




The Distribution of Wind Power Forecast Errors from Operational Systems

Bri-Mathias Hodge, Erik Ela, and Michael Milligan
National Renewable Energy Laboratory

Presented at the 10th International Workshop on Large-Scale Integration of Wind Power into Power Systems
Aarhaus, Denmark
October 25–26, 2011

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Conference Paper
NREL/CP-5500-52568
February 2015

Contract No. DE-AC36-08GO28308

NOTICE

The submitted manuscript has been offered by an employee of the Alliance for Sustainable Energy, LLC (Alliance), a contractor of the US Government under Contract No. DE-AC36-08GO28308. Accordingly, the US Government and Alliance retain a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for US Government purposes.

This report was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or any agency thereof.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Available electronically at http://www.osti.gov/scitech

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone: 865.576.8401
fax: 865.576.5728
email: mailto:reports@adonis.osti.gov

Available for sale to the public, in paper, from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
phone: 800.553.6847
fax: 703.605.6900
email: orders@ntis.fedworld.gov
online ordering: http://www.ntis.gov/help/ordermethods.aspx

Cover Photos: (left to right) photo by Pat Corkery, NREL 16416, photo from SunEdison, NREL 17423, photo by Pat Corkery, NREL 16560, photo by Dennis Schroeder, NREL 17613, photo by Dean Armstrong, NREL 17436, photo by Pat Corkery, NREL 17721.

NREL prints on paper that contains recycled content.

The Distribution of Wind Power Forecast Errors from Operational Systems

Bri-Mathias Hodge, Erik Ela, and Michael Milligan

***Abstract* — Wind power forecasting is one important tool in the integration of large amounts of renewable generation into the electricity system. Wind power forecasts from operational systems are not perfect, and thus, an understanding of the forecast error distributions can be important in system operations. In this work, we examine the errors from operational wind power forecasting systems, both for a single wind plant and for an entire interconnection. The resulting error distributions are compared with the normal distribution and the distribution obtained from the persistence forecasting model at multiple timescales. A model distribution is fit to the operational system forecast errors and the potential impact on system operations highlighted through the generation of forecast confidence intervals.**

***Index Terms*— forecasting, distribution functions, power systems, wind power generation, stochastic systems**

I. Introduction

Wind power is playing an increasingly important role in power system operations as the amount of energy supplied by wind continues to increase. Wind power, however, has a maximum generation limit that changes through time (variability) and cannot be predicted with perfect accuracy (uncertainty). One way system operators attempt to deal with the variability and uncertainty of wind power is through wind power forecasting. While forecasting can help to reduce the impact of wind power uncertainty, even the best forecasts are not perfectly accurate. Understanding how operational wind forecasting systems are inaccurate can lead to more efficient system operations. This may be in the form of a formal process for considering the forecast uncertainty, such as with a stochastic unit commitment process, or through informal decisions made to increase system flexibility when the consequences of an inaccurate forecast could significantly hinder system operations. One method of utilizing the information contained in previous forecast errors is through the production of forecast confidence intervals. Interval forecasts can complement the point forecasts normally used by establishing bounds on the expected value, within certain probabilities. Calculating the forecast confidence intervals usually relies on assuming an error distribution for the point forecast. The wind power forecast errors are often assumed to follow a normal distribution [1-4], despite the fact that this provides a poor fit for most forecasting methods and timescales. Weibull [5], beta [6] and Cauchy [7] distributions have also been utilized. Lange studied wind power forecast- error distributions for timescales between 6 and 48 hours ahead, focusing on the errors incurred while translating wind speed data to wind power output [8]. It was found that while numerical weather prediction (NWP) model wind speed error distributions were normal, the wind power error distributions were not. Focken et al. demonstrated that geographic diversity in wind plant location could provide smoothing of the forecast errors of aggregate wind power at longer timescales [9].

The remainder of the paper is organized as follows. In Section II, the methods and data used in this study are detailed. Section III reports on the results of analyzing wind power forecasting-error distributions from different forecasting methods over a number of timescales. Section IV shows the results of matching a model distribution to the forecast errors. Conclusions are then drawn and future areas for examination outlined in Section V.

II. Methods and Data

In this section, we describe some of the important methods and data utilized in the study. Section II-A contains information on the datasets analyzed. Section II-B provides some background on statistical distributions and methods of characterizing the same.

