UNIVERSIDADE TÉCNICA DE LISBOA
INSTITUTO SUPERIOR TÉCNICO

Managing Wind Power Forecast Uncertainty in Electric Grids

Brandon Keith Mauch

Co-supervisors: Doctor Jay Apt (Co-Chair)
Doctor Pedro M.S. Carvalho (Co-Chair)

Thesis approved in public session to obtain the PhD Degree in
Engineering and Public Policy

Jury Final Classification:

Jury

Chairman: Doctor Jay Apt
Doctor Pedro M.S. Carvalho

Members of the Committee:
Doctor Luis Marcelino Ferreira
Doctor Paulina Jaramillo
Doctor Paul Fischbeck

2012

Managing Wind Power Forecast Uncertainty in Electric Grids

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Engineering and Public Policy

Brandon K. Mauch

B.S., Mechanical Engineering, University of Kansas
M.S., Mechanical Engineering, University of Wisconsin

Carnegie Mellon University
Pittsburgh, PA

December 2012

Acknowledgements

I would like to thank the Fundação para a Ciência e a Tecnologia (FCT) for generously providing funding for me through the Carnegie Mellon – Portugal program. The collaborative experience with Portuguese students and faculty through classes and research gained through this program was an invaluable aspect of my doctoral studies. I hope the relationship between Carnegie Mellon and the FCT continues to provide similar opportunities for future students and faculty.

There were many people who helped me during my doctoral studies. First, I want to thank my co-advisors, Jay Apt from Carnegie Mellon and Pedro Carvalho from the Instituto Superior Técnico in Lisbon. Their knowledge of the electric industry as well as their research guidance enriched my experience and permitted me to have much more confidence in my work. I am very grateful to have had such wonderful support from them.

I also want to thank the remaining members of my committee, Paulina Jaramillo, Paul Fischbeck and Luis Marcelino Ferreira. Paulina and Paul provided valuable advice and endless conversations that helped me on multiple occasions. Marcelino provided considerable advice during my stay in Portugal. There are many other faculty members who provided assistance during my studies. Special thanks go to Fallaw Sowell, Francis McMichael and Mitchell Small.

I am very thankful for the entire EPP staff who provided incredible support with administrative issues, and my fellow students who made the whole process much more fun.

Finally, and most importantly, I want to thank my wife, Sarah, for her support over the past four and a half years. She made a lot of sacrifices for me during this process.

Abstract

Electricity generated from wind power is both variable and uncertain. Wind forecasts provide valuable information for wind farm management, but they are not perfect. Chapter 2 presents a model of a wind farm with compressed air energy storage (CAES) participating freely in the day-ahead electricity market without the benefit of a renewable portfolio standard or production tax credit. CAES is used to reduce the risk of committing uncertain quantities of wind energy and to shift dispatch of wind generation to high price periods. Using wind forecast data and market prices from 2006 – 2009, we find that the annual income for the modeled wind-CAES system would not cover annualized capital costs. We also estimate market prices with a carbon price of $20 and $50 per tonne CO_2 and find that the revenue would still not cover the capital costs. The implied cost per tonne of avoided CO_2 to make a wind-CAES profitable from trading on the day-ahead market is roughly $100, with large variability due to electric power prices.

Wind power forecast errors for aggregated wind farms are often modeled with Gaussian distributions. However, data from several studies have shown this to be inaccurate. Further, the distribution of wind power forecast errors largely depends on the wind power forecast value. The few papers that account for this dependence bin the wind forecast data and fit parametric distributions to the actual wind power in each bin. A method to model wind power forecast uncertainty as a single closed-form solution using a logit transformation of historical wind power forecast and actual wind power data is presented in Chapter 3. Once transformed, the data become close to jointly normally distributed. We show the process of calculating confidence intervals for wind power forecast errors using the jointly normally distributed logit transformed data. This method has the advantage of fitting the entire dataset with five parameters while also providing the ability to make calculations conditioned on the value of the wind power forecast.

