IET Generation, Transmission & Distribution

ORIGINAL RESEARCH OPEN ACCESS

Quantifying Distribution System Resilience From Utility Data: Large Event Risk and Benefits of Investments

Arslan Ahmad | Ian Dobson

Department of Electrical and Computer Engineering, Iowa State University, Ames, Iowa, USA

Correspondence: Arslan Ahmad (arslan@iastate.edu)

Received: 27 May 2025 | **Revised:** 20 August 2025 | **Accepted:** 12 October 2025

ABSTRACT

We focus on blackouts in electric distribution systems that have a large cost to customers. To quantify resilience to these events, we show how to calculate risk metrics from the historical outage data routinely collected by utilities' outage management systems. Risk is defined using a customer cost exceedance curve. The exceedance curve has a heavy tail that implies large fluctuations in large blackout costs, and this makes estimating the mean large cost in the usual way impractical. To avoid this problem, we use new resilience metrics describing the large event risk; these metrics are the probability of a large cost event, the annual log cost resilience index, and the average of the logarithm of the cost of large-cost events or the slope magnitude of the tail on a log–log exceedance curve. Resilience can be improved by planned investments to upgrade system components or speed up restoration. The benefits that these investments would have had if they had been made in the past can be quantified by "rerunning history" with the effects of the investment included, and then recalculating the large event risk to find the improvement in resilience. An example using utility data shows a 2% reduction in the probability of a large cost event due to 10% wind hardening and 6%–7% reduction due to 10% faster restoration in two different areas of a distribution utility. This new data-driven approach to quantify resilience and resilience investments is realistic and much easier to apply than complicated approaches based on modeling all the phases of resilience. Moreover, an appeal to improvements to past lived experience may well be persuasive to customers and regulators in making the case for resilience investments.

1 | Introduction

Overhead distribution systems are vulnerable to extreme weather such as extreme wind. For example, the August 2020 upper Midwest USA derecho caused ~11 billion dollars of damage and left more than one million customers without power. The frequency and intensity of such extreme weather events are gradually increasing [1, 2]. This motivates quantifying the resilience of distribution systems by calculating the risk of large-cost events, as well as quantifying the benefits of planned investments to reduce these risks, and finding ways to help justify these investments to customers and regulators.

Almost all of the literature quantifying distribution system resilience either optimizes expected (mean) losses [3, 4] or addresses the resilience for specific extreme events [5–9], or uses reliability indices such as SAIDI that address system reliability averaged over the year [2, 5, 10]. Expected or average losses are dominated by more routine outages and do not directly measure extreme event risk. The field is starting to move beyond specific resilience events and average metrics. For example, Carrington [11] extracts resilience events of all sizes from observed data and obtains the overall statistics of resilience metrics from the outage and restore processes. Moreover, papers led by Dubey [3, 12] have pioneered simulation models that assess value at risk (VaR) and conditional value at risk (CVaR) resilience metrics that directly quantify the risk of large events. While resilience quantification in distribution systems typically uses detailed models of a subset of resilience processes to simulate and assess resilience [2], excellent opportunities are opening up to assess distribution system resilience directly from observed utility data. Ahmad [13] uses utility data not only to quantify resilience with metrics but to "rerun history" with the effects of investments in resilience included to quantify the benefits of those investments. However, Ahmad [13] uses metrics for resilience events such as number of outages, duration, and customer hours not served, and does not use a metric directly describing risk.

In this paper, we aim to:

1. Formulate new metrics that use utility data to quantify distribution system resilience in terms of the customer risk of large-cost events, despite the problems of large fluctuations in blackout costs caused by heavy-tailed statistics, and
2. Extend the historical rerun method to quantify the effects of resilience investments on the large event risk. The resilience investments considered are hardening the distribution infrastructure to withstand higher winds, and faster restoration. We compare the effects of these investments on the risk metrics.

This paper extends and expands the methods and metrics of the conference paper [14] and also applies new ALEC and ALCRI metrics [15].

The historical rerun method quantifies the resilience improvement that a proposed resilience investment would have had if that investment had been made in the past [13]. Since it is driven by real data, this has the advantage of incorporating all the factors affecting system's resilience over the past period, such as weather, trees, human factors, operating procedures, equipment aging, system reconfigurations, and restoration practices. Thus, the historical rerun method has no modeling error from these factors. The historical rerun method does not predict the future, but the model-based methods of predicting the future with simulation must represent the considerable complexities of all the phases of resilience and are very complicated, whereas data drives the historical rerun method and it is much simpler and straightforward. Moreover, in communicating the benefits of a proposed resilience investment to stakeholders, the historical rerun method has some advantages. The benefits that would have applied to the lived experience of stakeholders in the past, both for particular large events and in general, may well be more persuasive than the benefits that are modeled and simulated for predicted events at some indeterminate time in the future.

2 | Literature Review

Reliability indices [16] such as SAIDI, SAIFI, CAIDI, etc. remain the most common and most reported performance measures by most utilities and regulators. These reliability metrics measure the frequency and customer impacts averaged over the year of regularly occurring outages. These reliability indices are sometimes extended to include extreme outage events together with the regularly occurring events [2, 5, 10, 17–20]. However, the reliability metrics often entirely exclude extreme events (major event days) from the calculations so that the large variability in the extreme events does not make the reliability metrics vary erratically from year to year [21] [16, Section 6.3]. Even if extreme events are included in the calculation, these reliability metrics do not directly characterize extreme event risk. This is because reliability indices focus on normal operating conditions and they value each lost kilowatt hour equally across time, when in fact the value of lost load can compound the longer it is lost [21]. With the increase in extreme weather events [1], the risk of power outages due to extreme weather events merits special treatment. Some investments to improve reliability using reliability metrics may also improve resilience, but this is not quantifiable unless resilience metrics are also used.

The traditional risk assessment techniques in distribution systems quantify system reliability. They calculate the steady state of Markov chain models and also use analytic methods such as when the components of a radial distribution feeder are modeled as a series network [17, 18]. Steady state probabilities averaged over the year are computed, primarily for the regularly occurring power outages, although a few classes of weather severity can also be considered [10].

In the wider context of risk analysis, Kaplan and Garrick [22] define risk as a combination of the probability of occurrence and the severity of consequences. Yet, most studies quantifying extreme events risk in power distribution systems emphasize only one of these components. For example: Xi Chen [23] calculates area failure probability (but not cost) to express power outage risk to residential customers in active distribution systems. Po-chen Chen [24] calculates the hazard and vulnerability of distribution systems, where hazard represents the probability of a weather event of a particular intensity and vulnerability is the probability of outage under the occurring weather conditions. Du [25] uses weather-related failure counts to express weather-related risk in distribution systems and trains a Bayesian neural network to predict different risk levels corresponding to predicted weather-related failures caused by wind, rain, and lightning. Guikema [26] suggests Bayesian data-based priors that can be used for failure probability estimation in power distribution systems.

There are other works that focus on the risk of extreme events, but they usually only consider a limited number of samples of one type of extreme event [27]. Han [7] develops models that predict damage, power out, and customers out to estimate the risk for hurricanes in distribution systems before they make landfall. Xu [8] proposes hazard resistance-based spatiotemporal risk analysis to predict the number of customers out on a distribution system feeder during a hurricane event and use this as a risk measure. Guikema [9] improves hurricane risk management by predicting the number of utility poles that will be damaged before the arrival of a hurricane. Reed [5] analyzes the occurrence of power outages during extreme winds and Liu [6] analyzes the same during ice storms and hurricanes. Watson [28] uses EAGLE-I outage data along with weather data, environmental data, and US census data to train a gradient boosted tree model to predict the impacts of tropical storms. These studies provide valuable insights into system vulnerabilities under particular conditions, but don't allow for a comprehensive risk analysis or an analysis easily generalizable to other event types or geographic contexts.

In contrast, our approach analyzes observed events of all sizes to enable a more comprehensive risk analysis.

Works that predict outages in distribution systems during multiple types of extreme events [27, 29, 30] can be used for risk analysis if they are supplemented with an impact assessment of the predicted outages.

