PUBLIC UTILITIES COMMISSION
STATE OF CALIFORNIA

FILED
1-29-16
11:35 AM

Attachment 1

"Intervenor Perspective Regarding an Improved Methodology to Promote Safety and Reliability of Electric and Natural Gas Service in California," prepared for the S-MAP Workshop January 25, 2016, by Charles D. Feinstein, Ph.D. and Jonathan A. Lesser Ph.D. on behalf of The Utility Reform Network/Indicated Shippers/Energy Producers and Users Coalition (Revised January 28, 2016)

INTERVENOR PERSPECTIVE REGARDING AN IMPROVED METHODOLOGY TO PROMOTE SAFETY AND RELIABILITY OF ELECTRIC AND NATURAL GAS SERVICE IN CALIFORNIA

Prepared for the S-MAP Workshop

January 25, 2016

Revised January 28, 2016

By:

Charles D. Feinstein, Ph.D.

and

Jonathan A. Lesser, Ph.D.

On behalf of

The Utility Reform Network/Indicated Shippers/Energy Producers and Users Coalition

TABLE OF CONTENTS

APPENDIX

I. INTRODUCTION: OVERVIEW OF PROPOSED METHODLOGY

The Utility Reform Network, Indicated Shippers, and Energy Producers and Users Coalition ("Joint Intervenors") appreciate this opportunity to provide an alternative perspective regarding an improved approach to managing risk and ensuring that electric and natural gas utilities in California provide safe, reliable, and affordable services to their customers. Like the utilities' current models, the approach proposed in this paper defines risk as the product of the likelihood of failure (LoF) and the consequences of failure (CoF). However, with respect to LoF, the proposed approach uses mathematical probabilities of failure events determined by relying on subject matter experts and other data regarding the condition of, and likelihood of threats to, the utility's system, and eliminates the utilities' current extra step of converting their frequencies to an artificial, nonlinear scale of values between 1 and 7. With respect to CoF, the approach relies on multi-attribute scaling of event consequences in a way that prioritizes safety and accounts for any other consequences the utilities and the CPUC may wish to include. Importantly, this improved approach enables the measurement of risk reduction from any proposed mitigation. This allows the utilities to select the optimal combination of risk mitigation actions given the constraints under which the utilities operate.

II. GOALS OF A SUCCESSFUL RISK MANAGEMENT METHODOLOGY

A successful risk management methodology will meet a number of key goals.

A. Ensures Public and Employee Safety are the Priority

Safety of the public and utility employees must be the top priority in assembling a risk reduction portfolio of actions, regardless of the processes and methodologies adopted by the utilities to determine how best to reduce risk. The goal of a risk management methodology should be to ensure electric and natural gas services are as safe as possible, given the funding and resource constraints the utilities face, while ensuring reliable and affordable service that satisfies customer needs and complies with regulatory and environmental requirements.

B. Promotes Cost-Effective and Optimized Risk Management

The methodology must be able to identify and implement the most *cost-effective* risk-reduction measures to improve safety and reliability, similar to how utilities secure needed electricity supplies for their customers at the lowest possible cost. <u>However, the cost-effectiveness of alternative risk management actions and programs can be determined only if risk reduction is measured</u>.

The methodology should enable a utility to propose an *optimal* portfolio of risk management alternatives subject to the constraints utilities face (e.g., budget, available labor and equipment, allowable number of service interruptions for repairs, and so forth). In other words, the methodology should determine a portfolio of decisions, programs, and actions that a utility can implement that will achieve the greatest risk reduction for a given set of resource expenditure levels or, alternatively, achieves a desired improvement in operating risk (measured by risk reduction) at the lowest expected cost in resources required to achieve the risk reduction target.

C. Is Transparent, Easy-to-Use, and Understandable

The processes and methodologies adopted by the utilities should be *understandable* to <u>all</u> parties, and it should be clear how each utility's risk management process works.

The processes and methodologies used by the utilities to design their risk management programs should be fully *transparent* to all parties, including regulators and intervenors. True transparency means that parties not only can understand each step in a utility risk management process, but also can replicate each step of the process, up to and including a utility's final selection of risk management actions and programs.

D. Allows for Common Application and Uniformity (Can Be Used by All Utilities)

The methodology should be able to be used by any utility, while still accounting for the utilities' individual characteristics. A common methodology ensures there is consistency in how failure events are measured and evaluated. It eases implementation, reduces costs of implementation, and enables easier comparisons of different utility risk management programs.

III. INTEVENORS' RECOMMENDED METHODOLOGY

In the next sections, we describe an alternative methodology and how it meets all of the four key goals. This methodology, and the software to implement it, was developed over a 10-year period, 1998 – 2008, with the support of the Electric Power Research Institute (EPRI), and EPRI-member utilities, including PG&E, Southern California Edison, and San Diego Gas and Electric. The methodology has been successfully deployed by many EPRI electric utilities, as well as non-EPRI utilities, and other entities, including PJM. The EPRI software is available to all utilities, but it is also straightforward for the utilities to develop the necessary software on their own. The appendix provides a list of references and additional discussion of the methodology, as well as a list of known utilities that have used the software.

A. Summary of How the Proposed Methodology Estimates Risk Reduction and Enables Utilities to Evaluate the Cost-Effectiveness of Risk Management Programs

A successful risk management methodology will measure risk reduction to promote cost-effective and optimized risk management. <u>If a utility cannot measure the amounts by which its contemplated risk reduction efforts would reduce risk, then that utility cannot determine whether that risk reduction strategy is cost-effective</u>.[1]

Like many utilities' approaches and ASME B 31.8s, the proposed methodology defines risk as the product of the likelihood of a failure event (LoF) and the consequences of that failure event (CoF). A risk mitigation action will change the values of LoF, CoF, or both. (An action

---

[1] The Utilities' Uniformity Report in this case states that "none of the implemented funding methods is currently capable of generating a risk reduction per dollar invested" although they further indicated that "with further discussions" a "potential approach" could allow them to "evaluate various mitigation strategies at the risk level to demonstrate how a collection of controls and mitigations contribute to reduction of a specific risk." They added that "cross-risk prioritization of mitigations is more challenging and will require more efforts to establish a methodology that can optimize spend based on risk reduction benefits across an entire organization." ("Combined Utilities S-Map Uniformity Report", December 1, 2015 ("Uniformity Report"), p. 14) Joint Intervenors read these statements as indicating that the utilities believe that: (i) they are not now able to calculate risk reduction; (ii) they may be able to make progress toward that goal at some future unspecified time; and (iii) the ability to compare and quantify risk reductions across an entire organization is a longer-term challenge.

that does not change either is not a risk mitigation action.) The reason stems from what we mean by a risky situation.

Specifically, a risky situation is one in which it is possible for an adverse event to occur and the *consequences* of that event would be sufficiently harmful that the utility would be willing to pay to avoid those consequences.[2] For example, it is possible that a cyber-attack on a utility customer database could succeed with the consequence that customer records are stolen. The *event* is the successful cyber-attack. The harmful *consequence* that the utility would be willing to pay to avoid is the theft of the customer records. Because there is some chance that a cyber-attack can both occur and succeed, and because some customer records can be stolen, the utility experiences some level of risk.

Risk *mitigation* is an action that could be taken to reduce risk. For example, the risk of a pipeline explosion could be mitigated by: (i) replacing existing pipe with new pipe that is less likely to suffer a rupture in the event of an earthquake; (ii) moving the existing pipe to a new location far away from populated areas; or (iii) doing both things. The first action reduces LoF, the second reduces CoF, and the third reduces both. (Of course, in reality, a single risk mitigation action could reduce the risks of multiple failure events or system-wide risk, such as a cyber-attack.)

In order to measure the amount of risk reduction associated with a risk management strategy, whether that strategy applies to a single asset or an entire utility system, both "before" and "after" LoF and CoF values must be estimated. The risk reduction is the difference between the (LoF *x* CoF) values before and after the mitigation action:

$$\text{Risk Reduction} = (\text{LoF x CoF})_{\text{BEFORE}} - (\text{LoF x CoF})_{\text{AFTER}}.$$

Suppose that a utility subject matter expert (SME) has estimated the LoF for a specific risk to be 50% this year. The SME has also estimated the consequences of a failure event (CoF)

---

[2] In general, it is important to note that if there is no uncertainty with respect to the occurrence of the event, then there is no risk; instead, there is a good or bad consequence associated with the (certain) occurrence of the event. Also, if, for all possible events, all consequences are beneficial, there is similarly no risk because the utility would not be willing to pay to avoid a beneficial outcome. Both uncertainty and at least one bad outcome must be present for there to be any risk.

to have a value of 100 units.[3] The risk associated with this particular failure event is therefore LoF *x* CoF = 50% x 100 = 50 units.[4]

Suppose a risk mitigation action (Action A) will reduce LoF from 50% to 10%. Then the post-mitigation risk will be 10% *x* 100 = 10 risk units. Therefore, this risk mitigation action will reduce risk by: $(\text{LoF x CoF})_{\text{BEFORE}} - (\text{LoF x CoF})_{\text{AFTER}} = (0.5\ x\ 100) - (0.1\ x\ 100) = 50 - 10 = 40$ risk units.

Alternatively, suppose a different risk mitigation action (Action B) will reduce CoF from 100 to 20, but will *not* affect LoF. Thus, this other risk mitigation action leaves LoF unchanged at 50%. A similar computation indicates that the risk reduction of this risk mitigation alternative is: $(\text{LoF x CoF})_{\text{BEFORE}} - (\text{LoF x CoF})_{\text{AFTER}} = (0.5\ x\ 100) - (0.5\ x\ 20) = 50 - 10 = 40$ risk units.

So, we have two ways to reduce the risk from the failure event from 50 risk units to 10 risk units. Because both risk mitigation actions reduce risk by the same amount, we should be indifferent between them if the resources required for implementation are the same.

Finally, suppose we undertook both Actions A and B, thus reducing LoF to 10% and CoF to 20 units. In that case, Risk Reduction = $(\text{LoF x CoF})_{\text{BEFORE}} - (\text{LoF x CoF})_{\text{AFTER}} = (0.5\ x\ 100) - (0.1 \text{ x } 20) = 50 - 2 = 48$ risk units.