The model present in Chapter 3 is applied in Chapter 4 to calculate increases in net load uncertainty introduced from day-ahead wind power forecasts. Our analyses uses data from two different electric grids in the U.S. having similar levels of installed wind capacity with large differences in wind and load forecast accuracy due to geographic characteristics. A probabilistic method to calculate the dispatchable generation capacity required to balance day-ahead wind and load forecast errors for a given level of reliability is presented. Using empirical data we show that the capacity requirements for 95% day-ahead reliability range from 2100 MW to 5600 MW for ERCOT and 1900 MW to 4500 MW for MISO, depending on the amount of wind and load forecast for the next day. We briefly discuss the additional requirements for higher reliability levels and the effect of correlated wind and load forecast errors. Additionally, we show that each MW of additional wind power capacity in ERCOT must be matched by a 0.30 MW day-ahead dispatchable generation capacity to cover 95% of day-ahead uncertainty. Due to the lower wind forecast uncertainty in MISO, the value drops to 0.13 MW of dispatchable capacity for each MW of additional wind capacity.

Direct load control (DLC) has received a lot of attention lately as an enabler of wind power. One major benefit of DLC is the added flexibility it brings to the grid. Utilities in some parts of the U.S. can bid the load reduction from DLC into energy markets. Forecasts of the resource available for DLC assist in determining load reduction quantities to offer. In Chapter 5, we present a censored regression model to forecast load from residential air conditioners using historical load data, hour of the day, and ambient temperature. We tested the forecast model with hourly data from 467 air conditioners located in three different utilities. We used two months of data to train the model and then ran day-ahead forecasts over a six week period. Mean square errors ranged from 4% to 8% of mean air conditioner load. This method produced accurate forecasts with much lower data requirements than physics based forecast models.

Contents

List of Figures

Chapter 1 - Introduction

Environmental concerns have led many to advocate for increased production of electricity from renewable resources. Due to its low cost relative to other forms of renewable energy, wind power is growing rapidly around the world. Over the past decade, installed capacity in the U.S. increased tenfold. It accounted for nearly 3% of the total electricity produced in the U.S during the year 2011 (EIA, 2012). However, just ten years earlier, the share of electricity generated by wind was below 0.5% (Figure 1.1). In Europe, 6% of the electricity consumed is provided by wind power (Wilkes et al., 2012).

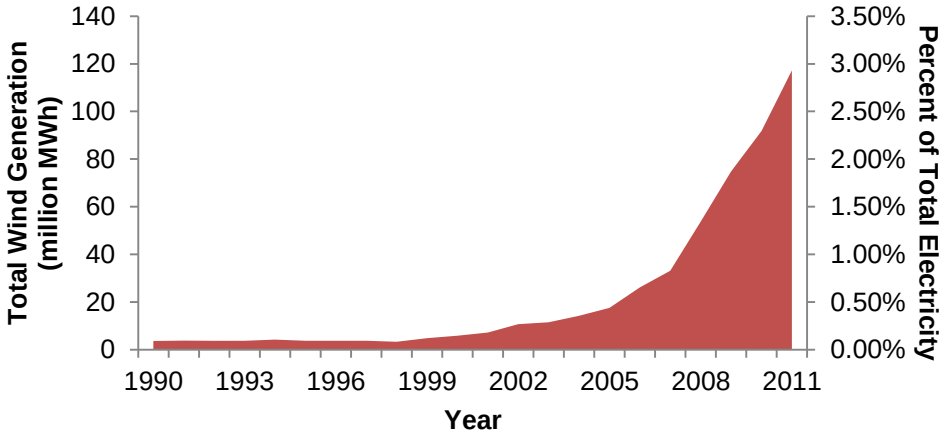

Figure 1.1: Total wind energy generation in the U.S. from 1990 to 2011.

