-policy-division/meeting-documents/risk-tolerance-simple-optimization-level4-final_branded4_110124.pdf.

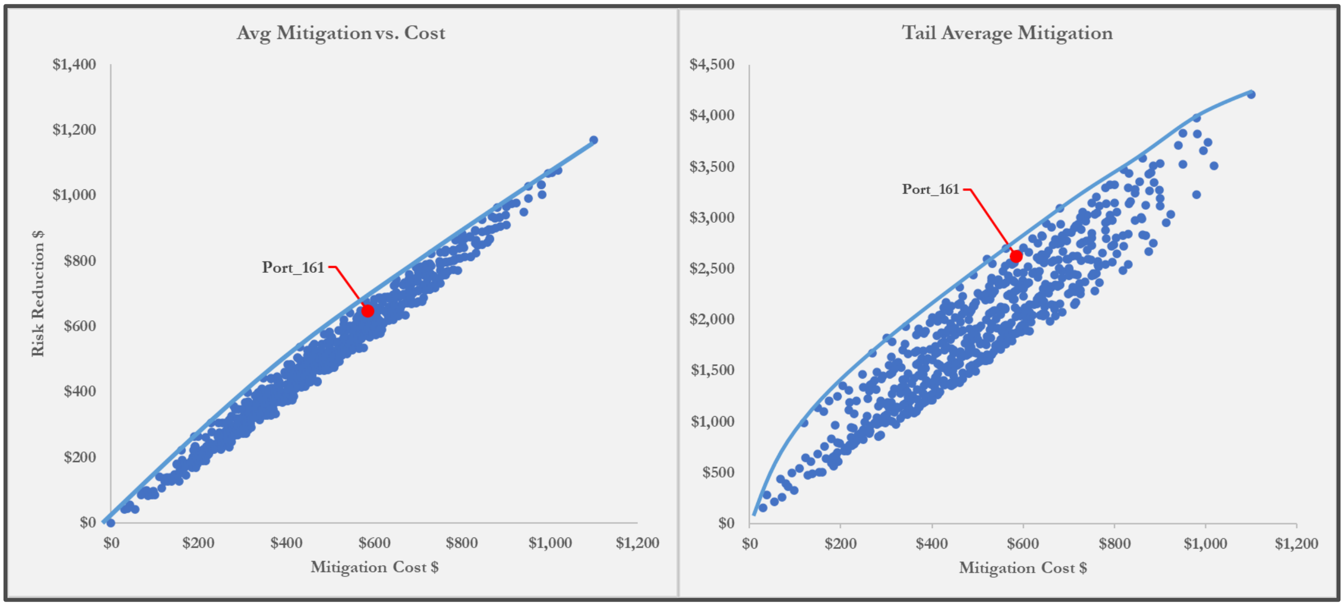

Figure 4-11 Hypothetical optimal mitigation selection example.

The evaluators of the utility's risk assessment are also interested in the safety versus reliability risk reduction of the chosen portfolio, which is represented in the herringbone diagrams in Figure 4-12.