Taking these examples a step farther illustrates that a well-designed methodology provides a clear, transparent way to evaluate the cost-effectiveness of the alternative risk management measures.

Suppose Action A (reducing LoF to 10%) costs $1,000. Suppose Action B (reducing CoF to 20 units) costs $1,500. Then Action A has a risk reduction cost-effectiveness of 40/$1,000 = 0.04 risk units per dollar. Similarly, Action B has a risk reduction cost-effectiveness of 40/$1,500 = 0.033 risk units per dollar. Finally, completing both Actions A and B, and assuming there are no economies of scale, has a risk reduction cost-effectiveness of 48/($1,000 + $1,500) = 0.019 risk units per dollar. These are summarized in Table 1 below.

---

[3] Currently, the utilities also use a 1 to 7 scale to measure the consequences of failure events. Although this scale can be used, in our projects we have found that individuals have a far easier time understanding consequences that are ranked on a scale of 0 to 100.

[4] The specific units don't matter because risk is typically dimensionless, that is, it is simply expressed as a number. Economists often use a fictitious unit called "utils," because they express consequences in terms of "utility." In this paper, we often use the term "risk units."

Table 1: Example Analysis of Risk Reduction and Cost-Effectiveness

Risk Action	Description	Risk Reduction Amount (Risk units)	Cost ($)	Risk Reduction (Risk units per dollar)
A	Reduce LoF from 50% to 10%	40	$1,000	0.040
B	Reduce CoF from 100 to 20	40	$1,500	0.033
A and B	Reduce LoF and reduce CoF	48	$2,500	0.019

If we had only $2,000 to spend, and the only important cost-effectiveness measure was risk reduction per dollar spent, then the most efficient risk mitigation strategy in this example would be Action A.[5]

It is important to note from this example that we could measure the cost-effectiveness of the different mitigation strategies only because we measured risk reduction in a way that allows meaningful comparisons of risk reductions from different potential mitigations.

B. The Proposed Methodology is a Probabilistic Approach Without the Need for Massive Amounts of Data

The utilities have acknowledged probabilistic modeling as a long-term goal, but have suggested that such models will require large amounts of data, collected over many years, in order to be developed.[6] In fact, a robust probabilistic model can be developed in the short term relying on subject matter expert (SME) judgment as a substitute for an initial lack of data. Of course, having additional data available can help the SME provide better judgment, but an initial lack of data is not a reason to delay implementation of probabilistic modeling techniques.[7]

The proposed methodology is a probabilistic methodology. It uses mathematical probabilities for LoF values. We specify LoF as a mathematical probability having a value between zero and one or, equivalently, between 0% and 100%. The fundamental reason for

---

[5] We note that this kind of ranking can be misleading when more than one resource is constrained, as is usually the case with the utilities. Nevertheless, the efficiencies with respect to each resource can be separately computed.

[6] See PG&E/Markland at 2-14:5-11 (May 1, 2015); San Diego Gas and Electric and Southern California Gas Company Presentation, "S-MAP Workshop 1" at Slide 3 (August 3, 2015).

[7] Even when data are available, SMEs are integral for determining the validity and applicability of those data to the specific risk analysis.

doing this is that it lets us compute the risk reduction associated with any risk mitigation action consistently and as accurately as possible. Doing so is a fundamental requirement if we are to select an optimal set of actions that provide the maximum risk reduction subject to all of the applicable constraints. Further, the methodology recognizes that risk (LoF x CoF) is an expected value, a number computed from a probability distribution.[8]

It is also important to note that our probabilistic methodology does not explicitly rely on value of life estimates to measure safety consequences. As we describe below, our methodology uses what is called a "multi-attribute utility" approach to estimate CoF values. Although the attributes are expressed in natural units (e.g., an electric system reliability attribute could be expressed in lost service measured in MW-minutes, financial attributes can be measured in terms of their dollar impact, and safety attributes can be measured in terms of lives lost or number of injuries, etc.), all of these attributes are weighted and scaled so that CoF values are ultimately expressed in dimensionless units. We refer to them as "risk units," but they can be called anything.

One of the concerns expressed in the S-MAP workshops has been the uncertainty surrounding the CoF values associated with failure events. The proposed methodology addresses these concerns. Specifically, the methodology uses estimates developed by SMEs, as well as any available data about the possible likelihoods or consequences of different failure events. The uncertainty in the estimates provided by SMEs is captured specifically.

For example, an SME might be asked to provide a 10% - 50% - 90% range for the estimated CoF associated with a failure event, such as a wildfire. These three estimates are then converted into an expected value for the CoF, which is then used to measure the risk reduction (which is itself the difference of two expected values) from mitigation.

The effects on risk mitigation decisions of any uncertainty about the specification of CoF are addressed through straightforward sensitivity analysis. For example, suppose the expected consequences of a failure event are estimated to be 50 units. As a result of that expected value, the utility selects a mitigation strategy. A sensitivity analysis can vary the expected consequence

---

[8] The simplest way to see this is to consider a single failure event. The single failure event occurs or it does not. The probability of the occurrence of the failure is LoF. If the failure occurs, the consequences are measured by CoF, a number. The probability that the failure does not occur is 1-LoF. The consequences that result if the failure does not occur may be conveniently set to 0. Then the expected value of the consequences of failure is LoF *x* CoF + (1 - LoF) *x* 0 = LoF *x* CoF = Risk.

values (e.g., 30 units, 70 units, etc.) to determine whether the risk mitigation strategy changes in response to changes in the expected consequences. In many cases, the mitigation strategy will not change. In some cases, however, even a slight change in the CoF or LoF will change the risk mitigation strategy. When that occurs, it is valuable to collect additional data about the failure event or refine estimates in order to reduce the uncertainty and increase confidence in the risk mitigation decision.

We use the expected value of CoF to compute risk and risk reduction because there is no other consistent way of measuring risk reduction and selecting an optimal portfolio of risk management actions.[9]

C. Estimating LoF, the Likelihood of Failure

1. LoF is Based on Condition Dependent Hazard (Failure) Rate

LoF measures the uncertainty of occurrence of a failure. LoF is the mathematical probability of an event, most often the probability that a failure occurs over a given time period, which in our methodology is one year. (The time scale is arbitrary, however, and need not be one year.) (Other events like cyber-attacks are also described by the applicable probabilities of occurrence). The probability of an event is a number between 0% (which is interpreted to mean that it is certain that the event will not happen) and 100% (which is interpreted to mean that the event is certain to happen). Engineers refer to the probability that an asset will fail over time as the *hazard rate,*[10] but it can just as well be called a *failure rate*.

---

[9] The utilities have also remarked about the uncertainty surrounding their arrival rate estimates as the reason why, as shown in Table 2 of the next section, they use ranges for arrival rates, e.g., an event might occur between once every 100 years and once every 30 years. The way to address such uncertainty about LoF is to convert a range of arrival rates into a mathematical probability using the Poisson distribution (see footnote 13, *infra*).

[10] The hazard rate is the *conditional probability* that an event, typically a failure, will occur in the next interval of time, given that the event has not yet occurred. For most mechanical and electrical equipment, the hazard rate is often represented as a so-called *bathtub curve* that defines three periods in the life of a single piece of equipment. Initially, there is a burn-in period of relatively large hazard that decreases over time until the steady-state is reached. The steady-state is the flat part of the bathtub curve, such that the hazard is constant over time until the burn-out period is reached. The burn-out period occurs when assets age sufficiently so that the probability of failure increases. The burn-in period may last for minutes or years. The steady-state period may last for several years. The onset of burnout depends on the particular equipment. The shape of the hazard curve over time resembles the cross-section of a bathtub, hence the name. This form of the hazard curve is based

Hazard rates are typically based on observed behavior of equipment. For example, as electric transformers age, they can leak fluid and eventually fail. The hazard rate for transformer failures would be based on observations of how often such failures occur for transformers of different ages. Transformers can also fail when they are struck by lightning. Thus, all other things equal, transformers located in regions with lots of lightning strikes are more likely to fail than transformers in regions where lightning is rare. The effect of lightning on transformer failure can be accounted for by adjusting the hazard rate. Similarly, all else equal, natural gas transmission pipe is more likely to rupture in locations where there is a lot of earthquake activity in comparison to locations where there is much less earthquake activity.

The LoF for any asset can depend on several factors. These can include outside events (e.g., earthquakes, wildfires, terrorism, etc.); the condition of the asset (e.g., natural gas transmission pipe that is known to have manufacturing defects, wooden utility poles that have wind and insect damage, etc.); and even operator errors (e.g., employees who may not operate equipment correctly). Figure 1 provides a simple illustration for a natural gas transmission pipe of age *t*, where the age is measured in years since the pipe was placed in service. (We limit the event and conditions, and do not show operator error in this figure to provide a clearer example.)

Figure 1: Determinants of LoF

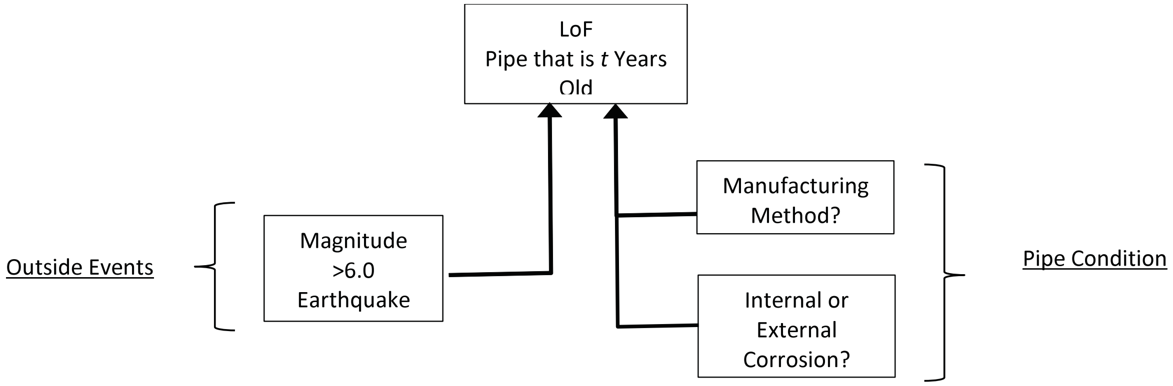

In Figure 1, the LoF for pipe that is *t* years old depends upon our knowledge of the occurrence of one outside event (earthquake) and the occurrence of two attributes that affect the pipe's condition: (i) the manufacturing method and (ii) the presence of internal or external

---

solely on the age of the equipment. If other information is known, such as the actual condition of the equipment, then the hazard function will take a different shape.

corrosion. (The question marks in each of the two boxes reflect whether or not these attributes are known to be present.)