A. Data Utilized

In this study, we have utilized data from two distinct areas of the United States. One dataset includes aggregated wind power output and forecasts from all of the wind power plants installed in the ERCOT interconnection for a 13-month period. This dataset includes all day-ahead forecasts, made once a day at 16:00 the day prior, for the 13-month period. Practically, this means that the data is for forecasts provided between 8 and 31 hours in advance. Another dataset used in this work includes the forecasts and actual production from a wind plant in the Xcel Energy Colorado territory with an approximate nameplate capacity of 300 MW. This dataset includes three months of data from the summer and fall seasons with hourly forecasts produced every 15 minutes for the next 72 hours. Both datasets provide useful information as the day-ahead ERCOT forecasts are those most likely to be used in the unit commitment process, while the Xcel forecasts provide information on how the forecasts improve with a shorter forecasting interval. It is important to note that the two datasets utilize methodologically similar forecasting systems, based on NWP models, but the forecasts come from different forecast providers.

B. Statistical Distributions

Probability density functions are used to describe the range of values that a random variable can take, and the likelihood of a sample falling in a particular range. A number of different probability distribution functions have been applied to wind power forecasting errors including the normal distribution [1-4], the Weibull distribution [5], the beta distribution [6] and the Cauchy-Lorentz distribution [7]. In addition to these model parametric distributions, we will also consider the Laplace and hyperbolic distributions. The Laplace distribution can help to better represent the ends of the distribution when there are semi-heavy tails, an important consideration in wind power forecasting. The hyperbolic distribution is similar to the Laplace distribution, except that it allows asymmetry between the two sides of the distribution. Both are subclasses of the generalized hyperbolic distribution. Fig. 1 shows examples of some of the different distributions thus far mentioned.

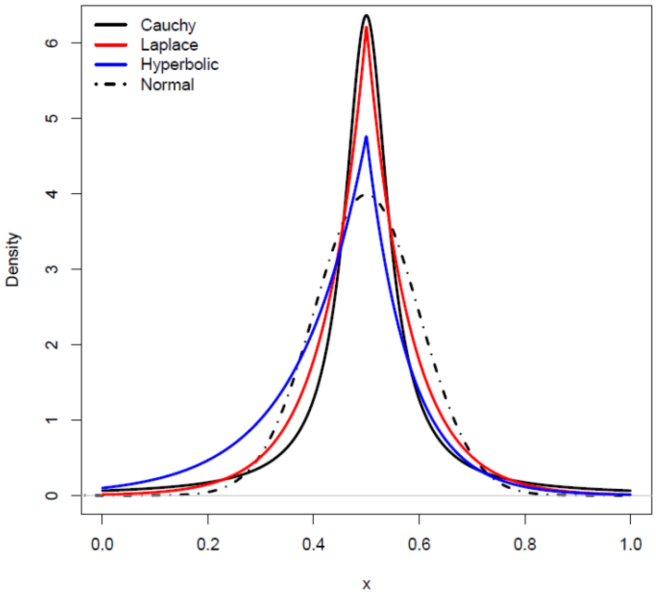

Fig. 1. Illustration of the different probability distributions on the x-interval [0,1]. The Cauchy distribution displayed has parameter values of $x_o = 0.5$ and $\gamma = 0.026$. The Laplace distribution has parameter values of $\mu = 0.5$ and $b = 0.055$. The hyperbolic distribution has parameters $\pi = -0.24$, $\zeta = 9.62\ E^{-5}$, $\delta = 9.79\ E^{-6}$, $\mu = 0.5$. The normal distribution is shown with mean = 0.5 and standard deviation = 0.1.

While the first two standardized moments are often applied to the characterization of wind forecasting error distributions, we will also consider the third and fourth moments. The third moment is known as skewness and is a measure of the probability distribution's asymmetry. A negative skew is one where the left side of the distribution has a longer tail, but the bulk of the values are on the right hand side. A positive skew indicates the opposite is true. The fourth moment is known as kurtosis and is a measure of the magnitude of the peak of the distribution. Alternatively, kurtosis is a measure of the thickness of the tails of the distribution. A distribution with a high kurtosis value is known as leptokurtic. Leptokurtic distributions possess a more pronounced peak, slimmer shoulders, and longer tails when compared to a normal distribution with the same mean and variance. The Cauchy and Laplace distributions shown in Fig. 1 demonstrate a few examples of leptokurtic distributions. In the rest of this work when we refer to the kurtosis of a distribution, we will mean specifically the excess kurtosis; that is the kurtosis above that of the normal distribution.

