What Can Reanalysis Data Tell Us About Wind Power?

Stephen Rose and Jay Apt

Abstract

Reanalysis data sets have become a popular data source for large-scale wind power analyses of because they cover large areas and long time spans, but those data are uncertain representations of "true" wind speeds. In this work we develop a model that systematically quantifies the uncertainties across many sites and corrects for biases of the reanalysis data. We apply this model to 32 years of reanalysis data for 1002 plausible wind-plant sites in the U.S. Great Plains to estimate variability of wind energy generation and the smoothing effect of aggregating distant wind plants. We find the coefficient of variation (COV) of annual energy generation of individual wind plants in the Great Plains is 8-17%, but the COV of all those plants aggregated together is 3.6%, Similarly, the average variability of quarterly cash flow to equity investors in a single wind plant is 37%, but that can be reduced to 26 – 29% by small creating portfolios of two wind plants selected from regions with low correlations of wind speed.

Introduction

Wind power is generating an increasing fraction of electricity in many countries and affecting electrical system operation and planning. Long-term wind data are important for predicting these effects. For developers and financers, long-term data reduce uncertainty about the expected revenues of a proposed wind plant. For electrical grid operators and planners, long-term data make it possible to estimate the probabilities of rare events, such as extreme low winds that necessitate conventional power plants as backup. Long-term data are also necessary to assess trends and cycles in wind resource.

Meteorological monitoring stations have collected data for many decades, but those data have several characteristics that make them problematic for wind power analyses (Brower 2012). First, meteorological stations measure wind speeds at 10-m height, which is far below the 60 – 100m hub heights of utility-scale wind turbines. It is possible to extrapolate measured wind speeds to those heights, but such extrapolations are uncertain because meteorological stations do not typically measure variables such as atmospheric stability and surface roughness that are required to calculate the vertical wind profile. Second, meteorological stations are often not located near areas well-suited for wind power development; in the U.S. most stations are located at airports. Third, observations contain errors and gaps in coverage, especially data collected manually before automated stations were deployed (Fiebrich 2009). Finally, measurement instruments, station locations, and surrounding land cover sometimes change, which make it difficult to compare measurements from a single site taken in different periods.

Because of these problems with historical data, many wind power researchers have turned to reanalysis data, which interpolate meteorological observations in space and time using numerical weather prediction models. Recent examples include NARR (Mesinger et al. 2006), ERA-40 (Uppala

et al. 2006), MERRA (Rienecker et al. 2011), and the Climate Forecast System Reanalysis (CFSR) we use in this work (Saha, Moorthi, Pan, Wu, Jiande Wang, et al. 2010a). Reanalysis data sets are attractive because they span several decades, contain observations for variables, locations, and times not recorded in historic data, and have uniformly-good data quality and no missing observations (Brower 2012). Researchers have used reanalysis data for wind resource assessment, long-term trends (Holt & Jun Wang 2012; Pryor et al. 2009), long-term variability (Henson et al. 2012), geographic smoothing (Giebel 2000; Fisher et al. 2013; Huang et al. 2014), and extreme winds (Larsén & Mann 2009).

Relatively few of the researchers who use reanalysis wind speeds have validated those data against historical data. Researchers who compare reanalysis-predicted wind speeds at 10-m height to historical measurements from meteorological stations find significant uncertainties: RMS error of 2.5-3 m/s for surface-level winds in NARR (Mesinger et al. 2006) and correlation coefficients of 80-90% and energy correction factors of 1.06-1.10 for MERRA and CFSR (Liléo & Petrik 2011). However, these validations do not capture errors and uncertainties introduced when wind speeds are extrapolated from 10-m to typical wind turbine hub heights using assumed vertical wind speed profiles. A few authors validate reanalysis data using wind speeds measured at heights closer to wind turbine hub height (50 – 100m). A comparison of daily average wind speeds from several reanalysis models to tall tower data calculates average R^2 values of 0.73 for CFSR and 0.67 for MERRA (Brower et al. 2013). A thorough analysis with offshore wind speed measurements in the UK finds MERRA under-predicts hourly wind speeds by an average of 7% and over-predicts the COV of annual wind speeds by an average of 17%. That study also calculates R^2 values of 64-93% for hourly speeds, 80-97% for daily averages, and 90-99% for monthly averages (The Crown Estate 2014). Henson finds correlation coefficients of 75-87% for hourly MERRA wind speeds with data from on-shore sites in Massachusetts (Henson et al. 2012).

In this work we present a model that corrects biases and quantifies the uncertainty in wind energy calculated from reanalysis data. Whereas previous studies estimate uncertainty for individual sites assuming a separate model for each site, the model we present quantifies the uncertainty attributable to between-site differences as well as within-site variability. We apply this model to generate 32 years of quarterly energy generation for individual wind plants, which we analyze to estimate inter-annual variability of wind energy generation and quarterly variability in cash flow to equity holders in a wind plant.

Method

We estimate the quarterly energy generated by each of 1,002 wind plants in the U.S. Great Plains for the period 1979 – 2010 using reanalysis wind speed data. We calculate the 80-m height hourly wind speed at each wind plant site by extrapolating data from the CFS reanalysis (Saha, Moorthi, Pan, Wu, Jie Wang, et al. 2010b). We then aggregate the energy for each site by quarter and apply a model we develop to correct for biases and quantify uncertainty in the CFS data. Finally, we simulate 10^3 probable realizations of quarterly energy generation at each site.

A1.1 Wind Plant Locations

We simulate the wind power at the locations of all wind plants from the Eastern Wind Integration and Transmission Study (EWITS) (Brower 2009) that are in the U.S. Great Plains (north and west of the Mississippi and Ohio rivers). We combine the few wind plants that are less than 5 km apart, which leaves 1,002 wind plants. We consider only sites in Great Plains for four reasons. First, most

wind power development in the U.S. is taking place in the Great Plains. Second, the terrain is generally flat, so we expect the reanalysis-predicted wind speeds to be more accurate. Third, the area has a good coverage of historical record of meteorological observations, which are assimilated into the reanalysis model. Finally, good empirical validation data were available for the Great Plains in the form of tall-tower wind speed data not assimilated into the reanalysis model.

The empirical data consists of several years of hourly wind speed measurements from tall towers at 78 sites in the Great Plains, measured at heights of 50 - 120m. At sites with anemometers at multiple heights, we selected the one closest to 80m. A table listing the locations, heights, date ranges, and mean wind speeds of the empirical data sets is given in the Supporting Information. These data were collected by economic development agencies in various states and then checked for quality control and compiled into a single database by the University of North Dakota Energy & Environmental Research Center (Simonsen & Stevens 2004).

A1.2 Extrapolating Hub-Height Wind Speed from CFSR Data

We estimate hourly wind speed at 80-m height $u(z = 80)$ at each location using 1 – 6 hour forecast data from the CFSR (Saha, Moorthi, Pan, Wu, Jie Wang, et al. 2010b) and the following formula for a logarithmic vertical wind profile given by Panofsky (Panofsky 1963):

$$u(z) = \frac{u_*}{\kappa}\left(\log(\frac{z}{z_0}) - \Psi(z/L)\right) \quad (1)$$

where:

u_* = friction velocity [m/s]
κ = 0.4 (von Kármán constant)
z = hub height [m]
z_0 = surface roughness length [m]
$\Psi(z/L)$ = correction for atmospheric stability, a function of the stability measure z/L
L = Obukhov length [m]

The friction velocity and surface roughness values are taken from the CFS reanalysis data. The Obukhov length is calculated by an expression given in the Supporting Information. The correction for atmospheric stability Ψ is given in (2) for unstable ($z/L < 0$), neutral ($0 \leq z/L < 0.5$), and stable ($z/L \geq 0.5$) atmospheric conditions:

$$\Psi\left(\frac{z}{L}\right) = \begin{cases} \frac{p_1(z/L)+p_2}{(z/L)^2+q_1(z/L)+q_2} & z/L < 0 \\ -5(z/L) & 0 \leq z/L < 0.5 \\ -3.76(z/L)^{0.45} & 0.5 \leq z/L \end{cases} \quad (2)$$

where p_1 = -2.0, p_2 = -0.36, q_1 = -0.26, and q_2 = 2.4. The expressions we present in (2) for stable and unstable conditions are new in this work because we find the expressions typically given in the literature (Panofsky 1963; Emeis 2013) are a poor fit for the CFS reanalysis data.