National and state level energy policies are driving the growth seen over the past decade. At the national level, wind benefits from the production tax credit (PTC) which currently provides $22 per MWh generated. Additionally, most states in the U.S. have a renewable portfolio standard (RPS) which sets a minimum level of electricity production from renewable resources.

States also provide tax incentives to renewable power generators. Installed wind capacity will likely continue to grow, although the rate of growth is uncertain if the PTC expires.

The growth of wind power is concentrated in areas with the most favorable resources and incentives. Six states now generate more than 10% of total electricity from wind: Iowa, South Dakota, Minnesota, North Dakota, Colorado and Oregon (Wiser and Bolilnger, 2012). At these levels, wind can no longer be treated as must run, and grids have started curtailing wind to ensure adequate reliability (Rogers, 2010).

Wind power uncertainty and variability pose substantial challenges to the integration of wind power in electric power systems. System operators of electric grids strive to match load levels with nearly equal generation to keep system frequency within acceptable limits. Adding intermittent and variable generation to electric grids increases the difficulty they face. The addition of wind power on the generation side of the grid means that conventional generation must be adjusted to balance wind power changes as well as changes in load.

In order to ensure reliable operation of an electric grid with significant levels of wind capacity, many system operators rely on wind power forecasts (Kehler et al., 2010)). Day-ahead forecasts are used for daily planning decisions, such as unit commitment, dispatch schedules and operating reserves. Figure 1.2 shows an example of one week of hourly day-ahead wind forecasts and actual wind power for the Electric Reliability Council of Texas (ERCOT). Forecast errors cause deviations in the balance between generation and load. Currently, load forecast uncertainty is more substantial than wind forecast uncertainty for most of the U.S., but future levels of wind power will likely create a situation where wind forecast uncertainty dominates that of load forecasts.

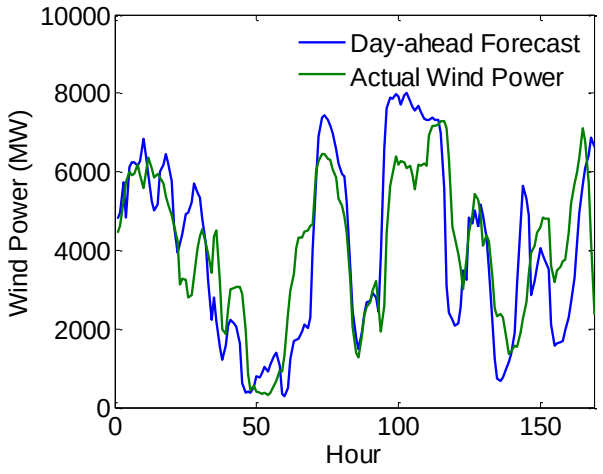

Figure 1.2: ERCOT hourly day-ahead wind power forecast (blue line) and actual wind power (green line) for the week of April 5-11, 2009.

One method to reduce the uncertainty of wind power is to combine wind farms with energy storage. We present a model in Chapter 2 to determine profit maximizing dispatch schedules from a wind farm with compressed air energy storage (CAES) in the case where wind exceeds RPS requirement and is thus not treated as must run. We also assume that tax benefits such as the PTC are expired. We selected CAES for energy storage due to its low cost relative to other large scale energy storage technologies. The addition of energy storage enhances wind energy by allowing limited control of dispatch from a wind farm and smoothing the fluctuations in wind generation. Coupling wind with energy storage allows firm commitments in the day-ahead energy market. At current market prices and capital cost estimates for wind and CAES, the wind-CAES system was found to be unprofitable. When considering the costs of operating the modeled wind-CAES system in the day-ahead market, we found the deficit to be roughly $100 per tonne of CO_2 avoided.