Apart from quantifying the historical risk of the distribution system, real-time/online risk assessment and short-term prediction of risk are also addressed in some works. For example, Leite [31, 32] uses historical outage management system data and weather data to model the failure probability and estimate interruption cost with a time series. These two factors are then used to develop a risk matrix to predict hourly risk levels for each distribution feeder section. Lin and co-authors [33] propose a data-driven online risk assessment method for distribution systems. They estimate risk in real-time using 25 different indices from five different categories: load, grid, resilient resources, emergency response and repair resources, and meteorology. Such predictive approaches can usefully guide operational decisions but are not designed for planning and evaluating investments.

A much smaller subset of literature combines probability and consequence estimation to give a risk analysis. In particular, Dubey and co-authors [3, 12] have pioneered value at risk (VaR) and conditional value at risk[1] (CVaR) to give a risk analysis for power distribution systems resilience. Related works that discuss the use of VaR and CVaR for risk analysis in power systems include references [34, 35]. VaR and CVaR are important measures used in financial risk management to assess potential losses in an investment. VaR estimates the maximum potential loss that can be incurred with a given probability over a specified period. In other words, it identifies a loss threshold exceeded with a fixed probability (e.g., the 95th percentile cost). Whereas, CVaR estimates the *mean* cost conditional on exceeding the VaR threshold. These measures, long established in finance and other infrastructure risk domains, are attractive because they explicitly capture tail risk rather than average performance. However, the application of CVaR to power outage cost data faces statistical challenges when the cost distribution is sufficiently heavy-tailed, because estimates of CVaR do not converge for the practically available amount of data [15]. As we discuss later, this prevents the use of CVaR in cases when the tail index is less than or equal to 2 and motivates our new metrics that do converge in practice.

The problem of high variability in distribution system blackouts associated with heavy-tailed probability distributions in observed outage statistics is well known. Calculations of SAIDI often exclude the highly variable major event days to get consistent year-to-year estimates of reliability [16, Section 6.3]. Median instead of mean can be used to get better estimates of a typical value under normal conditions when the data includes extreme values [36]. Logarithms are often used to handle high variability in the data. For example, Sullivan [37] discusses several forms of logarithmic transformation of outage data to improve regression models. Pandey [38] introduces the Area Index of Resilience AIR metric that averages the logarithm of event customer hours to reduce its variability. Carreras [39] discusses the corresponding heavy tail problem for blackout sizes in transmission systems.

In summary, the existing literature related to risk in distribution systems can be grouped into four broad streams:

1. Reliability-driven methods and indices (SAIDI, SAIFI, etc.) that describe the steady-state risk averaged over the year, capturing common outages but excluding or diluting the effect of extreme events.
2. Event-specific resilience studies that analyze or predict the impacts of particular hazard types, providing depth but limited generalizability.
3. Methods that partially address risk by calculating one of the components of risk (probability or consequence) rather than a complete risk calculation.
4. Risk-based approaches that combine probability and consequence, including VaR/CVaR methods, but face statistical limitations under heavy-tailed cost distributions.

Our work advances the risk analysis of extreme events in distribution systems by:

1. Simultaneously incorporating both the probability and consequence components of risk of large cost events in distribution systems.
2. Using actual observed outage data for events of all types and magnitudes, avoiding dependence on specific single-hazard scenarios.
3. Addressing the deficiencies of mean-based metrics in heavy-tailed settings by introducing risk metrics that can be practically calculated.
4. Use the new risk metrics to quantify the benefits of resilience investments from a risk-to-customer perspective.

This positions our contribution as a data-driven and statistically robust alternative to existing risk and resilience quantification methods, capable of both describing the risk of large cost events and evaluating the benefits of potential resilience investments.

3 | Outage Data and Extracting Events

We use 6 years of detailed outage data recorded by a US distribution utility. The dataset contains records of 32,278 individual power outages that occurred in the utility's network. Each outage entry corresponds to an outage of a component in the distribution system and includes the number of customers affected during the outage, the outage's start and end times, and its cause codes. We exclude the scheduled and planned outages and only consider the unscheduled outages in this analysis.

To analyze the wind resilience investments, we use NOAA weather data from weather stations available within the distribution network's geographic area. For each outage, we use the weather data from the closest available weather station. The overall distribution network is thus divided into multiple small areas based on the number of weather stations [13]. For this paper, we use two of the largest areas in the distribution network: area A and area B, which, respectively, contain 12,715 and 7876 unscheduled distribution system outages of at least

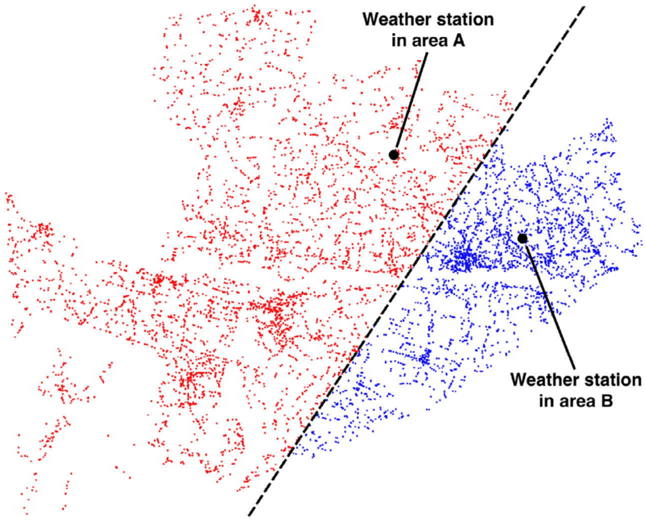

FIGURE 1 | Geographic locations of outages and the associated weather stations in two areas of a distribution system.

1-min duration. The areas and their corresponding outages are shown in Figure 1. The discussion and results in the subsequent sections are based on area A. All the results of area B are given separately in Section 8.

We group the outages into resilience events during data pre-processing. Resilience events are formed by overlapping outages. The start of an event is defined by an initial outage that occurs when all components in the distribution system are operational, and the end of the same event is defined by the first subsequent time when all the components are restored. Two example events are shown in Figure 2. More details about resilience events and their automatic extraction from the outage data are available in [13]. 3706 resilience events are formed from the outages in area A.

4 | Estimating Customer Cost

The normalized customer cost of a power outage event in a utility that serves n_{customer} customers can be described in terms of the total customer hours lost in that event as:

$$C = \frac{k\ A_{event}}{n_{\text{customer}}}$$

where A_{event} is the total customer hours lost in that event, which is equal to the area under the customer performance curve [40], and k is the average cost per customer per hour of an outage.

The value of k can be estimated in various ways [41] such as customer surveys or online tools like DOE's ICE calculator [42] and NREL's CDF calculator [43]. We use $k = \$370.2$ (2022 USD), based on the average proportions of customer classes (residential, commercial, industrial) in the utility and expert feedback from another utility. The normalization by the number of customers served by the utility n_{customer} spreads the blackout cost to the affected customers equally over all the customers served by the utility to give a blackout cost per served customer. This allows comparison of costs when the number of customers served by the utility changes.

The cost of a power outage to customers depends on different factors. These include the number of customers affected by the outage, outage duration, customer class, the affected customer's power outage risk level (houses with ill residents are at elevated risk), the criticality of services offered by the affected customer (hospitals, old homes, police, etc.), along with various other direct and indirect socio-economic factors [37, 44–51]. Incorporating all of these factors would give a more accurate yet complex model for the cost to customers. Different values of k can be used for different customer classes and multiplied with each outage individually as per its affected customer class to get a more accurate estimate of the customer cost.

While we address here a resilience event's cost to customers, there are costs to the utility as well, which could be similarly modeled and analyzed to quantify the risk to the utility.

5 | Estimating Large-Cost Events Risk

One basic definition of risk associates probabilities with costs of events or groups of events [22], and can be described by the probability distribution of the cost. One useful way to present the probability distribution of cost is the cost exceedance curve $\bar{F}_C(c) = \text{P}[C > c]$, which is the probability of the event customer cost C exceeding the value c as c varies [2].

Figure 3 shows the customer cost exceedance curve obtained from the utility data.

