Large Blackouts in North America: Historical trends and policy implications

Paul Hines[a,*], Jay Apt[b], Sarosh Talukdar[c]

[a] *School of Engineering, 301 Votey Hall, University of Vermont, 33 Colchester Ave., Burlington, VT 05405, USA*

[b] *Carnegie Mellon Electricity Industry Center, Department of Engineering & Public Policy and Tepper School of Business, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

[c] *Carnegie Mellon Electricity Industry Center, Department of Engineering & Public Policy and Department of Electrical and Computer Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

---

* Corresponding author. Tel: +18026569669; fax: +18026563358. *E-mail address:* paul.hines@uvm.edu (P. Hines)

Running Head: Trend Analysis of Large Blackouts in North America

Abstract

Using data from the North American Electric Reliability Council (NERC) for 1984-2006, we find several notable trends. We find that the frequency of large blackouts in the United States has not decreased over time, that there is a statistically significant increase in blackout frequency during peak hours of the day and during late summer and mid winter months (although non-storm-related risk is nearly constant through the year) and that there is strong statistical support for the previously observed power-law statistical relationship between blackout size and frequency. We do not find that blackout sizes and blackout durations are significantly correlated. These trends hold even after controlling for increasing demand and population and after eliminating small events, for which the data may be skewed by spotty reporting. Trends in blackout occurrences, such as those observed in the North American data, have important implications for those who make investment and policy decisions in the electricity industry. We provide a number of examples that illustrate how these trends can inform benefit-cost analysis calculations. Also, following procedures used in natural disaster planning we use the observed statistical trends to calculate the size of the 100-year blackout, which for North America is 186,000 MW.

Keywords: Electricity reliability, power system blackouts, cascading failures, policy analysis

1. Introduction

Recent large electricity disruptions in North America (e.g. 14 Aug. 2003, 12 Sept. 2005) and Europe (e.g. 28 Sept. 2003, 4 Nov. 2006) highlight the extent to which urban societies depend on reliable electricity infrastructure. Understanding what trends exist in the history of electricity infrastructure failures is essential to the process by which policy makers choose among priorities for improving this and other infrastructures. After a major blackout, policy makers typically commission detailed analyses of the specific events that preceded a failure in electricity service (UTCE, 2006; USCA, 2004). Sometimes these event-focused reports result in substantive policy changes aimed at forestalling future failures; the creation of a US "Electricity Reliability Organization" by the US congress (EPACT, 2005) is one example. However, a narrow focus on one particular blackout can obscure longer time-trends that can have important implications for policy and investment decisions. The goals of this paper are (1) to determine what trends exist (or do not exist) in the available historical record of large blackouts[1] in the United States and (2) to show how these trends are relevant to policy analysis within the electricity industry.

Data classification and research hypotheses

To clarify this discussion it may be helpful to define "blackout" and the measures used in this paper to classify blackouts. In this paper, a blackout is any unplanned disruption of electricity service to multiple customers that lasts more than 5 minutes. Shorter disruptions are commonly considered power quality events, and are therefore not considered here. We use "large blackout" to describe events that result in service disruptions to at least 50,000 customers or 300 MW of demand. This definition is based on the regulatory reporting requirements in the United States. Each blackout in our data is categorized in three dimensions: size, time, and cause. Size is measured in both MW and customers affected. Time is measured in the start time of the event and its duration. There are two types of blackout causes. External causes, such as storms, vandalism, or operator error, are the causes that initiate the blackout. In some blackouts, a set of external initiating events triggers a sequence of subsequent component outages, known as a cascading failure. Where the reports indicate that a blackout was propagated, at least in part, by cascading failures, these events are noted as caused by cascading failure. Some blackouts thus have multiple causes. Section 2 discusses the data and classification process in more detail.

After classifying our data in the above three dimensions we test the following trend-related hypotheses:

1. The data show an observable decrease in the frequency of large blackouts, after adjusting for demand growth and population increases.
2. There are no seasonal trends in the data. (Blackout probability does not change with time of year.)
3. There are no time of day trends in the data. (Blackout probability does not change with time of day.)
4. The fit between the blackout data and a power-law cumulative probability distribution is significantly better than the fit to an exponential distribution for large blackouts.
5. The data show a positive, significant correlation between blackout size and blackout duration.

The rationale for each hypothesis is discussed briefly in Section 3.

Robust trends in the history of large blackouts can be valuable inputs to the investment and policy decision-making process. A power-law relationship between event size and probability (Hypothesis 4), as previously observed by Carreras et al. (2000; 2004) and Talukdar et al. (2003), indicates that it is important to design electricity infrastructure to be robust to large failures rather than focusing narrowly on small failures. Section 4 provides several calculations illustrating how the observed trends can impact the benefits and costs of investment and policy options.

Related research results

Several recent papers note useful patterns in the North American blackout data. Carreras et al. (2000; 2004) show that large blackout sizes follow a power-law probability density function (pdf). Talukdar et al. (2003) show that the data fit a power-law statistic far better than they do to an exponential (Weibull) pdf. Carreras et al. (2004) argue that time-correlations in the blackout data (using the Hurst parameter, which measures auto-correlation over multiple time-scales) give evidence of self-organized criticality, providing a plausible explanation for the power-law tail. While it may be possible to find other explanations for the power-law frequency distribution (see Newman, 2005, for a discussion of power-law generating mechanisms), there is some agreement that the data fit well to a power-law statistic. However the existing research provides little evidence for the statistical significance of this relationship. Clauset et al. (2009) report only moderate statistical support for the power-law conclusion. Using a more comprehensive and carefully filtered, data set than what has been reported in past studies, the results reported here provide strong statistical support for the previously reported power-law frequency distribution.

Simonoff et al. (2007) study the blackout data available from NERC (the North American Electric Reliability Corporation) for 1990-2004 from the perspective of assessing the risk associated with a terrorist attack. The authors build several regression models based on these data, which provide some evidence relevant to the hypotheses proposed above. In related work, Greenburg et al. (2007) use results from an economic model of terrorist-initiated blackouts to show the value of targeted investments that improve the resiliency of a power grid.

We use a more extensive data set (1984-2006) than these existing studies and filter the data in several ways to remove artifacts that could lead to misleading conclusions. Specifically, we control for demand growth, supply shortages, extreme natural events and the spotty reporting of small events. The removal of small events—those less than the 50,000 customer or 300MW thresholds that trigger federally mandated reporting in the United States—is particularly important if we are to be confident that the observed trends are not artifacts of selective reporting. We then suggest applications of the revealed trends to decision-making and policy problems.

The largest blackouts tend to be the result of either extreme natural events (hurricanes, ice storms, etc.) or cascading failures (see Table 1). Cascading failures are the subject of a rapidly growing body of research. Baldick et. al. (2008) provide a thorough review of the research on cascading failures. Ren and Dobson (2008) study a nine-year time series of branch outage data for one utility to develop a risk model for a specific region. Dobson et al. (2005) and Chen et al. (2005) describe probabilistic models of cascading failure risk. Trends indentified in this paper may be useful to refine these blackout risk assessment tools.

This paper is organized as follows. Section 2 describes the data that are used in this study. Section 3 describes the trends that are found (or not found) in these data. Section 4 provides some example calculations illustrating how these trends affect decision-making and Section 5 presents our conclusions.

2. The NERC disturbance data for 1984-2006

Both the US Department of Energy (DOE) and the North American Electric Reliability Council (NERC) require that organizations submit reports when sufficiently large disturbances occur within their territories. DOE publishes the resulting data as "Form 417" reports, and NERC provides the data through the Disturbance Analysis Working Group (DAWG) database and “System Disturbance Reports.” By DOE regulations, utilities and other load serving entities must report all disturbances that interrupt more than 300 MW or 50,000 customers (DOE, 2007). Some smaller disturbances are also included in the DOE and NERC reports. Until the establishment of NERC as the US “Electricity Reliability Organization” (its rules went into effect in 2007), NERC had limited ability to enforce its reporting rules during the period that we study. However, because of the DOE reporting requirements and the resulting publically available reports, NERC has had access to the records that utilities submit to DOE in addition to the information collected directly by NERC. As a result there is substantial agreement between the two sources, particularly for blackouts larger than 300 MW. Of the 148 EIA records (2000-2006) with recorded sizes greater than or equal to 300 MW, 19 (13%) do not include a corresponding report in the NERC records. Sixteen of these are storm-related events that appear to be isolated to the distribution infrastructure of a single utility, and are thus outside of NERC’s primary jurisdiction. The three EIA records that do not fall into this category affected a small number of customers within a single service territory (4 customers on Sept. 7, 2003, 1 customer on Nov. 5, 2003, and 940 on May 3, 2004). The analysis presented on this paper is based on the NERC reports because they cover a longer time period and include more detailed reports, allowing for more careful classification of the event categories.