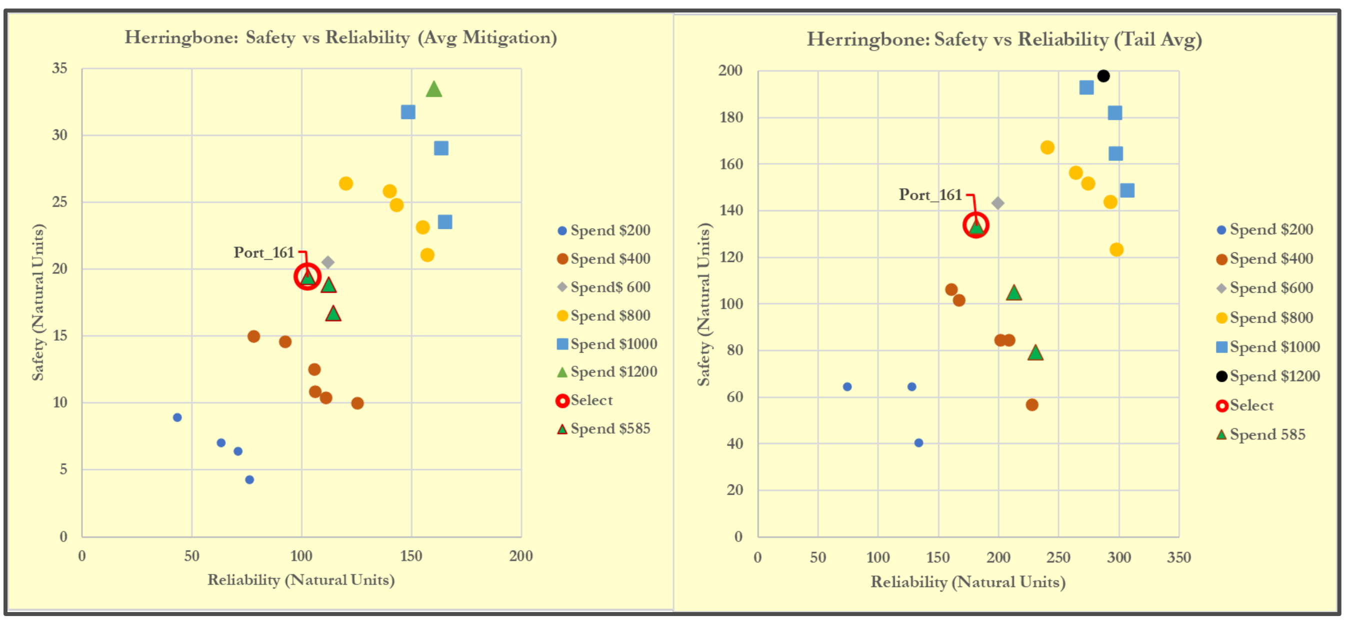

Figure 4-12 Herringbone representation of safety vs. reliability trade-off.

Port_161 has the highest safety impact of the alternative optimal portfolios at the given budget level of $585 (green triangles).

Table 4-3 calculates the key statistics from the optimizations, including applying the relevant quantification of risk tolerance. Residual 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	Overall Residual Risk	$718	$2,095
D	Risk Tolerance - Neutral	$800	
E	Risk Tolerance - Averse	$700	$1,800
	BCR of Mitigation Benefit	1.11	

Table 4-3 Simple optimization outcomes: Risk-averse tolerance unachieved

The key takeaways for whether this portfolio is acceptable or not depends on whether a utility is risk-neutral or risk-averse. This portfolio meets the threshold for a utility that chooses to be risk-neutral— the average overall residual risk of $718 (row C) is below the neutral risk tolerance level of $800 (row D).

For a utility that chooses to be risk-neutral, there is no comparison with tail average risk; by definition, risk neutrality does not distinguish between average and tail risk. For a utility that chooses to be risk averse, however, overall residual risk is above risk tolerance requiring further mitigation.[70]

The selected portfolio BCR is 1.11, which exceeds 1.0 by a healthy margin. It may warrant further conversation on whether the budget for this portfolio can be increased to reduce overall residual 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 Table 4-4. Tolerance for average risk is $800 for risk-averse and risk-neutral utilities, and average overall residual risk is within tolerance for risk-neutral and risk-averse utilities. In this example, tail average risk remains out of tolerance for a risk averse utility.

[70] The comparison for the average is $718 (row C) versus $700 (row E) and for tail risk $2,095 (row C) vs. $1,800 (row E)

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

Table 4-4 Simple optimization outcomes: risk-averse tolerance achieved for average risk

5. Modifying the RDF, and Ensuring a Manageable Transition

5.1 Risk Tolerance: Gaining Consensus

In this proposal SPD recognizes the importance of risk tolerance, but two key questions remain:

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

SPD argues 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. For this reason, SPD recommends that a forum of key stakeholders be established whose consensus on risk tolerance would represent the residents of California. This will be called the California Utility Risk Tolerance Stakeholder (CURTS) Working Group. The CURTS Working Group should be engaged by each utility during the preparation for filing its RAMP Report. SPD recommends that the CURTS Working Group include parties to the RDF Proceeding as well as the parties to the IOU's most recent GRC and RAMP Proceedings. Finally, due to the importance of the wildfire risk in RAMP and GRC Applications, SPD recommends that representatives from the Office of Energy Infrastructure Safety also join the CURTS Working Group.