5.1 | Probability of a Large Cost Event

To help describe and communicate the risk of large-cost events in Figure 3, we define large cost events as those events with cost $c \geq c_{\text{large}}$, where c_{large} is the threshold for the minimum large normalized cost. Suppose that n_{large} is the number of events with $c \geq c_{\text{large}}$ and n_{event} is the total number of events. Then the probability of an event having large cost is the probability of an event cost exceeding c_{large}:

$$p_{\text{large}} = \text{P}[C \geq c_{\text{large}}] = \frac{n_{\text{large}}}{n_{\text{event}}} \tag{1}$$

For our data we choose $c_{\text{large}} = 0.59$ USD per customer served, which corresponds to the 90$^{\text{th}}$ percentile of the observed normalized customer costs. Therefore $p_{\text{large}} = 0.10$. c_{large} is shown as the vertical dotted black line in Figure 3. Note that the event cost $c_{\text{large}} = 0.59$ USD is normalized per customer served and corresponds to an event cost of 1.8×10^5 USD.

The probability p_{large} that an event has large cost does not account for the frequency of events. Therefore it is also useful to define the annual frequency of large-cost events f_{large} by multiplying p_{large} by the average annual event rate E_{rate}:

$$f_{\text{large}} = p_{\text{large}}\, E_{\text{rate}}$$

For our utility data, $E_{\text{rate}} = 623$ events per year, so that $f_{\text{large}} = 62.3$ events per year.

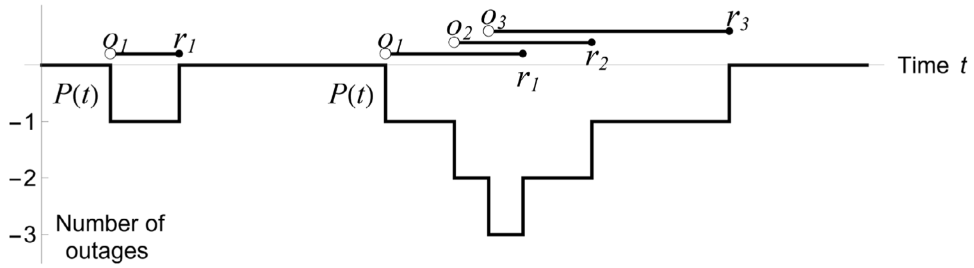

FIGURE 2 | Event with one outage and an event with 3 outages. Each outage's start time (open circle) and restore time (dot) are shown above the time axis. Below the time axis is the performance curve $P(t)$ for each event. The event ends when $P(t)$ returns to zero. Licensed under CC by 4.0 from [13].

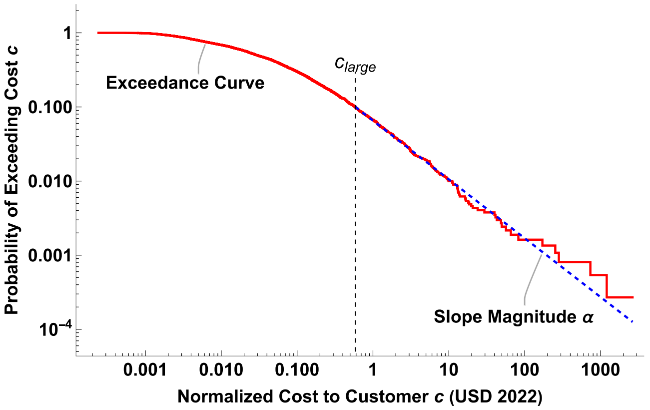

FIGURE 3 | Customer cost exceedance curve, fitted tail distribution, and large-cost threshold c_{large} on a log–log scale (Area A).

The probability p_{large} that an event has a large cost can be compared to VaR [12]. For p_{large}, one fixes the threshold cost c_{large}, and then p_{large} is the probability that the cost of an event exceeds c_{large}. For VaR, one fixes a probability p and then VaR is the minimum cost c such that the probability is p that the event cost over a fixed period of time, such as 1 year, exceeds c. So p_{large} differs from VaR, but they both encode information about the right hand portion of the customer cost exceedance curve.

5.2 | Slope of the Cost Exceedance Curve

The customer cost exceedance curve shown in Figure 3 exhibits an approximate straight-line behavior in the right-hand portion of the curve. c_{large} is chosen within this linear region. Since Figure 3 has a log–log scale, the linear tail has an approximate power-law behavior given by:

$$\overline{F}(c) = \left(\frac{c}{c_{\text{large}}}\right)^{-\alpha}, \qquad c \geq c_{\text{large}} \tag{2}$$

To verify the straight-line behavior on the log–log plot, take the logarithm of (2) to obtain $\log \overline{F}(c) = -\alpha \log c + \alpha \log c_{\text{large}}$. The tail slope index α is the absolute slope of the tail above c_{large} on the log–log plot of the exceedance curve, so that α describes a linear trend for the observed risk. A larger value of α gives a steeper tail and improved resilience.

α can be estimated using the maximum likelihood Hill estimator [52]:

$$\alpha = \left[\frac{1}{n_{\text{large}}} \sum_{c \in C_{\text{large}}} \ln \frac{c}{c_{\text{large}}}\right]^{-1} \tag{3}$$

Using (3), we estimate $\alpha = 0.79$, where α is the magnitude of the slope of the customer cost exceedance curve for large costs. (It follows that the corresponding large cost slope magnitude of the probability density function of C on a log–log plot is $\alpha + 1$.) α is calculated using the ratio of costs c/c_{large} and therefore does not depend on the multiplicative scaling of the costs or k or $n_{\text{customers}}$.

The 95% confidence intervals of α are given in Table 2. A formal goodness-of-fit test (using Clauset's [53, Section 4.1] method) gives $p = 0.611$ with 2500 samples, hence the plausibility of a Pareto (straight-line) tail is not rejected at $p = 0.1$. This further substantiates the linear approximation to the tail on the log–log plot in Figure 3. A similar goodness-of-fit test with c_{large} as the cutoff gives $p = 0.865$ with 10,000 bootstrap samples, and thus also fails to reject the null hypothesis of Pareto fit above c_{large} at $p = 0.1$.

5.3 | Average Log Event Cost ALEC

We define the average log event cost or ALEC metric for large cost events as [15].

$$\text{ALEC} = \frac{1}{n_{\text{large}}} \sum_{c \in C_{\text{large}}} \log_{10} c \tag{4}$$

The use of the logarithm to base ten shows that ALEC indicates the mean of the order of magnitude of the large event costs. For example, ALEC = 1 indicates costs of mean order of magnitude 1, or 10. Note that ALEC is also equal to $\log_{10}$ of the geometric mean of the large event costs. That is, 10^{ALEC} is the geometric mean of the large event costs.

Using (3) and (4), ALEC can be directly related to α:

$$\alpha = \frac{(\ln 10)^{-1}}{\text{ALEC} - \log_{10} c_{\text{large}}} \tag{5}$$

$$\text{ALEC} = \frac{1}{\alpha \ln 10} + \log_{10} c_{\text{large}} \tag{6}$$

ALEC and the frequency of large cost events f_{large} can be combined together to get another useful metric, the ALCRI Annual Log Cost Resilience Index for large cost events [15]:

$$\text{ALCRI} = \frac{1}{n_{\text{year}}} \sum_{c \in C_{\text{large}}} \log_{10} c = f_{\text{large}}\, \text{ALEC} \qquad (7)$$

The part of the customer cost exceedance curve to the right of c_{large} defines the large-cost event risk and is approximately linear on the log–log plot. This linear region can be described by p_{large}, the value of the exceedance curve at p_{large}, and its slope magnitude α. Since ALEC is directly related to α via (3) and (4), p_{large} and ALEC also describe the large-cost event risk. Changes in p_{large} and ALEC affect the risk of large-cost events differently: Reducing p_{large} moves the exceedance curve vertically downwards, whereas reducing ALEC increases the slope magnitude α to greatly reduce the risk of the very largest cost events and slightly reduce the risk of the events with cost just above c_{large}. ALCRI combines ALEC and f_{large} (which is directly related to p_{large}) to reflect both the changes in probability and slope. It is particularly useful in methods such as optimization, where a single metric reduces complexity compared to using two metrics.

