Is It Always Windy Somewhere? Occurrence of Low-Wind-Power Events over Large Areas

Mark A. Handschy,[1] Stephen Rose,[2] and Jay Apt[2,3]

Author Affiliations

[1] Enduring Energy, LLC, 5589 Arapahoe Ave., Suite 203, Boulder, Colorado, USA
[2] Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, USA
[3] Tepper School of Business, Carnegie Mellon University, Pittsburgh, USA

Corresponding Author

Mark A. Handschy
Enduring Energy, LLC
5589 Arapahoe Ave., Suite 203
Boulder, CO 80303 USA
markh@enduringenergyllc.com

Abstract

As wind power grows from its present 4% market share in the US, knowing how often the wind fails and power must be supplied by other generators becomes important. The statistics of these low probability events have "thin tails"; the wind fails less frequently than would be predicted by a Gaussian distribution. In order to investigate a future in which wind plants are geographically numerous, we examine the occurrence frequency of low wind-power levels for arrays of wind generators simulated from anemometer data at nine tall-tower sites spread across the contiguous United States. We find that the number of low-power hours per year declines exponentially with the number *N* of sites comprising the array. Power levels below 5% of total capacity, for example, drop by a factor of about 60, from 2140 h/y for the median single site to 36 h/y for the generation aggregated from all nine sites. The systematic dependence of the low-power duration on both *N* and on power threshold is in accord with an explanation based on the theory of Large Deviations. Combining this theory for tail behaviour with the normal distribution for behavior near the mean allows us to estimate the entire generation duration curve as a function of the number of sites in the array.

Broader context

Wind power variability must be considered when integrating this low-pollution electricity source onto the grid. It has long been recognized that this variability can be reduced by combining the outputs of wind generators spread over a large geographic area. The magnitude of the benefit has most often been understood conceptually in terms of the size of the typical deviation of wind power from its mean, which falls as $1/\sqrt{N}$ for a number *N* of independent wind sites. A different variability measure, one of critical importance to the reliability and economics of grids with large amounts of wind power, is the probability of widespread low-wind conditions. Explanations of how this probability might depend on geographic area or number wind sites are scant. Results reported here show that geographic diversity can reduce tail probabilities much more sharply than it reduces the size of typical variations. The theory of Large Deviations convincingly explains the exponential decrease we observe in the occurrence of low wind-power conditions with increasing *N*.

Introduction

The benefit of geographic diversity in reducing the variability of wind power has been investigated since integration of wind generation into the electric grid was first seriously considered. Based on his analysis of 5,000 wind speed data points recorded by the U.S. Weather Bureau at twenty cities east of the Mississippi River, Thomas speculated in 1945 that firm capacities of 50–60% of average generation could be obtained,[1] while shortly thereafter Putnam assessed the capacity value of geographic diversity to be worth less than the cost of transmission needed to achieve it.[2] The smoothing benefit provided by geographic diversity would have considerable economic importance if it allowed a grid system to meet reliability targets with less conventional "dispatchable" generating capacity than would otherwise be needed for a similar amount of unsmoothed wind power. In the terminology of grid reliability this is equivalent to asking to what extent geographic diversity increases wind power's effective load carrying capacity (ELCC).[3-6]

The probability that the aggregated power from an array of wind generators falls below some small generation level is of particular importance in determining ELCC, as pointed out by Kahn[3] and by Haslett and Diesendorf.[4] Characterizing such "tail" probabilities and modeling how they depend on factors such as the number and geographic layout of wind plants making up the array can be challenging. Conventional measures of variation around the mean, such as the variance or standard deviation, reveal little about tail probabilities. Even though the power statistics of large arrays of independent wind generators approach the normal distribution, as required by the Central Limit Theorem, they remain distinctly non-normal for small power levels near the hard lower bound at zero output.

Some previous studies have characterized the occurrence of low wind levels empirically by examining historical wind-speed records. In a 1978 study of data from 25 weather stations in what was then West Germany, Molly found that the times during which total generation of arrays of hypothetical wind plants was zero declined from 1500–7200 hours per year for single sites to less than 5 h/y for arrays of 18 sites within the 800-km (N-to-S) national region.[7] A more recent study by Archer and Jacobson using wind-speed data from meteorological stations in the U.S. Midwest found that the incidence of average afternoon wind speed less than a typical turbine cut-in speed (i.e. $v < 3$ m/s) dropped from 7.6% of the time for single sites to 2.6% for three sites spread over a 120×160 km area, to what they said was 0% for eight sites spread over a 550×700 km area.[8] From the duration curve they present in follow-on work[9] one can see that wind generation was below 5% of turbine capacity 21% of the time for a single site, 10% of the time for a 7-site array, and 1.6% of the time for

a 19-site array. In a study of the Nordic region using actual wind generation records, Holttinen found that while Denmark alone had production below 1% of capacity nearly 5% of the time during the years 2000–2002, the entire Nordic region never fell that low.[10] Using numerical-weather-model reanalysis data roughly corresponding to the territory of the Midwest Midcontinent independent system operator (MISO) Fisher *et al.* found that for a network of 108 sites the output level that could be counted upon all but 10% of the time was 7% of capacity during the winter and 3% of capacity during the summer.[11]

A number of studies have attempted to estimate the complete probability distributions of aggregated wind power, including the tails, in terms of parameters determined from the contributing generators. Justus and Mikhail[12] devised an interesting technique wherein they supposed that an entire array of wind generators was characterized by a single "effective" array wind speed having a Weibull distribution with its shape parameter chosen to make its standard deviation σ_N smaller than the standard deviation for a single site σ_1 according to the number of sites in the array and their average correlation $\bar{\rho}$: $\sigma_N = \sigma_1[(1 + \bar{\rho}(N - 1))/N]^{1/2}$. They then modelled the array output power distribution by transforming the array wind speed distribution through a new power curve that cut in at a lower wind speed and reached rated capacity at a higher speed than did the turbines supposed to be deployed at the individual sites. This produces a complete model output-power probability density function; the "narrowing" of the wind-speed distribution and the "widening" of the power curve act in concert to greatly diminish the probability of low array power. Kahn used their extensive data and pioneering method to calculate the ELCC for wind arrays in California.[3] Carlin and Haslett pursed an alternate approach where they assumed the square-root of site wind speed was approximately normally distributed, allowing the probability of zero array power to be calculated almost exactly from the characteristics of contributing individual sites.[13] They noted "that the effect of dispersal on the probabilities of zero and rated power is significantly more marked than on the coefficient of variation of windpower," a result we concur with and further quantify here. Hasche fit array output-power probability distributions to the Pearson-family Type I (beta) distribution by using empirical functions to match observed and model moments.[14] The chosen beta distribution has the advantage that it naturally accommodates the bounds on array output power at zero and total turbine capacity. Alternately, non-parametric distributions can be defined using kernel estimators.[15]