We determined these novel expressions for Ψ in stable and unstable conditions by substituting historical wind speeds from the empirical data described in Section 2.1 for $u(z)$, solving (1) for Ψ, and then fitting empirical curves to Ψ as a function of z/L using robust (least absolute squares) regression. This method is less sensitive to outlier data than are other forms of regression. We divide the empirical data randomly into two equally-sized subsets: a "training" set used to fit the curves and a "validation" set used to test the fit. Within each subset, data from all sites are pooled, so we cannot separate within-site variation from between-site variation. We exclude the small number of data points from hours with $z/L > 10$ or $z/L < -10$ from the curve fitting because variance of the residuals grows rapidly for $|z/L| > 10$.

A1.3 Correcting Biases and Quantifying Uncertainty in Quarterly Energy Generation

The hourly hub-height wind speeds extrapolated from reanalysis data using (1) have biases and uncertainties from the extrapolation procedure and from the reanalysis model. We develop a correction and uncertainty model to correct for the biases (a procedure known as Model Output Statistics (Potter et al. 2007)) and quantify the uncertainties. Our procedure is similar to the linear regression method described by Brower for determining the relationship between wind speeds at different locations (Brower 2012).

We develop this correction and uncertainty model by comparing the wind power calculated from reanalysis data with (1) to wind power calculated from historical data. First, we interpolate the raw data from the reanalysis model to the locations of the sites in the "validation" subset of the empirical data described in Section 2.1. We use that interpolated raw reanalysis data as inputs to (1) to calculate hub-height wind speeds. Finally, we convert the reanalysis and empirical wind speeds to wind power using the power curve for a generic 2-MW wind turbine (Brower 2009).

The correction and uncertainty model given in (3) is a hierarchical random-effects model (Ntzoufras 2009) that estimates "actual" quarterly energy $E_{i,j}$ for site *i* in quarter *j* as a function of reanalysis energy $R_{i,j}$ for the corresponding site and quarter. The slope β is fixed, the offset α_j for site *j* is drawn from a normal distribution with mean 0 and standard deviation σ_α (4), and the error term ε for each measurement is drawn from a normal distribution with mean 0 and standard deviation σ_ε (5).

$$E_{i,j} = \alpha_i + \beta\, R_{i,j} + \epsilon_{i,j} \quad (3)$$

$$\alpha_i \sim \mathrm{N}(0, \sigma_\alpha) \quad (4)$$

$$\epsilon_{i,j} \sim \mathrm{N}(0, \sigma_\epsilon) \quad (5)$$

We fit the model to the available data using Markov Chain Monte Carlo (MCMC) methods, as implemented in OpenBUGS version 3.2.3 (Thomas et al. 2006). Fitting the model parameters using MCMC methods yields distributions of probable values for each parameter (β, σ_α, σ), rather than point estimates; summary statistics for the distributions of values of β, σ_α, σ are given in Table 1. Figure 1 plots the data to which the mode is fitted, overlaid with the model using the mean parameter values in Table 1. The inset shows the model for a single site (thin red line) with 1-standard-deviation error bounds (dashed lines).

Table 1: Summary statistics of hierarchical model of quarterly energy for a 2-MW turbine described in (3) - (5). Units are MWh/quarter for a 2-MW turbine.

	Mean	Median	Std. Dev.	Pearson Corr. Coeffs. for β, σ_α, σ
β	0.85	0.85	0.023	$\begin{bmatrix} 1 & & \\ -0.51 & 1 & \\ 0.16 & -0.076 & 1 \end{bmatrix}$
σ_α	376	372	52	
σ	173	173	5.3	

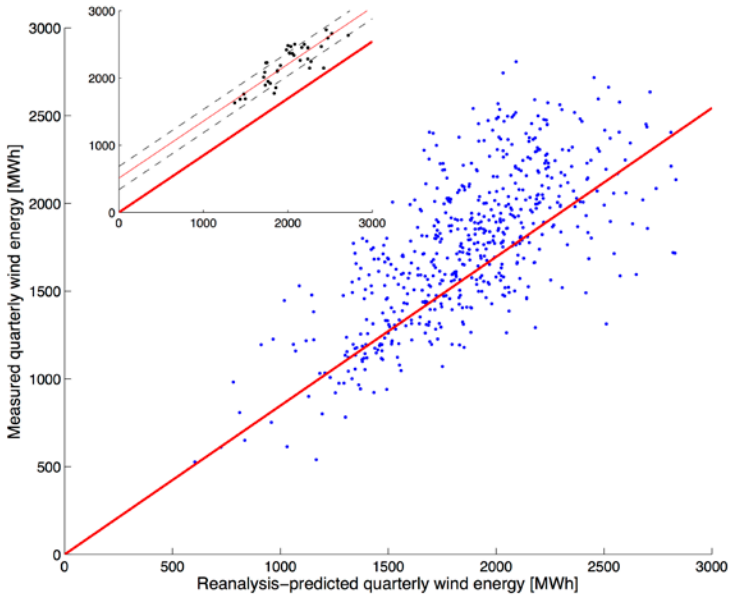

Figure 1: Comparison of measured energy generation to reanalysis-predicted generation. Each point represents the quarterly generation for a 2-MW turbine in a given site. The red line shows the nominal values of the model described by equation (3) and parameters in Table 1. The inset compares data from a single site to its corresponding model: the thin red line shows the site-specific model, which is offset from the nominal model by α_i, and the dashed lines show the error term ε.

We apply this model to quarterly wind generation from reanalysis data by first randomly drawing values of the parameters β, σ_α, and σ to simulate 10^3 probable realizations of quarterly energy for each site. Those values of σ_α, and σ drawn for each realization are then used as the parameters for normal distributions from which the offset for each site α_i and the measurement error for each quarterly energy value $\varepsilon_{i,j}$ are drawn. We subsequently refer to the resulting values of quarterly energy as "corrected" reanalysis data.

The hierarchical model in (3) - (5) estimates the bias (β), between-site variability (σ_α), and within-site variability (σ) of quarterly energy generation. The mean value of the bias parameter $\beta = 0.85 \pm 0.023$ shows that the "actual" quarterly energy production is slightly lower than predicted from the reanalysis data. The mean value of the within-site variability parameter $\sigma = 173$ MWh/quarter and between-site variability parameter $\sigma_\alpha = 376$ MWh/quarter (for a turbine that generates an average of 1800 MWh/quarter) show that between-site variability dominates the uncertainty in energy generation. These results suggest that additional research is needed to determine the source of the between-site variability and find additional inputs to the model to better explain that variability.

To the best of our knowledge, this is first model of the uncertainties and biases in reanalysis-predicted wind speed or energy. Previous work summarized in Section 1 calculated the R^2 values (or the related correlation coefficients) for the relationship between reanalysis and actual wind speeds but not applied those findings to estimating the uncertainty bounds on the. For comparison with previous work, we fit linear functions (unrelated to the model described in (3)) to quarterly wind energy for each of the 38 sites in the validation subset of the empirical data and find R^2 values in the range 0.04 – 0.997, with a mean of 0.72.

Results

We use the realizations of corrected wind energy developed in Section 2.3 to estimate several measures of long-term wind energy variability.