In this example, in order to determine the LoF for a pipe that is *t* years old*,* we need to know a few things. First, we need to know the probability that a magnitude 6.0 or larger earthquake will take place. Second, we need to know the probability that the pipe of age *t* will fail *if* a magnitude 6.0 or greater earthquake occurs. Third, we need to know how the condition of the pipe affects the probability of failure.

To determine the condition of the pipe, we would either need to inspect it directly or estimate its condition based on what we already know about it (e.g., age, type of manufacturing, etc.). To determine the LoF, we would need to know the probability that a pipe *t* years old has a manufacturing defect and, *if* it does, the probability of failure given that the manufacturing defect is present. Similarly, we need to know the probability that a pipe *t* years old has external or internal corrosion (or both) and, *if* it does, the probability of failure given either or both those forms of corrosion are present.

The probabilities of failure *if* an earthquake takes place, or *if* there is a manufacturing defect, or *if* there is corrosion present are called *conditional probabilities*., This means that the probability of failure is *conditioned* on (or depends upon) another event taking place – either an earthquake or a manufacturing defect or corrosion, or any combination of those three events.

Let’s consider the influence of pipe condition on LoF, illustrated on the right-hand side of Figure 1. Our knowledge of the condition of the pipe with respect to manufacturing defects and corrosion alters the LoF and creates what is called a *condition-dependent hazard rate*. A condition-dependent hazard rate reflects the probability an asset will fail over time and depends on the condition of the asset.

Estimating a *condition-dependent hazard rate* is straightforward, and always begins with SMEs. The SMEs should know whether available data regarding the asset are accurate, applicable to the situation, and so forth. And in the event there are no available data, the SME can use his knowledge to estimate the hazard rate. In our methodology, we typically develop condition-dependent hazard rates by specifying three alternative conditions: “good,” fair,” and “poor,” and asking an SME to describe what it means for an asset to be in “good,” fair,” or “poor” condition. Using these discrete descriptors makes it easier for SMEs to describe how asset condition affects the likelihood of failure.

A new piece of pipe, free of manufacturing defects and fully coated, for example, would be in "good" condition. An old, heavily corroded piece of pipe would be in "poor" condition. The condition-dependent hazard rate simply estimates the probability of failure for assets in these different conditions. The resulting condition-dependent hazard rates might look as shown in Figure 2. As shown, pipe in poor condition, regardless of age, always has a higher probability of failure than pipe in good condition.

Figure 2: Condition-Dependent Hazard Rates

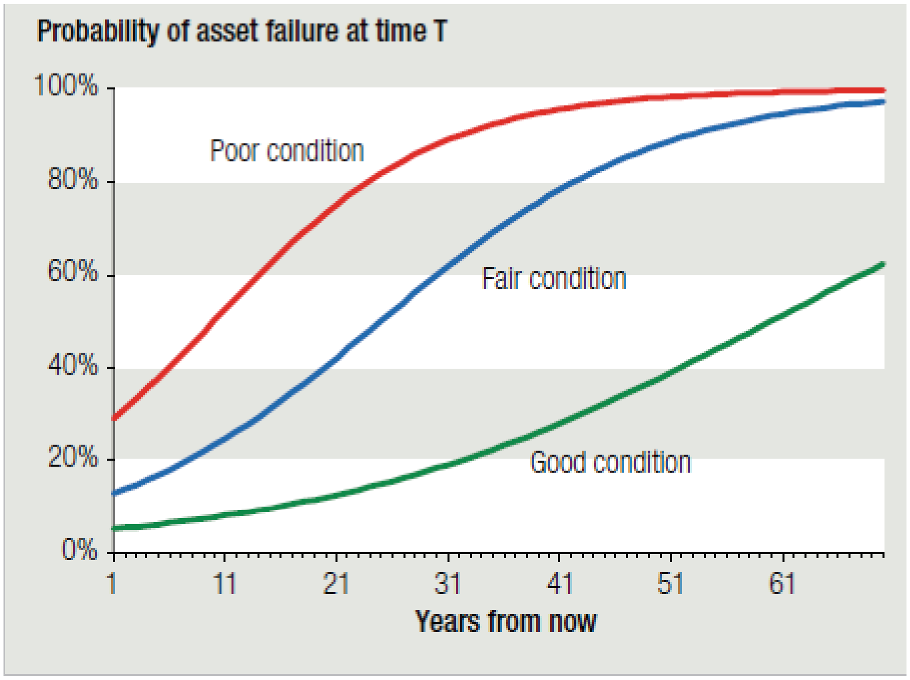

Once the condition of the pipe is determined, we still need to factor in the effect of an earthquake on the hazard rate. The simplest approach is to modify the condition-dependent hazard rates directly to reflect that the effect of the earthquake on pipe failure depends on the condition of the pipe. SMEs can provide estimates of these dependencies if data are not available. For example, in Figure 2, the occurrence of an earthquake would shift all of the hazard rate curves upwards, depending on the magnitude of the earthquake. Similarly, a wildfire could be more likely to cause failures in older, insect-damaged wooden utility poles than new ones. The same applies to other enterprise risk-level events, such as terrorism, floods and severe storms.

To summarize, the information required to develop condition-dependent hazard rates, and the effects of outside influences (e.g., earthquakes, wildfires, severe storms, etc.) is as follows:

- SMEs define what it means for assets to be in different conditions (e.g., good, fair, poor) and develop hazard rates for equipment in those conditions;

- SMEs also provide information about the types of outside events that can lead to asset failure and the likelihood of those outside events; and

- SMEs provide "multipliers" that are used to shift the hazard rate curves to account for the outside events. For example, in Figure 2, a magnitude 6.0 earthquake might shift the "good" condition hazard rate curve up by 10%, the "fair" condition rate curve by 20%, and the "poor" condition curve by 50%. (Of course, applying these multipliers recognizes that the hazard rate can never be greater than 100%.)

This information also lets us evaluate systemic failure events, that is, events that can cause widespread failures of equipment. In some cases, outside events (e.g., earthquakes, wildfires) will cause systemic failures by their nature. In other cases, a single failure event can cause cascading failures. For example, if a high-voltage transmission line fails, other lines can become overloaded and fail, too.

2. Contrast with the Utilities' Approach to LoF

As described in the utilities Uniformity Report, the utilities currently assign LoF values to risky events using a 1 to 7 scale.[11] This scale defines event frequencies, i.e., *arrival rates*, as shown in Table 2.

Table 2: Utility Arrival Rates and LoF Values

LoF Value						
1	2	3	4	5	6	7
Remote	Rare	Infrequent	Occasional	Frequent	Regular	Common
Once every 100+ years	Once every 30-100 years	Once every 10-30 years	Once every 3-10 years	Once every 1-3 years	1-10 times per year	>10 times per year

---

[11] Uniformity Report, December 1, 2015, p. 9.

The utility scales are similar to Richter scales, which are based on orders of magnitude. PG&E, for example, calls this LoF scale "quasi-logarithmic," and the resulting 1 to 7 scale for LoF seems to be nonlinear because the underlying arrival rates that the scale is supposed to represent change in ways that are similar to order-of-magnitude differences. Moreover, this scale means that an event occurring once every 35 years has the same LoF (2) as an event occurring once every 90 years, an event occurring every 3 years has the same LoF (4) as an event occurring every 9 years, and so forth, even though the mathematical probabilities of these events are all different.[12]

Consider the following example in which, instead of actual mathematical probabilities, we use the utilities' LoF scale. Consider two risky events that, as measured using the utilities' current 1 to 7 scale have the same CoF value, say CoF = C = 5. Suppose the initial LoF for the first event is 7 and the initial LoF of the second event is 2. And, suppose that the risk mitigation actions applied to each event will reduce the initial LoF values of each event by one unit. Thus, the risk mitigation action of the first event will reduce its LoF from 7 to 6, and the risk mitigation action of the second event in the second situation will reduce its LoF from 2 to 1. This means that for the first situation the risk reduction is:

$$(\text{LoF } x \text{ CoF})_{\text{BEFORE}} - (\text{LoF } x \text{ CoF})_{\text{AFTER}} = (7\ x\ 5) - (6\ x\ 5) = 1\ x\ 5 = 5 \text{ risk units.}$$

Similarly, for the second event, the risk reduction is:

---

[12] The utilities may suggest that their reliance on frequency ranges allows them to account for imprecision and uncertainty, but their approach does not yield better information in comparing risks and measuring risk reduction. Some may express the claim that "it is better to be approximately right than precisely wrong," but that is not the choice utilities face when selecting methodology and inputs. Instead, utilities should make the most informed decisions given the quality of the available inputs. The goal is to extract the useful information contained in the uncertain estimates by converting those estimates into numbers that facilitate analysis. Keeping wide ranges of actual levels of variables such as CoF and LoF does not facilitate analysis. Instead, it undermines the utilities' ability to calculate risk reduction. Although there may be uncertainty about the true (but unobserved) frequency of specific failure events, developing an optimal portfolio of risk mitigation actions requires using specific probabilities. The uncertainty in the specification of the CoF and LoF should be used to find the best point estimates of the CoF and LoF. And those point estimates should be used in the analysis. The impacts on the decisions of the selected risk mitigation actions to changes in the estimates of the probability of the failure event then can be evaluated using sensitivity analysis. As we have noted above, the uncertainty in any estimate is, in itself, not a fundamental issue. What is fundamental is the effect of the uncertainty on the decision. That effect cannot be understood without using an analytic methodology to measure it.

$$(\text{LoF } x \text{ CoF})_{\text{BEFORE}} – (\text{LoF } x \text{ CoF})_{\text{AFTER}} = (2\ x\ 5) – (1\ x\ 5) = 1\ x\ 5 = 5 \text{ risk units.}$$

That is, the amount of risk reduction in both cases is the same. Therefore, we should be indifferent between the two actions with respect to risk reduction.