Some recent investigations have focused on the effects of spreading arrays of wind generators over especially large distances. Kempton *et al.*[16, 17] considered an array of offshore wind plants distributed along the entire extent of the U.S. East Coast, while Fertig *et al.*[18] and Louie[19] evaluated the smoothing effect on wind generation of interconnections between independent system operators (ISOs) across the U.S. Huang *et al.* used reanalysis data to study the variability of coupled wind plants spread over the Great Plains of the U.S. from Montana to Texas.[20] A common feature of these studies is a sharp decline with increased geographic diversity of the fraction of time the aggregated wind power falls below small generation thresholds. Here, we attempt to provide a systematic explanation for the dependence of this behavior on the number of sites being aggregated and the characteristics of the individual wind-power distributions.

Methods

Simulated wind-power data

We investigate the number of hours per year that aggregate wind power is less than a chosen threshold by simulating the power output of arrays of widely separated wind plants using historic wind speed data. We select 9 wind plant sites in the continental United States, shown in Figure 1,

according to three criteria: they have publically-available data from towers taller than typical meteorological stations (instrument height h in Table 1), their mean wind speeds are similar (mean $\bar{v}$ in Table 1), and their simulated wind power outputs are poorly correlated (correlation coefficients in Table 1). We use 1-hour mean wind speeds for the period from January 2007 to December 2012, a total of 5.26×10^4 hours. Most sites have anemometers at multiple heights; we used the data from the height with mean wind speed closest to 6 m/s because that is the speed range for which data is available at the largest number of sites. We exclude any measurements identified as bad by the supplier, measurements less than 0 or greater than 40 m/s, or measurements inconsistent with those taken by other sensors on the same tower at the same time. This quality control excludes 9–38% of the data from the individual sites and leaves 1.46×10^4 hours when data are available from all sites simultaneously. The mean wind speeds from these nine sites, shown in Table 1, are lower than would typically be selected for commercial wind power development, but that is the range for which we could get the largest number of sites with similar wind speeds.

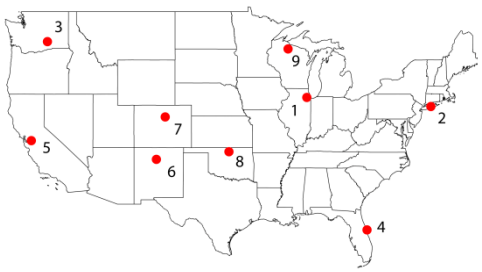

Figure 1. Locations of nine tall-tower sites. (See ESI for site details).

Wind power is simulated from the historical wind speed data using a turbine power curve based on the Vestas V110 2.0 MW turbine,[21] with a cut-in wind speed of 3 m/s, reaching rated power normalized to 1 at wind speed of 11 m/s, and cut-out wind speed of 25 m/s. Details of the functional form of the power curve are given in the Electronic Supporting Information (ESI).

We count the number of hours per year that the simulated power aggregated from combinations of N sites is less than a chosen threshold p_0, for N ranging from 1 (individual site) to 9 (all sites combined), and plot p_0 vs. the count in the form of a generation duration curve. This curve, which depicts the same information as a cumulative distribution function (cdf), is a plot with the power threshold p_0 on the y-axis and the number of hours per year that aggregate power is less than p_0 on the x-axis. (Here we plot *hours less than threshold* rather than the conventional hours greater than threshold, but on an x-axis with its origin at the right where values increase to the left. The curves thus retain their conventional form but allow the use of a logarithmic axis to portray small duration values.) For each combination of sites, the duration curve was calculated by averaging the individual sites' simulated power at each hour (excluding any hour for which data was missing from one or more of the sites comprising that array).

To characterize a single site with behavior "representative" of the nine sites we also pooled all

		h	$\bar{v}$	μ = cf	δ_0	δ_1	correlation coefficients							
	site	(m)	(m/s)				1	2	3	4	5	6	7	8
1	Argonne	60	5.4	0.27	0.11	0.02								
2	Brookhaven	88	5.8	0.32	0.13	0.03	0.062							
3	Hanford	122	5.0	0.27	0.36	0.08	−0.034	−0.026						
4	Kennedy	90	6.0	0.35	0.12	0.03	0.091	0.078	0.015					
5	LLNL	23	6.0	0.37	0.27	0.12	−0.025	−0.028	−0.041	−0.026				
6	Los Alamos	92	4.6	0.23	0.38	0.04	0.049	0.002	0.077	0.008	0.004			
7	NWTC	80	4.8	0.22	0.37	0.06	0.070	0.048	−0.077	0.052	−0.070	−0.007		
8	SGP	25	6.1	0.36	0.15	0.08	0.170	0.031	−0.014	0.071	−0.043	0.240	0.055	
9	WLEF	122	6.3	0.41	0.13	0.03	0.220	−0.010	0.029	0.029	−0.055	0.065	0.023	0.061
	"representative"			0.31	0.23	0.06								

Table 1. Data characteristics for each site: anemometer height h above ground, average wind speed $\bar{v}$; for simulated wind power, the capacity factor cf, the fraction of time the turbine is stalled δ_0, and the fraction of time it is at full power δ_1; Pearson cross-correlation between sites. The parameters of the representative site are those of the histogram in Figure 2.

the hourly simulated power values from the nine individual sites (about 4.2×10^5 samples) and calculated a single histogram, as shown in Figure 2, taking care to separately accumulate those simulated power values that were exactly 0 (δ_0) or exactly 1 (δ_1). We use a histogram with 72 bins (70 full-width and 2 zero-width for δ_0 and δ_1), but we show in the ESI that our results are not sensitive to the number of bins. For these nine sites the "representative" wind plant produces zero power approximately 2000 h/y, and full power about 520 h/y. Its capacity factor or average output power is $\mu = 31\%$. The 0.00–0.014 bin is empty because the minimum non-zero power output of the simulated turbine is 0.014, which is the power produced at the cut-in wind speed.