Now we summarize the considerations [15] in choosing the large-cost threshold c_{large}. c_{large} should be chosen in the approximately linear right hand portion of the log–log plot of the exceedance curve. Moreover, there is a tradeoff: a smaller c_{large} includes more large event data, giving a less variable estimates of α and ALEC, whereas a larger c_{large} better describes the trend indicated by the very largest cost events. Choosing c_{large} is a well-known delicate issue in estimating α with the Hill estimator (3) [52].

5.3.1 | ALEC and ALCRI in Plain Terms

Tracking the cost of power outages shows that most outages are small and don't cost much, but a few are massive, such as those caused by a major hurricane, and are very expensive. The costs of these rare, large events can be thousands of times greater than the more common ones. Due to this huge variation in costs, if we try to calculate the average cost of a large blackout over time, the result is erratic whenever one of the most extreme rare events occurs. The average does not converge, and does not give a meaningful value that is representative of the large event risk [15]. It is like trying to find the average wealth in a room that includes a billionaire; the billionaire's wealth makes the average misleading.

This is where the ALEC metric comes in. It handles these massive variations and gets a more stable and practical measure of the risk from large blackouts. Instead of averaging the costs directly, ALEC takes the average of the logarithms of the costs. Using a base-10 logarithm essentially looks at the number of zeros in the cost, or its "order of magnitude." ALEC answers the question: "On average, how many zeros are in the cost of a large outage?". For example, ALEC ≈ 2 implies that each large cost outage typically costs in the hundreds of dollars *per customer served*. An unnormalized ALEC value can be obtained by adding $\log_{10}(n_{\text{customer}})$ to ALEC. For example, if a utility serves 30,000 customers and ALEC $= 0.5$, then unnormalized ALEC will be ≈ 5 which implies that each large cost outage typically costs hundreds of thousands of dollars.

5.3.1.1 | ALEC versus ALCRI. While ALEC indicates the *cost* of one typical large outage, ALCRI indicates the *annual cost* of large outages. Think of ALEC as a "per-event" number, and ALCRI as a "per-year" number. ALCRI accounts for both the size and annual frequency of large cost outages. Higher ALCRI means higher annual burden from large cost outages, and lower ALCRI means lower annual burden from large cost outages. ALCRI can be compared across different years to see if the large-cost event impact on customers is getting better or worse.

5.3.1.2 | Step-by-step calculation of ALEC & ALCRI.

1. Normalize all event costs gathered over a number of years, such as 5 years, by the number of customers served.
2. Select only the large cost events—those above a threshold cost c_{large} (e.g., c_{large} could be fixed at the top 10% most expensive events in one of the years).
3. Take the base-10 logarithm of all the large event costs.
4. Average all the log cost values in step 3 to get ALEC.
5. Sum all the log cost values in step 3 and divide by the number of years to get ALCRI.

5.4 | Heavy Tails and Implications for Metrics

For our data, and as observed for two other distribution systems [15], the tail slope magnitude $\alpha < 1$, indicating a very heavy tail, which is associated with very large variability in the costs of large cost events and severe statistical problems in estimating a mean large events cost [15]. Therefore there is no typical or representative large event cost. Moreover, although the extreme events are rare, their cost is so high that the extreme event risk is greater than the regular event risk. These problems are indicated by the fact that the Pareto distribution (2) has infinite mean when $\alpha < 1$. It is not realistic that (2) has no upper bound of a maximum possible event cost. However, even if such a maximum possible event cost is assumed (such as the cost of a 1 month blackout of the entire distribution system) so that the mean is large and finite, it takes an impractical amount of data to estimate the mean [15].

There are uncertainties in fitting the tails of heavy-tailed distributions [52, 53] such as the customer cost exceedance curve, including whether it is decisively fit by the Pareto distribution (2) and the manner and extent to which it can be reasonably extrapolated beyond the maximum cost event observed. Moreover, further progress is needed in methods of estimating direct and indirect customer costs, beyond our assumption of direct customer costs as proportional to customer hours. Nevertheless, the available evidence so far indicates that, at least for some distribution systems, there is no workable mean cost of large cost events. That is, quantifying the large event risk in ways that use the mean large event cost is not viable. In particular, the optimization of mean event cost and the calculation of Conditional Value at Risk CVaR are not viable. Our new metrics of large event risk are chosen to account for the heavy tails and avoid these problems [15]. Further careful analysis of costs, catastrophic events, and heavy tails is indicated.

6 | Quantifying Investment Benefits by Rerunning History

Having established resilience risk metrics that capture the frequency and severity of large-cost outage events without relying on deficient mean-based estimates, we now explain in this section how these metrics can be used to quantify the benefits of different resilience investments. Specifically, we demonstrate how a utility can use its own historical data to retrospectively assess how resilience improvements, such as infrastructure hardening or faster restoration, would have changed the large event risk. This approach, which we refer to as the historical rerun method, provides a practical and persuasive framework for quantifying the value of investments, thereby supporting better-informed future planning decisions.

In the historical rerun method, we first assess the historical resilience performance of the system using the risk metrics explained earlier. We do this by extracting events from the outage data and calculating base-case metrics of the events. We then consider what would have happened if an investment to upgrade the system had been made several years ago. How would that investment have improved the system's performance over these past years? Using this concept, we modify the outage data to reflect the effects of such an investment (explained further in Sections 6.1 and 6.2) and then recalculate the risk metrics using the modified outage data. Comparing the results with and without the effect of the investment quantifies the impact of that investment. Suppose there are n_{event} events and m_{Bi} and m_{Ai} are the values of a particular metric of the ith event before and after incorporating the effects of an investment. Then the overall benefits of the investment are calculated using that metric as:

$$\text{Percentage Benefits} = \left(\frac{1}{n_{\text{event}}} \sum_{i=1}^{n_{\text{event}}} \frac{m_{Ai} - m_{Bi}}{m_{Bi}} \right) \times 100 \quad (8)$$

The historical rerun method is discussed in more detail in [13].

The historical rerun method is limited in that while it can maintain or decrease the number of outages, it cannot synthesize new outages. Nevertheless, it could still be used in cases where the extreme weather events are expected to become more intense, as long as the benefits of the proposed investments outweigh the expected increase in the severity of the events. For instance, if the average wind speeds in an area are projected to rise by x mph and an investment is proposed to harden the infrastructure in that area to withstand y mph higher wind speeds on average, then as long as $y > x$, the historical data and the rerun history method can be used because there will be an overall reduction in the number of outages rather than an increase. Moreover, an expected future increase in the frequency of extreme events is easily accommodated by increasing the average annual event rate E_{rate}.

6.1 | Modeling Benefits of Wind Hardening

We develop the area outage rate curve [13] using the outage data and weather data. The area outage rate curve gives the empirical average outage rate of an area of the distribution system as a function of wind speed. The empirical area outage rates for the utility data we analyze can be fit by an exponential function [13] which is shown in Figure 4.

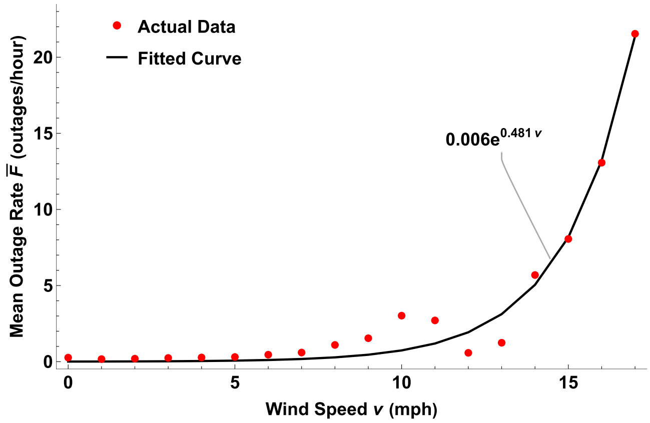

FIGURE 4 | Area outage rate curve of area A. Dots are the empirical mean outage rate at each wind speed, and the curve is an exponential fit of the empirical data.