There are 933 event reports in the 1984-2006 NERC records. In some of these reports, multiple entries from different organizations refer to a single large blackout. For example, the August 14, 2003 event spans six reports. In order to accurately record these blackouts, we combine these types of multiple reports into a single event record. After combining reports, 856 events remain. Of these, 418 events were both smaller than 300 MW and affected fewer than 50,000 customers. Since small event reporting is largely optional and spotty, this analysis focuses on the larger events. Table 2 lists descriptive statistics for these data with and without the smaller events.

Disturbance categories and locations

Disturbances recorded in the NERC data proceed from a wide variety of triggering events including natural disasters, storms, human error and mechanical failure. For this analysis we sort disturbances into the following primary cause categories: earthquakes, tornados, hurricanes or tropical storms, ice storms, lightning, wind or rain storms, other cold weather, fire, intentional attack, supply shortage, other external (not-human or equipment) event, equipment failure, operator error, voltage reduction and calls for voluntary demand reduction (the latter two are not blackouts, as discussed below). Some records indicate multiple causes, and are thus put into multiple categories.

Some events were initiated by a natural cause, such as lightning, but grew through a set of cascading failures. Unfortunately, it is difficult to isolate with certainty all cascading failures in these data. Some of the NERC reports give sufficient detail to identify which events are clearly cascading failures, but many lack sufficient detail to determine the extent to which dependant events precipitated from the natural- or human-initiated disturbances. This makes it difficult to calculate precisely the total impact of cascading failures. For example, many reports describing blackouts initiated by lightning do not describe the sequence of events in detail, which could obscure a dependant sequence of switching events worsening the resulting blackout. Still assuming that most large cascading failures are included in the NERC data, these data allow us to calculate a rough upper bound on the historical impact of cascading failures.

To better understand the geographic distribution of the events, we divide the records into those located in the West, Northeast (NE), Midwest (MW), Southeast (SE), Texas (ERCOT) and Hawaii (3 events). Figure 1 shows the relative frequency of blackouts in each cause category, and in each region. This figure shows that, for example and as expected, hurricanes and tropical storms are major sources of blackouts in the Southeast and Texas, but not in other regions. Table 3 provides the category results in tabular form.

Data filtering

In order to ensure the reliability of this trend-analysis, the data are filtered in several ways. First we remove the 102 events in the "voluntary reduction" and "voltage reduction" categories since these do not generally disrupt electricity service and reporting may be unreliable. We also remove the events that have no recorded size (207 events). The term blackout is hereafter used to refer to the remaining 547 events, which did result in service interruption. 258 of these have recorded sizes of 300 MW or greater. 304 are recorded to affect 50,000 or more customers. Second some of the records include event size in either MW or customers, but not both. Of the 258 blackouts that are listed as 300 MW or larger, 65 do not report the number of customers affected. From the text associated with these reports, there is no reason to believe that these events actually affected zero customers. In fact in several cases the DOE data do show the number of customers affected, where NERC reports do not, indicating that the empty value is not an indication of zero customers affected. Similarly, of the 304 blackouts specified to be larger than 50,000 customers, 73 do not include the size in MW. To estimate the missing size data, the missing MW or customer entries were filled according to the US average customers per MW, based on 2006 US DOE Energy Information Agency (EIA) data (300.4 customers per MW).[1] Replacing missing data brings the number of blackouts larger than 300 MW to 307 and those larger than 50,000 customers to 382.

Finally, to avoid underestimating the importance of older blackouts, measured as a percent of demand or customers, we adjust event sizes to control for demand growth and an increase in the total number of customers. This is done by scaling sizes in customers to reflect population growth (from US census data) and by scaling sizes in MW to reflect the annual electricity demand (net electric energy generation from DOE EIA data). Specifically, we scale the data to reflect year-2000 customers and year-2000 MW using the same approach that would be used to adjust for inflation in financial calculations. Given annual demand (D_y) and population (P_y) data, Equations (1) and (2) give the adjusted size in MW (S_{MW}') or customers (S_{cust}').

[1] The value of 300.3 customers per MW comes from taking the EIA reported number of electricity customers in 2006 (140.4 million) and dividing by the electricity demand in 2006 (4.0938 million MWh/8760 hours).

$$S_{MW}' = S_{MW} \frac{D_{2000}}{D_y} \qquad (1)$$

$$S_{cust}' = S_{cust} \frac{P_{2000}}{P_y} \qquad (2)$$

The choice of 2000 as the base year was arbitrary. Choosing a different year would have little to no effect on the outcomes in this paper, since the result is merely to change the numerator in Eqs. (1) and (2), thus affecting all event sizes equally.

This scaling had several minor effects on the data. The adjusted size of blackouts before 2000 increases and the sizes of blackouts after 2000 decreases. The adjusted size of a 300 MW blackout in 1984 becomes 498 MW, whereas the size of a 300 MW blackout in 2006 becomes 279 MW. After this adjustment the number of blackouts larger than 300 (year-2000) MW increases by 10 to 317, and the number affecting 50,000 or more year-2000 customers decreases 373. Table 2 describes the number of events before and after scaling in detail.

Customer interruption frequency (SAIFI) in the NERC data

Given the number of customers interrupted in each blackout, it is possible to calculate the total number of customer interruptions and thus infer the fraction of all blackouts reflected in these data. From the total number of customer interruptions, one can calculate the apparent System Average Interruption Frequency Index (SAIFI), a common reliability measure used in the electricity industry. After adjusting for demand growth, and dividing by the number of electricity customers in the US in the year 2000, the apparent SAIFI from the NERC data is:

$$\text{SAIFI} = \frac{\text{141 M interruptions}}{\text{(23 years)(128 M customers)}} = 0.05 \qquad (3)$$

According to La Commare and Eto (2005) SAIFI in the United States is approximately 1.2 or 1.3. Thus the disturbances in the NERC data represent about 4% of all customer interruptions recorded in the national SAIFI numbers. However, due to reporting requirements, the vast majority of large events is likely to be included in these data, so in the majority of the analysis that follows we separate out the large events.

Regarding the utility of the NERC data for trend analysis

The NERC DAWG data certainly contain inaccuracies and are far from a complete record of all blackouts between 1994 and 2006. While the EIA and NERC data agree in many respects, the reported blackout sizes frequently differ somewhat. Submission of disturbance data to NERC was largely voluntary until 2007, when mandatory reliability rules came into effect in the United States. Submission of disturbance data to DOE is required, but some argue that DOE has not sufficiently enforced its reporting requirements to ensure complete data accuracy. Recent discussions regarding the creation of mandatory reliability rules may explain an observable increase in the frequency of small event reports (see Figure 2). Regarding the completeness of the data, some data are certainly missing from the records. For example there are no records for March-November 1998, which is certainly a data collection error. The 1998 data are dropped from the analysis of Hypotheses 1 and 2, where a bias could result from the data collection error. The remaining hypotheses are unaffected by this issue.

Despite some inaccuracies, there are good reasons to believe that the data are sufficiently accurate to support the trend analyses described in this paper. Even if there has been an uptick in the reporting of small disturbances (those less than 300 MW, for example) very large blackouts are very public events. It is difficult for a utility to neglect the reporting of large blackouts. It is precisely for this reason that we filter the data for large events, a precaution that some past authors have not taken when studying these data (e.g., Simonoff et al., 2007). Also, even if there is a change in reporting over time, there is no reason to believe that this would have any effect on the results pertaining to Hypothesis 2-5. In these cases we are looking at size, monthly, hourly, and duration trends, which would be useful even if the data that exist represent a sparse subset of the actual history of all blackouts.

3. Trends in the blackout data

In the introduction we posit five hypotheses regarding trends in the blackout data. This section describes the methods and results that inform our conclusions regarding these trends. The first three sub-sections focus on time-trends in the data and the remaining two focus on the power-law trend and the relationship between blackout size and restoration time.

Hypothesis 1: Blackout frequency has decreased with time (reject)

The past 25 years have been a period of enormous technological growth. At least some of this growth has affected the electricity industry through increasing information technology in the control and data collection equipment that facilitate the uninterrupted flow of electricity between suppliers and consumers. It thus seems reasonable to hypothesize that there is an observable reduction in the overall frequency of large blackouts in the NERC data.

Figure 2 shows the number of blackouts per year in various size-categories for blackouts affecting more than 50,000 customers, after adjusting event sizes for population growth. The data show a fairly clear increase in the annual blackout frequency after 1998. It appears that some data are missing for 1998, as no events are recorded for February through November. The median frequency before 1998 is 11.5 events/year, whereas it is 23.5 events per year after 1998. Amin (2008) and Simonoff et al. (2007) both note this increase over time. But a good portion of this increase occurs within the smaller event categories. The observed increase could be the result of increased reporting of smaller blackouts. The increase in the frequency of events in the "wind/rain" category may indicate that utilities are increasingly likely to report smaller weather-related events in recent years. To control for this potential increase in reporting we focus on events that are greater than 300 MW after filling in missing values and adjusting for demand growth. Also we remove all events that result from extreme natural events (hurricanes, tropical storms, ice storms, earthquakes, and tornadoes). The resulting blackout frequency is shown in figure 3.