III. Results

Having established the importance of forecast error distributions and described the relevant statistical background, we now characterize the distributions for the datasets under consideration. Section III-A compares the distributions resulting from operational forecasting methods with the normal distribution. In Section III-B, we compare the wind power forecast error distributions from an operational system at a single wind plant with those obtained from the persistence model over multiple timescales.

A. Operational Distributions vs. the Normal Distribution

While it has been common to assume that wind power forecasting errors follow a normal distribution, this simplistic assumption can lead to inaccuracies in both wind integration studies and actual system operations. Forecast errors from the persistence model have been shown to deviate significantly from the normal distribution [7], and an examination of operational forecasting errors reveals that the same conclusion may be drawn for timescales relevant to power system operations. Forecasts that are an aggregation of multiple plants and multiple timescales would be the most likely to follow a normal distribution, following the central limit theorem. Yet even for the ERCOT day-ahead dataset, which is an aggregation of a number of wind plants (approximately 9000 MW of capacity) and 24 different forecasting periods, an examination of the forecast error distribution histogram reveals that the resulting errors are significantly non-normal, as seen in Fig. 2.

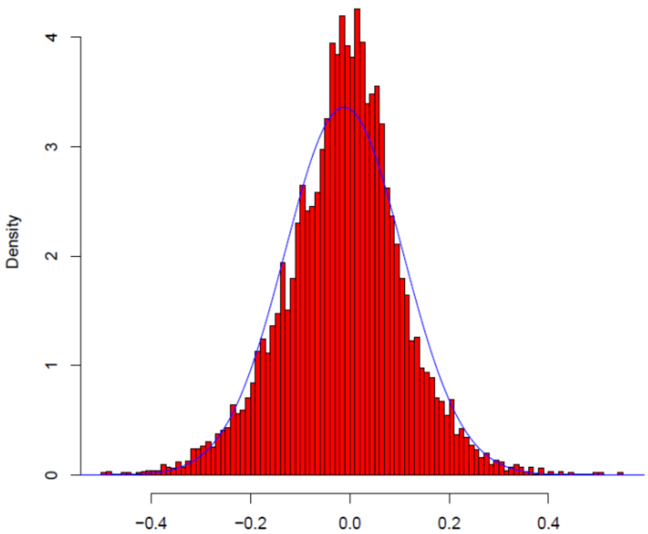

Fig. 2. Histogram of all of the day-ahead forecasts for the ERCOT system over a 13-month period. Note that this includes forecasts at different timescales, from 8 to 31 hours ahead. $\gamma = -0.62$; $\kappa = 1.03$. The blue line represents a normal distribution with the same mean and standard deviation.

The quantile-quantile (Q-Q) plot is a means by which two distributions can be graphically compared. Though the histogram shown in Fig. 2 seems to indicate that the forecast error distribution of the ERCOT day-ahead forecasts is not normal, the use of a normal Q-Q plot provides additional assurance. As is readily apparent in Fig. 3, the sample distribution deviates significantly from the normal distribution. A Shapiro-Wilk [10] test was run for the data, with the null hypothesis, that the sample ERCOT data comes from a normal distribution, being rejected at a significance level of $\alpha = 0.000001$, i.e., the 99.9999% confidence interval.

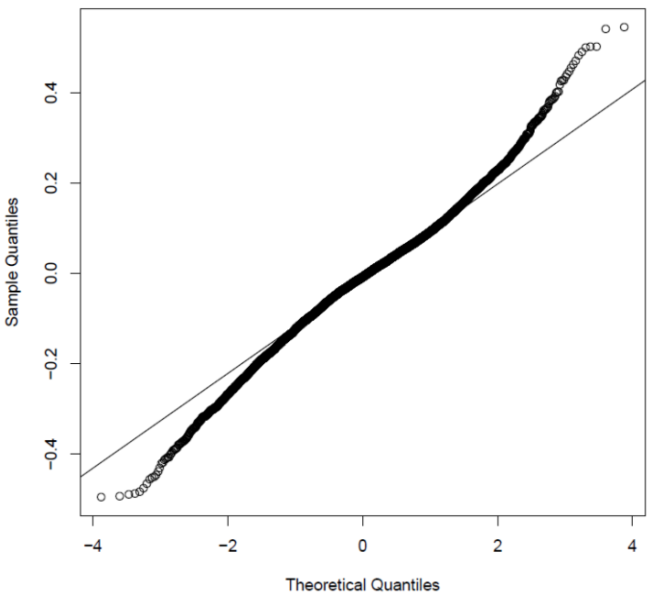

Fig. 3. Normal Q-Q plot for the ERCOT day-ahead forecasts over a 13-month period with forecasts made once per day. The line in the graph passes through the first and third quartiles of the observed data and should run through all of the data points if the distribution is normal.