Chapter 3 presents a method to model wind forecast errors with a logit-normal distribution. Wind forecast errors are highly dependent on the wind forecast level. Estimating wind forecast uncertainty by fitting a distribution to wind forecast errors misses this important aspect of wind

forecast uncertainty. The model presented represents forecast uncertainty as a function of the wind forecast level. It also represents forecasts' errors across a wide range of forecasts with only five parameters.

The model in Chapter 3 is used in Chapter 4 to investigate conventional generation levels required to balance wind uncertainty. From a grid perspective, wind power uncertainty will require more operating reserves in future systems with high levels of wind capacity. While current levels of wind capacity add to the overall uncertainty in grid operations, load forecasts are the dominate source of uncertainty. In future grids with much higher levels of wind capacity, this will no longer be true. We present an analysis of reserve requirements for systems with high levels of wind capacity. This analysis is intended to support the decision of system operators in the procurement of operating reserves.

Finally, Chapter 5 changes course a bit and presents a forecast model for residential air conditioner loads. Direct load control (DLC) is a type of demand response (DR) that will provide more flexibility in the grid. This will likely play a large role in wind power integration. DR has received increasing attention in recent years. Traditionally, DLC was used for reducing peak loads and emergency situations. However, today, many utilities can sell load reductions into wholesale power markets. The challenge of using residential loads is the ability to estimate and verify the amount of load reduction when residential loads are called upon. We present a censored regression model of air conditioner loads based on historical load data and ambient temperature and then test it on hourly data from 467 air conditioners. The results indicate that air conditioner load can be accurately forecast for use in DLC programs without data intensive physical models.

1.1 References

1. EIA, 2012. Electric power annual 2012, http://www.eia.gov/electricity/data/state/
2. Kehler, J., Ming, H., McMullen, M., Blatchford, J., 2010. ISO perspective and experience with integrating wind power forecasts into operations, 2010 IEEE Power and Energy Society General Meeting, pp. 1-5.
3. Rogers, J., Fink, S., Porter, K., 2010. Examples of wind energy curtailment practices, NREL Subcontract Report, NREL/SR-550-48737.
4. Wilkes, J., Moccia, J., Drangan, M., 2012. Wind in power: 2011 European wind statistics, European Wind Energy Association Technical Report, http://www.ewea.org/ fileadmin/ewea_documents/documents/publications/statistics/Stats_2011.pdf
5. Wiser, R., Bolinger, M., 2012. DOE 2011 wind technologies market report, Department of Energy, http://www.windpoweringamerica.gov/pdfs/2011_annual_wind_market_report.pdf.

Chapter 2 - Can a wind farm with CAES survive in the day-ahead market?[1]

2.1 Abstract

We investigate the economic viability of coupling a wind farm with compressed air energy storage (CAES) to participate in the day-ahead electricity market at a time when renewable portfolio standards are not binding and wind competes freely in the marketplace. In our model, the CAES is used to reduce the risk of committing uncertain quantities of wind energy and to shift dispatch of wind generation to high price periods. Other sources of revenue (capacity markets, ancillary services, price arbitrage) are not included in the analysis. We present a model to calculate profit maximizing day-ahead dispatch schedules based on wind forecasts. Annual profits are determined with dispatch schedules and actual wind generation values.

We find that annual income for the modeled wind-CAES system would not cover annualized capital costs using market prices from the years 2006 to 2009. We also estimate market prices with a carbon price of $20 and $50 per tonne CO_2 and find that revenue would still not cover the capital costs. The implied cost per tonne of avoided CO_2 to make a wind-CAES profitable from trading on the day-ahead market is roughly $100, with large variability due to electric power prices.