Several statistical tests were employed. Table 4 shows the results of Kolmogorov-Smirnov (K-S) t-tests, which test the hypothesis that two sets of non-Gaussian data (event frequency before and after the mid-point year) come from the same probability distribution. Table 4 also shows the results of correlation tests between blackout frequency and time. Each test is run on the data before and after adjusting for demand and population growth. It is important to note that this treatment decreases the number of recent blackouts above a given threshold and increases the number of older blackouts, thus increasing the likelihood that we would find an apparent decrease in the frequency of large blackouts. Also, because of the apparently missing reports for 1998, the 1998 data are not included in the statistical tests. Adding the 1998 data back in to the data set does not substantially change the conclusions for any of the tests described in Table 4. In most cases a statistically significant (at the α=0.05 level) increase in event frequency is apparent. In no case is a statically significant increase apparent, even after removing blackouts caused by natural disasters (earthquakes, ice storms, tornadoes, and hurricanes/tropical storms).

There are a number of plausible explanations for this observation. Firstly the observed increase is certainly not an artifact of the data filtering that we performed (filling in missing customer/MW counts and scaling the data to adjust for increases in demand and population increases). Table 4 shows these data before and after these treatments, both of which show roughly the same trend (no apparent increase in frequency). Filling in missing values increases the number of events per year slightly (from 10.2 to 12.3 blackouts ≥300 MW per year), but does not result in a decrease by any measure. Scaling to adjust for demand growth reduces the number of recent large blackouts, and increases the number of older large blackouts. This scaling ensures that observed frequency changes are not the result of increases in demand and population growth. Secondly, it is important to note that there is some evidence that reporting of small events may have increased in recent years (since 2000), perhaps as a result of increasing federal oversight of electricity reliability. Figure 2 shows a notable increase in the number of events in the 50,000 to 300,000 customer range, which could be the result of increased reporting. This may explain the dramatic increase in blackout frequency reported by Amin (2008). However, as discussed earlier, it is more difficult to avoid reporting very large blackouts. Given that it is difficult to not report large blackouts and given that DOE reporting requirements have not changed significantly for the largest events, it is less likely that reporting of large blackouts has increased over time.

To conclude we do not find sufficient evidence to conclude that blackout frequency is increasing in time. On the other hand there is sufficient evidence to reject our initial research hypothesis that blackout frequency is decreasing in time. Despite imperfections in the data, Hypothesis 1 can be safely rejected.

Hypothesis 2: Blackout frequency does not change seasonally (reject)

In this section we test the hypothesis that blackout frequency does not change seasonally. To do so, we estimate monthly blackout frequencies using 3-month centered rolling-average windows. Blackout frequency increases substantially during the late summer and mid-winter months, though due to the large variance among years, only some of the statistical tests show that the seasonal trends are significant (Table 5). Removing hurricanes and tropical storms increases the significance of the results, since these have a different seasonal pattern relative to other seasonal patterns. Since the differences between the July and October averages are significantly different, after removing hurricanes, we can reject the hypothesis and conclude that blackout risk does change with time-of-year, though with the caveat that the trend appears to be well correlated with the seasonal nature of storms. Non-storm-related risk appears to be nearly constant through the year.

Hypothesis 3: Blackout probability does not change with time-of-day (reject)

In this section we test the hypothesis that blackout probability does not change with the time-of-day. Figure 5 shows a centered rolling 3-hour hour average blackout frequency for all 24 hours in the day, after excluding the events that do not have a start time recorded. As with the seasonal data, a rolling average is used to smooth out some of the noise that results from the relatively small data set. Blackout probability increases substantially during the peak hours (Figure 5). Blackout frequency is 3.9 times higher during 15:00-19:00 compared to 23:00-2:00. A K-S test shows that this difference is significant ($P < 0.01$). Storm activity typically increases during the mid-afternoon hours, which clearly accounts for some of the increase. Figure 5 shows a substantial increase in weather related events during the mid-afternoon hours. Alternatively (or perhaps additionally), the increase may be the result of power networks being more stressed during mid-afternoon hours, indicating proximity to critical points at which blackout probability increases sharply (Carreras et al., 2004; Liao et al., 2004; Dobson et al., 2007).

Hypothesis 4: Blackout sizes follow a power-law probability distribution (support)

The sizes of large blackouts in the United States follow a power-law probability distribution (Carreras et al., 2000; Talukdar et al., 2003; Carreras et al., 2004). International blackout data also show a power-law size-frequency relationship (Holmgren and Molin, 2006; Dobson et al., 2007), indicating that this relationship is fundamental to the structure of power grids in developed countries. (Power grid failures in many less developed nations have substantially different dynamics. The results and methods discussed in this paper may not be directly applicable to less developed nations.) Using a goodness of fit measure, based on the KS statistic, Clauset et al. (2007) find that the blackout data fit a power-law distribution with "moderate" confidence.

Power-law probability distributions come in a number of forms, but one of the most common is the Pareto distribution. Because the Pareto distribution naturally accounts for data with a fixed minimum value, it is a logical distribution for fitting the blackout data. The cumulative distribution function (cdf) for a Pareto-distributed random variable x with minimum value $x_{\min}$, can be written as follows:

$$P(x \leq X) = 1 - \left(\frac{x_{\min}}{X}\right)^k \qquad (2)$$

where k is the scaling exponent and X is any real value in $[x_{\min},\infty]$. The pdf is:

$$P(x = X) = \frac{k x_{\min}^{k}}{X^{k+1}} \tag{3}$$

and the expected value (E) is:

$$E[x] = \begin{cases} \dfrac{k x_{\min}}{k-1}, & k > 1 \\ \infty, & k < 1 \end{cases} \tag{4}$$

To find the parameters, $x_{\min}$ and k, that provide the best fit to the blackout data we use the method described in Clauset et al. (2007), which uses least squares estimation to find the best fit exponent and the KS statistic to find $x_{\min}$. For the blackout data with size in year-2000 MW, we find a scaling exponent (k) of 1.2 (±0.1) with $x_{\min}$ = 1012 (±340) MW. The data with sizes in year-2000 customers give a scaling exponent of 1.14 with $x_{\min}$ = 291,000 customers. Both fits give KS P values of 0.84, indicating that the fit between the Pareto distribution and the blackout data is quite good. This P value is higher than that reported by Clauset et al. (P=0.62), perhaps as a result of our larger data set. Comparing this to a minimum-value Weibull distribution whose cdf has the form:

$$P(x \le X) = 1 - e^{-\left(\frac{X - x_{\min}}{a}\right)^{b}} \tag{5}$$

shows the clear superiority of the power-law fit. As shown in Table 6, when the small events are included in the Weibull fit the KS P value is less than 0.05. In all cases the power-law provides a vastly superior fit than the Weibull, indicating that one can reject the hypothesis that the data are exponentially distributed with confidence. As shown in Figure 6, the Weibull assumption fails to predict the size of the larger events, thus indicating that the data have a power-law tail.

While the size data in MW or customers fits well with a power-law probability distribution, blackout durations do not (see Table 6). The power-law fit for the duration of the 223 events with known durations, one hour or longer, and sizes over 300 MW (after scaling) has a KS P value of 0.001, whereas the Weibull fit provides a good fit to the data with $P = 0.83$. The size distribution for the longest events (the events over 100 hours in duration, N=24) falls off sharply, as indicated by the higher exponent in the power-law fit to the duration data (Table 6).

We conclude that there is strong statistical support for the power-law relationship between blackout size and frequency. There is virtually no chance that the data come from a process with an exponential probability distribution function. This is true, even after controlling for demand and population growth. Also, the scaling exponent is near the critical value, $k = 1$, at which the expected value of the distribution becomes infinite, indicating that large events contribute substantially to overall blackout risk.

Hypothesis 5: Blackout sizes and restoration times are positively correlated (reject)

Due to difficulties associated with starting large power plants without off-site power, restoration can take many hours, even if no equipment damage has occurred. Many large blackouts require extended restoration periods. It can take days or weeks to restore customers after the losses from a natural disaster. It thus seems reasonable to expect to find correlation between the size of a blackout and its duration.

To estimate blackout duration we identified those blackout reports that included both the time at which the blackout began and the time it took to restore service (595 of 865 events). Blackout restoration is typically an incremental process. Some of the reports indicate this by giving a time at which restoration began and a time at which restoration completed. When both initial and final restoration times are given, we record the duration by the final time.