Considering the operational forecast distributions for the single Xcel plant demonstrates that the normal distribution is also a poor representation of the observed forecast errors for smaller systems and shorter timescales. This is shown in the normal Q-Q plot of hour-ahead forecast errors in Fig. 4. The observed Xcel data deviates very strongly from the normal assumption throughout the distribution. A Shapiro-Wilk test on the hour-ahead forecast errors confirms the visual analysis and rejects the normal distribution hypothesis at the 99.9999% confidence interval. Further visual evidence is provided in the histogram of Fig. 6.

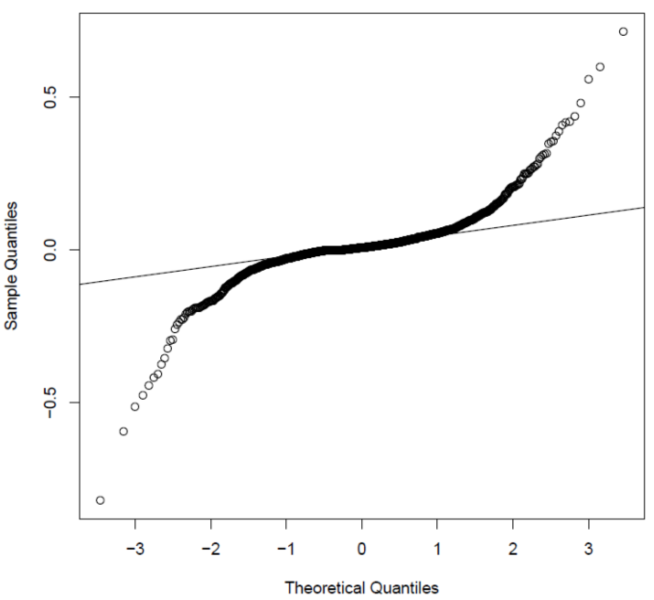

Fig. 4. Normal Q-Q plot for the Xcel plant one-hour forecasts over a 3-month period with forecasts made once per hour. The line in the graph passes through the first and third quartiles of the observed data and should run through all of the data points if the distribution is normal.

B. Operational vs. Persistence Forecasting Models

The persistence model is often used as a naïve method for short-term wind forecasting, especially in wind integration studies without access to NWP models. For this reason, it is interesting to compare the forecast error distributions that result from the persistence model with those from an operational model. One method used to examine the error distributions of the two forecasting methods is to graphically represent the distributions with histograms. The number of bins used exceeded the number needed according to Scott's rule [11] in all cases, and a value of n = 100 was found to work well over all of the timescales for the Xcel data.

First we will examine the hour-ahead timeframe, which is commonly used in the economic dispatch process. As may be seen in Fig. 5, the persistence model is fairly accurate at this timescale, producing many very accurate forecasts with a few large errors. However, the operational forecasting method significantly outperforms persistence, and creates a much different error distribution in the process. Fig. 6 shows the histogram from the operation forecasting system, with a significantly higher kurtosis value and a smaller skew; this indicates more accurate forecasts and less forecasting bias.

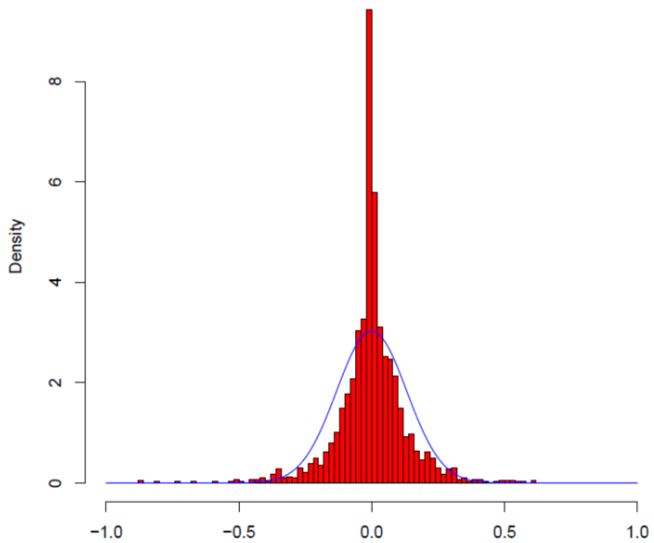

Fig. 5. Histogram of the one-hour-ahead persistence forecasts for one Xcel wind plant over a three-month period. γ = -0.51; κ = 5.97. The blue line represents a normal distribution with the same mean and standard deviation.