But on the face of the utilities' approach, this is not true because of the correspondence of the scaled LoF numbers and the underlying arrival rates. The reason for this becomes more apparent when the scaled LoF values are converted to actual mathematical probabilities. In this example, using the Poisson distribution[13] to convert arrival rates to probabilities, the change in LoF from 7 to 6 corresponds to a reduction in probability of about 0.37 or 37%, which implies a risk reduction of 0.37 x 5 = 1.85 risk units. However, the change in LoF from 2 to 1 corresponds to a change in probability that is less than 0.03 or 3%, which implies a risk reduction less than 0.03 x 5 = 0.15 risk units. Thus, the risk reduction achieved by mitigating the first event is more than ten times greater than the risk reduction achieved by mitigating the second event (1.85 / 0.15 = 12.33). The reason that the risk reductions are not the same is that the 1 to 7 scaled LoF values do not correspond to mathematical probabilities.

For very infrequent events, the differences between arrival rates and mathematical probabilities are small. For example, if one estimates that something will take place, on average, once every 1,000 years, the true probability of occurrence over a single year is approximately 0.001,[14] or 0.1%, almost the same as the arrival rate. But as the arrival rate of an event becomes greater, say once per year, or 10 times per year, then the difference between an arrival rate and probability becomes significant. This is especially problematic for events the utility classifies as having arrival rates greater than one per year (LoF values of 6 and 7, in Table 2). An arrival rate of 10 times per year cannot be the same as a probability of 10, because probability cannot be greater than one (100%).

---

[13] The Poisson distribution is the most common way to convert an expected rate of occurrence of an event (e.g., once every five years, once every 100 years, etc.) into a mathematical probability. The conversion simply substitutes the expected rate of occurrence into an equation. For a discussion, see Sheldon Ross, *Introduction to Probability Models*, 6th ed. (New York: Academic Press 1997), pp. 249-257. For example, suppose we expect an event to occur once every 5 years. Then, the probability the event will occur within one year from now is $1 - e^{-1/5} = 0.181$ = 18.1%. In general, if an event X has an expected rate of occurrence of once every T years, then the probability X will occur within a time period of one year from now is $1 - e^{-1/T}$.

[14] This is estimated using the Poisson distribution, as described previously. For an arrival rate of once every 1,000 years, the corresponding probability is $1 - e^{-1/1000} = 0.001$.

The Commission should require the utilities to express LoF as a more understandable and straightforward mathematical probability. Further, LoF should be based on condition dependent hazard rates.

3. How Implementing Mathematical Probabilities for LoF Meets the Goals of a Successful Risk Management Methodology

The benefits of relying on mathematical probabilities include:

- The proposed change supports the S-MAP goal of prioritizing Public and Employee Safety by ensuring that the risk of failure events is estimated as accurately as possible. Failure to rely on mathematical probabilities may overestimate some risks and underestimate others, to the potential detriment of safety.

- The proposed change supports the S-MAP goal of Cost-Effectiveness by making it possible to compute both risk reduction and cost effectiveness. Expressing LoF as the probability of an event makes the computation of risk reduction possible.

- The proposed change supports the S-MAP goal of Understandability because it relies on actual probabilities and makes it possible to understand how a risk mitigation action changes the likelihood of occurrence of the risky event.

- The proposed change supports the S-MAP goal of Transparency by focusing attention on the likelihood of occurrence of a risky event. In particular, arguments about the conclusions reached by the utility can be based on disagreements about probabilities of risky events rather than on the way those event likelihoods are treated in the algorithm.

D. Estimating CoF, the Consequence of Failure

1. CoF Is Based on a Well-Constructed "Multi-Attribute Utility Function"

Utilities measure the benefits of their risk reduction actions with respect to multiple objectives. In effect, these objectives are designed to capture the benefits of improved performance with respect to *impact dimensions* including, but not limited to, *safety*, *system reliability*, *customer satisfaction*, *regulatory compliance*, *environmental consequences, and financial consequence*. A risk management action may achieve risk reduction benefits with respect to any or all of the value attributes. For example, deciding to add a safety awareness training class for field employees may have risk reduction benefits with respect to customer and employee *safety*, service *reliability*, *regulatory compliance,* and *financial consequences*.[15]

A failure event may affect one or more of these dimensions: for example, injuries may occur, energy demand may be unserved, customer satisfaction may decrease, and expenses may be incurred all as the result of a single failure event. These separate consequences of failure can be forecast (using either data or SMEs) and reasonable estimates of these consequences can be found. Any methodology that is used by utilities to measure the risk associated with utility operations and risk reduction associated with utility decisions must ensure that *safety* is given the top priority among those different value attributes.

Because it is inconvenient to consider multiple dimensions separately, economists created what are known as *multi-attribute utility functions*.[16] Multi-attribute utility functions combine the impacts in each consequence dimension into a single numerical value, measured in risk units (or treated as a number with no unit attached to it).

Each consequence dimension is called an *attribute,* in the sense that the measureable behavior of each dimension caused by the failure event is an attribute of the consequences of the failure event.

The measurement of the attribute in terms of its natural unit is called the *level* of the attribute. For example, a reliability dimension might be measured in terms of how many customers lost service and for how long because of an event.

The numerical value that a multi-attribute utility function provides measures the consequences of a failure event. Specifically, the value provides a single, numerical measure associated with the joint occurrence of different levels of the identified attributes that occur as a result of a failure event. Thus, a wildfire that resulted in pole failures could cause safety, environmental, reliability, and financial consequences. The multi-attribute utility function

---

[15] Each of these impact dimensions is an observable and measureable variable that can be described in terms of its own *natural unit*. For example, financial consequences can be naturally described in dollars but system reliability is more naturally described in terms of frequency of outage, duration of outage, service hours lost, unserved energy, or any other natural descriptor.

[16] Here, the word *utility* means an economic quantity, not a regulated entity that provides electricity or natural gas service.

determines an overall CoF value by converting the levels of these different attributes associated with the wildfire into a single score.

Our proposed methodology includes an approach to creating and testing multi-attribute utility functions that was developed and implemented in software for EPRI. (The EPRI methodology is described in reference [18]. An excellent reference to the theory of multi-attribute utility functions is [19]. Other reports and references are available from the authors of this white paper.)

In order to use a multi-attribute utility function and combine the individual attribute scores into a single overall score that permits estimation of risk reductions and cost-effectiveness, the multi-attribute utility function must be properly designed and have two special properties. First, the attributes themselves must be *value-independent*. This simply means that the contribution of a single attribute to the overall score does not depend on the level of any other attribute. This property is important to avoid double-counting consequences. For example, if a failure event occurs, then the financial consequences of the failure to the overall score cannot depend on the level of customer satisfaction. Conversely, the customer satisfaction consequences of a failure would not depend on the amount of money required to settle any lawsuits resulting from the failure. Second, the multi-attribute utility function is *additive.* This means that the contributions of each attribute are added together to determine the overall CoF score.

In order to implement the multi-attribute function, a utility must specify the following:

- *Attribute range*, which is the span of natural units over which the attribute will be measured. The span must include both the most benign level of the attribute and the most harmful level of the attribute. For example, a financial consequence scale could start at $0 (assuming a failure event does not improve the utility's finances) and might extend to billions of dollars. Constructing the attribute range is important because the outcome of all possible failure events must be included in the range. Defining the lower- and upper-bound levels of each attribute is typically done by interviewing utility experts, other SMEs, utility management, regulators, and other stakeholders.

- *Attribute scale,* which measures the relative value of each of the levels of each attribute. The attribute scale converts an attribute level measured in natural units (for example, 10 outages per year) into a value or numerical score (e.g., 38). This can be expressed in risk units or is treated as a dimensionless number. The attribute scale is itself arbitrary. Moreover, the scales will differ depending on the nature of the attributes. Thus, some scales always will be linear, whereas others need not be. When we have performed studies for utilities and others, we have found that expressing the scale from 0 to 100 is easiest to understand. With that scale, the worst (most harmful) level in the attribute range—representing the worst possible outcome—is set to a value of 100. The best level in the range is set to 0 on the scale. In the example above, a $0 financial consequence would be assigned a score of 0, while the worst, multiple billion-dollar financial consequence would be assigned a scale value of 100. The scaled values of intermediate attribute levels are found by comparing the values of attribute levels within the range. The simplest way to specify the intermediate levels of the attribute scale is to decide the relative value of two different changes in attribute levels. For example, eliminating the possibility of a single death might be viewed by SMEs as having the equivalent value as eliminating the possibility of 12 major injuries. Therefore, the scale value of one death is 12 times greater than the scale value of one major injury. Again, however, the scale value for any level of a given attribute does not depend on the scale values of any of the other attributes because of value independence.

- *Attribute weight*, which measures the relative importance of the attributes as compared to one another. As noted, ranges and scales for each attribute are determined without reference to any other attribute. Weights, on the other hand, require comparison and direct tradeoffs among attribute levels for different attributes. These comparisons provide ratios that can be converted into a set of attribute weights (which are all greater than zero) and which sum to 100%.

One way to ensure that public and employee safety is given the highest priority is to set the weight on safety before any other weights. For example, the CPUC could require all of the utilities to assign at least a 50% weight for safety. The remaining weights can then be determined by the utility so that they are all consistent with the attribute ranges and scales.

> If the attribute weight for safety is pre-determined, such as set to 50%, then the simplest way to develop the remaining attribute weights is to compare the values associated with changing each attribute level from its worst to its best against similar changes in level of other attributes. (This step is illustrated in an example below.) The conversion process is a simple algebraic exercise. Typically, as the weight on an attribute increases, the greater the contribution of that attribute to the overall score.