Surprisingly we do not find a statistically significant correlation between event size and duration (see Table 7 and Figure 7). In a few cause-categories, there is a significant positive correlation between size and duration (lightning, wind/rain, and "other external cause"), but in some other categories a weak (insignificant) negative correlation exists resulting in no significant correlation for the data set as a whole.

It is possible that the lack of a significant correlation is at least in part a result of the difficulties associated with measuring, and incomplete reporting of, duration. To accurately measure blackout duration it would be best to perform a weighted average of the restoration time over all customers. The existing reports are not sufficiently detailed to support this measurement. While we can conclude that the data do not show a significant correlation, we cannot describe in detail, from these data, the relationship between blackout size and duration.

Additional Hypotheses

Numerous additional questions about blackout trends remain. Some of the most important questions regard patterns in overall blackout risk (frequency times size) and the frequency of cascading failures. In particular, it would be useful to investigate the following additional hypotheses:

Hypothesis 6: At least one half (or some other percentage) of large blackouts are caused primarily by cascading failures.

Hypothesis 7: The sizes of the largest blackouts are increasing with time.

Unfortunately, due to the nature of the data that are available, it is difficult to address Hypotheses 6 and 7 with confidence. While NERC provides fairly detailed reports on the largest events in its records, most of the reports do not provide sufficient detail to determine the extent to which a given blackout was caused by cascading failures. Hypothesis 7 is of particular importance, as many have conjectured that the size of the August 14, 2003 event was larger than past events, at least in part, because the electricity industry is pushing the existing transmission capacity closer to its physical limits over time. It is hard to support this conclusion from the data. The large variance in blackout sizes that results from the power-law probability distribution makes it difficult to discern trends in event sizes. Figure 8 shows the sizes of the largest annual blackouts over the period of analysis. It is perhaps useful to note that the 1965 blackout is nearly as large as the 2003 blackout, after adjusting for population growth. According to the Federal Power Commission report (1967), the 1965 blackout affected about 30 million people, or 15% of the US population. The 2003 blackout affected 17% of the US population.

Summary of trends found in the blackout data

To summarize, we find (A) that blackout frequency has not decreased from 1984 to 2006, (B) that blackouts are substantially more frequent in the summer and winter and (C) during mid-afternoon hours, (D) that large blackouts occur much more frequently than would be expected from a exponential statistics, and (E) that there is no apparent correlation between blackout size and restoration time.

4. Relevance of trends to policy problems

Assuming that the trends observed in the NERC blackout data accurately represent actual trends in the electricity industry, they can have a significant impact on investment and policy choices in the electricity industry. This section aims to show that a careful consideration of observed blackout trends can lead to better decisions, which in turn should increase reliability and efficiency. Ignoring observed trends could result in significant mis-allocation of scarce resources.

In this section we make four primary modeling assumptions. Firstly we assume that the trends observed in the NERC data represent trends that actually exist in the time-series of blackouts in North America. As discussed in Section 2, these data are not without error, but the trends that we have found show sufficient statistical power to be useful for the types of analysis described in the following. Assuming that industry members collect more reliable data going forward, the utility of the types of decision-making procedures that follow will increase. Secondly we assume that blackout size and blackout duration are independent random variables. While we cannot sufficiently observe the processes that generate blackouts (at least not from the data reported on in this paper) to verify this assumption with certainty, the lack of a significant correlation between blackout size and duration provides some evidence that this assumption is a reasonable approximation. Again it is important to note that blackout duration is a difficult quantity to measure or use effectively in decision analysis because blackout restoration is incremental and good records are often not kept or reported regarding the amount of load restored over time. Thirdly we assume that blackout costs scale linearly with blackout size in MW. If blackout cost is a linear function of unserved energy (MWh) and duration is an independent random variable, then assumption three follows from assumption two. Finally, we assume that blackout risk (probability times cost) remains constant over time. As reported in Section 3, the data indicate that the frequency of large blackouts is not decreasing in time. Unless emerging reliability standards, or other industry changes significantly change the underlying processes that generate blackouts, this assumption is supported by the NERC data.

It is important to note that the calculations that follow are not intended to recommend specific, immediate policy changes, but rather to illustrate how blackout trends can be used in the analysis of policy trends.

Relevance of the time-of-day and time-of-year trends

Because blackout risk changes with time-of-day and time-of-year, it is rational to focus the bulk of blackout risk reduction efforts on peak periods. For example, consider a utility that wants to compare two policy options regarding an increase in operator staff to reduce blackout risk[2]. Under Option 1 it deems that by doubling the number of on duty operators during all hours it can reduce the blackout probability at all hours by 50%, without changing the blackout size distribution. Let the annual cost of this option be C_1. Under Option 2 it deems that it can reduce blackout frequency during the highest risk hours by doubling is operator staff during only these hours. Let us assume that it can increase its staff during 50% of all hours for 50% of C_1. The cost of Option 2 is thus $C_2 = 0.5C_1$. Let $\Pr(B|h)$ and $\Pr'(B|h)$ represent the probability of a blackout at hour h before and after the policy change, c_b be the average per MW blackout cost and $\mathrm{E}[S]$ be the expected value for the utility's blackout size distribution. Assuming that blackout size and blackout probability are independent, the expected one-year benefit of the policy P is:

$$\mathrm{E}[V_P] = c_b \sum_{h=1}^{8760} \left(\Pr(B \mid h) - \Pr'(B \mid h) \right) \mathrm{E}[S] \qquad (6)$$

Assuming that c_b and $\mathrm{E}[S]$ remain unchanged by the increase in operating staff, Option 1 will reduce blackout costs in the utility's service area by 50% ($\mathrm{E}[V_1] = 0.5$

[2] It is notable that most electric utilities, and other system operators, increase operations staff during peak periods and daytime hours, largely as a result of increased blackout risk. The calculations provided here illustrate that this policy is reasonable given trends observed in the data. This result is not intended to argue that increasing operating staff beyond an already increased daytime workforce would significantly decrease blackout probabilities. The calculation given here illustrates how the available trends can be combined with utility-specific information (eg. assumptions about affects of increased operating staff) to compare the benefits and costs of decisions.

$E[V_0]$). Option 2 however will not reduce risk as much, but does so during peak periods only. If Option 2 reduces blackout probability by 50% during the 12 highest risk hours (7am-7pm) given the time-of-day blackout frequency trend shown in Figure 5, Option 2 would reduce overall blackout costs by 34% ($E[V_2] = 0.34\ E[V]$). The utility can obtain most of the desired risk reduction with half the cost by doubling staff during only peak periods. Similarly, the monthly blackout data can be used to focus risk reduction efforts on seasons with the highest risk.

Relevance of the power-law relationship between blackout size and frequency

The existence of a power-law probability distribution in the blackout size data is important because it indicates that large events are substantially more common than one would predict from exponential distributions such as a Gaussian or Weibull, which are commonly used in engineering reliability analysis. The end result is that a blackout of any size (up to the extent of the entire network) has a significant, non-zero probability. More practically, this result indicates that, if costs scale linearly or super-linearly with blackout size in MW, blackout mitigation efforts should focus on the largest events in nearly equal proportion to the smaller events.

Another effect of the power-law distribution is apparent when calculating the size of a 100-year blackout, using methods commonly applied to storm impact assessment. Given that the sizes of very large blackouts (with sizes greater than 1000 MW, where the power-law frequency distribution fits the data well) follow a Pareto distribution, the probability of obtaining at least one blackout with size S or larger being among a set of n randomly selected very large blackouts is:

$$\Pr(\max(s) \geq S \mid n) = 1 - \Pr(s < S \mid n) = 1 - \left(1 - \left(\frac{1000}{S}\right)^{k}\right)^{n} \qquad (7)$$

Assuming the number of blackouts in a given year follows a Poisson distribution and that blackout frequency and blackout size are independent, the probability of obtaining a blackout of size *S* or larger in a given year (year *y*) is

$$\Pr(\max(s) \geq S \mid y) = \sum_{n=1}^{\infty} \left(1 - \left(1 - \left(\frac{1000}{S} \right)^{k} \right)^{n} \right) \left(\frac{\lambda^{n} e^{-\lambda}}{n!} \right) \qquad (8)$$

To obtain the size of the 100 year blackout (S_{100}) we find *S* such that $\Pr(\max(s) \geq S|y) = 1/100$, given the observed arrival rate of $\lambda = 5.3$ very large blackouts per year, and the scaling exponent ($k = 1.2$) obtained in Section 3 (Hypothesis 4). Solving (8) numerically with these parameters gives $S_{100} = 186{,}000$ MW. By comparison, according to DOE/EIA data, the peak demand (EIA: "Net Internal Demand") for the continental US in 2000 (the base year for the size measures) was 681,000 MW. Thus, if the observed statistical pattern holds for very large blackouts, and if the US were to see a 100-year blackout next year, it would interrupt about one quarter of all electricity service in the continental US. This result is very sensitive to the exponent on the power-law distribution (*k*). Decreasing *k* to 1.15 gives $S_{100} = 233{,}000$ MW.