A1.4 Variability of Single-site Annual Energy Generation

The energy generated by a single wind plant varies from year to year due to weather and climate. We quantify that variability for each EWITS site using the corrected reanalysis data described in Section 2.3. The mean coefficient of variation (COV = σ/μ) of annual energy generation for individual sites range from 5.9% to 15%, with a mean value of 7.9%; the mean COV values are plotted in Figure 2. We refer to "mean COV" because we calculate 10^3 realizations of COV for each site from the uncertain reanalysis data described in Section 2.3. The median 1-standard-deviation confidence interval for the COV of those realizations at an individual site is ±2.9 percentage points and 90% of the confidence intervals are smaller than 7.7 percentage points.

These values are similar to previous estimates of COV of annual energy: Milligan calculates 10% from historical weather data at one low-wind-speed site in North Dakota (Milligan 1997), Baker calculates 12 – 13% from historical weather data for 3 sites in the Pacific Northwest (Baker et al. 1990), and Wan calculates 8% - 13% from historical wind power production data at 4 sites in the Great Plains (Wan 2012). The COV for individual sites shows a geographic trend that is the inverse of the geographic trend in wind resource (Elliot et al. 2010; Katzenstein & Apt 2012): sites with better wind resource (average annual wind speed) have lower COV (plotted in the Supporting Information).

This long-term variability of annual energy at an individual site is important to wind plant developers because it sets an upper limit on the allowable debt load for the wind plant. Most plants sell their energy on a fixed-price contract, so revenue variability is proportional to variability in energy generation. The expected revenue in a bad year determines the amount of debt financing a wind plant can obtain. Typically, the debt payments are set to some multiple of the plant's revenue in the 1st percentile ("P99") or 10th percentile ("P90") year (Tindal 2011). If two plants have identical mean generation but different year-to-year variability in generation, the plant with less variability will be able to take a larger amount of debt and have a higher debt-to-equity ratio. In practice, uncertainties about future revenue will be larger than the results we give here because wind plant developers estimate the distribution of energy generation from much shorter periods of data: 1 – 2 years, compared to the 32 years we use in our analysis.

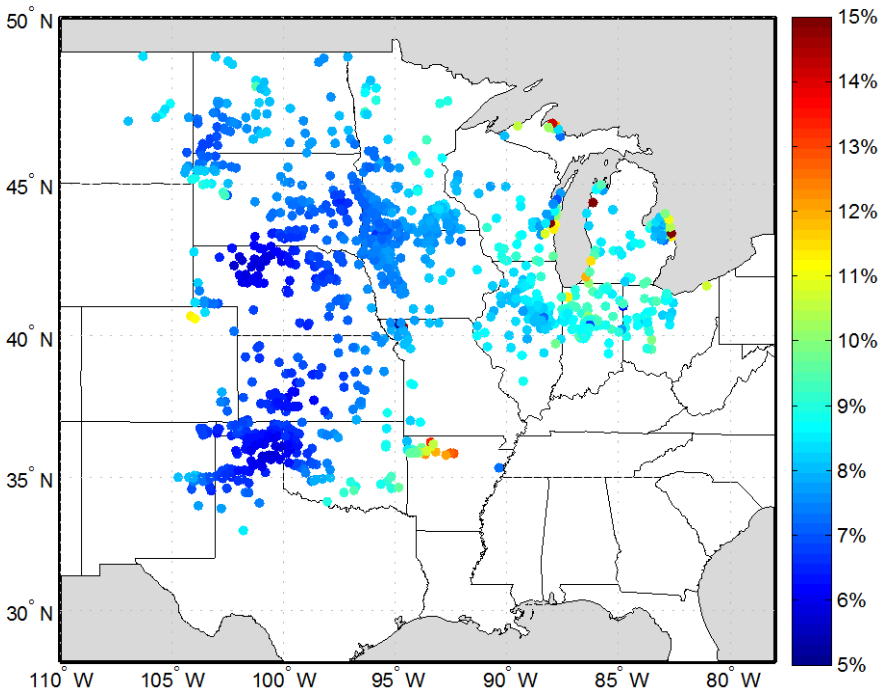

Figure 2: Mean Coefficient of Variation (COV) of annual energy generation for EWITS wind plants in the U.S. Great Plains and the Eastern Interconnect. Each COV value has a 1σ confidence interval of ±2.9 percentage points.

A1.5 Variability of Aggregate Annual Energy Generation

The inter-annual variability of wind generated in large regions is of interest to power system planners. We calculate the aggregate annual energy generation for all the EWITS sites, using the corrected reanalysis data weighted by capacity. The mean COV of energy for the aggregated sites, 3.6%, is much smaller than the COVs calculated for individual sites, which illustrates the smoothing effects of aggregating many wind plants. Figure 2 plots a time series of aggregate annual energy (with 1 standard deviation error bars) and Table 2 presents summary statistics.

We also present results from previous studies for comparison with our results. Giebel used reanalysis data to calculate a COV of 6.4% for aggregate annual energy for 83 sites in northern Europe; see Table 2 (Giebel 2000). Katzenstein used historical data to calculate a COV of 5.4% for 16 sites in the Great Plains (Katzenstein et al. 2010); see Table 2 and the red line in Figure 2. The time series of our results in Figure 2 is qualitatively similar to Katzenstein's, but our results show less variability (3.6% compared to Katzenstein's 5.4%). To test whether the lower variability of our results is caused by the much larger number of sites we aggregate (1002 compared to 16) or the turbine power curve we use, we use our corrected reanalysis data to calculate the aggregate energy for the 16 sites Katzenstein analyzes with the same power curve. We calculate a COV of 3.8% for those sites. The summary statistics for our analysis of those 16 sites is labeled "This work (compare to Katzenstein)" in Table 2 and a plot of the time series is given in the Supporting Information. For comparison, we also include statistics for aggregate U.S. hydroelectric generation in the last row of Table 2 (EIA 2012). These statistics show that aggregate annual hydroelectric is much more variable than aggregate annual wind generation.

The year-to-year variability of aggregate generation is important for the financing of wind plants because these results show that aggregating many wind plants distributed across a large area

significantly reduces the variability of energy generation and corresponding revenue. Grid operators already manage the inter-annual variability of hydroelectric generation, which has a COV of 12%, approximately three times as large as the variability we estimate for wind generation. Long-term planning for generation capacity may benefit from understanding the size of this inter-annual variability. However, grid operators typically estimate the contribution of wind power to peak generation capacity based on the correlation between hourly wind power and electricity demand.

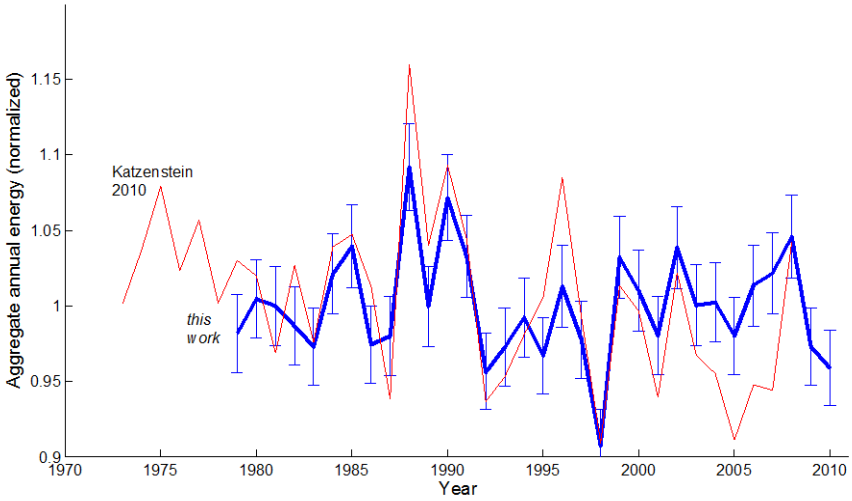

Figure 3: Annual wind energy the aggregate of all EWITS wind plants in the U.S. Great Plains (blue) with error bars showing 1σ confidence intervals. Red line plots aggregate annual wind energy for 16 sites in the Great Plains from Katzenstein (Katzenstein et al. 2010).