The resulting multi-attribute utility function will measure the utility of a set of attribute levels as the weighted sum of the scaled values of the attribute levels. For example, if there are four attributes—say, safety, reliability, financial, and environmental—the overall consequence score for a failure event X, which we can write as "CoF(X)," would be calculated using the following equation:

$$\begin{aligned} \text{CoF}(X) = {} & w_{\text{SAFETY}}\text{SafetyScore}(X) + w_{\text{RELIABILITY}}\text{ReliabilityScore}(X) \\ & + w_{\text{FINANCIAL}}\text{FinancialScore}(X) + w_{\text{ENVIRONMENT}}\text{EnvironmentalScore}(X) \end{aligned}$$

$$\text{where: } w_{\text{SAFETY}} + w_{\text{RELIABILITY}} + w_{\text{FINANCIAL}} + w_{\text{ENVIRONMENT}} = 100\%$$

CoF can be a dimensionless number, or expressed in terms of some arbitrary units, such as "risk units.

2. Properly Designing a "Multi-Attribute Utility Function"

The first step in constructing a multi-attribute utility function is to identify the risks that are present. Then, it is natural to consider the mitigation actions that can address those risks. Given the risks and the actions, the next question is how the actions will change the system in order to mitigate the risk. This leads naturally to the definitions of the attributes. Thus, the definition of the attributes follows from consideration of risks and actions. (One favorite question is "Why are you considering this action; what do you get if you do it?") The definition of the attributes requires defining the following aspects: (i) How is the attribute measured (natural units)? (ii) What levels can the attribute take on (attribute range)? and (iii), Are there subordinate attributes (attribute structure)?

After completing the attribute definitions, all of the attributes are presented in a logical hierarchy to facilitate further analysis.[17] At this point, the attributes are checked for value independence. If dependencies exist, attributes must be redefined until all attributes are independent with respect to value.

Next, the scale is created for each attribute. The typical scale sets the worst attribute level at 100 and the best attribute level at 0. The interior levels are assigned scale values by asking questions about the benefit of different changes in level. One interesting question, to set the midpoint, is: "Beginning at the worst level, how much improvement must be made in order to achieve half the benefit of moving from worst to best?" Other similar questions permit the specification of the attribute scales. This transparent scaling process is immediately checked and reviewed.

Next, the set of weights is specified. The fundamental weight-setting tradeoff is straightforward, but not simple. The weights are determined by specifying the relative importance of changing each attribute level from its worst level to its best level. These relative importance questions are answered with reference to the attribute hierarchy, starting at the bottom and working upward. At each stage of the hierarchy, the most important lower-level attribute is promoted, so that every tradeoff compares attribute level changes at the same hierarchical stage and every tradeoff compares actual changes that can be made to the real system. Rather than asking vague questions such as, "How important is reliability?" the questions are designed to compare the relative values of specific changes in different attribute levels. For example, a question might ask to compare avoiding a loss of $2 million compared, with avoiding an outage that would result in 1,000 customers losing power for three days. Such relative importance tradeoffs provide the ratios of the weights. Because the weights sum to one, such ratios are sufficient to solve for the weights.

The final step is to check the weights, scales, and ranges for internal consistency by creating many examples. Because the entire process is transparent, it is straightforward to discover and correct both logical and mathematical errors.

---

[17] The attribute structure is represented hierarchically. The structure contains all logical relationships among subordinate attributes. The structure ends when there are no lower-level attributes. For example, a complex attribute like *reliability* may have a structure that includes outage rate, outage duration, number of customers affected, types of customers affected, unserved energy, and other attributes.

The most critical aspect for determining the CoF values of different risk events is ensuring that the attribute ranges, scales, and weights are internally consistent.[18] This is where the multi-attribute utility process often breaks down. What frequently happens in practice is that attribute ranges and scales are first set by utility operational personnel and then the weights are specified by upper management personnel <u>independently</u> of the attribute ranges and scales. This leads to consistency problems, because the weights should be established by comparing and trading off specific changes in attribute levels. The weights, scales, and attribute ranges must all work together.

If the attribute weights are set without reference to the specific changes in the attribute levels, then there is no reason to believe that the relative values of the changes in the levels are measured correctly. For example, suppose that one of the attributes is *safety*. Suppose further that the worst conceivable safety outcome is some widespread catastrophe involving hundreds of deaths and hundreds of injuries. The best conceivable outcome is no death and no injury. These outcomes define the highest and lowest scores over the *safety* attribute range.

Next, suppose that another attribute is *financial consequences*. This attribute is measured in dollars. The worst conceivable outcome is a financial impact of a large number of dollars, say $5,000,000, and the best outcome is $0. Similarly, we can define the worst conceivable outcomes for the *reliability* and *environmental* attributes.

Now, suppose that the attribute weights are set without reference to the specific attribute ranges. For example, because safety is to be the most important attribute, that attribute is given a weight of 0.50. Financial consequences are not as important, so that attribute is given a weight of 0.10. This means that the score for changing the safety attribute level from worst to best is $0.50 \, x \, 100 = 50$ and the score for changing the financial consequences attribute from worst to best is $0.10 \, x \, 100 = 10$.

However, in practice, most changes in the level of the safety attribute accomplished by risk mitigation projects will involve relatively few deaths and relatively few injuries. Hence, the *changes* in the safety attribute scale will be nearer the low end (say 1) rather than the high end. But the changes in the financial consequences accomplished by risk mitigation projects could be

---

[18] It is not clear from the utilities' S-MAP filings that their current application of multi-attribute utility functions are internally consistent.

up to hundreds of thousands of dollars, which corresponds to changes in the financial consequences attribute scale that are nearer 10 rather than 1 or 100. This means that the typical change in the safety score would be on the order of 0.50 *x* 1 = 0.50 and the typical financial consequences change in score would be on the order of 0.10 *x* 10 = 1. This result contradicts the requirement that safety should be the most important attribute.

The error results from specifying attribute weights without simultaneously considering the tradeoffs among the attribute ranges. Such mistakes occur because the weights are not assigned based on the attribute range, and the attribute range is not specified based on real-world project considerations. As noted above, the attribute ranges, scales, and weights must work together and they must describe the effects of real projects.

The simplest method to ensure prioritization of safety is for the Commission to set a prioritized weight for safety and allow the utilities to identify all other attributes and develop all other weights. To the extent that the Commission elects to specify all consequence dimensions and the relative weights, it would similarly be required to set the attribute levels. While a public process could be established to identify attributes and set weights, this process would be time and labor intensive, and would significantly delay implementation of proper multi-attribute utility functions. To the extent that the Commission decides it is desirable to identify and weight all attributes, either on its own or within a public process, that process is best suited for the next S-MAP proceeding.

3. Application of COF: A Case Study Example

The following example is based on a recent study we performed for an electric utility. In this example, for the sake of simplicity, consider three high-level attributes: Money (which measures the value provided by incremental cash flows or cost savings), Reliability, and Safety.

The natural unit for the Money attribute is dollars. The range for the Money attribute is $0 to −$500,000. The scale for the Money attribute is linear, with a value of 100 at −$500,000 and a value of 0 at +$0. The limits were present for weighting purposes only; the scale of money can always be extended linearly.

The natural unit for the Reliability attribute is unserved energy, measured in MW-minutes. The range for the Reliability attribute is 0 MW-minutes to 19,000 MW-minutes. The scale for the Reliability attribute is piecewise linear, with 100 at 19,000 MW-min, 10 at 5,000

MW-min, and 0 at 0 MW-min. This means that 90% of the benefit (going from a scaled value of 100 to 10) is gained in reducing the MW-min lost from 19,000 to 5,000.

The natural unit for the Safety attribute is based on the health consequences for a single individual. (That will be extended additively when more than one person is exposed.) The range for the Safety attribute is (No injury to Death). Note that this is a discrete range that consists of four levels (No injury, Minor injury, Major injury, Death). The scale for the Safety attribute is a set of four numbers, where Death is 100, Moderate injury is 30, Minor injury is 10, and No Injury is 0. This means that eliminating the possibility of death for a single individual (100) is worth the same as eliminating the possibility of minor injury (10) for ten individuals. It is important to note that the Safety attribute is not expressed in dollars. Thus, there is no need to specify a statistical value of life.

The weights for Money, Reliability, and Safety are 6.25%, 18.75%, and 75%, respectively. These weights were derived by making the following tradeoffs: It is four times as valuable to change the safety level from Death to No Injury as it is to change the reliability level from 19,000 MW-min unserved to 0 MW-min unserved. And it is three times as valuable to change the reliability level from 19,000 MW-min unserved to 0 MW-min unserved as it is to change the cash flow level from -$500,000 to +$0. (In the actual study, utility personnel were asked to rank these changes.)

For example, suppose that, if a specific failure event occurs, then the consequences are the following:

- The expected dollars (lost) in incremental cash flow are –$300,000;
- The expected number of MW-min unserved is 2,000;
- There is a 50% chance of No injury, a 25% chance of Minor injury, a 15% chance of Major injury, and a 10% chance of Death for 3 people exposed to the consequences of the failure.

Graphically, the scales look like Figure 3 below. Notice that the money scale is a straight line, the reliability scale is piecewise linear (with small outages having less impact than larger outages), and the safety scale is a series of discrete values because the most natural way to describe a health outcome is to state the actual situation rather than use a continuous variable.