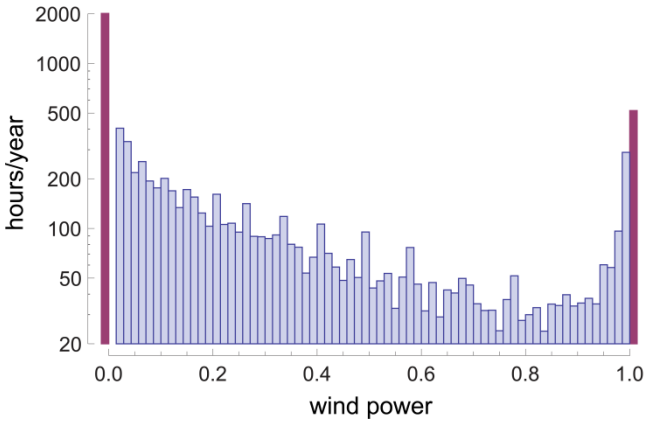

Figure 2. Histogram of simulated wind power for wind-speed records from all nine sites. Dark bars at each end represent instances of power being exactly 0 or exactly 1; bars are widened and offset for clarity. The distribution has a mean μ = 0.31 and standard deviation σ = 0.34.

Large Deviations Theory model

Although the simulated power outputs of the sites in Figure 1 are neither independent nor identical we nevertheless model the simulated aggregate wind power of an N-site array as the mean $\overline{P_N}$ of N independent identically distributed (i.i.d.) copies of a "representative" random variable X having the probability distribution given by the histogram shown in Figure 2. Large Deviations Theory (LDT) gives, under quite broad conditions, tight bounds on the probability that this mean is less than some small value (see Lewis and Russell[22] for an accessible introduction to LDT). According to LDT, $\Pr(\overline{P_N} < p_0)$ falls with N as $e^{-Q(p_0)N}$. The "rate function" $Q(p_0)$ is given as the Legendre transform of the random variables' cumulant generating function $\lambda(\theta)$:

$$Q(p_0) \equiv \sup_\theta [p_0\theta - \lambda(\theta)]\,;\ \lambda(\theta) \equiv \ln\langle e^{\theta X}\rangle, \tag{1}$$

with $\langle\cdot\rangle$ denoting expectation. Although the exponential of the rate function captures the leading asymptotic dependence of probability on N, following Rozovsky[23] a more complete expression is given by:

$$\Pr(\bar{P} < p_0) = \frac{-1}{\vartheta\sigma(\vartheta)\sqrt{2\pi N}} e^{-Q(p_0)N}(1 + o(1))\,, \tag{2}$$

where ϑ is the value of θ at which the supremum in equation (1) is found, and the second derivative of the cumulant generating function gives $\sigma(\vartheta) = [\lambda''(\vartheta)]^{½}$. With bin heights y_k for the histogram in Figure 2 normalized to represent the total fraction of samples in each of the 70 bins, we calculated the cumulant generating function as:

$$\lambda(\theta) \equiv \ln\left[\delta_0 + \delta_1 e^\theta + \sum_1^{70} y_k e^{\left(\frac{2k-1}{140}\right)\theta}\right]. \tag{3}$$

Finding the maximum of equation (1) by numerical search gives Q and ϑ; evaluating the second derivative of $\lambda(\theta)$ at ϑ let us calculate the desired probability in equation (2).

In this simple model, the variability-reducing benefit of aggregating wind power from N sites is largely determined by the magnitude of rate function Q. As seen in Figure 3, Q rises as power threshold p_0 is decreased below mean μ, indicating that modeled variability reduction through geographic

diversity becomes more effective at smaller thresholds. The limiting value of Q at small thresholds can be understood by considering the probability that the output wind power of an N-site array is zero. This requires that its N sites are all simultaneously at zero power, which under our model's i.i.d. assumption has probability $(\delta_0)^N$; hence the LDT rate function saturates for small power thresholds p_0 at $-\ln(\delta_0)$: $e^{-QN} \sim (\delta_0)^{-N}$. The importance of the model's probability distribution being bounded to zero outside the [0,1] limits on wind power generation can be seen by comparing the rate function calculated above with the rate function for a normal distribution having the same mean and variance. This can easily be expressed analytically,[22] and is shown in Figure 3 as well. Both go to zero at μ, but for thresholds ("deviations") away from the mean, the tail of the mean of normally-distributed variables, which extends to $-\infty$, declines at a slower rate than does the tail of the mean of our bounded wind-power variables. That is, for bounded distributions like wind power, LDT shows the distribution's tails are thinner than those of the normal distribution. The effect of changes in δ_0 and in the output-power distribution on the rate function is further illustrated by the dashed curve, which shows the rate function calculated from the simulated wind power distribution for a site in Sweetwater, TX (see ESI) where higher average wind-speeds (7.9 m/s) result in the modeled turbine being above cut-in all but 6.3% of the time.

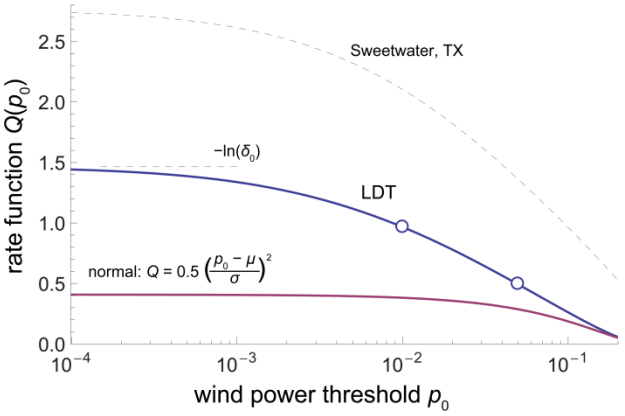

Figure 3. Rate functions for the distribution of Figure 2 computed according to LDT and the normal approximation. Symbols show the rate function values used for the LDT model curves in Figure 4. Also shown is the rate function for a site (Sweetwater, TX) with wind resource quality more typical of a commercial wind plant.