2.2 Introduction

Wind energy in the Unites States has experienced rapid growth as a result of aggressive energy policies at multiple levels of government. In twenty-nine U.S. states, renewable portfolio standards (RPS) place mandates on the amount of electricity production from renewable resources (DSIRE, 2011). RPS mandates and penalties for non-compliance vary from state to

[1] This chapter is based on the published paper: Mauch, B., Apt, J., Carvalho, P., 2012. Can a wind farm with CAES survive in the day-ahead market?, Energy Policy, 48, pp. 584-593

state. At the federal government level, the primary incentive for electricity production from wind is the production tax credit (PTC). For each unit of energy produced from wind, the generator receives a tax credit during the first ten years of generation. Additionally, many state and local governments offer tax incentives for renewable energy projects such as accelerated depreciation and reduced or waived property taxes (DSIRE, 2011).

Wind power comprised only 2.9% of the electricity generated in the US during 2010 (EIA, 2011a). However, due to the policies mentioned above, the Energy Information Administration (EIA) of the U.S. expects wind generation to double from 2009 to 2035 (EIA, 2011b). As the share of electricity from wind energy grows, grid stability will become an important issue. Unforeseen drops or increases in wind generation must be balanced in real time with fast ramping generation such as natural gas or hydro power plants.

The growth of wind energy has spurred interest in coupling wind farms with energy storage in order to alleviate these problems to some extent, and allow wind farms to readily participate in the day-ahead market. Energy storage could enhance wind energy by allowing limited control of dispatch from a wind farm and smoothing fluctuations in wind generation. This would allow less reliance on expensive reserve generation for balancing wind forecast errors. It also provides a wind farm the ability to shift a portion of dispatch from low price periods to periods of peak prices. This is an attractive feature since wholesale electricity prices tend to be low at night when U.S. onshore wind production is stronger than during the day.

Previous work by Garcia-Gonzalez et al (2008), Castronuovo and Lopes (2004) and Greiner et al. (2009) proposed models to determine optimal dispatch schedules for a wind farm with energy storage participating in the day-ahead market. In each model, the stochastic problem was solved by averaging deterministic results obtained from a set of possible wind generation profiles. We take an approach similar to Kim and Powell (2009) by creating an optimization

model based on dynamic programming. In this algorithm, optimal dispatch quantities are calculated for each hour based on the expected state of the energy storage system and wind forecast at that particular hour. Our model differs from Kim and Powell in three ways (1) we do not assume a probability distribution for wind generation, but rather use real wind data; (2) electricity is not sold in the regulation market; and (3) available stored energy each hour is limited by the power output of the storage facility. In order to characterize wind forecast uncertainty, we use historical data from a wind farm to create empirical probability distributions of the wind forecast errors. Since forecast uncertainty is dependent upon location, this analysis is specific to the data we acquired. Furthermore, it should be mentioned that aggregating wind farm forecasts reduces uncertainty (Focken, 2002).

Electrical energy is sold either through bilateral contracts or wholesale markets (day-ahead or real time). Currently, most wind generation in the U.S. is sold through bilateral power purchase agreements (PPA) to utilities that pay an agreed per-unit price for all electricity produced over a 15 to 25 year period (Harper et al., 2007; Windustry, 2010). PPAs guarantee that all energy will be sold, and remove the risk of price fluctuations inherent in wholesale power markets. Utilities benefit by securing renewable energy requirements for RPS mandates. This arrangement essentially means wind generation is treated as "must run" except for times when grid stability is at risk or transmission is constrained.

Here we analyze the economic viability of selling energy on the day-ahead market from a wind farm with a compressed air energy storage (CAES) plant. Electricity generated by the wind turbines can either be dispatched directly to the grid or stored in the CAES plant for later dispatch. We created a model to determine dispatch quantities that maximized expected hourly profits in the day-ahead market given the uncertainty of wind power forecasts (derived from forecast and power output data from a large wind farm). Dispatch quantities from the model

were then used with actual wind power data to determine the profit realized from optimal dispatch. Finally, annualized capital cost estimates were subtracted from the summed hourly profits to determine annual profits.