Figure 3: Attribute Scales

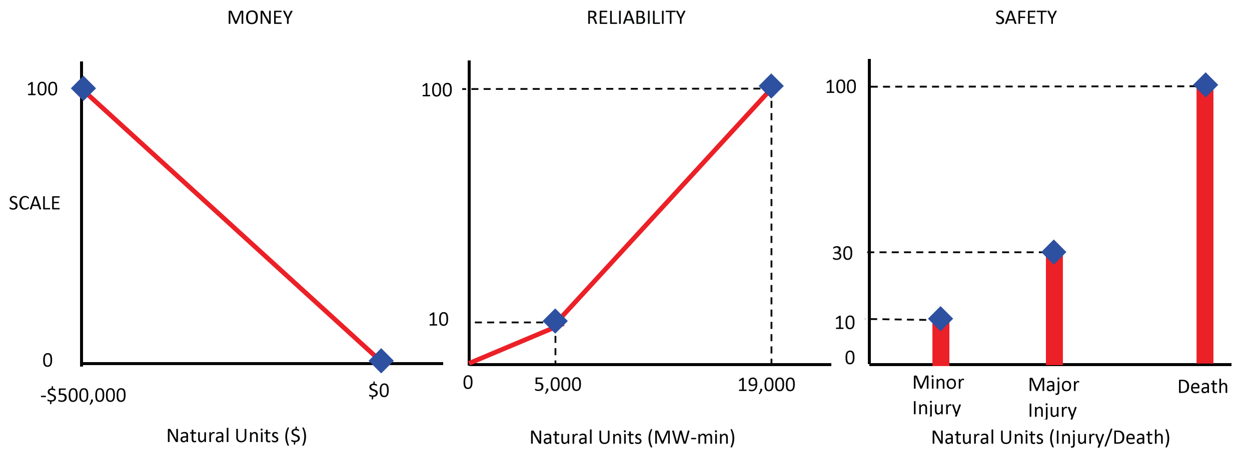

The scaled outcomes are the following:

(1) The scaled value of a -$300,000 loss is 60 units, because we move from the top of the attribute range $0) to the actual event (-$300,000), which is 60% of the way down, thus moving 60% of the way to 100. (That is, ($0 - $300,000) / ($0 - -$500,000) = 60% of 100 maximum risk units, or 60 risk units.)

(2) The scaled value of 2,000 MW-min of unserved energy is 4, because 5,000 MW-minutes has a value of 10 units and 0 MW-minutes has a value of 0 units, thus 2,000 MW-min / 5,000 MW-min = 40% of 10 risk units, or 4 risk units.

(3) The scaled value of Safety per person exposed is the expected scaled value of the consequences of the failure, or: (0.50 *x* 0) + (0.25 *x* 10) + (0.15 *x* 30) + (0.10 *x* 100) = 17. Because there are three people, the scaled value of Safety is 3 x 17 = 51 risk units.

The overall consequence score for this failure event is just the weighted sum of these individual scaled attribute values, or

CoF = (6.25% *x* 60) + (18.75% *x* 4) + (75% *x* 51) = 38.25 risk units.

If a risk mitigation action decreases the LoF for this failure event, then the risk reduction is equal to 38.25 multiplied by the change in LoF. For example, if we reduce the LoF from 50% to 40%, the reduction in risk is 10% *x* 38.25 = 3.825 risk units.

Similarly, suppose a risk mitigation action could not reduce the LoF, but could reduce the number of persons exposed to the consequences of failure to just two. In that case, the risk reduction caused by that action equals LoF *x* (0.75 *x* 17) = LoF *x* 12.75, because the weighted scaled safety value for one person has been removed by the action. And if the LoF is 50%, reducing the number of persons exposed from 3 to 2 reduces risk by 50% *x* 12.75 = 6.375 risk units.

This example demonstrates that the computation of CoF is straightforward. Moreover, it can be used to find the risk reduction associated with any risk mitigation action.

4. Contrast with the Current CoF Methodologies

The utilities currently assign CoF values using a nonlinear 1 to 7 scale. The impacts of different failure events are described by the utilities as shown in Table 3.[19] The attribute range is the discrete set that identifies the attribute levels *Negligible, Minor, Moderate ... Catastrophic.* The attribute scale is the set of numbers 1, 2, 3 … 7.

Table 3: Utility CoF Categories and Scores

Utility CoF Scores						
1	2	3	4	5	6	7
Negligible	Minor	Moderate	Major	Extensive	Severe	Catastrophic

The utilities have adopted different impact ranges to describe the consequences of different failure events. The utilities have not identified common attribute categories, but all agree to include a safety category. While the utilities have adopted common words (*e.g.*, “catastrophic,” “severe,” etc.) for the safety category (Figure 4), the criteria themselves are vague.

---

[19] Uniformity Report, pp. 8, 10-12.

Figure 4: Utilities' Uniform Safety Impact Criteria

Impact Level	Description
Catastrophic (7)	**Fatalities**: Many fatalities and life threatening injuries to the public or employees.
Severe (6)	**Fatalities:** Few fatalities and life threatening injuries to the public or employees.
Extensive (5)	**Permanent/Serious Injuries or Illnesses:** Many serious injuries or illnesses to the public or employees.
Major (4)	**Permanent/Serious Injuries or Illnesses:** Few serious injuries or illnesses to the public or employees.
Moderate (3)	**Minor Injuries or illnesses:** Minor injuries or illnesses to many public members or employees.
Minor (2)	**Minor Injuries or illnesses:** Minor injuries or illnesses to few public members or employees.
Negligible (1)	No injury or illness or up to an un-reported negligible injury.

Two things follow from this combination of attribute levels and scales. First, as with the 1 to 7 LoF scale the utilities have adopted, failure events with different consequences still can be assigned the exact same scores. For example, a "Severe" CoF value of 6 corresponds to a "few" fatalities and life-threatening injuries. Few means one or more, but how many more? Are 10 public fatalities a "few" or "many?" This kind of ambiguity always accompanies the specification of an attribute that is naturally measured numerically (*e.g.*, number of injuries, etc.) in terms that are not numerical. Second, as we discussed with respect to LoF, this seven point scale used for CoF makes it impossible to compute the risk reduction of different mitigation strategies. For example, using the utilities' CoF scale and holding the LoF for a specific failure constant, the risk reduction achieved by changing CoF from 7 to 6 will not be the same as the risk reduction achieved by changing CoF from 2 to 1, because the CoF scale is not additive.

Furthermore, currently the utilities do not appear to set the weights for the different attributes in a consistent manner. Instead, it appears that the weights are set independently of the attribute ranges by individuals who did not specify the attribute ranges. Therefore, there is no reason to believe that either risk or risk reduction will be measured correctly using such attribute weights.

The required changes to the utilities' method to estimate the consequences of failure events are straightforward: (1) replace the weights they now use with weights that are derived from the attribute ranges; and (2) replace the seven-point scale with an actual consequence scale. Whether the attribute scale ranges from 0 to 100 or 1 to 7 is not crucial, although, as we said, we have found using a continuous scale from 0 to 100 is easier to understand and facilitates the tradeoffs that are necessary to specify the attribute scale over the attribute range.

What the utilities seem to have done is combined the attribute range and attribute scale so that CoF is reported as an integer value between 1 and 7. This simplification does not have sufficient flexibility to define attribute scaling for different kinds of attributes. Also, because this is an attribute *scale,* equal changes in scale value should be equally valuable, *i.e.*, changing CoF from 7 to 6 should equal the value of changing CoF from 6 to 5, and so forth. But this is not the case with the existing utility methodologies. Hence, the utilities' CoF scale does not measure the value of a change in attribute level accurately or consistently. We suggest that the scale be continuous,[20] in other words, all values within the range are possible, rather than just discrete values of 1, 2 … 7. This *discrete* scale should be abandoned.

5. How Improving the Computation of CoF Meets the Goals of a Successful Risk Management Methodology

The benefits of the proposed methodology for determining CoF values include:

---

[20] According to the Uniformity Report, PG&E states its consequence scale is continuous. *See* Uniformity Report, p. 9. The report also states that, "The ultimate objective is to use continuous values, but the data available today and the quality of the data is not conducive to using continuous values." In fact, the scaling approach we suggest meets the utilities' ultimate objective and can be implemented. The scaling will most likely be done through the combined efforts of SMEs, who can best identify potential consequences, and utility management, who can best evaluate the various tradeoffs among those consequences.

- The proposed methodology prioritizes *Public and Employee Safety* and allows the CPUC and the utilities to establish an understandable and transparent weight for safety.

- The proposed changes support the S-MAP goal of *Cost-effectiveness* by making it possible to compute both risk reduction and cost effectiveness in a straightforward manner.

- The proposed changes support the S-MAP goal of *Understandability* by eliminating the 1 to 7 discrete nonlinear value CoF scale. The change makes it possible to understand the value-based results of a risk mitigation action that changes the levels of one or more attributes. The change also motivates planners to think in terms of measuring risk reduction by attribute.

- The proposed changes support the S-MAP goal of *Transparency* by establishing a direct link between mitigation actions that result in changes in the levels of the consequence attributes and risk reduction. Furthermore, by reporting the attribute ranges, scales, and weights, all CoF results can be easily replicated. In particular, disagreements over the conclusions reached by a utility will focus on how a utility's risk management programs change the levels of different attributes, and not about the way the attributes have been defined and measured.

E. Optimization Techniques Can Identify the Portfolio of Risk Mitigation Actions that Achieve the Greatest Risk Reduction at a Given Budget

1. Optimization Considers All of the Constraints Utilities Face in Carrying Out a Risk Management Program

The most cost-effective collection of risk mitigation actions is the portfolio of actions that achieves the greatest level of risk reduction, subject to the constraints the utilities face. Utility constraints likely vary among utilities and over time and may include, but are not limited to, budget (including rate affordability considerations), available labor and equipment, and operational constraints that limit what can be done (because service needs to be maintained), even if the budget and labor were available.

If Intervenors' recommendations with respect to LoF and CoF are accepted in the short term, then the risk reduction of every risk mitigation action will be known. Each utility should have the tools to characterize each proposed mitigation by reference to the utility's perceived constraints, such as cost, labor required, and potential service interruptions when a risk mitigation project is underway. With these inputs, determining the optimal set of risk mitigation actions to undertake is a straightforward and well-defined optimization problem. Of course, intervenors in a rate case or the Commission itself may make different judgments about, and set different levels for, these constraints. For example, determining the affordability of rate increases ultimately will be decided by the Commission based on input from all parties, which will affect the utilities' budgeted amounts devoted to risk management activities. While constraints may ultimately be adjusted by the Commission, the utilities should transparently identify and describe the constraints they face when developing their risk mitigation portfolio.