How consensus can be achieved should be the first decision made within CURTS Working Group. SPD recommends that the Commission allow the Working Group to determine its own order of operations. However, SPD does recommend that the Working Group consider well established approaches to building consensus, such as rank choice voting and/or the Delphi Method.[71]

In the following sections, SPD will propose three perspectives that the utilities, in consultation with the CURTS Working Group, should consider as they begin implementing the Risk Tolerance requirements set out in this proposal.

[71] The Delphi method for instance has been used in solving environmental decision-making, see Armour, Carl L., and Samuel C. Williamson. *Guidance for modeling causes and effects in environmental problem solving.* US Department of the Interior, Fish and Wildlife Service, Research and Development, 1988. For a brief overview and critique of the Delphi Method, see Pill, Juri. "The Delphi method: substance, context, a critique and an annotated bibliography." *Socio-economic planning sciences* 5, no. 1 (1971): 57-71. The Working Group could also consider using a Decision Theater environment to support building consensus. For an example see White, Dave D., Amber Y. Wutich, Kelli L. Larson, and Tim Lant. "Water management decision makers' evaluations of uncertainty in a decision support system: the case of WaterSim in the Decision Theater." *Journal of Environmental Planning and Management* 58, no. 4 (2015): 616-630.

5.1.1 Risk Tolerance Considerations

Several issues must be considered even before the challenging work of quantifying tolerance begins.

- Should tolerance be set at the overall residual risk level in dollars, as was done in the example in Section 4.8?
- Alternatively, should tolerance be set at the attribute level, in natural units? This would mean setting individual risk tolerances for safety, reliability, and financial consequences.
- 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 given year total risk was within the utility's tolerance, but wildfire risk accounted for 99% of the total risk that year?

SPD recommends that to start, risk tolerance should be set both in aggregated dollars for overall residual risk and at the consequence attribute level for each utility.

5.1.2 Approaches for Setting Risk Tolerance

Optimizing risk mitigations based on risk tolerance requires risk tolerance to be set at some level for the utility. That level can be integrated into the utility's risk assessment, for every combination of risk events and attributes, or somewhere in between. Figure 5-1 below shows the range of options.