Wind hardening has the effect of shifting the area outage rate curve in Figure 4 to the left, which means we would observe lower outage rates at each wind speed as compared to the outage rates before the wind hardening. Based on this understanding, we shift the area outage rate curve to the left to represent a wind-hardening investment which gives a 10% decrease in outage rates. To capture the variation in results as different outages are removed to decrease the outage rate, we take 2000 samples of reduced outages randomly from the outage data according to the reduced outage rates. We compute the risk metrics for each sample, and then take the average over the 2000 samples to obtain the average risk metrics resulting from the investment.

6.2 | Modeling Benefits of Faster Restoration

While investments can be made for hardening the infrastructure, investments can also be made to improve the restoration of outages. We also model such investments using the historical rerun technique [13]. If investments had been made to acquire more repair crews, better stocks of spare parts, and better route scheduling, then the restoration rates of the outage events would have improved, resulting in the earlier completion of the restorations. Let $c_{\text{faster}} < 1$ be the factor by which the restoration duration of outage events is reduced after the implementation of an appropriate investment. Then the restore duration of the kth restore $(r_k - r_1)$ in an event is reduced by a factor of c_{faster}, as long as the new restore time occurs after its corresponding outage. The new restore time of the kth restore r_k^{new} is then calculated by:

$$r_k^{\text{new}} = \max\{r_1 + (r_k - r_1)c_{\text{faster}}, o_{\pi(k)}\},\ k = 1, \ldots, n. \quad (9)$$

Here, r_1 is the time of the first restore, and $o_{\pi(k)}$ is the occurrence time of the outage that is restored in the kth restore. To demonstrate the effects of faster restoration investments on the risk metrics, we assume an investment that would have resulted in a 10% faster restoration. We thus assume $c_{\text{faster}} = 0.9$ and rerun history by updating the restoration times of outages in the data

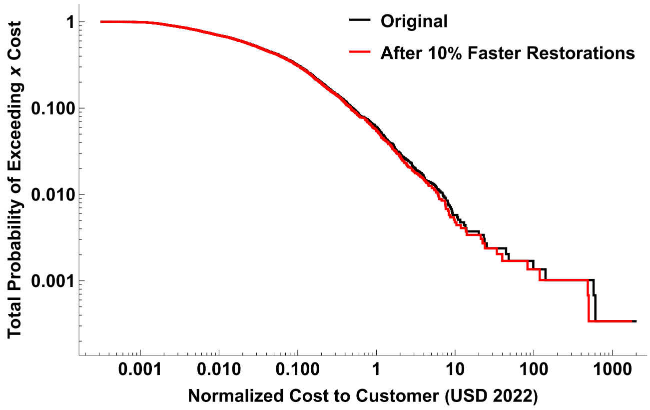

FIGURE 5 | Effect of 10% faster restoration on the customer cost exceedance curve (Area B).

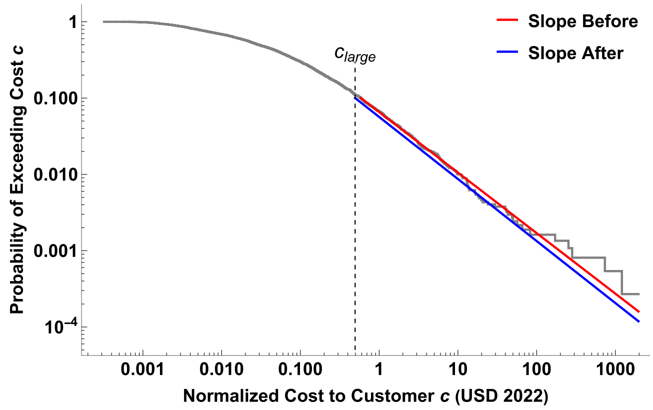

FIGURE 6 | Comparison of tail slopes of customer cost exceedance curve before and after 10% faster restoration (Area A).

using (9). Faster restorations in events decrease the cost of events and thus decrease the risk, as shown in Figure 5 for area B.

6.3 | Discussion of Rerunning History

The historical rerun method has several advantages:

- It is driven by real historical data, which includes all the factors that are very difficult to capture in models, such as weather, human factors, emergency system reconfigurations, equipment aging, and restoration practices.
- Model-based approaches can be used for predicting future behaviors, but have considerable uncertainties, such as uncertainties in parameters and modeling omissions and approximations. Using historical data removes many of these uncertainties.
- It is easy to communicate the benefits of the proposed resilience investment to stakeholders, as those benefits would have applied to the already lived experience of stakeholders, particularly for large events.
- It is computationally inexpensive and easy to implement, and the outage data is already available to utilities.
- In addition to distribution systems, it can easily be used in transmission systems to quantify the benefits of investments for both individual events and groups of events in a region [54].

The main limitation of the historical rerun method relative to detailed models relates to predicting the future. Detailed infrastructure models have many assumptions, but can describe with more detail and flexibility changes to the system and its environment to predict the future. Whereas historical rerun has more limited capabilities: it can only account for more frequent wind events and, to a limited extent explained in Section 6, more severe wind events.

The historical rerun requires some standard engineering practice specifying the effectiveness and cost of proposed investments. The 10% decrease in outage rates or the 10% faster restoration considered in this paper can be achieved in different ways with different types of investments. Some of the investments can even yield both a decrease in outage rates and faster restorations. Types of investments available vary significantly from one utility to the other and are influenced by factors like cost, geographical location, dominant vulnerabilities, reliability and maintenance policies, and regulations. Utilities have extensive expertise in this engineering, which allows them to translate specific investments into estimated improvements in outage rates and restoration times. This paper shows how to take those estimates of percentage improvements and calculate the resilience benefits of the specific investments with metrics by rerunning history.

7 | Results of Resilience Investments

The effects of wind hardening investments on risk metrics are shown in Table 1. As a result of the 10% wind hardening investment, the probability p_{large} that an event has a customer cost more than 1.8×10^5 USD is reduced by almost 2%. Since the wind hardening decreases the number of outages, the annual event rate E_{rate} is also decreased by more than 6%. Consequently, the expected annual frequency f_{large} of large-cost events also decreases by nearly 8%.

The effects of faster restoration are also shown in Table 1. Investments made for improving the restoration rate of events by 10% would have resulted in a significant 7% decrease in the probability p_{large} that an event has a large customer cost and the same 7% decrease in the annual frequency of large-cost events f_{large}. In other words, there would have been 7% less chance that an outage event costing customers 1.8×10^5 USD or higher would occur if such an investment had been made. The percentage reductions in p_{large} and f_{large} are the same because faster restoration does not affect the number of events n_{event} or the annual rate of events E_{rate}. As explained earlier, the exceedance curve shifts downward when p_{large} decreases, and this can be seen in Figure 6 comparing the slope of the exceedance curve before and after 10% faster restoration. Although the downward shift is more at the right end of the exceedance curve because the slope magnitude α has also increased by 3.68%.

One detail about the faster restoration modeling is that events with only one outage are not affected by it as their restoration

TABLE 1 | The effects of historical rerun with 10% wind hardening and 10% faster restoration in Area A.

Metric	Base value	After 10% hardening		After 10% faster restoration	
		Value	Difference	Value	Difference
ALEC	0.315	0.319	1.04%	0.296	−6.16%
α	0.793	0.789	−0.58%	0.823	3.68%
p_{large}	0.10	0.098	−1.90%	0.093	−7.01%
E_{rate}	622.8	583.7	−6.28%	622.8	0%
f_{large}	62.35	57.32	−8.07%	57.98	−7.01%
ALCRI	19.64	18.29	−6.90%	17.16	−12.62%

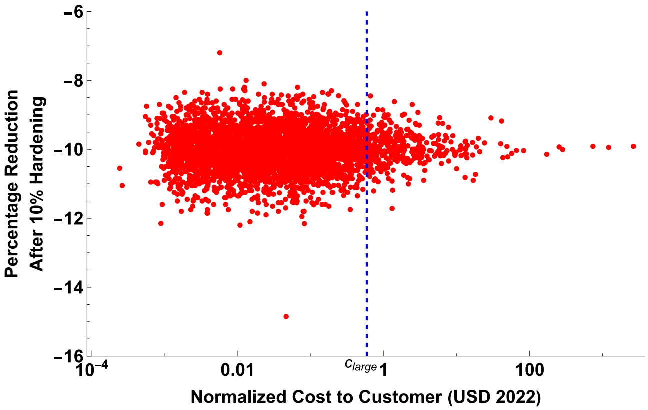

FIGURE 7 | Percentage reduction in the cost of each event after 10% Wind Hardening (Area A).