Table 2: Summary of annual variability of aggregate wind power from this work and two previous studies. Results from this work give ±1σ confidence intervals in parentheses. The last row (labeled "EIA 2012") summarizes annual aggregate hydroelectric generation in the U.S. for comparison.

	Data source	Sites	COV	Max year	Min year	Max year-year change
This work	Reanalysis	1002	3.6% (3.5–3.7)	+9.2% (6.3-12)	-9.2% (6.8-12)	12.4% (12.0-12.9)
This work (compare to Katzenstein)	Reanalysis	16	3.8% (3.4-4.1)	+8.3% (6.2-10)	-7.9% (5.8-10)	11.8% (9.8-13.9)
Katzenstein 2010	Historical	16	5.4%	+15%	-9.6%	22%
Giebel 2000	Reanalysis	83	6.4%	+13%	-12%	18%
EIA 2012	Historical (Hydro)		12%	+26%	-23%	21%

A1.6 Variability in quarterly cash flow to equity investors

The variability of wind generation affects not only debt financing for a wind plant, as we describe in Section 3.1, but it also equity financing. Variations in cash flow to equity investors are significantly larger than variations in wind generation because the equity investments are leveraged by debt financing. We use the corrected reanalysis data to estimate the COV of quarterly cash flow to an equity investor for single wind plants or portfolios of two wind plants selected to reduce variability.

We create a simple financial model of a wind plant based on typical financing terms in the U.S. in 2013 (Wiser & Bolinger 2014). The installed capital cost of a wind plant is \$2 x 10^6 /MW, a portion of which is financed with debt at 6% interest for 15 years. The debt load is determined as 1.2 times the 10-year average annual 10^{th} percentile revenue (P90), which is calculated from the reanalysis data described in Section 2.3 with a fixed energy price of \$25/MWh, operating cost of \$24/MWh, and federal production tax credit of \$23/MWh. This yields debt financing of approximately 35% of the capital cost, which is significantly lower than the historical averages for wind plants because the capital cost in 2013 was higher than the historical average and the energy price was significantly lower.

Given those financing terms, we calculate the COV of quarterly cash flow for individual wind plants and pairs of wind plants and plot the results in Figure 4. The median COV for individual sites is 32% with an inter-quartile range from 25 – 42%. Figure 4 also plots the COV values for various portfolios composed of two wind plants. This shows that aggregating two plants reduces variability of quarterly cash flow, similar to the effects of aggregating many plants we discuss in the previous section. The median COV of randomly-chosen pairs of sites is 28%, a decrease of 4 percentage points from the COV for individual sites. The variability can be reduced further by carefully selecting pairs of sites. The median COV for optimally-chosen pairs (chosen to minimize average correlation of quarterly energy generation) is 23%. However, we find it is possible to achieve COV values nearly as low by selecting pairs of sites from the regions shown in Figure 5. The median COV values for pairs of sites selected from regions A, B, and C are 21 – 25%. The Supporting Information gives additional details on how the regions were determined.

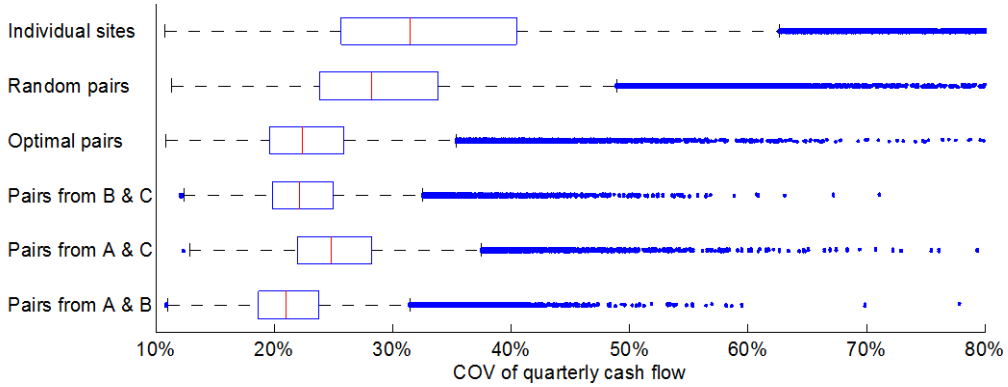

Figure 4: Coefficient of Variation (COV) of quarterly cash flow to equity investors for individual sites and sites paired according to various criteria. The red lines denote the median, boxes span the 25th and 75th percentile values, the whiskers extend to 1.5 times the inter-quartile range.

To test how much of the variability in our results is random rather than seasonal, we subtract the seasonal means from the quarterly cash flow and calculate a median COV of 23% for the same individual sites (a boxplot of these results is given in the Supporting Information). This shows that most of the variability is random, rather than seasonal. For comparison, Dunlop estimates a COV of 93% for an individual plant, which is significantly higher than the COV values we calculate (Dunlop 2004).

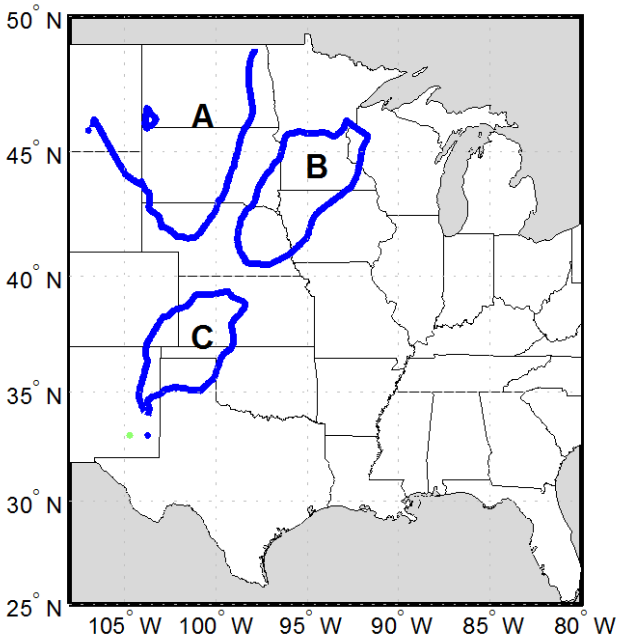

Figure 5: Regions that have high correlations in quarterly wind energy generation. The correlations between regions are low.

Discussion

In this work we develop a model to correct for biases and quantify uncertainties in wind energy calculated from Climate Forecast System Reanalysis (CFSR) data. This model is the first application of model output statistics to reanalysis wind data and the first to systematically quantify the uncertainties of the reanalysis data across many sites. We find the reanalysis data has a small positive bias: measured quarterly energy is 85% of the predicted value from CFSR data for matching locations and time periods. More importantly, we find energy predicted from CFSR data has significant uncertainties, dominated by between-site variability.

In spite of the between-site variability in the reanalysis data, we find robust results for measures based on wind energy aggregated over large areas. We estimate the COV of aggregate wind energy from 1,002 EWITS wind sites in the Great Plains to be 3.6% ± 0.1%. This inter-annual variability is much smaller than the variability at individual sites (5.9% - 15% ± 2.9%), which demonstrates the smoothing effect of aggregating wind plants spread across a large area. We also show robust reductions in the variability of quarterly cash flow to equity investors when pairs of wind plants from certain regions are combined into portfolios.