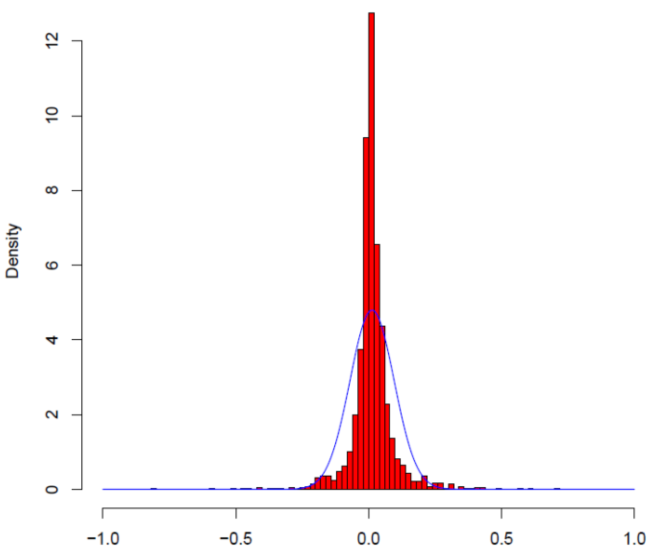

Fig. 6. Histogram of the one hour-ahead forecasts for one Xcel wind plant over a three month period. γ = -0.01; κ = 17.62. The blue line represents a normal distribution with the same mean and standard deviation.

Another important timeframe in power system operations is the day-ahead forecasting that occurs for the unit commitment process. While day-ahead forecasts span a number of hours, usually from 12-36 hours in advance, here we will highlight the 24-hour-ahead forecast as a representative sample. At the 24-hour-ahead forecast interval the performance of the persistence method is significantly worse, and has a noticeably different error distribution. Forecast errors that span the entire plant capacity are relatively common, as seen in Fig. 7.

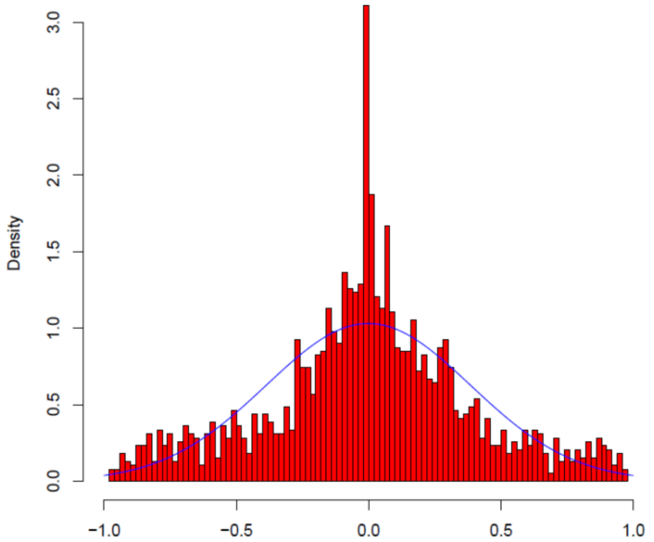

Fig. 7. Histogram of the 24-hour-ahead persistence forecasts for one Xcel wind plant over a three-month period. γ = -0.03; κ = 0.15. The blue line represents a normal distribution with the same mean and standard deviation.

The operational performance of the operational forecasting system at the 24-hour-ahead timeframe is also significantly worse than for the hour-ahead interval, as seen in Fig. 8. However, the performance of the operational method is still significantly better than the persistence method. An interesting trend in the error distributions from the operational method is that a significant skewness starts to develop at longer time frames with many small overproduction forecasts, but underproduction forecasts that are larger in magnitude.