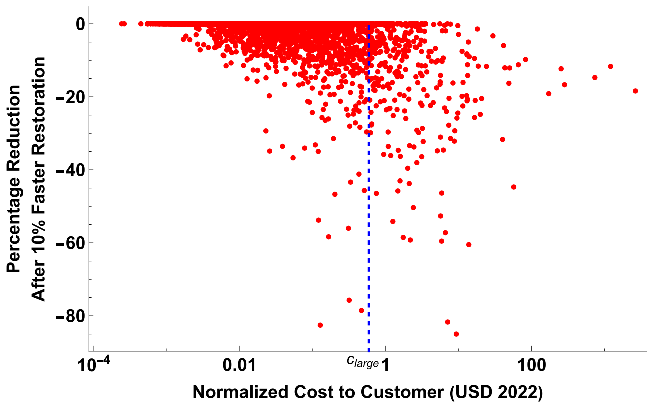

FIGURE 8 | Percentage reduction in the cost of each event after 10% Faster Restoration (Area A).

process, which starts and ends with the first restore, has zero duration. There are 2142 such events in the utility data and their costs remain the same after the faster restoration modeling.

We note in Table 1 that 10% wind hardening slightly increases ALEC. This is because ALEC is the mean of the log of large cost events, as given in (4). ALEC decreases when either the sum of the log of large costs $\sum_{c\in C_{\text{large}}} \log_{10} c$ decreases or the number of large cost events n_{large} increases. The overall effect of any type of investment on ALEC is determined by how much that investment affects each of these two factors. In the case of wind hardening investments, both the sum of large costs and the number of large cost events decrease, but the decrease in the number of large cost events (−8% on average), is more than the decrease in the sum of their costs (−7% on average). Therefore, the value of ALEC slightly increases. However, we see in Tables 1 and 3 that 10% hardening results in a significant decrease in the frequency of large events f_{large} because hardening reduces the number of events overall. In the case of faster restoration investments, the number of large cost events decreases less than the decrease in the sum of large cost events (−7% vs. −13%), which results in an overall decrease in ALEC, as shown in Table 1.

Both the hardening and faster restoration decrease the customer cost of outage events as shown in Figures 7 and 8. However, they do it in different ways: hardening decreases the cost by decreasing the number of outages in events, whereas faster restoration decreases the cost by decreasing the duration of outages in events.

TABLE 2 | Parameters of Area A and Area B.

Parameter	Area A	Area B
n_{event}	3706	2944
c_{large}	0.586	0.495
n_{large}	371	295
ALEC	0.315	0.209
95% CI of ALEC	(0.260, 0.371)	(0.151, 0.268)
α	0.793	0.844
95% CI of α	(0.715, 0.876)	(0.750, 0.943)

Due to this, 10% hardening results in a relatively uniform 10% decrease on average in the cost of all large-cost events (as shown in Figure 7) as compared to 10% faster restoration which decreases the cost of many (41%) of the large cost events by more than 10%, as shown in Figure 8.

8 | Results of Area B

Area B has 7876 outages that are grouped into 2944 outage events. Table 2 compares different details of area A and area B, while Table 3 contains the area B results of rerunning history with 10% wind hardening and 10% faster restoration. The exceedance curve of area B is shown in Figure 9. Goodness-of-fit test for Area B

TABLE 3 | The effects of historical rerun with 10% wind hardening and 10% faster restoration in Area B.

Metric	Base Value	After 10% hardening		After 10% faster restoration	
		Value	Difference	Value	Difference
ALEC	0.209	0.212	1.13%	0.202	−3.46%
α	0.844	0.840	0.44%	0.856	1.43%
p_{large}	0.10	0.098	−2.08%	0.094	−5.76%
E_{rate}	495.1	461.5	−6.79%	495.1	0%
f_{large}	49.61	45.28	−8.73%	46.75	−5.76%
ALCRI	10.37	9.60	−7.42%	9.44	−8.92%

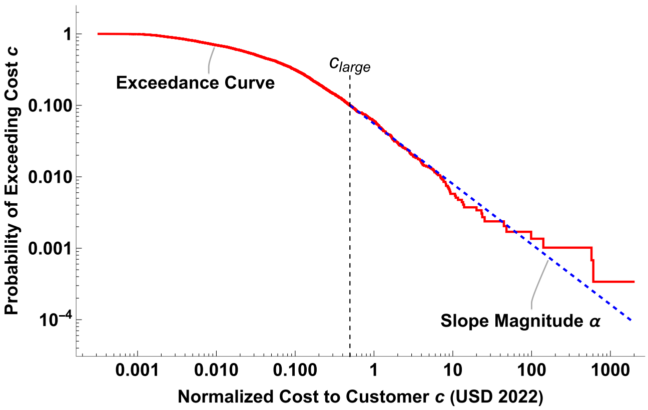

FIGURE 9 | Customer cost exceedance curve of Area B, along with fitted tail distribution, and large-cost threshold on a log–log scale.

data with c_{large} as the cutoff gives $p = 0.077$ with 10,000 bootstrap samples, and thus fails to reject the null hypothesis of Pareto fit above c_{large} at $p = 0.05$.

A comparison of the results of area A and area B (Table 1 vs Table 3) shows similar trends in the improvements in the risk metrics due to 10% wind hardening and 10% faster restoration. Particularly, the percentage improvements due to 10% wind hardening are approximately the same in both areas.

9 | Conclusions

Fundamental to our resilience analysis of distribution system utility data is grouping observed outages into events in which outages accumulate before they are restored. Events of all sizes are easily extracted from utility outage data [11]. It is straightforward to evaluate the customer hours lost for each event. Then we calculate a customer cost for each event that is proportional to the customer hours and normalized by the number of customers served by the utility.

We define risk by a customer cost exceedance curve [22]; that is, the probability that the customer cost of an event exceeds a given amount. We define large cost events as events with costs exceeding a threshold value c_{large}. Then p_{large} is the probability that an event has cost greater than c_{large} as well as the value of the exceedance curve at c_{large}. We also obtain the annual frequency of large-cost events f_{large} by multiplying p_{large} by the average annual event rate E_{rate}.

For the larger costs in our utility data, the exceedance curve is approximately linear with slope magnitude α on a log–log plot. It follows that the exceedance curve above c_{large} and hence the risk of large events is described by its value p_{large} together with the slope magnitude α. Directly related to the slope magnitude α is the ALEC metric, which is simply the Average of the Logarithm of the large Event Costs. We propose the probability of large cost events p_{large} or the annual frequency of large cost events f_{large} together with the slope magnitude α or ALEC as novel large event risk metrics. Moreover, the annual frequency f_{large} and ALEC can be multiplied together to obtain the novel ALCRI Annual Log Cost Resilience Index, which is simply the yearly sum of the logarithms of the large event costs. This new formulation of extreme event risk in distribution systems incorporates both the probability and cost of extreme events and solves the statistical problems of heavy-tailed data.

Our utility data shows customer cost exceedance curves with heavy tails, showing that large cost events will occasionally happen and have substantial risk. Further, the heaviness of the tail (slope magnitude $\alpha < 1$) and the consequent expected occurrence of catastrophic events with highly variable and large costs strongly suggest that using average or mean values to characterize these large costs in distribution systems is not workable [15]. The costs of the large cost events are so variable that there is no typical or representative large cost. These problems are avoided by the new risk metrics p_{large} or f_{large} together with α or ALEC. ALEC takes the logarithm of large costs before taking the mean. Taking the logarithm transforms the heavy tail of the exceedance curve into a light-tailed distribution with a mean that can be practically calculated in the usual way.