Future electric grids may have very high levels of wind power capacity. Once a significant amount of electricity generation comes from wind power, it is reasonable to expect that at least a small portion of it will be sold on the day-ahead market. While we use current market prices in our model, we feel that this analysis will provide insight into the problem of coupling wind power with CAES. It should be noted that we do not consider additional revenue from capacity markets or ancillary services. Only day-ahead market income is considered without the benefits of a PTC or RECs, or the obligation to purchase wind required by an RPS. We also exclude the possibility of using CAES for energy arbitrage through the purchase and sale of grid electricity. These additional forms of revenue have been treated in previous papers (e.g. Fertig and Apt, 2010) and are not in the scope of this paper. This paper is organized into 4 sections. Section 2.3 describes the model used to determine the hourly day-ahead dispatch quantities and annual profit. Results from the model are presented in Section 2.4. Finally, conclusions are presented in Section 2.5.

2.3 Wind-CAES model

2.3.1 Overview

The model we present in subsequent sections uses wind forecast and market price data to determine the profit maximizing dispatch schedule from a wind-CAES system for the day-ahead market. For each hour of the following day, there are two decisions to be made; the amount of energy to dispatch and the amount of energy to store. Since the supply of energy from the wind farm is stochastic, we maximize expected values for hourly profit. The expected profit is a non-

linear function of available wind energy, and is calculated from a Monte Carlo routine using randomly drawn values of possible wind generation. The wind samples come from empirical probability density functions that were created with historical wind forecast and wind generation data. We found that a sample size of 100 values provided adequate results without requiring too much computational time. We programmed the model with FORTRAN code using dynamic programming to solve the optimization problem by cycling through many different energy storage levels for each hour of the dispatch period. For each storage level, maximum expected hourly profits were determined considering the present hour and all remaining hours in the time horizon.

2.3.2 Storage

Large scale energy storage exists in many forms including pumped hydroelectric, compressed air energy storage (CAES), batteries and flywheels. Currently, the least expensive options are pumped hydro and CAES. Nearly 21 gigawatts of pumped hydro storage exist in U.S. grids (EIA, 2010a), but only two CAES facilities exist worldwide, with others under development (Succar, 2011).

Due to the low capital costs and flexibility in location, we chose CAES as the energy storage technology to use in our model. However, this method is generally applicable to any utility-scale storage. CAES facilities store energy in the form of compressed air in underground caverns. A compressor pushes air into the cavern during the charging process. Air is allowed to escape through an expander and natural gas turbine when the stored energy is used to generate electricity. In stand-alone natural gas turbines, half of the energy contained in the gas is used to compress the air prior to combustion. A CAES facility connected to a wind farm uses electricity generated from wind energy to compress air resulting in a heat rate that is roughly half of that compared to standalone gas turbines (Succar, 2011).

Round-trip efficiency is an important parameter used in quantifying an energy storage system. For the case of CAES plants, energy input comes from two different types of energy sources. This complicates the round-trip efficiency metric. Succar (2011) showed that using commonly accepted values for CAES plants gives a round-trip efficiency around 82%. However, in this model we are interested in the ratio of energy output to energy input. For every 1 MWh of wind energy stored in a CAES facility, roughly 1.2 to 1.8 MWh of energy can be supplied due to the addition of natural gas energy (Succar, 2011).

2.3.3 Wind forecasts

Dispatch decisions in the day-ahead market are made a day before wind generation occurs meaning that participation for wind farms requires good wind forecasts. While dispatch quantities need to be calculated for a twenty-four hour period, a wind farm with CAES will use longer forecast look-ahead times to optimally manage energy storage levels over a multi day period. Our model uses forty-eight hours of prediction values to make dispatch decisions for a twenty-four hour period.

Forecasts are received and used for dispatch calculations the day before the dispatch is carried out. Our model assumes that dispatch quantities are submitted at noon on the day prior to dispatching electricity. Forecast values for wind power starting at midnight are used for dispatch calculations, and the first twelve hours are discarded. Figure 2.1 illustrates the timeline and wind forecast period for day-ahead dispatch used in the model described later.