We said above that LDT provides good estimates of the probability that mean aggregate power $\overline{P_N}$ is less than some small value of p_0, but we did not define "small value." For the purposes of LDT, p_0 is considered "small" if there is a small probability that $\overline{P_N}$ is less than p_0, i.e. if p_0 is far from the centre of the distribution of $\overline{P_N}$. For the cases we investigate in this paper with $N \leq 9$ and the distribution shown in Figure 2 with $\mu = 0.31$, p_0 values less than 0.07 can be considered "small." When p_0 is not small, i.e. when it is close to the mean, the Central Limit Theorem tells us that, at least for large N, $\Pr(\overline{P_N} < p_0) \sim \Phi[(\mu - p_0)/(\sigma/\sqrt{N})]$, where Φ is the cdf of the unit normal distribution and μ and σ are the mean and standard deviation of the distribution in Figure 2. We illustrate the ranges of p_0 for which LDT and a normal distribution provide good estimates of the probabilities with Figures S1 and S2 in the ESI.

Results and discussion

The number of hours per year that simulated wind power aggregated from an N-site array is less than a given threshold p_0 decreases essentially exponentially with the number N of aggregated sites, as shown in Figure 4. The box-plot shows the range of durations for all possible combinations of N of the 9 sites shown in Figure 1. The durations of low-wind events for our nine sites are significantly higher than would be expected at a typical commercial wind site as shown by the red crosses. Using the rate-function values shown in Figure 3 we calculated the dependence of $\Pr(\bar{P} < p_0)$ on N, with the results shown by the solid curves in Figure 4. The correspondence of these results to the simu-

lated power data, achieved without the use of any adjustable parameters, encouraged us to compare LDT results to the complete forms of the *N*-site generation duration curves.

Figure 5 plots duration curves for *N*-site arrays. For both $N = 3$ and $N = 6$ there are 84 different combinations possible; the solid-color curve shows the median of their durations while the dashed lines and shading encompass the 5–95% range. Each duration curve plots the fraction of hours (reversed *x*-axis) that the aggregate wind power is *less* than the value p_0 (*y*-axis). The inset is provided to show on familiar linear axes that this unconventional plot style nevertheless gives a conventional-looking duration curve, as explained above in Methods. The thin black curves show the LDT model, again calculated without adjustable parameters from the representative distribution in the histogram of Figure 2. Horizontal "slices" of Figure 5 for chosen power thresholds ($p_0 = 0.01, 0.05, 0.15$) are equivalent to the curves shown above in Figure 4. According to the correlation coefficients listed in Table 1 our sites are not completely independent, with two of the site-pairs having correlation coefficients in excess of 0.2. Nevertheless, the partial correlation seems not to prevent a close correspondence of the LDT curves to the median simulated wind-power data in Figures 4 and 5. Additional insight into the degree of correlation between sites can be had by comparing the variances of the *N*-site arrays to the variances of the underlying individual sites over which the array averages; if the individual sites were uncorrelated the array variance should be equal to $1/N$ times the sum of the individual variances. The variance of the 9-site array simulated wind power (0.0154) is actually 7.6 times smaller than the average single-site variance of the distribution in Figure 2 (0.118), indicating that its behaviour might be closer to that of an array of 8 uncorrelated sites: partial correlation effects are modest but definitely not negligible.

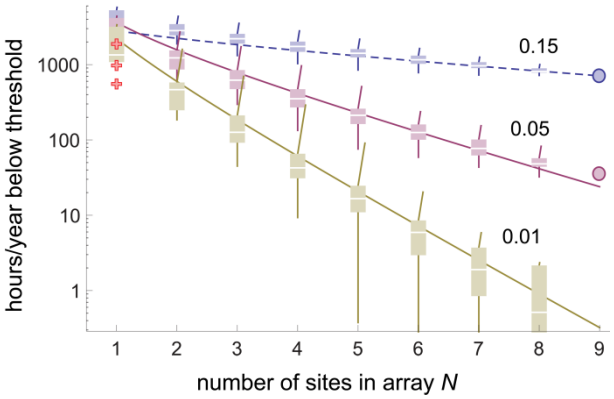

Figure 4. Fraction of time wind power for an array of *N* sites is less than $p_0 = 15\%, 5\%$, or 1% of total capacity. Low-power occurrence decreases approximately exponentially with *N*. Box and whisker symbols show the spread of the various possible combinations of *N* sites with the whisker spanning minimum to maximum and the box the central two quartiles. Circles plot value for a unique combination. For $p_0 = 1\%$ the missing circle and cut-off whiskers and box indicate values of 0. Dashed curve shows normal distribution: $\Phi[(\mu - 0.15)/(\sigma/\sqrt{N})]$. Solid curves show LDT: $1.69e^{-0.49N}/\sqrt{2\pi N}$, and $1.67e^{-0.97N}/\sqrt{2\pi N}$), respectively. LDT predicts far fewer hours below a given threshold than would a normal distribution. Red crosses plot durations for a site with mean wind speeds (7.9 m/s) more typical of a commercial wind power plant.

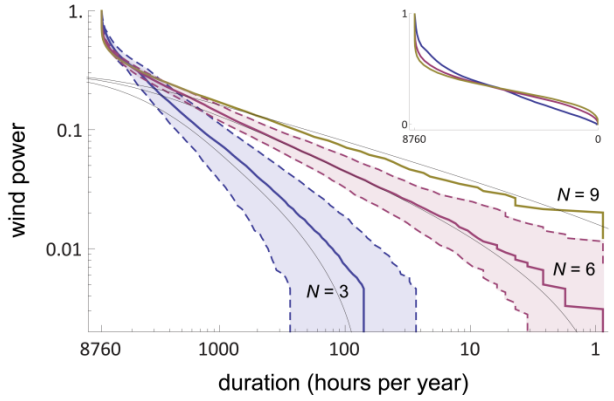

Figure 5. Duration curves for aggregate power from *N* sites: empirical (colored) and LDT (black). Dashed lines and shading enclose 5–95% range of durations from different combinations of *N* sites; solid line is the median. Inset: median duration curves on linear axes.[18]

We now return to the question posed in the title: "Is it always windy somewhere?" A more useful question to answer is "How windy is it *almost always*?" This is the question asked by the system planner who decides how much non-wind reserve generation must be available. To provide an answer we measure the aggregate simulated wind power capacity available except for allowed outages of 5, 50, and 500 h/y ("firm capacity") as a function of the number of sites *N*. These results, plotted

in Figure 6, unsurprisingly show that the capacity available at a given reliability level increases with the number of sites. For example, doubling the number of sites from $N = 4$ to $N = 8$ increases the median capacity available at least 50 h/y from 0.01 to 0.047. The results in Figure 6 also show that increasing the desired reliability level (i.e. tightening the definition of "almost always") decreases the available capacity. For example decreasing the allowed outage rate from 50 h/y to 5 h/y for $N = 8$ decreases the available wind power capacity from 0.047 to 0.02. For large N (not shown), the LDT curves approach a horizontal asymptote equal to average capacity factor μ.