While the customer cost distribution must be individually checked in each proposed application, we are optimistic that outage data from other distribution systems will show similar very heavy tails in customer cost so that the methods and conclusions of this paper can apply more broadly. For example, reference [15] shows four other distribution systems in the USA with very heavy tails in customer cost and provides further justifications for the methods and metrics of this paper.

The main reason we develop risk-based metrics in this paper is that risk-based decision-making is a sound and widely accepted approach for evaluating investments. Investments in resilience should account for both the probability of extreme events and the impact of the investment in reducing their impact. In particular, we develop risk metrics to evaluate investments from a risk-to-customer perspective, while ensuring that these risk metrics can be practically calculated in the presence of heavy tails in the cost data.

To summarize our advances in quantifying the risk of large customer cost events:

- We describe resilience risk by processing utility outage data into events, finding the customer cost of each event, and forming the exceedance curve of these event costs.
- The tail of the exceedance curve describes the risk of large cost events. In our data, the exceedance curve has a very heavy tail that makes conventional metrics that depend on calculating the mean unworkable.
- Practical metrics can describe the tail of the exceedance curve. These new metrics are the probability p_{large} that an event has large cost or the annual frequency f_{large} of large events together with either the slope α of the exceedance curve on a log–log plot or, equivalently, the ALEC metric. ALEC is simply the average value of the logarithm of the costs of the large cost events. The ALCRI metric describes both the frequency and size of large cost events with the annual log cost of the large cost events.

Having quantified the risk of large events from utility data, we investigate how investments can mitigate this risk. We use the historical rerun method to quantify the resilience improvement that a proposed resilience investment would have had if the investment had been made in the past.

Future work will include analyzing outage and cost data from additional distribution systems, extending the analysis to utility costs and risk, more specific case studies of hardening and faster restoration, and quantifying other resilience investments such as undergrounding.

Author Contributions

Arslan Ahmad: conceptualization, data curation, formal analysis, investigation, methodology, software, validation, visualization, writing – original draft, writing – review & editing. **Ian Dobson:** conceptualization, formal analysis, funding acquisition, investigation, methodology, project administration, software, supervision, validation, writing – original draft, writing – review & editing.

Acknowledgements

Support in part from Iowa State University EPRC, National Science Foundation (Grants 2153163 and 2429602), PSerc project S110, Argonne National Laboratory, and U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Hydrogen Fuel Cell Technology Office (Award Number DE-EE0010724) is gratefully acknowledged. The views expressed herein do not necessarily represent the views of US DOE or the United States Government.

Open access funding provided by the Iowa State University Library.

Conflicts of Interest

The authors declare no potential conflicts of interest.

Data Availability Statement

The data supporting the findings of this study are confidential.

Endnotes

[1] CVaR is also known as expected shortfall, average value at risk (AVaR), expected tail loss (ETL), and superquantile.

[2] The cost exceedance curve is also known as the survival function or complementary cumulative distribution function (CCDF) or risk curve of C.

References

1. United States Environmental Protection Agency, "Climate Change Indicators in the United States (Fifth ed., EPA 430-r-24-003)," Tech. Rep., accessed May 13, 2025, https://www.epa.gov/system/files/documents/2024-09/climate_indicators_2024.pdf.

2. S. Paul, A. Poudyal, S. Poudel, A. Dubey, and Z. Wang, "Resilience Assessment and Planning in Power Distribution Systems: Past and Future Considerations," *Renewable and Sustainable Energy Reviews* 189 (2024): 113991.

3. A. Poudyal, S. Poudel, and A. Dubey, "Risk-Based Active Distribution System Planning for Resilience Against Extreme Weather Events," *IEEE Transactions on Sustainable Energy* 14, no. 2 (2023): 1178–1192.

4. S. Ma, B. Chen, and Z. Wang, "Resilience Enhancement Strategy for Distribution Systems Under Extreme Weather Events," *IEEE Transactions on Smart Grid* 9, no. 2 (2018): 1442–1451.

5. D. A. Reed, "Electric Utility Distribution Analysis for Extreme Winds," *Journal of Wind Engineering and Industrial Aerodynamics* 96, no. 1 (2008): 123–140.

6. H. Liu, R. A. Davidson, and T. V. Apanasovich, "Spatial Generalized Linear Mixed Models of Electric Power Outages due to Hurricanes and Ice Storms," *Reliability Engineering & System Safety* 93, no. 6 (2008): 897–912.

7. S. R. Han, "Estimating Hurricane Outage and Damage Risk in Power Distribution Systems" (Ph.D. dissertation, Texas A&M University, 2008), https://www.proquest.com/dissertations-theses/estimating-hurricane-outage-damage-risk-power/docview/304333596/se-2.

8. L. Xu, N. Lin, D. Xi, K. Feng, and H. V. Poor, "Hazard Resistance-Based Spatiotemporal Risk Analysis for Distribution Network Outages During Hurricanes," *IEEE Transactions on Power Systems* 40, no. 3 (2025): 2143–2152.

9. S. D. Guikema, S. M. Quiring, and S.-R. Han, "Prestorm Estimation of Hurricane Damage to Electric Power Distribution Systems," *Risk Analysis* 30, no. 12 (2010): 1744–1752.

10. R. Billinton and G. Singh, "Application of Adverse and Extreme Adverse Weather: Modelling in Transmission and Distribution System Reliability Evaluation," *IEE Proceedings-Generation, Transmission and Distribution* 153, no. 1 (2006): 115–120.

11. N. K. Carrington, I. Dobson, and Z. Wang, "Extracting Resilience Metrics From Distribution Utility Data Using Outage and Restore Process Statistics," *IEEE Transactions on Power Systems* 36, no. 6 (2021): 5814–5823.

12. S. Poudel, A. Dubey, and A. Bose, "Risk-Based Probabilistic Quantification of Power Distribution System Operational Resilience," *IEEE Systems Journal* 14, no. 3 (2020): 3506–3517.

13. A. Ahmad and I. Dobson, "Towards Using Utility Data to Quantify how Investments Would Have Increased the Wind Resilience of Distribution Systems," *IEEE Transactions on Power Systems* 39, no. 4 (2024): 5956–5968.

14. A. Ahmad and I. Dobson, "Quantifying Distribution System Resilience From Utility Data: Large Event Risk and Benefits of Investments," in

CIRED Workshop on Resilience of Electric Distribution Systems (IEEE, 2024).

15. A. Ahmad and I. Dobson, “Logarithmic Resilience Risk Metrics That Address the Huge Variations in Blackout Cost,” *IEEE Transactions on Power Systems* 40, no. 6 (2025): 5507–5510.

16. T&D Committee of IEEE PES, “IEEE Guide for Electric Power Distribution Reliability Indices,” IEEE Std 1366-2022 (Revision of IEEE Std 1366-2012), 1–44 (2022).

17. R. Billinton and R. N. Allan, *Reliability Evaluation of Power Systems*, 2nd ed. (Plenum Press, 1996).

18. W. Li, *Risk Assessment of Power Systems: Models, Methods, and Applications* (John Wiley & Sons, 2014).

19. R. Billinton and P. Wang, “Teaching Distribution System Reliability Evaluation Using Monte Carlo Simulation,” *IEEE Transactions on Power Systems* 14, no. 2 (1999): 397–403.

20. R. Billinton and Z. Pan, “Historic Performance-Based Distribution System Risk Assessment,” *IEEE Transactions on Power Delivery* 19, no. 4 (2004): 1759–1765.

21. M. Keogh and C. Cody, *Resilience in Regulated Utilities* (National Association of Regulatory Utility Commissioners, 2013), https://pubs.naruc.org/pub/536f07e4-2354-d714-5153-7a80198a436d.

22. S. Kaplan and B. J. Garrick, “On the Quantitative Definition of Risk,” *Risk Analysis* 1, no. 1 (1981): 11–27.

23. X. Chen, J. Qiu, L. Reedman, and Z. Y. Dong, “A Statistical Risk Assessment Framework for Distribution Network Resilience,” *IEEE Transactions on Power Systems* 34, no. 6 (2019): 4773–4783.