The power-law relationship between size and probability can significantly affect the expected size of large blackouts, which often factors into the evaluation of benefits associated with risk management decisions, such as in the example decision-analysis calculations that use Eq. 6. Specifically, the power-law relationship can lead to a significant under-estimation of the value of policies or technology that could reduce the likelihood of very large blackouts. Let us assume that one would like to evaluate a technology that is expected to reduce the number of blackouts 1000 MW or larger by one half to $\lambda' = 5.3/2 = 2.7$ per year. We can estimate the annualized value of this decision by finding:

$$\mathrm{E}[V_P] = c_b(\lambda - \lambda')\mathrm{E}[S_{S \geq 1000}] \qquad (9)$$

For data that come from exponential distributions, such as the Gaussian, it is reasonable to estimate an expected value by calculating the mean. The mean of power-law distributed data does not provide a good estimate of the expected value. For one, the variance in a power-law system is so large that random samples from a power-law process are likely to produce widely varying statistics. The mean size of the 122 blackouts larger than 1000 MW is 3700 MW, however the uncertainty in this sample mean is very high. Bootstrap methods provide a 95% confidence interval for the sample mean of [2900, 5500]; the high variance being a result of the variance in the power-law distribution. If we use the power-law statistic to get the expected size by integrating the pdf up to a max blackout size of 700 GW, we get $E[S] = 4400$ MW, which is substantially higher than the sample mean (though within the bootstrap confidence interval). If one merely uses the sample mean to estimate the value of the technology change, one obtains $E[V_p] = 10{,}000\ c_b$. From the power-law distribution, the expected value is $E[V_p] = 12{,}000\ c_b$. By using the sample mean as an estimate of the process mean one would undervalue this technology by about 20%. If one ignored the power-law altogether, and assumed that blackout sizes follow exponential statistics, one would undervalue this decision by much more than 20%.

5. Conclusion

From the available records of large blackouts in North America between the years 1984 and 2006 we find (1) that the frequency of large blackouts in the United States has not decreased over time, that there is a statistically significant increase in blackout frequency during peak hours of the day (2) and seasons of the year (3), (4) that there is strong statistical support for a power-law statistical relationship between blackout size and frequency, and finally (5) that blackout sizes and blackout durations are not correlated. We find these trends to hold even after controlling for increasing demand and population and after eliminating small events, for which the data may be skewed by spotty reporting. Several example calculations show that these trends can have a significant effect on the net benefit of decisions within the industry.

Important unanswered questions remain. For example, the large annual variance in the size of the largest blackouts makes it infeasible to draw conclusions regarding changes in blackout sizes over time. Also, it would be useful to know the percentage of blackouts that are cascading failures. Unfortunately, most midsize events have insufficient detail in the reported characterizations to determine whether cascades played a role. Considering the importance of the implication for blackout mitigation techniques and technologies, we recommend that FERC and NERC place priority on collecting more detailed, accurate blackout data in North America, particularly focusing on identifying the extent to which particular events were exacerbated by cascading failure.

Acknowledgments

The authors gratefully acknowledge Edward Curtis at the US DOE National Energy Technology Laboratory for assisting with the blackout data analysis. Eduardo Cotilla Sanchez at the University of Vermont provided additional assistance with the analysis of power-law statistics in the data. P.H. acknowledges support from the DOE Office of Electricity Delivery and Energy Reliability, under RDS contract DE-AC26-04NT41817 and from the National Science Foundation under grant ECCS-0848247. J.A. acknowledges support from the US National Science Foundation under grants SES-0345798 and CNS-0428404.

References

Amin, S. M., 2008, "For the good of the grid," *IEEE Power and Energy Magazine*, 6(6): 48-59.

Apt, J., Lave, L.B., Talukdar, S., Morgan, M.G., and Ilic, M., 2004. "Electrical blackouts: A systemic problem," *Issues in Science and Technology*, 20: 55-61.

Baldick, R., et al., 2008, "Initial review of methods for cascading failure analysis in electric power transmission systems: IEEE PES CAMS task force on understanding, prediction, mitigation and restoration of cascading failures," in *Proc. of the IEEE Power and Energy Society General Meeting*.

Billinton, R., Allan, R.N., 1996. *Reliability evaluation of power systems, second edition*, Chapter 13, Plenum Press, New York.

Billinton, R., Otengadjei, J., Ghajar, R., 1987. "Comparison of 2 alternate methods to establish an interrupted energy assessment rate," *IEEE Transactions Power Systems*; 2(3): 751-757.

Blumsack S., Lave, L. and Ilic, M., 2007. "A Quantitative Analysis of the Relationship Between Congestion and Reliability in Electric Power Networks," *Energy Journal*; 28: 101-128.

Carreras, B. A., Newman, D. E., Dobson, I. and Poole, A. B., 2000. "Initial evidence for self-organized criticality in electric power system blackouts," in *Proceedings of Hawaii International Conference on System Sciences*.

Carreras, B. A., Newman, D. E., Dobson, I. and Poole, A. B., 2004, "Evidence for Self-Organized Criticality in a Time Series of Electric Power System Blackouts," *IEEE Transactions on Circuits and Systems I: Regular Papers*, 51: 1733-1740.

Chen, J., Thorp, J.S., Dobson, I., 2005. "Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model," *International Journal of Electrical Power and Energy Systems*, 27: 318-326.

CIGRE Task Force 38.06.01, August 2001. "Methods to consider customer interruption costs in power system analysis," CIGRE report 191.

Clauset, A; Shalizi, C; and Newman, M, 2009, "Power-law distributions in empirical data." *SIAM Review*, forthcoming, available from arxiv.org (arXiv:0706.1062).

Dobson, I., Carreras, B. A., Newman, D. E., 2005. "A loading-dependent model of probabilistic cascading failure," *Probability in the Engineering and Informational Sciences*, 19: 15-3.

Dobson, I., Carreras, B.A., Lynch, V.E., Newman, D.E., 2007. "Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization," *Chaos*, 17: 026103.

DOE, 2007, "Form OE-417, (revised 4/2007). Electric Emergency Incident and Disturbance Report," U.S. Dept. of Energy, Office of Electricity Delivery and Energy Reliability, online: http://www.eia.doe.gov/oss/forms.html.

EPACT, 2005, "The Energy Policy Act of 2005," United States Public Law 109-58, 109th Congress, online: www.gpo.gov.

FPC, 1967, "Prevention of Power Failures: An Analysis and Recommendations Pertaining to the Northeast Power Failure and the Reliability of the U.S. Power Systems." Federal Power Commission, July 1967.

Greenberg, M., Mantell, N., Lahr, M., Felder, F., Zimmerman, R., 2007, "Short and intermediate economic impacts of a terrorist-initiated loss of electric power: Case study of New Jersey," 35: 722-733.

Hines, P., Apt, J., Liao, H., and Talukdar, S., 2006. "The frequency of large blackouts in the United States electrical transmission system: an empirical study," in *Proc. of the Second Carnegie Mellon Conference in Electric Power Systems: Monitoring, Sensing, Software and Its Valuation for the Changing Electric Power Industry*.

Hirst, E., Aug. 2004"U.S. Transmission Capacity: Present Status and Future Prospects," Edison Electric Institute and Office of Electric Transmission and Distribution, U.S. Department of Energy.

Holmgren, A; Molin, C, 2006, "Using Disturbance Data to Assess Vulnerability of Electric Power Delivery Systems," *Journal of Infrastructure Systems*, 12(4): 243-251.

Kirschen, D., Strbac, G., 2004. "Why investments do not prevent blackouts," *Electricity Journal*, 17: 29-34.

LaCommare, K. H. and Eto, J. H., 2005. "Cost of Power Interruptions to Electricity Consumers in the United States," *Energy*, 31(12): 1509-1519.

Lawton, L., Eto, J. H., Katz, A. and Sullivan, M.,2003. "Characteristics and Trends in a National Study of Consumer Outage Costs," *Proc. of the Center for Research in Regulated Industries 16th Annual Western Conference.*

Liao, H., Apt, J. and Talukdar, S., 2004. "Phase Transitions in the Probability of Cascading Failures," in *Proc. of the First Annual Carnegie Mellon Conf. on the Electricity Industry, Electricity Transmission in Deregulated Markets: Challenges, Opportunities, and Necessary R&D Agenda*.

NERC "DAWG Database: Disturbances, Load Reductions, and Unusual Occurrences," North American Electric Reliability Council, Disturbance

Analysis Working Group, online: http://www.nerc.com/ dawg/, accessed Nov., 2007.

Newman, M.E.J. "Power laws, Pareto distributions and Zipf's law." *Contemporary Physics*, 46(5): 323-351.