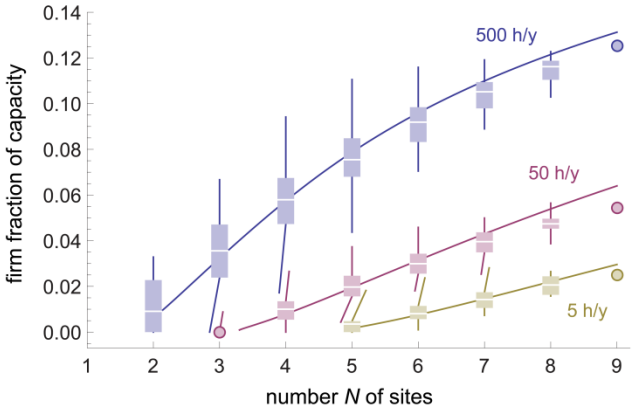

Figure 6. Simulated wind power available at least 5, 50, and 500 hr/yr from *N*-site arrays. Box and whisker symbols show simulated power from *N*-site arrays. Solid lines show LDT predictions based on the distribution in Figure 2.

Conclusions

The model presented here provides a quantitative basis for understanding the increase in firm capacity with geographic diversity. However, relating the predictions of our model to the variability of real wind power plants depends critically on the extent to which the number of statistically independent sites is a good proxy for geographic diversity. The empirical results presented here evidence good agreement with model results based on the assumption that the sites are independent. We presume that this agreement is a consequence of the weak correlation of our widely spread sites. The data we analyze inform speculation neither about the performance of arrays of more-closely spaced wind plants nor about achieving more than nine effectively independent sites within the contiguous U.S. This work may be useful in addressing planning for future modifications to the electric power grid. It demonstrates that aggregating wind plants can decrease the occurrence frequency of low-power events more dramatically than it decreases the magnitude of typical of variations around the mean. For weakly correlated sites we find the occurrence of low-power events in fact declines exponentially with N, in accord with Large Deviations Theory. For comparison, according to the Bienaymé Formula, the standard deviation of the mean decreases as $1/\sqrt{N}$. Thus, to decrease the odds of aggregated wind power falling below 1% of capacity by a factor of 20 for a 3-site array requires an increase in the number of aggregated sites from 3 to 6, at least for sites with characteristics similar to those investigated here. Cutting the standard deviation by a similar factor would require increasing the number of independent sites from 3 by a factor of 400 to 1200—almost certainly unfeasible.

For our data year-to-year variations in the simulated wind power distribution have only modest effects on the rate function (see ESI, Fig. S4). This may allow a grid planner to extrapolate the firm wind power capacity available with a given reliability from limited historical data. With regard to smoothing benefits in general, it is important to note the important caveat that our results do not calculate the time duration of individual low-wind-power events, i.e. they do not distinguish between ten one-hour periods and one ten-hour period of low power.

Additional work is needed to determine to what extent the methods and results presented above apply to situations with the higher inter-site correlations that typically result from clustering wind plants in areas with the best wind. The Gärtner-Ellis theorem[24] allows generalization of the Large Deviations approach to correlated non-identical random variables. Alternately, an array of partially-

correlated sites may be characterized as comprising a smaller *effective number*[25] N^* of independent sites, with array statistics obeying LDT in terms of N^*.[26]

The methods and results we present do not yet allow us to estimate the capacity value or ELCC of wind power as they neglect the correlation of wind generation with electrical load. However, the deviations of regional loads from their daily and seasonal cycles (i.e. the load "anomaly") might yield variables that could be combined with regional wind generation in a way suitable for a treatment like the one employed here.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 1332147 and by the Doris Duke Charitable Foundation, the RK Mellon Foundation, the Electric Power Research Institute, and the Heinz Endowments through the RenewElec project. The authors thank Prof. Julie Lundquist for directing them to many of the public sources of historical wind speed records used here. The authors also thank Argonne National Laboratory, Brookhaven National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, National Renewable Energy Laboratory, the U.S. Department of Energy's Hanford Site, and its Atmospheric Research Measurement facility, NASA Kennedy Space Center, and the NOAA Earth System Research Laboratory WLEF site for providing wind speed data. They are acknowledged in more detail in the Supporting Information.

References

1. P. H. Thomas, *Electric Power From the Wind: A Survey*. (United States Federal Power Commission, Washington, DC, 1945).
2. P. C. Putnam, *Power From the Wind*. (Van Nostrand Reinhold Company, 1948).
3. E. Kahn, "The reliability of distributed wind generators," Electric Power Systems Research **2**, 1-14 (1979).
4. J. Haslett, and M. Diesendorf, "The capacity credit of wind power: A theoretical analysis," Solar Energy **26**, 391-401 (1981).
5. M. Milligan, and K. Porter, "The Capacity Value of Wind in the United States: Methods and Implementation," The Electricity Journal **19**, 91-99 (2006).
6. A. Keane, M. Milligan, C. J. Dent, B. Hasche, C. D'Annunzio, K. Dragoon, H. Holttinen, N. Samaan, L. Soder, and M. O'Malley, "Capacity Value of Wind Power," Ieee Transactions on Power Systems **26**, 564-572 (2011).
7. J. P. Molly, "Balancing power supply from wind energy converting systems," Wind Engineering **1**, 57-66 (1977).
8. C. L. Archer, and M. Z. Jacobson, "Spatial and temporal distributions of U.S. winds and wind power at 80 m derived from measurements," Journal of Geophysical Research **108**, 4289 (2003).
9. C. L. Archer, and M. Z. Jacobson, "Supplying Baseload Power and Reducing Transmission Requirements by Interconnecting Wind Farms," Journal of Applied Meteorology and Climatology **46**, 1701-1717 (2007).
10. H. Holttinen, "Hourly wind power variations in the Nordic countries," Wind Energy **8**, 173-195 (2005).
11. S. M. Fisher, J. T. Schoof, C. L. Lant, and M. D. Therrell, "The effects of geographical distribution on the reliability of wind energy," Applied Geography **40**, 83-89 (2013).
12. C. G. Justus, and A. S. Mikhail, "Energy statistics for large wind turbine arrays," Wind Engineering **2**, 184-202 (1978).
13. J. Carlin, and J. Haslett, "The Probability Distribution of Wind Power From a Dispersed Array of Wind Turbine Generators," Journal of Applied Meteorology **21**, 303-313 (1982).
14. B. Hasche, "General statistics of geographically dispersed wind power," Wind Energy **13**, 773-784 (2010).
15. R. A. Sobolewski, and A. E. Feijóo, "Estimation of wind farms aggregated power output distributions," International Journal of Electrical Power & Energy Systems **46**, 241-249 (2013).
16. W. Kempton, F. M. Pimenta, D. E. Veron, and B. A. Colle, "Electric power from offshore wind via synoptic-scale interconnection," Proceedings of the National Academy of Sciences **107**, 7240-7245 (2010).