Implementation of optimization can help the utility demonstrate the cost-effectiveness of its proposed risk management portfolio in a transparent and understandable fashion. An optimization problem is a mathematical problem that is solved using an algorithm. For example, linear programming models, such as those that are often used to develop least-cost portfolios of generation and efficiency resources that meet future electricity demand, use a solution method (*i.e.*, algorithm) called the "simplex" method. The simplex method is the most popular solution algorithm for linear programming problems. Numerous optimization algorithms have been developed and are commercially available as software. For example, the "Solver" application in Microsoft Excel™ contains several algorithms, including the simplex method, which can be used to solve optimization problems.

Rather than applying optimization algorithms, the utilities now rely on other methods, called <u>heuristics</u>. Heuristics are measures that are supposed to give approximate solutions to optimization problems. For example, benefit-cost ratios are a common heuristic that, as discussed below, seldom identify optimized solutions.

2. <u>Application of Optimization: An Example</u>

As an alternative to optimization, the utilities may rank projects by their benefit/cost ratios, and select projects in rank-order until the budget is exhausted. This heuristic is sometimes called *prioritization.* If there is only a single resource constraint, the benefit/cost ratio can be a useful approximation. But, if there is more than one constraint, as there typically are in the utilities' risk management determinations, using benefit/cost ratios would not likely result in an optimal portfolio.

An example will illustrate the difference between a portfolio based on prioritization and one based on optimization. Suppose that there are seven risk reduction projects, each having a specific cost, an estimated benefit (expressed in "risk units"), and a specific labor requirement, as shown in Table 4.

Table 4: Pipeline Risk Reduction Strategy

Project	Risk Mitigation Project Cost (1,000$)	Risk Mitigation Benefit (Risk Units)	B/C Ratio	Labor Units Required	Benefit/Labor Ratio
	(A)	(B)	(C)	(D)	(E)
1	**$100**	300	3.00	5.0	60.00
2	**$20**	55	2.50	0.8	68.75
3	$150	350	2.33	2.5	140.00
4	**$50**	110	2.20	**1.0**	110.00
5	$50	100	2.00	**0.75**	133.33
6	$150	250	1.67	**1.5**	166.67
7	$150	200	1.33	**1.0**	200.00

In Table 4, the seven risk mitigation projects are ranked in order from highest B/C ratio to lowest (column **C**), as shown in the fourth column of the table. Column **E** shows the "Benefit/Labor Ratio," that is, the amount of risk reduction achieved per unit of available labor.

Suppose the utility has a budget constraint of $200,000. In that case, choosing projects solely by their B/C ratio would result in selecting projects [1, 2, and 4]. (These projects are highlighted in Column **A**.) Together, these projects cost $170,000 and provide 465 units of risk reduction (Column **B**). The remaining $30,000 the utility has to spend is too little to pay for any other projects.

Suppose, instead, that the only constraint is on available labor. For example, suppose the utility has only 4.75 units of labor available to the pipeline. If the projects were ranked based on their "benefit-to-labor ratio," *i.e.*, determining the most beneficial projects that could be completed given the labor constraint, then the pipeline would select projects [7, 6, 5, and 4]. The total labor used is 4.25 units and the total risk reduction benefit is 660 units. The remaining 0.5 labor units are too few to undertake any other project.

Therefore, with two constrained variables, the selection process based on benefit/cost ratio prioritization is not well defined. It is worth noting that the two prioritized selections have only one common member.

Next, we consider the results of optimization. We can solve three different optimization problems in the preceding example (using an Excel implementation of the zero-one integer programming algorithm in the Solver utility):

(1) The optimal group of projects (*i.e.*, the group of projects providing the highest risk reduction benefit) when there is only the $200,000 cost constraint is [1, 4, and 5]. These three projects provide a total of 510 units of risk reduction benefits and use the full $200,000.

(2) The optimal group of projects with just the 4.75 unit labor constraint is [3, 5, and 6]. These three projects use the full 4.75 units of labor and provide a total of 700 units of risk reduction benefits.

(3) Finally, the optimal risk reduction group of projects with *both* the $200,000 cost and 4.75 unit labor constraints in place is [3 and 4].

Thus, in this example the optimal group of projects to select when constraints are present (*i.e.*, the group of projects providing the optimal risk reduction benefit) may be significantly, or even completely, different than the sets of projects chosen when ranked by benefit/cost of benefit/labor ratios. Furthermore, the optimal group of projects selected depends on which constraints are present.

3. Implementation of Optimization

The example above demonstrates that prioritization is not the same thing as optimization, and that ranking projects based on benefit/cost ratios may not result in the most cost-effective portfolio when more than one constraint is present. Given the availability of computer software and relative ease of defining an optimization problem it is in the best interest of ratepayers to instead move towards optimization as a goal.

Implementing an optimization approach requires the utilities to identify and explain all of the constraints they face, such as the number of line crews available to undertake electric transmission and distribution projects at any given time (just to give one example). Once the constraints are identified, the optimization problem can be formulated as a simple zero-one integer programming problem and solved using widely available software. Moreover, the utilities can easily test the impacts of other uncertainties. For example, if the anticipated labor constraint changes because of additional hiring, an analysis can be rerun easily. Similarly, uncertainty regarding the benefits of mitigation of different projects can be evaluated to determine whether those uncertainties affect the choice of mitigation projects and, if so, how.

4. Optimization Still Requires Utility, Intervenor, and Regulator Judgment

No modeling approach should replace judgment of the utility, stakeholders and the regulator. Models are tools that can provide insights, but they should not be applied without proper oversight. Thus, the results of an optimization model should be viewed as advisory and reviewed carefully by all parties — utilities, regulators, and intervenors.

However, if the results of the optimization model are not adopted, then for the sake of transparency, it is incumbent upon utility management to explain why when it presents its rate case proposal. One of the benefits of an optimization model is that it identifies surprising and different results compared with a heuristic approach. What first might look like errors in the optimization model can actually provide new insights about the most cost-effective solutions. If an optimization model did nothing but reward intuition and pre-conceived solutions, the modeling approach would provide little added value. It is the unexpected results that make the modeling approach worthwhile.

The intervenors' proposed approach provides a clear roadmap for the utilities. If utilities, or other parties, believe that the results of a transparent and understandable analytical process must be adjusted to reflect other factors, then those factors can be identified clearly and their impacts can be evaluated.

5. Benefits of Optimization

The benefits of adopting optimization in the short run include:

- The proposed change supports the S-MAP goal of *Cost-Effectiveness* by making it possible to find the optimal set of risk mitigation actions. Optimality implies cost-effectiveness.
- The proposed change supports the S-MAP goal of *Understandability and Transparency* because a clear process for developing a project portfolio is implemented relying on analytic problem solving.

F. Strengths of the Proposed Methodology

The Intervenors' proposed methodology for risk management has a number of strengths, including:

- The methodology is straightforward to implement. The methodology is based on straightforward, intuitive definitions and principles, clearly identifies and defines the problem to be solved; and is solved using well-known and available mathematical methods.
- The results of the methodology rely on mathematical probabilities, not relative comparisons. The methodology: (i) uses mathematical probabilities of events to address uncertainty about the occurrence of risky events, and uses those probabilities to determine by how much different risk management strategies reduce the risks of failure events; and (ii) addresses uncertainty about the true conditions of assets that are vulnerable to failure events.
- The methodology can be implemented while the utilities gather additional data, initially relying primarily on subject-matter experts (SMEs) for the inputs necessary to use the methodology. Furthermore, the methodology can guide and focus future data collection, thereby avoiding time- and resource-consuming wasted efforts in gathering unimportant data.
- The methodology can be applied to both electric and natural gas utility operations, including transmission and distribution functions. In fact, the proposed methodology can be applied to almost any investment problem involving long-lived assets. The methodology has a proven track-record, having been used successfully by many electric utilities, as well as the PJM Regional Transmission Organization.

- The methodology does not require assigning a value to human life and health.
- The methodology complements the goals of the As Low As Reasonably Practicable (ALARP) framework. Commission staff published a paper discussing a potential role for ALARP in California. Specifically, staff highlighted ALARP's ability to "determin[e] how much risk mitigation is needed in a way that balances safety with cost."[21] Similarly, the paper encourages greater reliance on quantitative data and optimization techniques.[22] Consistent with these goals, the proposed methodology provides the Commission a quantitative means by which to address safety and affordability goals.

G. Implementation

1. Short Run Changes

In the short run, Intervenors recommend the following:

- Eliminate the unnecessary step of converting arrival rates of failure events into scaled 1 to 7 LoF values, and instead express LoF as a mathematical probability. LoF should be based on a condition-dependent hazard rate. The resulting LoF scale will be between 0% and 100%, linear, additive, and capable of measuring risk reductions associated with different mitigation strategies.
- Eliminate the existing discrete 1-7 CoF scale for failure events and replace it with a continuous scale. Also, the CPUC should consider implementing a more intuitive 0 to 100 unit scale. As part of this, the utilities or the Commission can establish a specific weight for the Safety attribute to ensure it is weighted most heavily. With that weight set, the utilities can also implement a multi-attribute approach that correctly defines weights and attribute scales together.
- Begin to implement additional optimization techniques by first requiring the utilities to clearly identify and quantify the key constraints affecting the utilities. Ideally, the

---

[21] "Safety and Enforcement Division Staff White Paper on As Low as Reasonably Practicable (ALARP) Risk-informed Decision Framework Applied to Public Utility Safety," December 24, 2015 at 3.

[22] *Id*. at 3-4.

utilities should rely on commercially available software to identify the optimal sets of risk management activities given those constraints. As an interim step, however, the utilities should prioritize risk mitigation activities based on risk reduction per dollar cost.

2. Long Run Changes

Over the longer term, the utilities should implement the full EPRI methodology, which has been rigorously tested and used by numerous electric utilities, and can be applied to both electric and natural gas operations.