Simonoff, J. S., Restrepo, C.E. and Zimmerman, R., 2007. "Risk-Management and Risk-Analysis-Based Decision Tools for Attacks on Electric Power," *Risk Analysis*, 27(3): 547-570.

Talukdar, S., Apt, J., Ilic, M., Lave, L. and Morgan, M, 2003. "Cascading failures: survival versus prevention," *The Electricity Journal*, 16,: 25-31.

Treaster, J. B., 2003. "The blackout: Business losses; insurers say most policies do not cover power failure," *The New York Times*, August 16, 2003, p.6.

USCA, 2004, "Final Report on the August 14, 2003 Blackout in the United States and Canada." US-Canada Power System Outage Task Force (USCA), online: www.oe.energy.gov.

UTCE, 2007, "Final Report System Disturbance on 4 November 2006," Union for the Co-ordination of Transmission of Electricity (UTCE), online: www.utce.org.

Tables and Figures for "Large Blackouts in North America: Historical trends and policy implications"

TABLE 1: THE 15 LARGEST NORTH AMERICAN BLACKOUTS, 1984-2006 (NERC)

	Date	Location	MW	Customers	Primary cause
1	14-Aug-2003	Eastern US, Canada	57,669	15,330,850	Cascading failure
2	13-Mar-1989	Quebec, New York	19,400	*5,828,000*	Solar flare, cascade
3	18-Apr-1988	Eastern US, Canada	18,500	2,800,000	Ice storm
4	10-Aug-1996	Western US	12,500	7,500,000	Cascading failure
5	18-Sep-2003	Southeastern US	10,067	2,590,000	Hurricane Isabel
6	23-Oct-2005	Southeastern US	10,000	3,200,000	Hurricane Wilma
7	27-Sep-1985	Southeastern US	*9,956*	2,991,139	Hurricane Gloria
8	29-Aug-2005	Southeastern US	9,652	1,091,057	Hurricane Katrina
9	29-Feb-1984	Western US	7,901	3,159,559	Cascading failure
10	4-Dec-2002	Southeastern US	7,200	1,140,000	Ice/wind/rain storm
11	10-Oct-1993	Western US	7,130	*2,142,000*	Cascading failure
12	14-Dec-2002	Western US	*6,990*	2,100,000	Winter storm
13	4-Sep-2004	Southeastern US	*6,018*	1,807,881	Hurricane Frances
14	25-Sep-2004	Southeastern US	6,000	1,700,000	Hurricane Jeanne
15	14-Sep-1999	Southeastern US	*5,525*	1,660,000	Hurricane Floyd

Italics indicate an estimated value, based on a US average of 300 customers per MW (EIA, 2006).

TABLE 2. DESCRIPTIVE STATISTICS FOR THE NERC DISTURBANCE DATA, 1984-2006

	All events ≥ 0 cust./MW	≥ 300 MW	≥ 50k cust.	≥300 MW or ≥50k cust.
Total # of events	856	278	321	438
# of blackouts	547	258	304	406
# after filling missing data	547	307	382	419
# after adjusting for growth	547	317	373	413
Mean size in MW	524	1,508	947	987
Median size in MW	86	634	300	385
Standard deviation in MW	2,396	4,034	3,648	3,285
Mean size in cust.	164,483	321,984	430,585	317,372
Median size in cust.	1,323	85,228	149,500	94,643
Standard deviation in cust.	689,815	1,106,958	1,075,888	939,638

TABLE 3. DESCRIPTIVE STATISTICS FOR EVENT CATEGORIES

	% of events	Mean size in MW	Mean size in customers
Wind/rain	31.4	679	235,840
Equipment failure	19.9	767	248,643
Ice storm	11.1	1,664	431,184
Hurricane or Tropical Storm	10.1	2,684	912,870
Other cold weather	8.8	1,045	271,924
Lightning	8.8	794	200,617
Operator error	8.5	1,226	358,440
Fire	5.6	972	294,994
Voltage reduction	3.9	437	1,162,860
Other external cause	3.6	1,518	823,691
Tornado	3.6	721	227,073
Supply shortage	2.3	600	896,432
Voluntary reduction	2.3	239	966,645
Earthquake	1.6	1,124	526,260
Intentional attack	0.7	2,154	165,000

This table includes only events larger than 50,000 customers. Some of the event sizes were interpolated from the size in customers or MW.

TABLE 4. STATISTICAL TESTS FOR THE HYPOTYESIS THAT BLACKOUT FREQUENCY IS DECREASING WITH TIME.

Data	N	Correlation[b] ρ	P	'84-'95 Median[d]	'96-'06 Median	P from K-S test
≥50k cust. (raw)[a]	321	0.68	0.000	11	20	0.012
≥50k y2k[b] cust.	373	0.78	0.000	11	23	0.003
≥100k cust. (raw)	214	0.62	0.002	7	12	0.012
≥100k y2k cust.	265	0.69	0.000	9	16	0.047
≥300 MW (raw)	278	0.60	0.003	10	15	0.047
≥300 y2k MW	317	0.67	0.001	11	17	0.003
≥300 y2k MW*	248	0.60	0.003	9	13	0.147
≥500 MW (raw)	180	0.57	0.006	7	8	0.147
≥500 y2k MW	214	0.51	0.015	8	12	0.047

[a] Data marked as "raw" are presented before any scaling or post-processing on the data, aside from combining records from the same event.
[b] Correlation measures the relationship between the year and the number of blackouts during that year. Data from 1998 were dropped due to an apparent reporting error.
[c] "y2k" sizes indicate that the data were scaled to account for demand or population growth using 2000 as a base year.
[d] Medians indicate the median events per year, not including the 1998 data.
* Not including extreme natural events (ice storms, hurricanes/tropical storms, earthquakes, tornadoes)

TABLE 5. STATISTICS FOR SEASONAL TRENDS

	Median blackouts per month. P values from K-S test comparing each season to the one previous, shown in parenthesis			
Data set	Dec. – Feb.	Mar. – May.	Jun. – Aug.	Sep. – Nov.
≥300 y2k MW, all blackouts	1.30 (0.59)	1.00 (0.98)	1.42 (0.36)	0.87 (0.10)
≥300 y2k MW, non-hurricanes	1.30 (0.20)	1.00 (0.98)	1.29 (0.36)	0.65 (0.04)
≥50k y2k cust. all blackouts	1.48 (0.10)	1.29 (0.84)	1.57 (0.84)	1.07 (0.20)
≥50k y2k cust. non-hurricanes	1.48 (0.00)	1.29 (0.84)	1.42 (0.84)	0.77 (0.01)

TABLE 6. BLACKOUT SIZE FIT STATISTICS

		Power-law fit		Weibull fit		
Data	x_{min}	*k*	KS *P*	*a*	*b*	KS *P*
y2k MW	300	0.90	0.13	907	0.625	0.04
y2k MW	500	0.97	0.12	1284	0.652	0.03
y2k MW	1016*	1.19	0.76	1775	0.635	0.19
y2k cust.	50000	0.71	<0.001	245103	0.643	0.03
y2k cust.	100000	0.89	0.07	302801	0.633	0.12
y2k cust.	291071*	1.12	0.68	559498	0.661	0.32
duration (hrs)	1	0.58	0.001	73.5	0.844	0.83
duration (hrs)	100	1.77	0.47	93.9	1.012	0.84

* x_{min} fit that results from the method described in Clauset et al. (2007)

TABLE 7. CORRELATION TESTS FOR THE HYPOTHESIS THAT LARGE BLACKOUTS RESULT IN LONG RESTORATION TIMES

Data	N	Corr. coef. (ρ)	P-value
≥300 y2k MW.	223	-0.019	0.78
≥50k y2k cust.	267	0.040	0.52

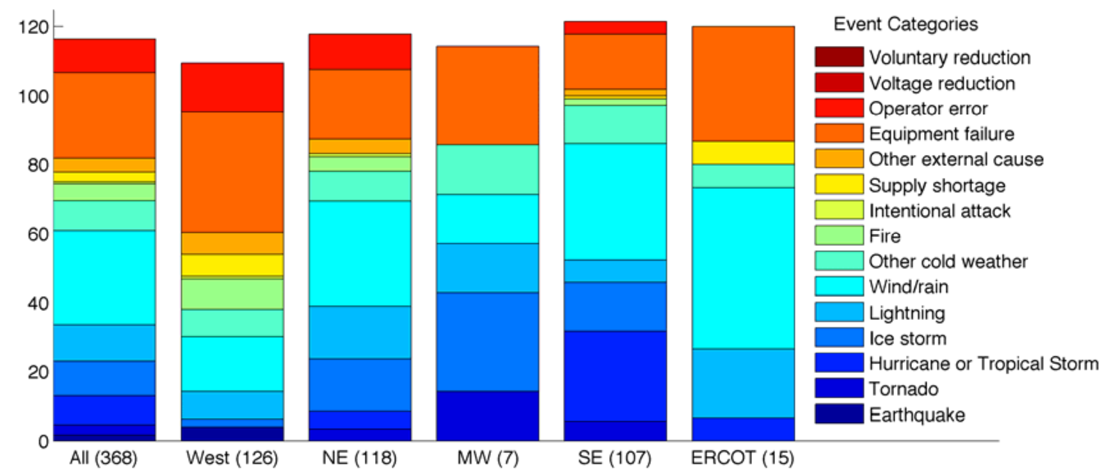

Figure 1. Percent of disturbances that fall into various initiating-event categories, for events larger than 50,000 customers, 1984-2006. The number of events is shown in parenthesis. The totals are greater than 100% because some records fall into multiple categories.