17. M. J. Dvorak, E. D. Stoutenburg, C. L. Archer, W. Kempton, and M. Z. Jacobson, "Where is the ideal location for a U.S. East Coast offshore grid?," Geophysical Research Letters **39**, L06804 (2012).

18. E. Fertig, J. Apt, P. Jaramillo, and W. Katzenstein, "The effect of long-distance interconnection on wind power variability," Environmental Research Letters **7**, 034017 (2012).

19. H. Louie, "Correlation and statistical characteristics of aggregate wind power in large transcontinental systems," Wind Energy **17**, 793-810 (2014).

20. J. Huang, X. Lu, and M. B. McElroy, "Meteorologically defined limits to reduction in the variability of outputs from a coupled wind farm system in the Central U.S.," Renewable Energy **62**, 331-340 (2014).

21. "Vestas V110-2.0 MW turbine," http://www.vestas.com/en/products_and_services/turbines/v110-2_0_mw, accessed June 25,

22. J. T. Lewis, and R. Russell, "An introduction to large deviations for teletraffic engineers," Computer Laboratory, University of Cambridge (1997).

23. L. V. Rozovsky, "A Lower Bound of Large-Deviation Probabilities for the Sample Mean under the Cramer Condition," Journal of Mathematical Sciences **118**, 5624-5634 (2003).

24. A. Dembo, and O. Zeitouni, *Large Deviations Techniques and Applications*, 2nd ed. (Springer, 2009).

25. C. S. Bretherton, M. Widmann, V. P. Dymnikov, J. M. Wallace, and I. Bladé, "The Effective Number of Spatial Degrees of Freedom of a Time-Varying Field," Journal of Climate **12**, 1990-2009 (1999).

26. M. A. Handschy, "Reduction of wind power variability through geographic diversity," in *Variable Renewable Energy and the Electricity Grid*, edited by J. Apt and P. Jaramillo (Taylor & Francis, New York, 2014), p. 176-188.

Electronic Supplementary Information:

Is it always windy somewhere? Occurrence of low-wind-power events over large areas

Mark A. Handschy,*[a] Stephen Rose[b] and Jay Apt[b,c]

Contents

Large Deviations Theory vs. Central Limit Theorem 1
Sensitivity analysis for "representative" power distribution 2
Turbine power curve 4
Historical wind-speed data sites 4

Large Deviations Theory vs. Central Limit Theorem

Figure S 1 below plots the empirical duration curve, the Large Deviations Theory (LDT) curve, and the normal distribution duration curve for $N = 6$ to illustrate how LDT provides a better model for the tail and the normal distribution provides a better model near the mean. In this case with $N = 6$ and a mean power of $\mu = 0.31$, the two models cross at a power threshold of $p_0 \approx 0.1$: LDT is a better model for lower thresholds and a normal distribution is a better model for higher thresholds. The normal distribution curve plots duration = $8760\Phi[(\mu - p_0)/\sigma_6]$. Figure S 2 plots the duration and power thresholds p_0 at which the LDT curve crosses the normal-distribution curve for various values of N.

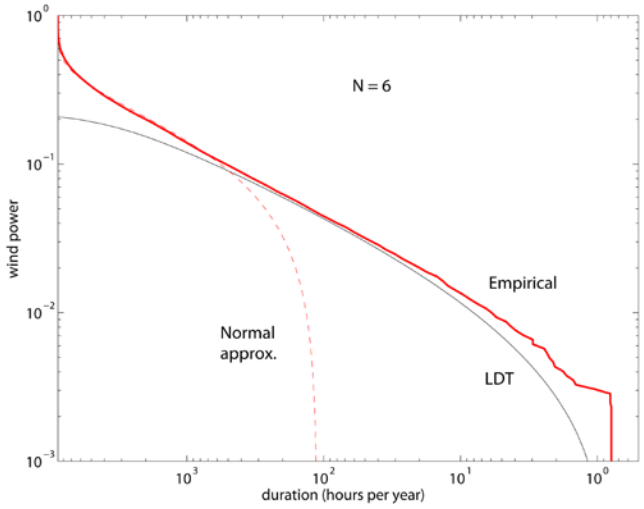

Figure S 1. Comparison of empirical generation duration curve for median simulated wind power from 6-site arrays with LDT model and with normal distribution of mean $\mu = 0.31$ and standard deviation $\sigma_6 = 0.34/\sqrt{6}$.

[a] *Enduring Energy, LLC, 5589 Arapahoe Ave., Suite 203, Boulder, CO 80303 USA.*
[b] *Department of Engineering and Public Policy, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, Pennsylvania 15213 USA.*
[c] *Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, Pennsylvania 15213 USA.*

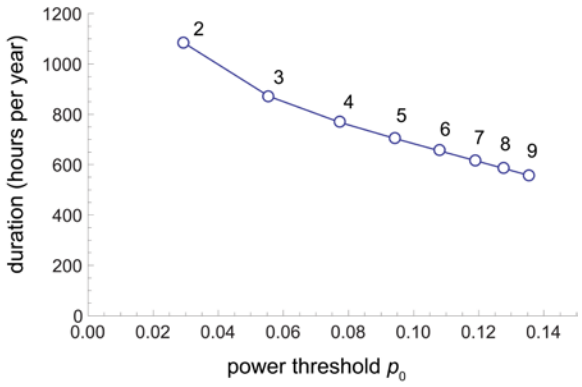

Figure S 2. Locus of cross-over between LDT model and normal distribution curves for various values of *N*.