A fundamental characteristic of utility systems is that they are dynamic, that is, the system's characteristics change over time. Accordingly, in the long run, the utilities should implement more complex approaches to determine the optimal set of risk mitigation actions given the dynamic nature of the utility system. Conditions of the assets change over time (equipment ages, corrosion increases, etc.) and the events that must be controlled (with respect to arrival and consequences) include asset failure. When asset failure will occur, however, is uncertain and utility assets must be operated over the indefinite future. Therefore, there are tradeoffs that must be made with respect to when to act and which actions to take. Depending on budget and other constraints, it may be optimal to defer action until a later date because it is safe to do so and use the budget available for more immediate purposes. A problem that has these characteristics — a dynamic system the condition of which is changing; risks that vary over time; risks that are based on the occurrence of random events; risk mitigation actions that can be deferred depending on need; constraints that vary over time; an operating horizon that extends into the indefinite future — can be formulated as a well-known problem called an *optimal control problem under uncertainty*.

We have formulated the risk management problem exactly that way and the references [1]-[7], [14] – [17] describe how the proposed methodology responds to that problem formulation. In particular, implementing an optimal control formulation will also identify the most valuable data to be gathered. The description of the system, specifying the uncertainty of failure, and the optimization algorithms used to solve the problem are all based on the optimal control formulation. Making these changes in the long run will further increase the benefits of the short run changes.

3. Successful Implementation Depends In Part on the Quality of Inputs

Adoption of the methodology proposed in this paper alone does not guarantee effective risk management. The success of utility risk management also depends on the quality of model inputs. Inputs may be sourced from historical utility records, external industry statistics and subject matter experts. Inaccurate inputs to the methodology, however, will result in inaccurate outputs.

Improving the quality of utility input data will be an evolving process. The state of data collection varies among utilities, but more methodical and continuous data collection will improve the quality of all utility risk management programs. An important benefit of implementing the proposed methodology is that it will signal the data-gathering required to provide the inputs to the methodology, eliminating wasted time and expensive efforts.

While data collection and retention procedures are under improvement, the utilities will be required to rely to a greater extent on SMEs. Experienced SMEs can be a valuable source of information, provided that their knowledge is transparently, systematically and accurately translated into objective inputs. Well-designed procedures for encoding expert judgment should be included in the implementation of the methodology. A well-designed methodology will also provide the utilities insights into the adequacy of the data they have collected, whether empirical data or SME estimates.

APPENDIX

The structure and implementation of the methodology has been described in a series of EPRI reports, which demonstrate how the methodology was applied to such electric systems and specific types of assets (*e.g.*, transformers, underground cable, wood poles, and breakers). (*See* references [1], [2], [3], [4], [5], below). Additional EPRI reports describe the development of the methodology ([6], [7], [8]), methods for estimating the reliability of system components ([9], [10], [11], [12]), and research into methods for converting customer needs into customer values [13] for purposes of ranking alternative strategies. The methodology has been peer- reviewed and presented at various conferences, including Institute of Electrical and Electronics Engineers (IEEE) meetings ([14], [15]). Non-technical descriptions have been published in the trade press ([16], [17]). These represent only a sample of the reports available from the authors of this whitepaper.

[1] "Guidelines for Intelligent Asset Replacement, Volume 1" (J. Bloom, C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2003. 1002086.

[2] "Guidelines for Intelligent Asset Replacement, Volume 2" (J. Bloom, C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2004. 1002087.

[3] "Guidelines for Intelligent Asset Replacement, Underground Distribution Cables. Volume 3" (J. Bloom, C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2005. 1010740.

[4] "Guidelines for Intelligent Asset Replacement, Wood Poles. Volume 4" (J. Bloom, C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2006. 1012500.

[5] "Substation Transformer Asset Management and Testing Methodology (J. Bloom, C.D. Feinstein, P.A. Morris) EPRI, Palo Alto, CA: December 2006. 1012505.

[6] "Cable Reliability Management Strategies: Research Status Report" (C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2003.

[7] "Asset Population Model with Testing for Managing Aging Power Delivery Assets" (C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: November 2004. 1008562.

[8] "Equipment Failure Modeling for Underground Distribution Cables (J. Bloom, C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2006, 1012498.

[9] "Reliability of Electric Utility Distribution Systems: EPRI White Paper" (C.D. Feinstein, P.A. Morris and R. Cedolin). EPRI Technical Report 1000424, October 2000.

[10] "A Review of the Reliability of Electric Distribution System Components: EPRI White Paper" (C.D. Feinstein, P.A. Morris and G.L. Hamm). EPRI, Palo Alto, CA: December 2001. 1001873.

[11] "Estimating Reliability of Critical Distribution System Components" (C.D. Feinstein, G.L. Hamm and P.A. Morris). EPRI, Palo Alto, CA: December 2003. 1001704.

[12] "Medium Voltage Cable Failure Trends: Research Status Report" (C.D. Feinstein, P.A. Morris). EPRI, Palo Alto, CA: December 2003.

[13] "Customer Needs for Electric Power Reliability and Power Quality: EPRI White Paper" (C.D. Feinstein, P.A. Morris and C. Downs). EPRI Technical Report 1000428, October 2000.

[14] "Optimal Replacement of Underground Distribution Cables" (J. Bloom, C.D. Feinstein, and P.A. Morris). *IEEE Power Systems Conference and Exposition*. Tampa, FL: 2006. (PDF version of this paper is available upon request.)

[15] "The Role of Uncertainty in Managing Aging Assets In Electric Utility Systems" (C.D. Feinstein and P.A. Morris). *IEEE PES Transmission and Distribution.* New Orleans: April 2010. (PDF version of this paper is available upon request.)

[16] "Spare Transformers and More Frequent Replacement Increase Reliability, Decrease Cost" (C.D. Feinstein and P.A. Morris), *Natural Gas and Electricity, Wiley Periodicals,* November 2008, 18-27.

[17] "Opening the Black Box: A New Approach to Utility Asset Management" (C.D. Feinstein and Jonathan A. Lesser). *Public Utilities Fortnightly,* 36-42, January 2014.

[18] "Project Prioritization System: Methodology Summary" (S.W. Chapel, C.D. Feinstein, P.A. Morris, V. Longo), EPRI, Palo Alto, CA: December 2001. 1001877.

[19] Value-Focused Thinking (Ralph Keeney). Harvard University Press, 2009.

Utilities That Have Used the EPRI Software

- American Electric Power (AEP)
- American Transmission Company (ATC)
- Baltimore Gas and Electric (BGE)
- British Columbia Hydro (BCH)
- Commonwealth Edison (Com Ed)
- Consolidated Edison (Con Ed)
- South Africa Electric Utility (ESKOM)
- Great Lakes Power (GLP)
- Green Mountain Power (GMP)
- Hawaii Electric Company (HECO)
- Kansas City Power and Light (KCP&L)
- Long Island Lighting (LILCO)
- MidAmerican Electric (MAE)
- Nashville Electric Company (NEC)
- Ontario Hydro (OH)
- PECO (Philadelphia Electric)
- PJM Interconnection (PJM)
- Public Service Electricity and Gas (PSE&G)
- Quebec Hydro (QH)
- Salt River Project (SRP)
- Southern Company (SoCo)
- Texas Utility (TXU)
- United Illuminating (UI)
- Wisconsin Electric (WI)

QUALIFICATIONS

Charles D. Feinstein, Ph.D.

Charles D. Feinstein is Associate Professor of Operations Management and Information Systems at the Leavey School of Business, Santa Clara University. Dr. Feinstein is cofounder of VMN Group LLC, a quantitative consulting company. He also teaches and has taught in the Department of Management Science and Engineering at Stanford University and in the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. He received his B.S. in Mechanical Engineering from Cooper Union. His advanced degrees (M.S. Aeronautics and Astronautics, M.S. Mathematics, PhD. Engineering-Economic Systems) are from Stanford. Dr. Feinstein has more than thirty years of experience in research, teaching, development and application of mathematical methods and mathematical modeling. His areas of expertise include optimization, decision analysis, probability, statistics, system dynamics, and systems analysis. His courses include operations research, operations management, investment science, systems analysis and design, linear and nonlinear programming, dynamic optimization & optimal control, and probability & statistics, at both the undergraduate and graduate levels. His previous employment includes positions as Senior Decision Analyst at Applied Decision Analysis, Inc. and Research Engineer at Xerox PARC (Palo Alto Research Center). He has been active in the academic and professional communities and has published more than fifty technical papers and reports as well as presented many lectures on both theoretical and applied research. His current interests include investment planning and risk analysis in the electric power industry. Dr. Feinstein and his colleagues at the Electric Power Research Institute (EPRI) developed state-of-the-art methodologies for managing the risks of failure and outage in the electric power system. He has written and presented extensively on managing aging infrastructure, project prioritization methodologies, distribution system risk analysis, and the application of distributed resources to distribution planning.

Jonathan A. Lesser, Ph.D.

Dr. Jonathan Lesser is the President of Continental Economics, Inc., and has over 30 years of experience working for regulated utilities, governments, and as an economic consultant. He has analyzed numerous economic policy and regulatory issues affecting the energy industry, including environmental policy, cost-benefit analysis, industry market structure, market power and market manipulation, the cost of capital, cost allocation and rate design, investment under uncertainty, risk management, incentive regulation, and economic impact analysis.

Dr. Lesser has prepared expert testimony and reports in cases before utility commissions in numerous U.S. states; before the Federal Energy Regulatory Commission (FERC); before international regulators in Latin America and the Caribbean; and in state and federal courts in commercial litigation cases. He has also testified before the U.S. Congress, and legislative committees in numerous states on energy policy and market issues. Dr. Lesser has also served as an independent arbiter in disputes involving regulatory treatment of utilities and valuation of energy generation assets.

Dr. Lesser is the coauthor of three textbooks: Environmental Economics and Policy, Fundamentals of Energy Regulation, and Principles of Utility Corporate Finance., as well as many academic and trade press articles. He has taught undergraduate and graduate courses in economics and business, and is currently a Lecturer in the Department of Economics at the University of New Mexico, where he is teaching a class on energy regulation and policy. He recently finished serving a three-year term as one of the Energy Bar Association "Deans" overseeing education programs on regulatory and ratemaking concepts for new attorneys.

(END OF ATTACHMENT 1)