The significant between-site variability (σ = 376 MWh/quarter for a 2-MW turbine) suggests two possible sources of error. First, the reanalysis data poorly models terrain features because of the low resolution of the reanalysis model relative to the size of features that affect wind power. We compared the reanalysis data from empirical data selected from only flat regions, but it is possible that we did not detect subtle terrain features (e.g. ridges, upwind vegetation) that affect wind energy production. In our future research we will validate these CFSR data in complex terrain, similar to Henson et al. (Henson et al. 2012), and at offshore sites similar to The Crown Estate (The Crown Estate 2014). Second, the method we use to extrapolate hourly wind speed to hub height may introduce uncertainties. Our future work will test whether uncertainties in the uncertainties in the

Obukhov length calculated from reanalysis data introduce uncertainties in the extrapolated wind speed. Our ultimate goal is to reduce the uncertainties for hub-height wind speeds so that reanalysis data can accurately model wind power from short time periods (hours) and small areas. This would allow reanalysis data to be used to estimate measures such as the capacity value of wind power, which depends on the coincidence between wind generation and demand for electricity.

To review, this work demonstrates a model that corrects biases in the CFSR data and quantifies its uncertainties. We find that CSFR data over-predicts wind plant generation output by ~15% and that year-to-year variability of Great Plains wind is likely to be less than half that of aggregate U.S. hydropower.

Acknowledgements

This work was funded by the R. K. Mellon Foundation through a grant to The RenewElec Project at Carnegie Mellon University. The authors would like to thank Bob Dattore, John Zack, George Young, Paul Knight, Fallaw Sowell, and Christopher Wikle for their insights and assistance.

A1.7 References

Arya, S.P., 1988. *Introduction to Micrometeorology* 1st ed., San Diego: Academic Press.

Baker, R.W., Walker, S.N. & Wade, J.E., 1990. Annual and seasonal variations in mean wind speed and wind turbine energy production. *Solar Energy*, 45(5), pp.285–289.

Brower, M., 2009. *Development of Eastern Regional Wind Resource and Wind Plant Output Datasets*, Golden, CO: National Renewable Energy Laboratory.

Brower, M.C. ed., 2012. Chapter 12. In *Wind Resource Assessment*. Hoboken: John Wiley & Sons, Inc.

Brower, M.C. et al., 2013. *A Study of Wind Speed Variability Using Global Reanalysis Data*, AWSTruepower.

Dunlop, J., 2004. Modern portfolio theory meets wind farms. *The Journal of Private Equity*, 7(2), pp.83–95.

EIA, 2012. *Annual Energy Review 2011*, Washington, DC: U.S. Energy Information Administration. Available at: http://www.eia.gov/totalenergy/data/annual/pdf/aer.pdf.

Elliot, D. et al., 2010. 80 and 100 Meter Wind Energy Resource Potential for the United States. *WINDPOWER 2010*.

Emeis, S., 2013. *Wind Energy Meteorology*, Berlin, Heidelberg: Springer-Verlag Berlin Heidelberg.

EnerNex, 2010. *Eastern Wind Integration and Transmission Study (EWITS)*, Golden, CO: National Renewable Energy Laboratory.

Fiebrich, C.A., 2009. Earth-Science Reviews. *Earth Science Reviews*, 93(3-4), pp.77–84.

Fisher, S.M. et al., 2013. The effects of geographical distribution on the reliability of wind energy. *Applied Geography*, 40(0), pp.83–89.

Fleagle, R.G. & Businger, J.A., 1980. *An Introduction to Atmospheric Physics* 2nd ed., New York: Academic Press.

Giebel, G., 2000. *Equalizing Effects of the Wind Energy Production in Northern Europe Determined from Reanalysis Data*, Roskilde, Denmark: Risø National Laboratory.

Henson, W.L.W., McGowan, J.G. & Manwell, J.F., 2012. Utilizing Reanalysis and Synthesis Datasets in Wind Resource Characterization for Large-Scale Wind Integration. *Wind Engineering*, 36(1), pp.97–110.

Holt, E. & Wang, Jun, 2012. Trends in Wind Speed at Wind Turbine Height of 80 m over the Contiguous United States Using the North American Regional Reanalysis (NARR). *Journal of Applied Meteorology and Climatology*, 51(12), pp.2188–2202.

Huang, J., Lu, X. & McElroy, M.B., 2014. Meteorologically defined limits to reduction in the variability of outputs from a coupled wind farm system in the Central US. *Renewable Energy*, 62(C), pp.331–340.

Katzenstein, W. & Apt, J., 2012. Energy Policy. *Energy Policy*, 51(C), pp.233–243.

Katzenstein, W., Fertig, E. & Apt, J., 2010. The variability of interconnected wind plants. *Energy Policy*, 38, pp.4400–4410.

Larsén, X.G. & Mann, J., 2009. Extreme winds from the NCEP/NCAR reanalysis data. *Wind Energy*, 12(6), pp.556–573.

Liléo, S. & Petrik, O., 2011. Investigation on the use of NCEP/NCAR, MERRA and NCEP/CFSR reanalysis data in wind resource analysis. In EWEA 2011. Brussels, pp. 1–18.

Mesinger, F. et al., 2006. North American regional reanalysis. *Bulletin of the American Meteorological Society*, 87(3), pp.343–360.

Milligan, M., 1997. Wind plant capacity credit variations: a comparison of results using multiyear actual and simulated wind-speed data. In Windpower'97 Proceedings; June 15–18, 1997. Austin, TX.

Ntzoufras, I., 2009. *Bayesian modeling using WinBUGS*, Hoboken, N.J.: Wiley. Available at: http://search.library.cmu.edu/vufind/Record/1554586.

Panofsky, H.A., 1963. Determination of stress from wind and temperature measurements. *Quarterly Journal of the Royal Meteorological Society*, 89(379), pp.85–94.

Potter, C.W., Gil, H.A. & McCaa, J., 2007. Wind Power Data for Grid Integration Studies. In 2007 IEEE Power Engineering Society General Meeting. IEEE, pp. 1–6.

Pryor, S.C. et al., 2009. Wind speed trends over the contiguous United States. *Journal of Geophysical Research: Atmospheres*, 114(D14).

Rienecker, M.M. et al., 2011. MERRA: NASA's Modern-Era Retrospective Analysis for Research and Applications. *Journal of Climate*, 24(14), pp.3624–3648.

Saha, S., Moorthi, S., Pan, H.-L., Wu, X., Wang, Jiande, et al., 2010a. The NCEP Climate Forecast System Reanalysis. *Bulletin of the American Meteorological Society*, 91(8), pp.1015–1057.

Saha, S., Moorthi, S., Pan, H.-L., Wu, X., Wang, Jie, et al., 2010b. NCEP Climate Forecast System Reanalysis (CFSR) 6-hourly Products, January 1979 to December 2010. Available at: http://dss.ucar.edu/datasets/ds093.0/.

Simonsen, T.K. & Stevens, B.G., 2004. Regional Wind Energy Analysis for the Central United States. In Global WINDPOWER 2004. Chicago, pp. 1–16.

The Crown Estate, 2014. *UK MERRA Validation With Offshore Meteorological Data*, London: The Crown Estate.

Thomas, A. et al., 2006. Making BUGS Open. *R News*, 6(1), pp.12–17.

Tindal, A., 2011. Financing wind farms and the impacts of P90 and P50 yields. In EWEA Wind Resource Assessment Workshop. Brussels.

Uppala, S.M. et al., 2006. The ERA-40 re-analysis. *Quarterly Journal of the Royal Meteorological Society*, 131(612), pp.2961–3012.

Wan, Y.-H., 2012. *Long-Term Wind Power Variability*, Golden, CO: National Renewable Energy Laboratory.

Wiser, R.H. & Bolinger, M., 2014. *2013 Wind Technologies Market Report* G. L. Barbose et al., eds.,

Supporting Information

A2 Power Curve

We use the piecewise power curve of a generic 2-MW wind turbine designed for IEC Class II wind conditions, used in the EWITS study (Brower 2009).