Sensitivity analysis for "representative" power distribution

In order to use LDT to model the distribution of the aggregate power of several non-identical wind plants, we developed a distribution for a hypothetical single site intended to be "representative" of all the individual sites. This "representative" distribution, shown in the main body of the paper as Figure 2, was calculated by pooling all the hourly simulated wind-power samples and then binning them into a single histogram. We show below that our results are not particularly sensitive to the data used to calculate the representative distribution.

Figure S 3 shows the results of LDT calculations using the power distribution of individual sites and using the representative distribution. It is clear that LDT using the "representative" power distribution matches the empirical duration curve (blue) far better than LDT using any single-site power distribution.

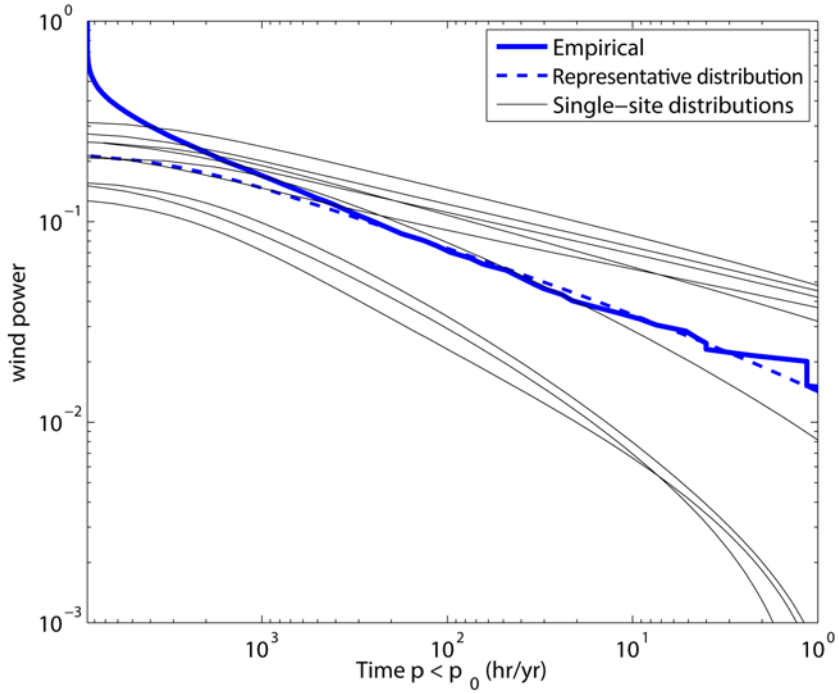

Figure S 3. LDT duration curves calculated using power distributions of each of the nine single sites along with the "representative" power distribution derived from all sites (Fig. 2 in main body), and 9-site aggregate empirical curve.

Figure S 4 shows the results of LDT calculations carried out as before, but instead of using the all-years "representative" distribution we substituted a histogram calculated from all sites but only for individual years. Each thin black line represents the LDT prediction from the pooled distribution for a different year. This figure shows that the distribution changes very little from year to year. This also suggests that a single year of data is probably sufficient to extrapolate the probabilities of events that occur less frequently than once per year.

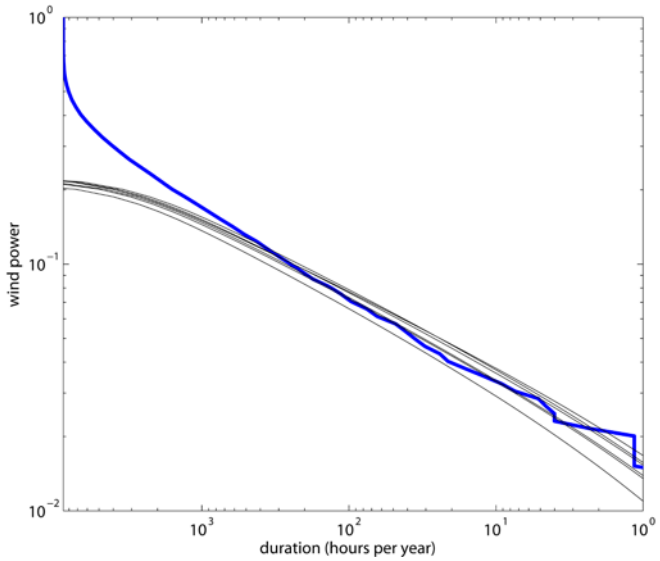

Figure S 4. LDT duration curves calculated using power distributions derived from single years of the all-sites-combined data, and the 9-site aggregate (all-years) empirical curve.

Figure S 5 shows the results of LDT calculations based on representative distributions calculated using different numbers of histogram bins (see histogram in Fig. 2 of the main body). The fit of the LDT duration curve to the empirical duration curve is not sensitive to the number of bins in the histogram.

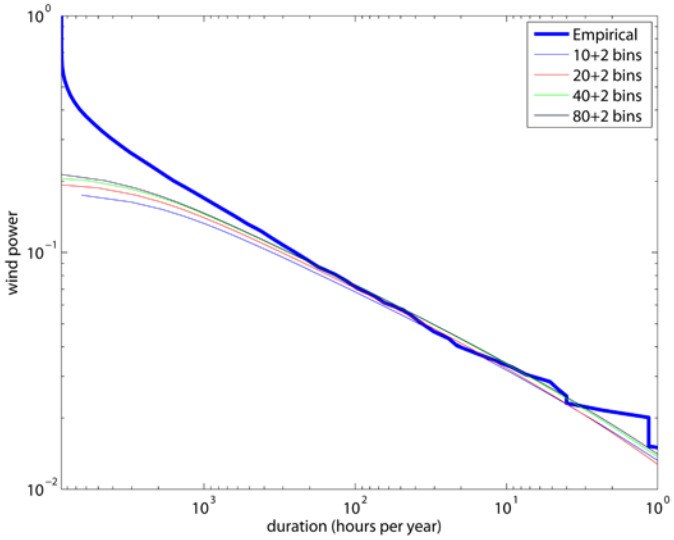

Figure S 5. LDT duration curves calculated by pooling data from all nine sites, but using different bin widths than were used in Fig. 2 (main body), and 9-site aggregate empirical curve. Each histogram retains zero-width bins to accumulate incidence of output power levels of exactly zero and exactly one.