24. P.-C. Chen and M. Kezunovic, “Fuzzy Logic Approach to Predictive Risk Analysis in Distribution Outage Management,” *IEEE Transactions on Smart Grid* 7, no. 6 (2016): 2827–2836.

25. Y. Du, Y. Liu, X. Wang, J. Fang, G. Sheng, and X. Jiang, “Predicting Weather-Related Failure Risk in Distribution Systems Using Bayesian Neural Network,” *IEEE Transactions on Smart Grid* 12, no. 1 (2021): 350–360.

26. S. D. Guikema, “Formulating Informative, Data-Based Priors for Failure Probability Estimation in Reliability Analysis,” *Reliability Engineering & System Safety* 92, no. 4 (2007): 490–502, https://www.sciencedirect.com/science/article/pii/S0951832006000160.

27. P. L. Watson, A. Spaulding, M. Koukoula, and E. Anagnostou, “Improved Quantitative Prediction of Power Outages Caused by Extreme Weather Events,” *Weather and Climate Extremes* 37 (2022): 100487, https://www.sciencedirect.com/science/article/pii/S2212094722000664.

28. P. L. Watson, D. Pasqualini, and E. Anagnostou, “A Data-Driven Decision Support Tool for Anticipating Tropical Storm Impacts to the United States Power Grid,” *IEEE Access* 12 (2024): 112 905–112 923.

29. D. Wanik, E. Anagnostou, B. Hartman, M. Frediani, and M. Astitha, “Storm Outage Modeling for an Electric Distribution Network in Northeastern USA,” *Natural Hazards* 79 (2015): 1359–1384.

30. D. Cerrai, M. Koukoula, P. Watson, and E. N. Anagnostou, “Outage Prediction Models for Snow and Ice Storms,” *Sustainable Energy, Grids and Networks* 21 (2020): 100294, https://www.sciencedirect.com/science/article/pii/S2352467719302668.

31. J. B. Leite, J. R. Sanches Mantovani, T. Dokic, Q. Yan, P.-C. Chen, and M. Kezunovic, “The Impact of Time Series-Based Interruption Cost on Online Risk Assessment in Distribution Networks,” in *2016 IEEE PES Transmission & Distribution Conference and Exposition-Latin America (PES T&D-LA)* (IEEE, 2016), 1–6.

32. J. B. Leite, J. R. S. Mantovani, T. Dokic, Q. Yan, P.-C. Chen, and M. Kezunovic, “Resiliency Assessment in Distribution Networks Using GIS-Based Predictive Risk Analytics,” *IEEE Transactions on Power Systems* 34, no. 6 (2019): 4249–4257.

33. C. Lin, F. Liu, L. Zhang, G. Li, C. Chen, and Z. Bie, “An Online Data-Driven Risk Assessment Method for Resilient Distribution Systems,” *CPSS Transactions on Power Electronics and Applications* 6, no. 2 (2021): 136–144.

34. R. Moreno, M. Panteli, P. Mancarella, H. Rudnick, T. Lagos, A. Navarro, F. Ordonez, and J. C. Araneda, “From Reliability to Resilience: Planning the Grid Against the Extremes,” *IEEE Power and Energy Magazine* 18, no. 4 (2020): 41–53.

35. E. D. Vugrin, A. R. Castillo, and C. A. Silva-Monroy, “Resilience Metrics for the Electric Power System: A Performance-Based Approach,” Tech. Rep. (Sandia National Lab., 2017).

36. K. Murray and K. R. W. Bell, “Wind Related Faults on the GB Transmission Network,” in International Conference on Probabilistic Methods Applied to Power Systems (PMAPS) (2014), 1–6.

37. M. Sullivan, “Estimated Value of Service Reliability for Electric Utility Customers in the United States,” Tech. Rep. (Lawrence Berkeley National Laboratory, 2009), https://escholarship.org/uc/item/5hg5h4zw.

38. S. Pandey, G. Kandaperumal, A. Ahmad, and I. Dobson, “Resiliency Metrics Quantifying Emergency Response in a Distribution System,” in *IEEE Power and Energy Society General Meeting* (IEEE, 2025), https://arxiv.org/abs/2501.06030.

39. B. A. Carreras, D. E. Newman, and I. Dobson, “North American Blackout Time Series Statistics and Implications for Blackout Risk,” *IEEE Trans. Power Systems* 31, (11 2016): 4406–4414.

40. I. Dobson, “Models, Metrics, and Their Formulas for Typical Electric Power System Resilience Events,” *IEEE Transactions on Power Systems* 38, no. 6 (2023): 5949–5952.

41. M. Sullivan, M. T. Collins, J. Schellenberg, and P. H. Larsen, “Estimating Power System Interruption Costs: A Guidebook for Electric Utilities,” Lawrence Berkeley National Laboratory, Tech. Rep., report #LBNL-2001164 (Lawrence Berkeley National Laboratory, 2018), https://escholarship.org/uc/item/0xn4r200.

42. “ICE Calculator,” accessed May 5, 2025, https://icecalculator.com/.

43. “Customer Damage Function Calculator,” accessed May 5, 2025, https://cdfc.nrel.gov/.

44. M. Macmillan, K. Wilson, S. Baik, J. P. Carvallo, A. Dubey, and C. A. Holland, “Shedding Light on the Economic Costs of Long-Duration Power Outages: A Review of Resilience Assessment Methods and Strategies,” *Energy Research & Social Science* 99 (2023), 103055.

45. T. Schröder and W. Kuckshinrichs, “Value of Lost Load: An Efficient Economic Indicator for Power Supply Security? A Literature Review,” *Frontiers in Energy Research* 3 (2015): 55.

46. M. Sullivan, J. Schellenberg, and M. Blundell, “Updated Value of Service Reliability Estimates for Electric Utility Customers in the United States,” Tech. Rep., report #LBNL-6941E (Lawrence Berkeley National Laboratory, 2015), https://escholarship.org/uc/item/10r8d6xg.

47. K. H. LaCommare, J. H. Eto, L. N. Dunn, and M. D. Sohn, “Improving the Estimated Cost of Sustained Power Interruptions to Electricity Customers,” *Energy* 153, (2018): 1038–1047.

48. P. H. Larsen, K. S. Kristina S. H., J. Eto, and A. Sanstad, “Frontiers in the Economics of Widespread, Long-Duration Power interruptions: Proceedings From an Expert Workshop,” Tech. Rep. (Lawrence Berkeley National Laboratory, 2019), https://escholarship.org/uc/item/8c8280md.

49. A. Ghodeswar, M. Bhandari, and B. Hedman, “Quantifying the Economic Costs of Power Outages Owing to Extreme Events: A Systematic Review,” *Renewable and Sustainable Energy Reviews* 207 (2025): 114984.

50. S. Baik, A. L. Davis, J. W. Park, S. Sirinterlikci, and M. G. Morgan, “Estimating What us Residential Customers are Willing to pay for Resilience to Large Electricity Outages of Long Duration,” *Nature Energy* 5, no. 3 (2020): 250–258.

51. D. Lineweber and S. McNulty, “The Cost of Power Disturbances to Industrial & Digital Economy Companies,” Electric Power Research Institute, Tech. Rep. (Electric Power Research Institute, 2001), https://www.epri.com/research/products/3002000476.

52. J. Nair, A. Wierman, and B. Zwart, *The Fundamentals of Heavy Tails: Properties, Emergence, and Estimation*, ser. Cambridge Series in Statistical and Probabilistic Mathematics (Cambridge University Press, 2022).

53. A. Clauset, C. R. Shalizi, and M. E. J. Newman, "Power-law Distributions in Empirical Data," *SIAM Review* 51, no. 4 (2009): 661–703.

54. A. Ahmad, I. Dobson, S. Ekisheva, C. Claypool, and M. Lauby, "The Improvement in Transmission Resilience Metrics From Reduced Outages or Faster Restoration can be Calculated by Rerunning Historical Outage Data," in *IEEE Power and Energy Society General Meeting* (IEEE, 2025), https://arxiv.org/abs/2501.06042.