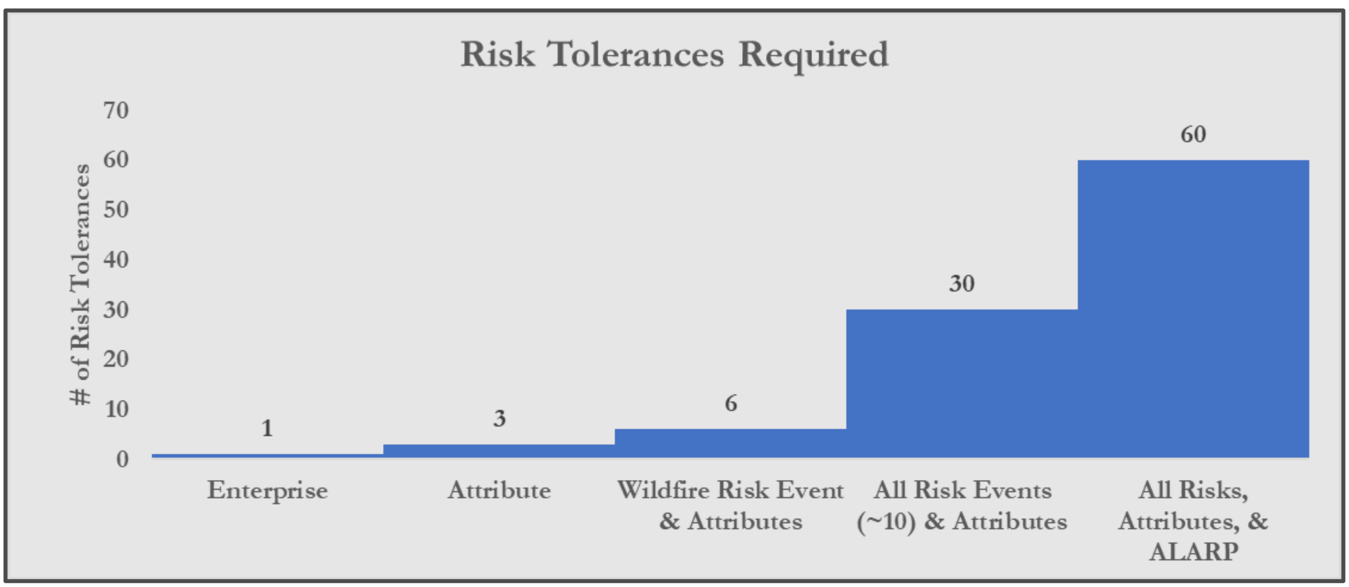

Figure 5-1 Range of risk tolerances required.

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 the utility includes ten risk events, three attributes and two ALARP tolerances. While requiring a lower number of tolerances would be desirable—at least to begin with—this does not 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 worth $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, SPD recommends the Commission require the utilities to start by setting risk tolerance at the attribute level for wildfire risk and a bucket for all other risk events, which would require 6 tolerances (3 attributes x 2 risk event categories). Each additional risk event added would increase the number of tolerances by 3. For example, setting tolerances for 3 attributes on the wildfire, loss of containment on gas transmission pipeline, LGUWR, and all other risk events, would require a total of 12 tolerances (24 for ALARP).

As the utilities, evaluators and stakeholders gain experience working with risk tolerance and optimization, more tolerances should be added.

5.1.3 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. Public workshops and technical sessions held in conjunction with the CURTS Working Group might include:

- Training on probabilistic risk assessments, LoRE and CoRE, how to understand a probability distribution for overall residual risk, and the difference between average and tail risk.
- Deciding on which tolerances are needed, as discussed in Section 5.1.2.
- Determining how to set risk tolerance levels for average risk and tail average risk. 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 needed 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 of codification could be a lengthy one. Thus, 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 5.1.2, 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 more transparent.

5.2 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 SPD confident that the fundamentals are in place. In a future Phase of the RDF Proceeding or its successor, SPD recommends a series of workshops on the technical requirements of simple optimization to ensure consistency and proper methodology across the utilities.

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. Table 5.1 is an example from PG&E's 2024 RAMP, which presents Monte Carlo trials for safety, reliability, and financial outcomes.

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

Trial	Safety				Reliability				Financial			
	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												

Table 5-1 From PG&E 2024 RAMP Monte Carlo Simulation.[72]

[72] A.24-05-008, PG&E 2024 RAMP (PG&E-2) at 2-50.

Table 5-1 shows data for each trial in a Monte Carlo simulation, which is a useful technique for building risk models that uses probability distributions to run the simulation. A future workshop could cover in detail the proper use of these stored probability distributions, including how to not collapse them into averages for input into other calculations as was done in Table 5-1.

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

2. *Data storage*. 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 for wildfire risk.[73] Advanced tools for efficient storage of simulation data such as metalogs[74] and sparse Monte Carlo[75] may be explored.

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.

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 it is done following the arithmetic of uncertainty and the proper "order of operations."

Topics 1-4 can be covered in greater detail in a workshop on Interrelationships of Risks in a future phase of this proceeding or its successor.

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.

---

[73] SCE 2023-2025 Wildfire Mitigation Plan at 131.

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

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

6. Simple optimization techniques. The use of average risk and tail average makes linear optimization possible, which is consistent with "simple" optimization and improves computational efficiency. Nonetheless, it will be necessary 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.

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 Workshop #3 on RMAR.

These last three topics have been discussed at length in this proposal. SPD believes that with the right adjustments to the RDF, it should be possible for the utilities to place these into practice. The Staff Proposal turns now to the specific changes that would need to be made to the RDF.

6. SPD 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): *Require reporting of overall residual risk*. Decision-makers need to know how much overall residual risk remains for every enterprise risk faced by the utility to determine if the risk tolerance of Californians has been met. This will also help the Commission determine whether the utility has properly designed its mitigation programs in a RAMP and GRC filing to quickly and cost-efficiently address overall residual risk. For this reason, SPD recommends that the Commission require the utilities to report on their progress of reducing overall residual risk for each enterprise risk addressed in a RAMP or GRC filing.