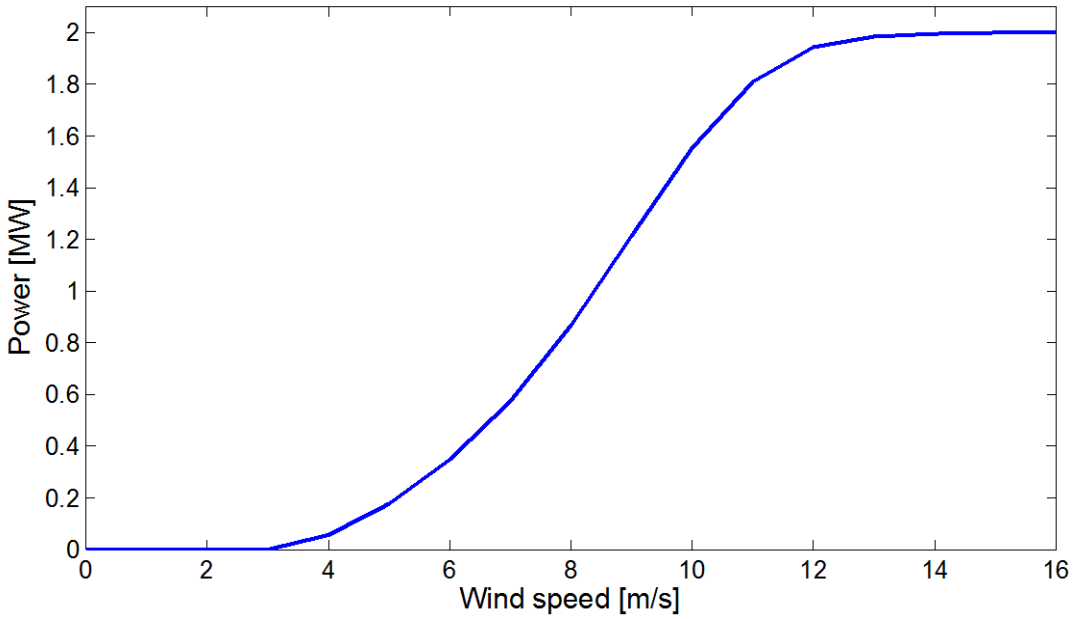

A3 Capacity Factor Comparisons

The first figure below plots our calculation of the capacity factor of the EWITS sites as a function of the cumulative capacity of the wind plants. We present these results for comparison to the similar curve calculated for the EWITS study (Figure 2-2 in (EnerNex 2010)). Although we exclude EWITS sites in the Northeast and Mid-Atlantic states, our curve below matches Figure 2-2 in the EWITS study well within our error bounds.

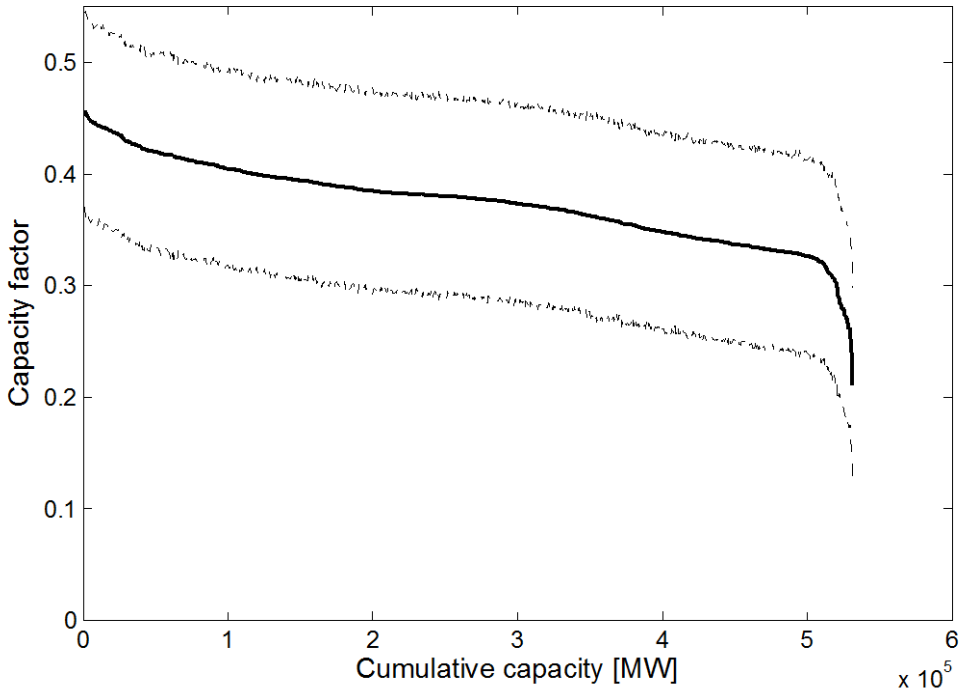

Figure 6: Cumulative capacity factor for the wind plants used in this study.

The figure below plots mean COV of annual energy against mean capacity factor for each site. There is a clear relationship between the COV and capacity factor: the annual generation is less variable at plants with higher capacity factors.

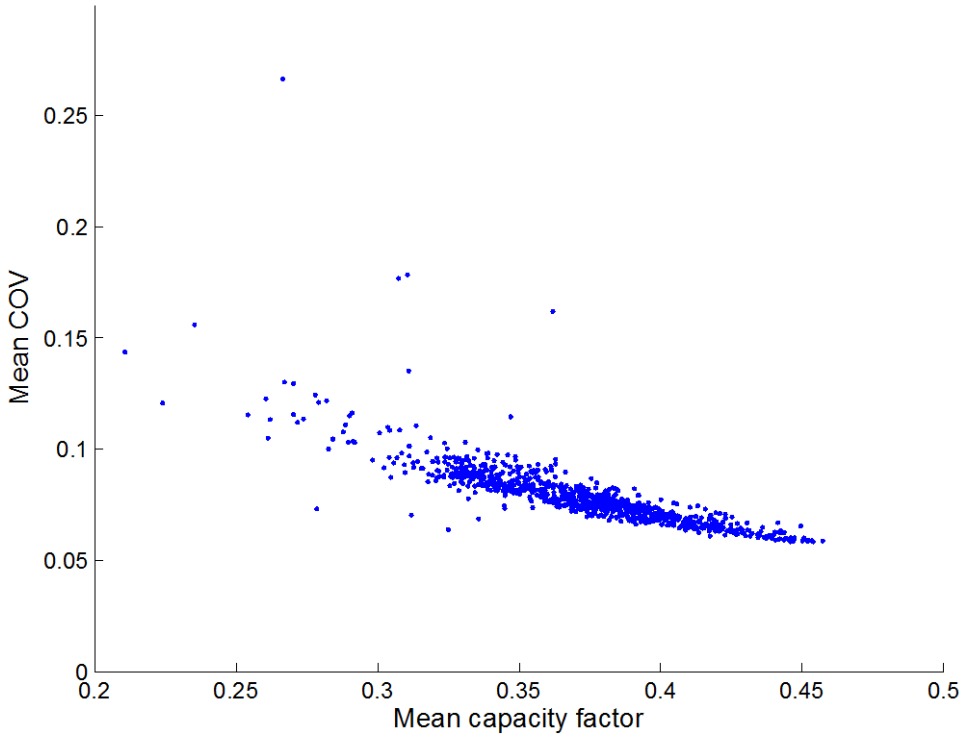

Figure 7: Mean COV vs. mean capacity factor for individual EWITS sites in the U.S. Great Plains (calculated from reanalysis data described in the body of the paper)

A4 Comparison of Aggregate Annual Generation

The figure below recalculates Figure 2 using only the 16 sites analyzed by Katzenstein (Katzenstein et al. 2010) and the same turbine power curve used b Katzenstein. Our results match well with Katzenstein's within our confidence bounds.