Turbine power curve

We transformed wind-speed samples to simulated wind power using the turbine power curve shown in Figure S 6, which approximates the manufacturer's curve for the Vestas V110 2.0 MW turbine. The turbine cuts in at v = 3 m/s, at which point its output steps discontinuously from 0 to 1.6% of capacity. For wind speeds greater than 11 m/s the output normalized by turbine capacity is unity, until the turbine cuts out at 25 m/s.

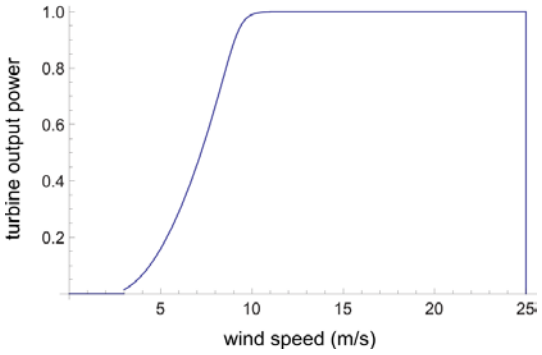

$$P(v) = \begin{cases} 0 & v \le 3 \text{ m/s} \\ \left[a + \dfrac{(1-a)u(v)}{(u_0^{\beta} + u(v)^{\beta})^{1/\beta}}\right]^2 & 3 \text{ m/s} < v < 11 \text{ m/s} \\ 1 & 11 \text{ m/s} \le v < 25 \text{ m/s} \\ 0 & v \ge 25 \text{ m/s} \end{cases}$$

$u(v) \equiv (v - 3 \text{ m/s})/(11 \text{ m/s} - 3 \text{ m/s})$
$a = \sqrt{0.0144}$
$u_0 = 0.79$
$\beta = 20$

Figure S 6. Model turbine power curve and parametric functional form used to transform historical wind-speed samples to simulated wind power.

Historical wind-speed data sites

Table S 1. Wind-speed data-site locations and characteristics.

		Location	Measurement height (m)	Mean wind speed (m/s)
1	Argonne National Laboratory	41.702º N, 87.995º W	60	5.4
2	Brookhaven National Laboratory	40.871º N, 72.889º W	88	5.8
3	ARM Southern Great Plains Central Facility	36.606º N, 97.489º W	25	5.0
4	Lawrence Livermore Natl. Laboratory, Tower 300	37.675º N, 121.541º W	23	6.0
5	Hanford Site, Station 21	46.563º N, 119.600º W	122	6.0
6	Los Alamos National Laboratory	35.861º N, 106.320º W	92	4.6
7	Kennedy Space Center Tower 313	28.626º N, 80.657º W	90	4.8
8	National Renewable Energy Laboratory(NWTC)	39.810º N, 105.235º W	80	6.1
9	WLEF TV	45.945º N, 90.273º W	122	6.3
	Sweetwater, TX, 51 Tall Tower South	34.412° N, 99.646° W	75	7.9

1. **Argonne:** Atmospheric and Climate Research Program, Environmental Science Division, Argonne National Laboratory. Downloaded on 2014 Jan 30 from:
   http://www.atmos.anl.gov/ANLMET/numeric/

2. **Brookhaven:** Personal communication with Scott Smith, Meteorological Services, Environmental and Climate Sciences Department, Brookhaven National Laboratory, 2014 Jan 30.

3. **Hanford:** Personal communication, Kenneth Burk, Hanford Site, Hanford Weather Station. 2014 Feb 18.

4. **Kennedy:** Kennedy Space Center Spaceport Weather Archive. Downloaded on 2014 March 17 from: http://kscwxarchive.ksc.nasa.gov/WeatherTower

5. **LLNL:** Lawrence Livermore National Laboratory. "LLNL Weather, Custom Weather Report Tool." Downloaded on 2014 March 27 from: http://www-metdat.llnl.gov/cgi-pub/reports/report.pl

6. **Los Alamos:** Los Alamos National Laboratory, Meteorology Team. Downloaded on 2014 Jan 30 from: http://environweb.lanl.gov/weathermachine/data_request_green_weather.asp

7. **NWTC:** Jager, D.; Andreas, A. (1996). NREL National Wind Technology Center (NWTC): M2 Tower; Boulder, Colorado (Data). ; NREL Report No. DA-5500-56489. Downloaded on 2014 March 26 from: http://dx.doi.org/10.5439/1052222.

8. **SGP:** Atmospheric Radiation Measurement (ARM) Climate Research Facility. 2001, updated hourly. Carbon Dioxide Flux Measurement Systems (30CO2FLX60M). 2001-01-01 to 2013-01-27, 36.605 N 97.485 W: Southern Great Plains (SGP) Central Facility, Lamont, OK (C1). Compiled by S. Biraud and M. Fischer. Atmospheric Radiation Measurement (ARM) Climate Research Facility Data Archive: Oak Ridge, Tennessee, USA. Data set accessed 2014 Mar 27 at: http://dx.doi.org/10.5439/1025038

9. **WLEF/Park Falls:** Davis, K.J., Bakwin, P.S., Yi, C., Berger, B.W., Zhao, C., Teclaw, R., and Isebrands, J.G., 2003. The annual cycles of CO2 and H2O exchange over a northern mixed forest as observed from a very tall tower, Global Change Biology, 9, 1278-1293. supported by Department of Energy, Ameriflux Network Management Project Support for UW ChEAS Cluster and National Science Foundation grant DEB-0845166. Data downloaded on 2014 Feb 2 from: http://flux.aos.wisc.edu/twiki/bin/view/Main/ChEASData

Sweetwater, TX: Alternative Energy Institute, West Texas A&M University: 51 Tall Tower South, June 1, 2003 to October 1, 2008. (Data from this site was used only to illustrate how a higher-wind single site compared to the 9 sites above, but was not used in any of our multi-site aggregations.) Downloaded on 2012 Dec 4 from: http://www.windenergy.org/datasites/51-talltowersouth/