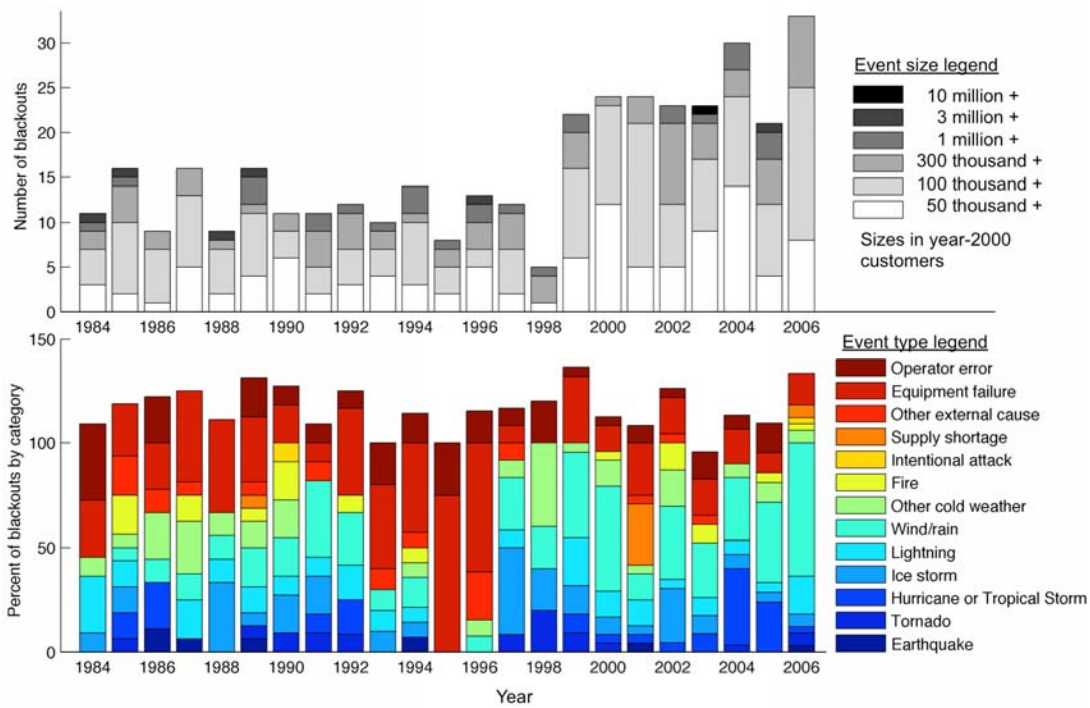

Figure 2. The annual number of blackouts affecting more than 50,000 customers, after adjusting for population growth. The lower graph shows the relative frequency of event categories over time, which shows an increase in the Wind/Rain category. Note that the increase in overall frequency is coincident with an increase in the relative frequency of naturally-caused events. This may be indicative of an increase in the reporting of smaller, weather-related events.

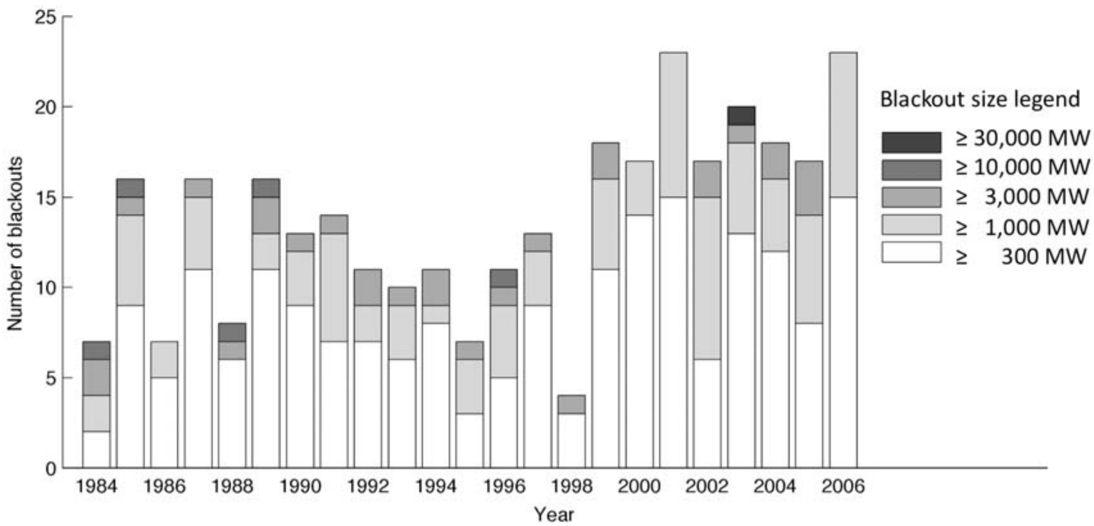

Figure 3. The number of large blackouts per year after removing small events, adjusting for demand growth, and removing extreme natural events. Event sizes are measured in year-2000 MW.

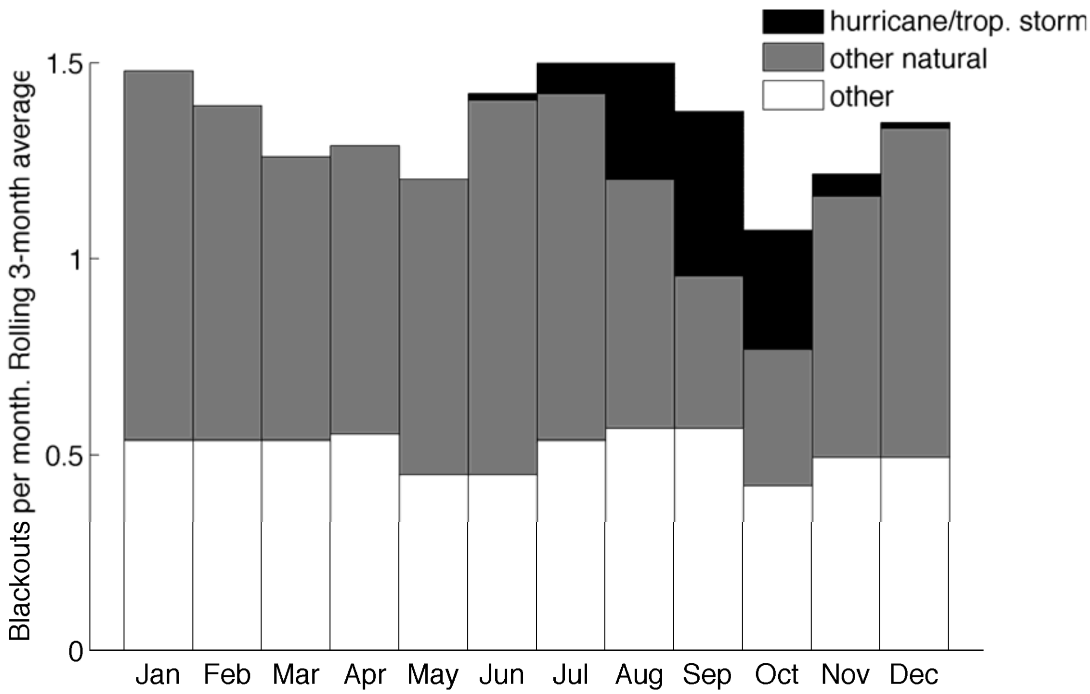

Figure 4. The seasonal frequency of large blackouts. Includes only events larger than 300 MW after adjusting for demand growth.