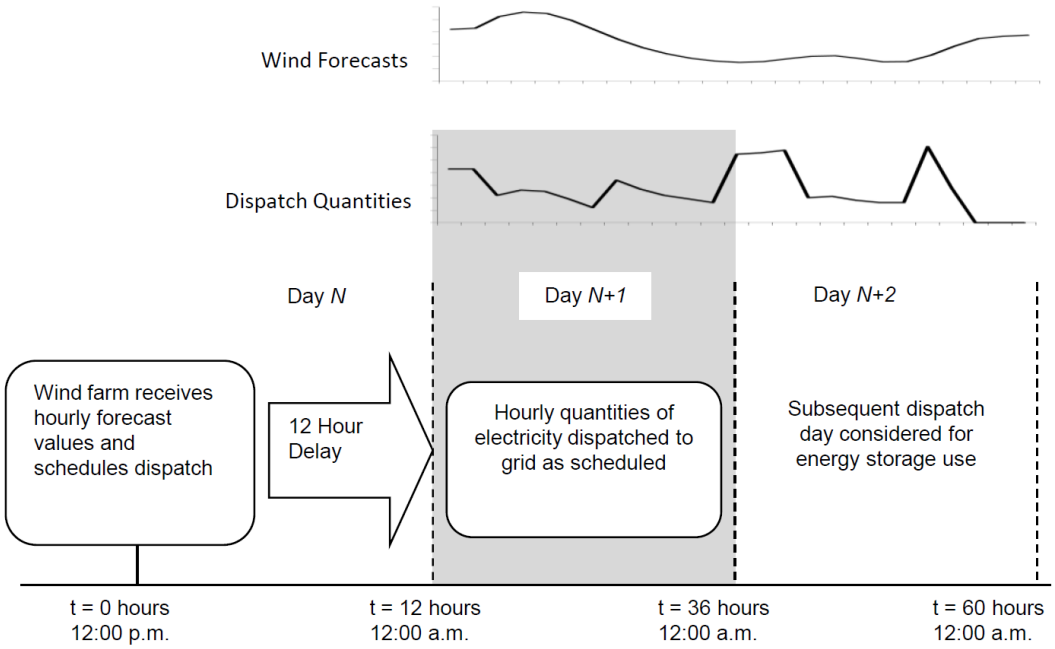

Figure 2.1: Wind forecast and dispatch quantity periods used along with timeline for participation in the day-ahead market. The market closes at noon each day. Dispatch quantities are calculated for two days with the first day values submitted to the market.

We obtained hourly wind forecast and generation data from a single wind farm for this study. The data span a period of two years starting at the beginning of 2008 and ending at the end of 2009. Data collection problems resulted in several gaps in the data. Since most gaps spanned a time period of several hours, all gaps were discarded from the dataset. Two hundred eighty-five days of wind data remained in the 2008 period, and 325 days remained in the 2009 period after removing bad data.

In order to characterize forecast uncertainty, we used the 2008 data to create empirical probability density functions of wind generation conditioned on the forecast value. For a given forecast value, the expected hourly profits are calculated using the wind generation probability density functions. The 2009 data were used as inputs for the dispatch model and assumed to have similar uncertainty as the 2008 data.

Forecast uncertainty depends on several factors including look-ahead time and the forecast values. As the look-ahead time moves further into the future, uncertainty increases. Common metrics to quantify forecast uncertainty are the mean absolute error (MAE) and root-mean-square-error (RMSE). Figure 2.2 shows the MAE and RMSE plots for the 2008 wind farm data.