SPD recommends adding the following definition:

Overall Residual Risk: all the risk on the utility's assets or systems after taking account of the historical progress of risk reduction for every GRC cycle to date.

Residual Risk: Risk remaining after application of Mitigations, including Mitigations classified as Controls for a given GRC cycle.

Additionally, SPD recommends the following update to Row 9 of the RDF:

9.	Risk Assessment	Using the Cost-Benefit Approach developed in accordance with Step 1A, for each Risk included in the Enterprise Risk Register, the utility will compute a monetized Safety Risk Value using only the Safety Attribute. The utility will sort its ERR Risks in descending order by the monetized Safety Risk Value. For the top 40% of ERR risks with a Safety Risk Value greater than zero dollars, the utility will compute a monetized Risk Value using at least the Safety, Reliability and Financial Attributes to determine the output for Step 2A. The output of Step 2A, along with the input from stakeholders described in Row 12 below, will be used to decide which risks will be addressed in the RAMP. The output of Step 2A must include a calculation of Overall Residual Risk, along with a diagram and supporting workpapers demonstrating the change of Overall Residual Risk since the utility's first RAMP filing. The Risk Assessment in preparation for RAMP will follow the steps in Rows 10 and 11.

Recommendation 2 (R2): *Require 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[76]).
- 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 SPD recommends 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 represented as a probability distribution.

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

---

[76] 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.

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 probability distribution representing Likelihood of a Risk Event (LoRE) by the probability distribution of Consequences of a Risk Event (CoRE) and be 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 3 (R3): *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.

SPD 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 R3, SPD 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 can also be described using the tail average above a specified percentile of the distribution if the utility so desires.

		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 4 (R4): *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 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*[77] 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. R4 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 R4, SPD recommends adding the following definitions to the RDF:

Constant Risk Exceedance Curve: the curve that results in the same Expected Value of Overall Residual 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.

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.

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

Risk Tolerance: Maximum amount of Overall Residual Risk that an entity or its stakeholders are willing to accept after application of risk Control or Mitigation. Risk tolerance can be influenced by legal or regulatory requirements.

Based on R4, SPD recommends 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 establishing 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. Additionally, if the utility chooses to present tail average risk as stipulated in Row 5, then the utility shall also present the tail average risk value for the probability above the specified percentile on 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) Working Group. Within 21 days of the CURTS Working Group 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 5 (R5): *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.

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, require each utility to declare and justify a risk tolerance, and evaluate risk reduction based on this risk tolerance.
5. *Phased approach*. SPD recommends the Commission require the utilities to start by setting risk tolerance at the attribute level for wildfire risk and a bucket for all other risk events, which would require 6 tolerances (3 attributes x 2 risk event categories).
6. *Long-term vision*. With some experience working with risk tolerance and simple optimization after each utility has completed one GRC cycle, the Commission could consider moving ahead with more sophisticated frameworks such as ALARP.

Based on R5, SPD recommends that in 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.

Recommendation 6 (R6): *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 R7 below).

SPD recommends adding the following definitions related to R6:

Mitigation Portfolio: a collection of one or more risk mitigations with a specified budget constraint for reducing the risk of a given enterprise risk. 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 R6, SPD 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 with a specified budget constraint. 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 (re: 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, budget constraint, 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 budget 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 transparent and 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 7 (R7): *Portfolio selection could be 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.[78]
- *Two scenarios*: Two efficient frontiers can be created, one for each of two stochastic optimization scenarios:
  - *Scenario 1*. Minimize average overall residual risk for various mitigation cost levels.
  - *Scenario 2*. Minimize tail average overall residual risk for various mitigation cost levels.

SPD recognizes that it may be necessary for the utilities 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. SPD also recognizes that various approaches to optimizing portfolios of mitigations for each risk event could be available as long as the goal is minimizing overall residual risk towards Californian's risk tolerance according to various affordability constraints.

[78] Alternative Analysis is defined in D.24-05-064, Appendix A Page A-3 and the requirements of Alternative Analysis can be found in D.18-12-014 at 34.