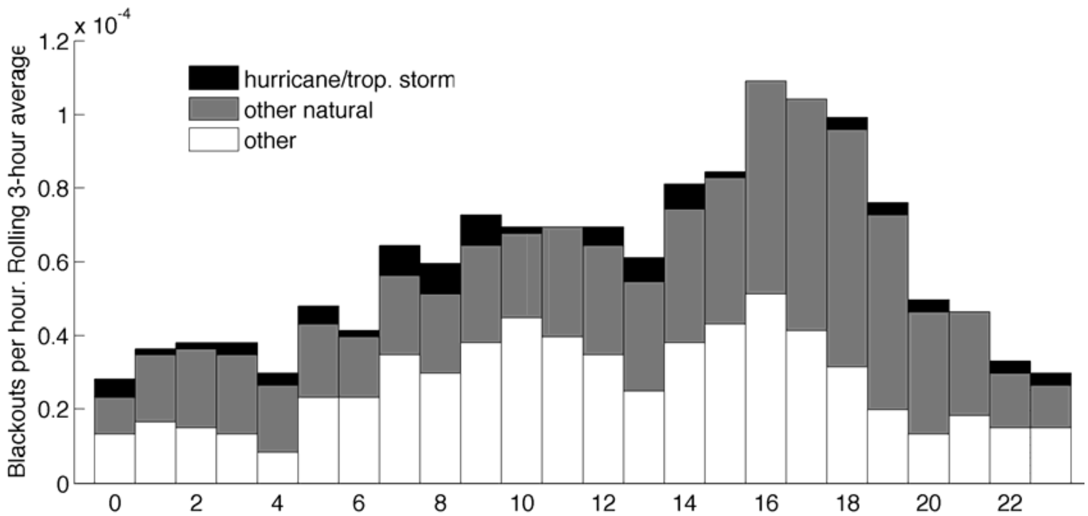

Figure 5. Blackout frequency as a function of time-of-day. The vertical axis shows the average number of events per hour, using 3-hour (centered) rolling averages to smooth out some of the noise in the data.

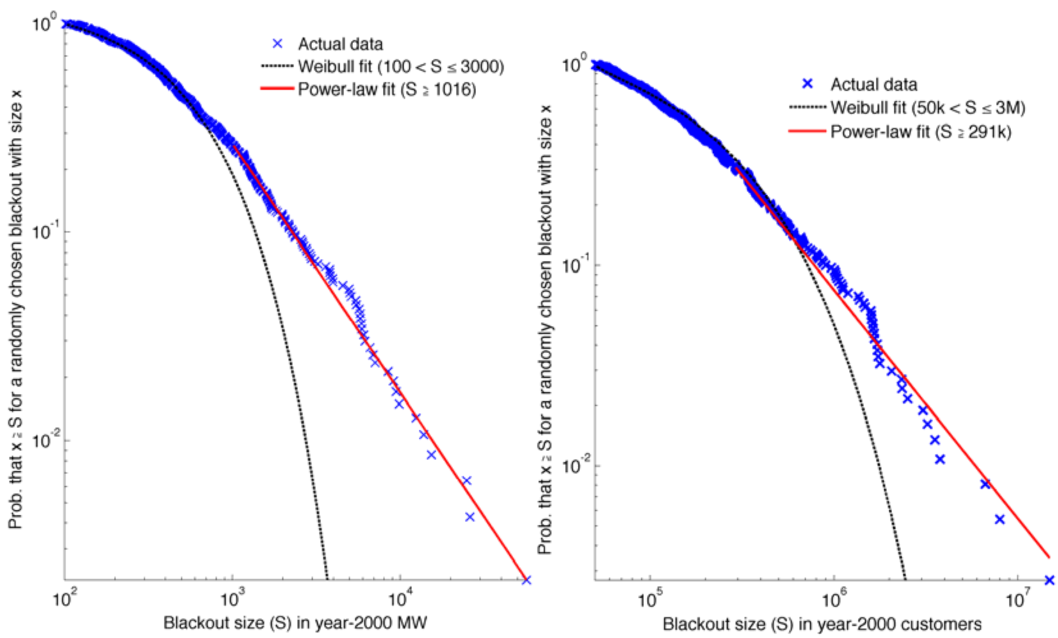

Figure 6. The complementary cumulative probability distribution of blackout sizes in MW (left) customers (right). While the Weibull distribution provides a good fit for predicting the frequency of small events, it grossly underestimates the probability of large blackouts.

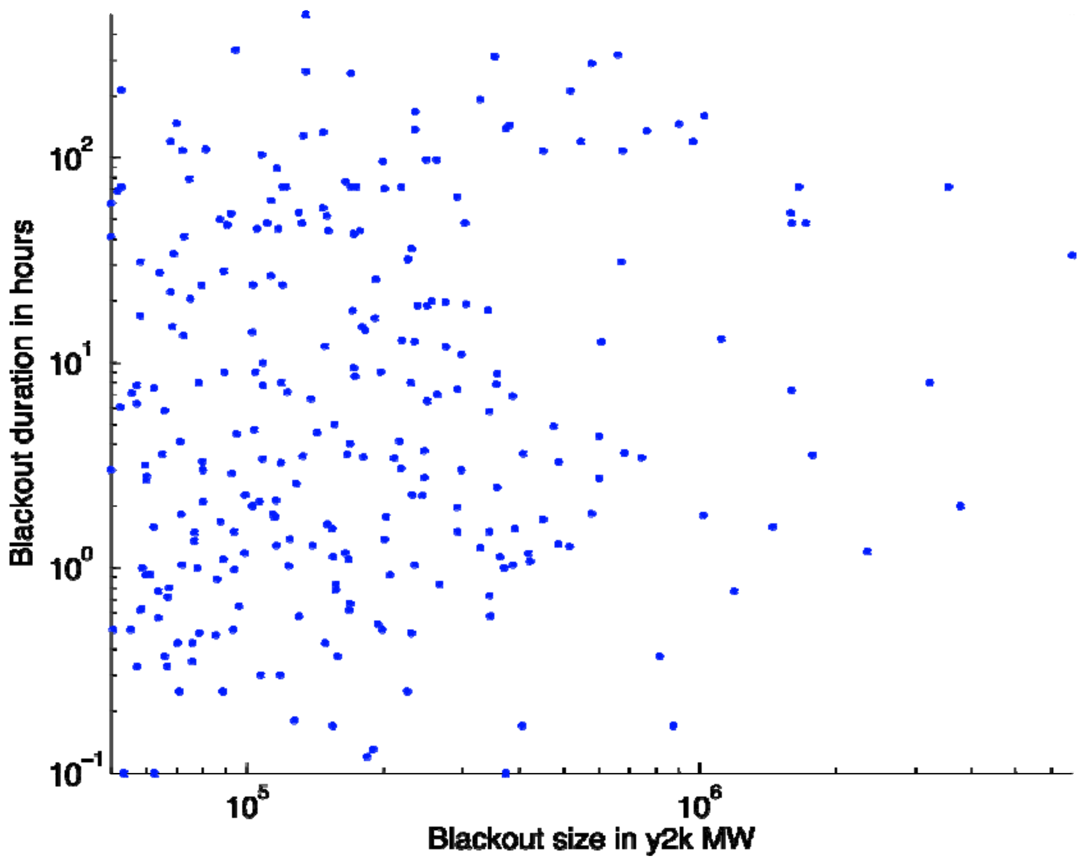

Figure 7. Blackout duration vs. blackout size in MW for events larger than 300 MW (after scaling). The two variables are almost perfectly uncorrelated (see Table 7 for statistics). The log-log scale is used here for clarity—the linear-scale figure shows a similar relationship.

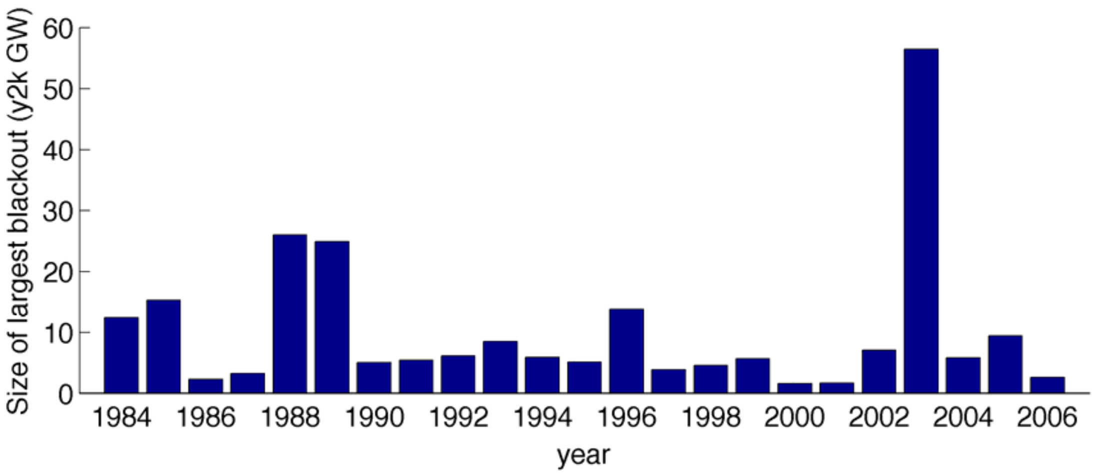

Figure 8. The sizes of the largest blackouts annually. As a point of reference, the 1965 and 1977 blackouts were 64 GW and 17 GW after scaling for demand growth.