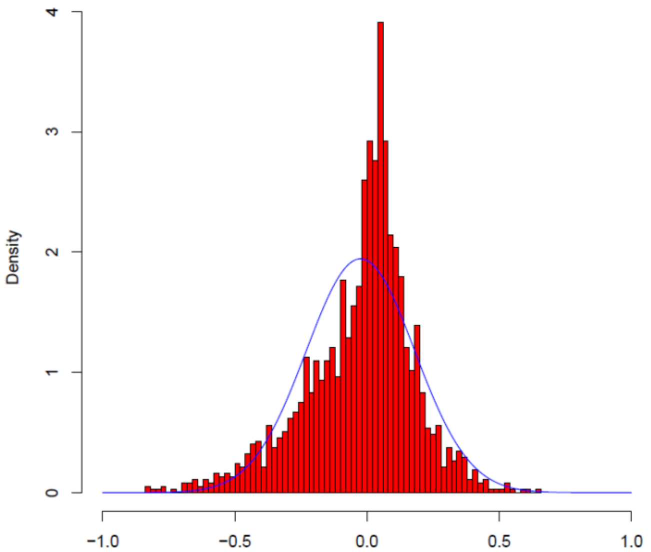

Fig. 8. Histogram of the 24-hour-ahead forecasts for one Xcel wind plant over a three-month period. γ = -0.65; κ = 1.05. The blue line represents a normal distribution with the same mean and standard deviation.

A comparison of the skewness and kurtosis values for both the persistence and operational forecasts over multiple forecasting intervals gives an idea of the trends that occur at increasing timescales. At increasingly large timescales the persistence model tends to trend toward an almost uniform distribution, though with significant mass remaining at the center. The operational model has significant skewness over most of the timescales considered, indicating a bias toward overproduction. The kurtosis values also indicate that the accuracy of the forecasts for both models decrease significantly as the forecasting period becomes longer, as would be expected.

TABLE I
SKEWNESS AND KURTOSIS VALUES FOR PERSISTENCE AND OPERATIONAL FORECAST ERRORS FOR THE XCEL DATASET AT DIFFERENT TIMESCALES

	Persistence		Operational	
Timescale	**Skewness**	**Kurtosis**	**Skewness**	**Kurtosis**
15 Minutes			-0.68	30.64
30 Minutes			-1.36	30.88
1 Hour	-0.51	5.97	-0.01	17.62
3 Hour	-0.21	1.81	-0.88	1.97
24 Hour	-0.03	0.15	-0.65	1.05
72 Hour	-0.03	0.06	-0.56	0.84

IV. DISTRIBUTION MODELING

Having established that the normal distribution is a poor representation of operational wind power forecasting errors over a number of timescales, we now examine how we may use a model distribution to better represent the observed error distributions. One of the most important considerations when examining the forecast error distributions is accurately modeling rare events, as represented by the tails of the error distribution. While operational forecasting systems produce very many accurate forecasts, they also regularly produce forecasts that are quite inaccurate, and could potentially lead to operational issues, depending on the state of the rest of the system. Understanding the frequency and magnitude of large forecast errors will help the system operator prepare for these contingency-like events.

To better represent the distributions obtained from operational forecasting systems, we have chosen the hyperbolic distribution as a model distribution. Distribution parameters were fit to the data using a maximum likelihood method implemented in the *hyperbFit* function of the *HyperbolicDist* package [12] in the *R* statistical computing environment [13]. Visual inspection of the histogram reveals that the fitted hyperbolic function provides a better representation for the observed error distribution for the ERCOT day-ahead forecasts than does the normal distribution, as shown in Fig. 9. The improvement is even more noticeable for the case of the single Xcel wind plant at the one-hour forecasting interval, displayed in Fig. 10. Both fitted distributions are able to capture the leptokurtic nature of the observed errors, as well as the skewness of the distribution. The fitted distributions tend to more accurately represent the taller peaks and slimmer shoulders of the error distributions when compared with the normal distribution. Additionally, the semi-heavy tails of the forecast error distributions are more accurately accounted for with the model distributions.

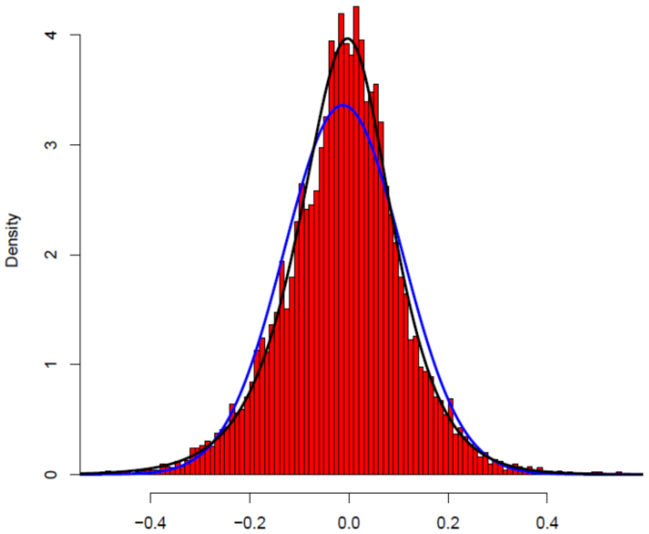

Fig. 9. Histogram of all of the day-ahead forecasts for the ERCOT system over a 13-month period. The blue line represents a normal distribution with the same mean and standard deviation. The black line represents a hyperbolic distribution fit to the data with: $\pi = 0.083$, $\zeta = 1.601$, $\delta = 0.105$, $\mu = 0.006$.