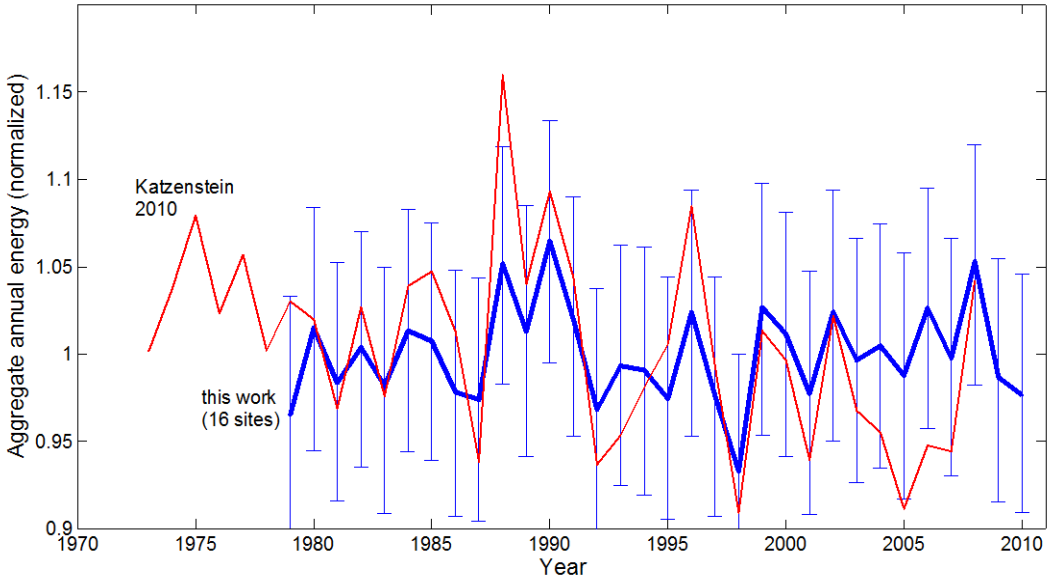

Figure 8: Aggregate annual energy calculated from reanalysis data for the 16 ASOS sites used by Katzenstein (Katzenstein et al. 2010)

A5 COV of quarterly cash flow with seasonal means subtracted

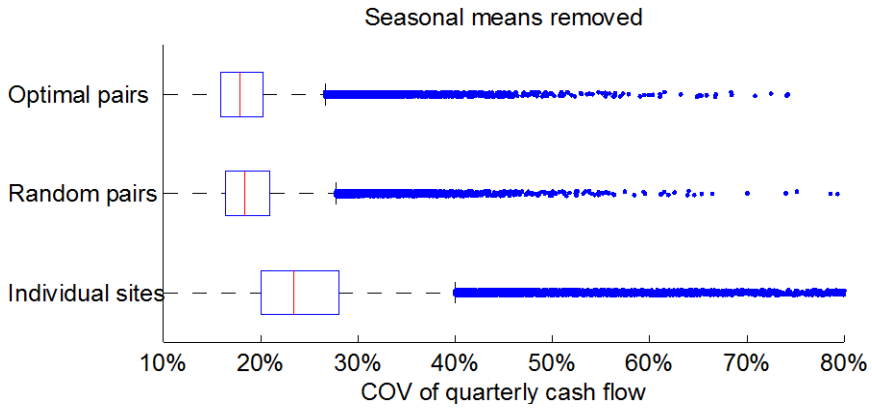

Figure 9: COV of quarterly cash flow to equity investors after seasonal means are removed. Compare to Figure 4 in the body of the paper.

A6 Defining Regions in Figure 5

The regions A, B, and C in Figure 5 of the body of the paper are determined by clustering the EWITS sites based on quarterly wind energy generation. We apply a k-means clustering algorithm tusing correlation of quarterly energy as the distance metric. The clustering is repeated 10^3 times, once for each probable realization of quarterly energies. Sites that are assigned to the same cluster in at least 70% of the realizations are considered robustly clustered, though our results and the boundaries of those clusters are not very sensitive to that threshold. Finally, the regions are defined by fitting contour lines around the sites that are robustly assigned to each cluster.

We selected $k = 5$ clusters heuristically as a balance between cluster size and between-cluster discrimination. More clusters increase the smoothing effect of combining sites from different

clusters, but reduces the number of sites that can be selected to construct portfolios. Clusters D and E are not shown because pairs of wind plants selected from those clusters did not show significant reduction in the COV of quarterly cash flow.

A7 Calculation of Obukhov Length *L*

	Description [units]	**Source**
$L = \frac{-u_*^3 T_{v0}}{\kappa\, g \left(\frac{H_{v0}}{\rho\, C_p}\right)}$	Obukhov length [m]	eq. 12.11 in (Arya 1988)
u_*	friction velocity [m/s]	from reanalysis data
$T_v = (1+0.61q)T$	virtual temperature	eq. 2.95 in (Fleagle & Businger 1980)
q	specific humidity [kg/kg]	from reanalaysis data
T	temperature 2m above ground [K]	from reanalysis data
$\kappa = 0.4$	von Kármán constant	
$g = 9.8$	gravity [m/s^2]	
$H_v = H + a_\theta H_L$	virtual heat flux [W/m^2]	eq. 12.10 in (Arya 1988)
H	sensible heat flux [W/m^2]	from reanalysis data
$a_\theta = 0.61 C_p \Theta / L_e$	dimensionless coefficient	
$C_p = C_{pd}\,(1+0.84*q)$	specific heat of moist air at const. press. [J/kg*K]	
$C_{pd} = 1004.67$	specific heat of dry air at const. press. [J/kg*K]	
$\Theta = T\left(\frac{p_0}{p}\right)^{\kappa_p}$	potential temperature [K]	
$p_0 = 10^5$	standard pressure [Pa]	
p	atmospheric pressure at ground level [Pa]	from reanalysis data
$\kappa_p = 2/7$	Poisson constant	
$L_e = 2.257 \times 10^6$	latent heat of evaporation [J/kg]	
H_L	latent heat flux [W/m^2]	from reanalysis data
$\rho = \rho_d((1+q)/(1+q(R_w/R_a)))$	Density of moist air [kg/m^3]	
$\rho_d = p/(R_a T)$	Density of dry air [kg/m^3]	
$R_a = 287.058$	Specific gas constant for dry air [J/kg*K]	
$R_w = 461.5$	Specific gas constant for water vapor [J/kg*K]	

A8 MCMC Procedure and Code

The OpenBUGS code we use to model the biases and uncertainties is given below. We simulate 10^4 samples. The "thin" parameter is set to 20 to record every 20th sample, which means OpenBUGS actually simulates 20 x 10^4 samples but does not display them to the user. We exclude the first 5 x 10^3 for burn-in.

A8.1 Code

```
## Inputs
# X: reanalysis quarterly wind energy for each site, sorted by site num
# Y: historical quarterly wind energy for each site, sorted by site num
# N_obs: total number of observations
# N_sites: number of sites
# meas_this_site: number of measurements available for each site
# my_site: site number corresponding to each mesaurement
# start_idx_this_site: index of first measurement for this site

model{
   for(i in 1:N_sites){
      for(j in 1:meas_this_site[i]){
         Y[start_idx_this_site[i]+j] ~ dnorm(mu[i,j], tau)
         mu[i,j] <- m*X[start_idx_this_site[i]+j] + a[i]
      }
      a[i] ~ dnorm(0, tau.a)  # offset, this site (between-site variability)
   }

## Prior Distributions (vague priors)
   m ~ dnorm(0, 0.001)  # slope
   tau.a ~ dgamma(0.001, 0.001)
   sigma.a <- 1/sqrt(tau.a)
   tau ~ dgamma(0.001, 0.001)
   sigma <- 1/sqrt(tau)
}
```

A8.2 Initial Values

```
m = 1
a = c(0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0,
0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)
tau.a = 0.2
tau = 0.1
```

A9 Empirical Data

The empirical data we use in this work for validation was provided by the University of North Dakota Energy & Environmental Research Center (Simonsen & Stevens 2004). We exclude sites within 1 km of forested areas or buildings. The data are divided randomly into a training subset and a validation subset.