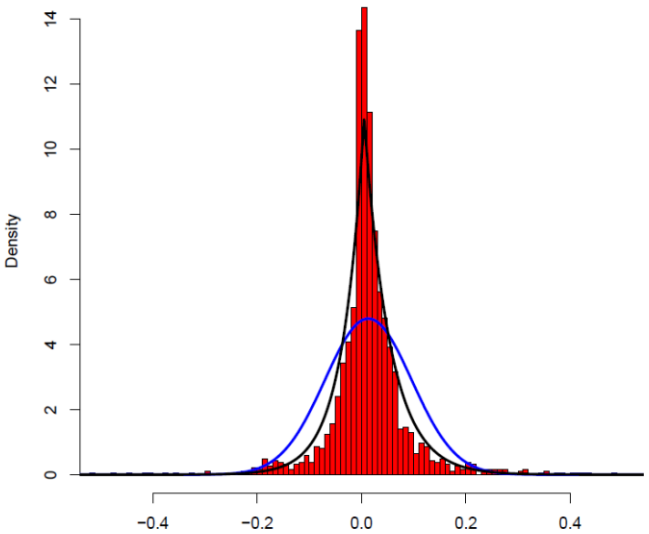

Fig. 10. Histogram of all of the one-hour forecasts for the Xcel plant over a 3- month period. The blue line represents a normal distribution with the same mean and standard deviation. The black line represents a hyperbolic distribution fit to the data with: $\pi = 0.087$, $\zeta = 3.88\ E^{-5}$, $\delta = 1.76\ E^{-6}$, $\mu = 0.005$.

One way that fitting a model distribution to the observed forecast errors may be useful is that it allows for the production of forecast bounds at different confidence levels to be computed, and subsequently applied to system operations. By comparing the forecast confidence intervals obtained with the fitted distribution with those generated using a normal distribution assumption, we can get a glimpse of the impact that this assumption would have on system operations.

Fig. 11 provides an example from the Xcel wind plant at the one-hour forecasting interval for a one-day period. The forecast and actual power outputs are displayed alongside the 95% confidence intervals of the forecast, based on a hyperbolic distribution in blue and a normal distribution in red. The confidence intervals are based on 10,000 samples from the respective distributions and reflect physical generation limits such as the maximum and minimum capacities. Since the one-hour forecast interval error distribution for this plant is extremely leptokurtic, the fitted hyperbolic distribution produces a tighter confidence interval due to its larger mass in the peak of the distribution. The total difference between the intervals produced using the different distributions reaches a maximum of 18 MW. This is an important difference as it represents approximately 6% of the total plant capacity. In a small balancing area, or during certain system conditions, this difference could require a change in the economic dispatch to increase system flexibility and ensure that any wind forecasting error could be easily accommodated. It is also interesting to see that the actual output does lie outside of the confidence intervals produced by both distributions once during the day, approximately the rate that would expected over larger sample sizes.

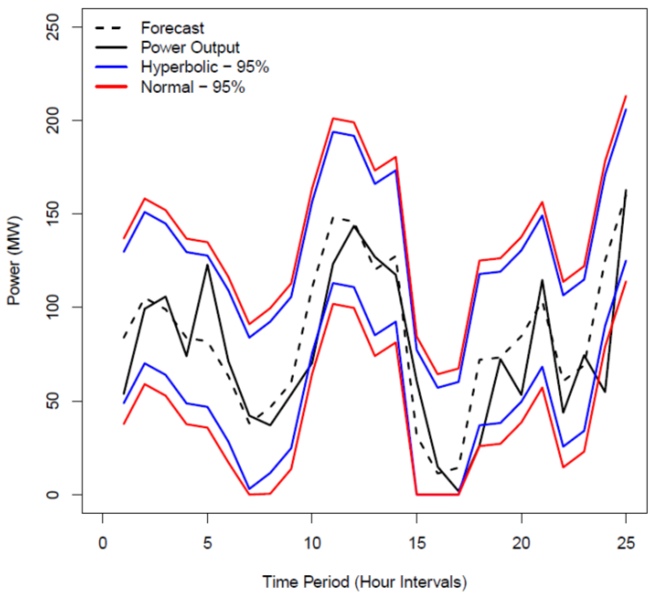

Fig.11. Plot of the wind power forecasts, actual power output, and associated forecast confidence intervals for the Xcel plant during a single-day interval. The solid black line represents the actual power output and the dashed black line the forecast output. The normal distribution confidence intervals at 95% are shown in red. The Cauchy distribution confidence intervals at 95% are shown in blue.

V. Conclusion

In this work, we have examined the shapes of wind power forecasting distributions from operational forecasting systems at multiple timescales through a statistical analysis. An important result of this work is the recognition that operational wind power forecasting errors may be poorly represented by the normal distribution, an assumption common to wind integration studies. It is also important to note that the distributions obtained from operational forecasting systems also differ significantly from those created by the persistence method, even at relatively short timescales. Wind power forecast error distributions can differ greatly depending on a number of factors, such as size of the wind plant or balancing area, forecasting method used, and timescale considered. Therefore, it is important to examine the shape of the distribution for the system being considered, and to keep in mind the goal of the study, before making any assumptions on the wind power forecast error distribution shape. For example, differences that may be critical in an operational study may be unimportant in a planning study. We have also suggested a distribution that may represent some types of wind power forecasting errors better than previously utilized parametric distributions. Further work will focus on a more detailed examination of the performance of the operational systems under different environmental states and power system conditions.

VI. Acknowledgment

The authors gratefully acknowledge David Maggio at ERCOT and Keith Parks at Xcel Energy for supplying the data used in the study. The authors would also like to thank Dr. Chris Dent of Durham University for helpful discussions on the modeling of rare events.

VII. References

[1] K. Methaprayoon, W. J. Lee, C. Yingvivatanapong, and J. Liao, "An integration of ANN wind power estimation into UC considering the forecasting uncertainty," in *IEEE Industrial and Commercial Power Systems Technical Conference*, Saratoga Springs, NY, 2005.

[2] V. Pappala, I. Erlich, K. Rohrig, and J. Dobschinski, "A Stochastic Model for the Optimal Operation of a Wind-Thermal Power System," *IEEE Transactions on Power Systems,* vol. 24, pp. 940 - 950, May 2009.

[3] E. Castronuovo and J. Lopes, "On the Optimization of the Daily Operation of a Wind-Hydro Power Plant," *IEEE Transactions on Power Systems,* vol. 19, pp. 1599 - 1606, August 2004.

[4] R. Doherty and M. O'Malley, "A New Approach to Quantify Reserve Demand in Systems With Significant Installed Wind Capacity," *IEEE Transactions on Power Systems,* vol. 20, pp. 587 - 595, May 2005.

[5] K. Dietrich, J. Latorre, L. Olmos, A. Ramos, and I. Perez-Arriaga, "Stochastic unit commitment considering uncertain wind production in an isolated system," in *4th Conference on Energy Economics and Technology*, Dresden, Germany, 2009.

[6] H. Bludzuweit, J. A. Dominguez-Navarro, and A. Llombart, "Statistical Analysis of Wind Power Forecast Error," *IEEE Transactions on Power Systems,* vol. 23, pp. 983 - 991, August 2008.

[7] B.-M. Hodge and M. Milligan, "Wind Power Forecasting Error Distributions over Multiple TImescales," in *IEEE Power & Energy Society General Meeting*, Detroit, MI, 2011.

[8] M. Lange, "On the Uncertainty of Wind Power Predictions - Analysis of the Forecast Accuracy and Statistical Distribution of Errors," *Journal of Solar Energy Engineering,* vol. 127, pp. 177-184, 2005.

[9] U. Focken, M. Lange, K. Monnich, H.-P. Waldl, H. Beyer, and A. Luig, "Short-term prediction of the aggregated power output of wind farms - a statistical analysis of the reduction of the prediction error by spatial smoothing effects," *Journal of Wind Engineering and Industrial Aerodynamics,* vol. 90, pp. 231 - 246, 2002.

[10] S. Shapiro and M. Wilk, "An analysis of variance test for normality (complete samples)," *Biometrika,* vol. 52, 1965.

[11] D. Scott, "On optimal and data-based histograms," *Biometrika,* vol. 66, pp. 605-610, 1979.

[12] D. Scott, "HyperbolicDist: The hyperbolic distribution," ed, 2009.

[13] "R: A Language and Environment for Statistical Computing," ed. Vienna, Austria: R Foundation for Statistical Computing, 2010.