A9.1 Training Data

Site ID	Latitude	Longitude	Height [m]	Start date	End Date	Mean Wind Speed [m/s]
14	46.2646	-96.5312	70	01-May-1996	17-Apr-2006	7.1
21	43.7828	-93.3472	70	07-Jun-1996	14-Mar-2002	7.3
27	44.0984	-95.5636	90	17-Aug-2000	26-Sep-2004	8.2
32	46.3861	-98.8645	51	13-Jul-2000	14-Aug-2005	7.8
59	45.0680	-96.0527	50	16-Jan-1998	26-Jun-2002	5.8
71	44.0416	-94.8477	70	01-Sep-2007	05-Oct-2007	8.5
76	48.7800	-98.0378	55	05-Jul-1995	30-May-2002	7.3
82	47.9870	-98.0097	55	05-Jul-1995	01-Jan-1998	7.7
89	46.7578	-96.6038	70	02-Jun-1995	08-Aug-2000	6.8

100	47.5003	-97.8727	85	07-Aug-2003	01-May-2005	8.4
117	41.1900	-92.5700	50	01-May-1995	29-Dec-1999	6.2
118	43.2700	-94.8600	50	01-May-1995	01-Jan-2000	8.0
120	43.2800	-93.6300	50	01-May-1995	08-Nov-1999	7.0
121	43.2500	-96.4700	50	01-May-1995	14-Dec-1999	6.9
122	41.4500	-91.1100	50	01-Jan-1996	16-Nov-1997	6.1
123	42.3000	-93.4400	50	01-May-1995	02-Jan-2000	7.5
125	43.4000	-95.6800	50	01-May-1995	02-Jan-2000	7.5
126	43.0000	-95.5200	50	01-May-1995	30-Dec-1999	7.3
127	41.9700	-95.5100	50	01-May-1995	29-Dec-1999	6.4
129	41.0000	-95.1700	50	01-May-1995	05-Dec-1999	6.9
141	43.7315	-88.2927	60	09-Jul-1998	11-Jul-2001	7.0
143	44.3588	-87.9857	83	09-Nov-1999	01-Oct-2001	7.3
147	43.3092	-89.3232	60	05-Oct-2000	01-Oct-2001	6.2
161	42.6742	-99.0972	50	05-May-1995	18-Nov-1996	7.2
189	46.3140	-101.8812	50	14-Jun-2002	03-Mar-2006	7.8
229	48.7646	-101.6947	80	02-Dec-2004	07-Aug-2006	8.2
257	38.9403	-104.3083	70	26-Dec-2002	31-Jan-2003	8.3
267	43.9050	-98.7533	70	26-Dec-2001	01-Jan-2006	8.8
274	44.1292	-99.4364	70	21-Nov-2001	14-Jan-2007	7.7
277	38.8469	-98.1889	80	18-Apr-2003	03-Sep-2005	9.0
278	39.7842	-98.1192	80	23-Apr-2003	04-Sep-2005	8.3
280	38.9592	-100.8526	80	01-May-2003	04-Sep-2005	8.5
289	37.2733	-97.5653	80	11-Jun-2003	03-Sep-2005	7.8
293	37.9957	-88.5579	70	12-Dec-2005	08-Jun-2006	5.6
294	41.3681	-90.1552	50	02-Apr-2005	07-Mar-2006	6.3
301	40.8160	-90.3954	50	03-Apr-2005	27-Feb-2006	6.1
306	35.4211	-99.5392	70	30-Oct-2003	27-Jan-2005	8.4
344	48.0508	-101.3900	50	04-Oct-2006	01-Jan-2008	8.3
353	46.0858	-97.9465	60	01-Dec-2004	01-May-2007	7.6

A9.2 Validation Data

Site ID	Latitude	Longitude	Height [m]	Start date	End Date	Mean Wind Speed [m/s]
2	45.5841	-96.0476	70	01-May-1995	23-Jul-2005	7.0
15	43.7287	-95.3566	70	22-Jun-1995	20-Dec-2007	7.4
16	44.7454	-94.3950	70	21-Jun-1995	15-Apr-2000	6.5
19	48.7052	-97.2657	50	19-Oct-2001	22-Apr-2007	7.1
20	43.8876	-95.9285	70	02-May-1996	18-Apr-2006	8.5
23	47.7563	-96.6697	70	01-Jun-1995	26-May-2005	6.7
35	47.5004	-97.8166	70	01-Jul-2000	31-May-2004	7.5
42	47.0681	-102.9272	55	05-Jul-1995	01-Jan-1998	7.4
46	48.8105	-96.9384	70	01-Jun-1995	20-Dec-2007	6.8
48	43.9877	-96.2080	90	21-Oct-1998	19-Dec-2007	7.5
51	35.0497	-99.0971	70	29-Mar-2002	05-Jan-2004	7.6
58	46.2988	-98.8660	70	13-Jul-2000	14-Aug-2005	8.5
60	44.3746	-92.2353	50	05-Jul-2001	06-Aug-2004	5.6
64	44.4139	-95.2928	70	14-Jul-2007	19-Dec-2007	6.9
67	44.4455	-96.0122	90	17-Aug-2000	13-Feb-2002	8.0
90	43.7823	-95.8735	90	16-Sep-1998	05-Dec-2007	7.6
102	47.1392	-100.7058	55	05-Jul-1995	01-Jan-1998	7.8
103	43.7041	-94.0293	70	22-Jun-1995	20-Dec-2007	7.4
108	43.7106	-96.0694	70	21-Jun-1995	02-Jul-2005	7.1
112	43.0400	-94.1400	50	01-Jan-1996	02-Jan-2000	7.2

113	42.8400	-95.3400	50	01-May-1995	02-Jan-1997	7.5
116	41.4700	-94.5400	50	01-Jan-1998	30-Dec-1999	7.3
139	42.9527	-90.4306	60	14-Nov-1997	01-Feb-2001	6.3
157	40.4417	-101.9917	50	01-May-1995	01-Apr-1997	7.7
158	41.0944	-103.6094	50	01-May-1995	01-Apr-1997	7.4
191	48.8925	-99.5606	50	14-Aug-2003	01-Nov-2005	7.6
196	43.7331	-92.0858	90	08-Aug-2003	26-Mar-2005	6.6
203	44.7586	-94.7158	70	30-Jul-2003	28-Oct-2006	7.2
220	46.1822	-103.0489	65	26-Sep-2002	07-Jul-2004	9.2
236	40.4173	-93.0390	50	14-Jan-2005	06-May-2007	6.3
244	45.1789	-97.0158	70	20-Apr-2002	01-Jan-2006	8.8
269	45.2058	-97.9053	70	19-Jan-2002	17-Feb-2005	8.2
275	45.6886	-99.2539	70	04-Dec-2001	01-Dec-2006	9.1
279	37.7373	-101.2091	80	29-Apr-2003	03-Sep-2005	8.7
281	38.4517	-99.8151	80	04-Jun-2003	04-Sep-2005	8.0
291	39.5120	-87.9659	50	17-Apr-2005	31-May-2006	6.3
342	46.9468	-99.1086	65	08-Dec-2003	01-Apr-2007	8.5
346	47.7337	-102.6580	50	03-Aug-2006	05-Jul-2007	7.2
350	48.6167	-98.4369	60	06-Jan-2006	07-May-2007	8.1