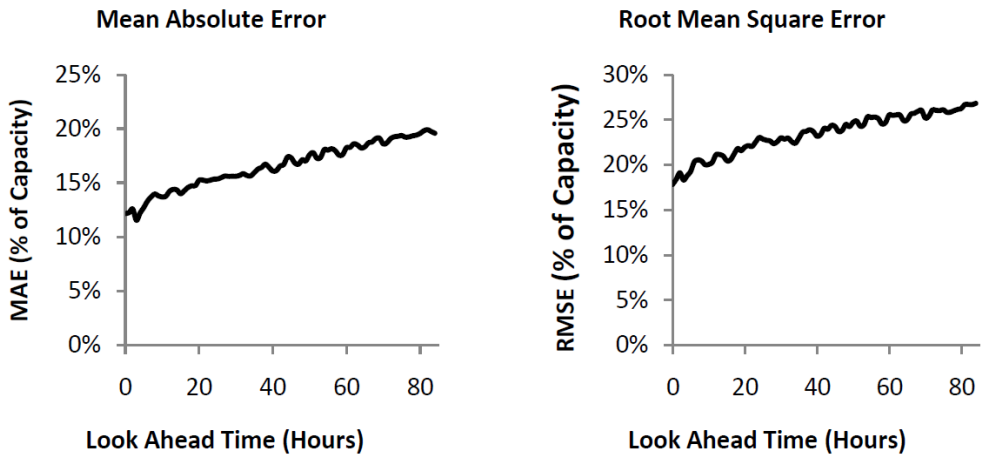

Figure 2.2: MAE and RMSE graphs for wind power forecasts with an 84-hour look-ahead time. The graphs were made with 2008 data from a single wind farm. Note that all values are normalized by the wind farm capacity.

If a prediction is made near the maximum output of the wind farm, then actual wind generation is more likely to be below the predicted value than above it. Alternatively, for a forecasted value near zero the actual wind generation will likely be above the predicted value. For this reason, treatment of the forecast uncertainty depended on the hour within the forecast horizon and the value of the forecast. We separated the forecast values into eleven power classes for each hour of the forecast time horizon. Figure 2.3 shows the mean absolute error and RMSE values for each power class eighteen hours after the forecast was made. Uncertainty bars indicate the shape of the two graphs is very similar.

Within each power class the forecast error probability density function was calculated empirically using 2008 forecast data. This produced a total of 11x48=528 different probability

distributions. This method is based on Bludszuweit et al. (2008) who divided forecast values into fifty power classes and fit a beta distribution to the forecast errors within each class. We chose to use empirical probability density functions of forecast errors for each class of forecast values.

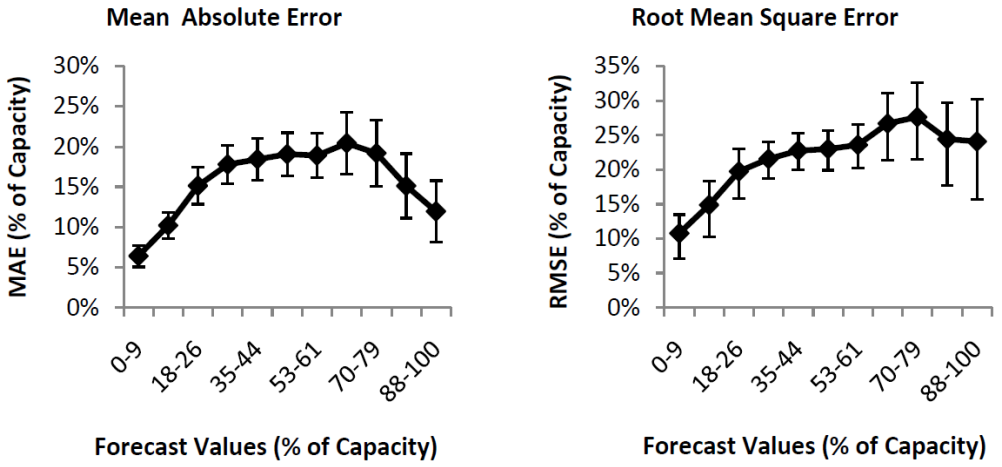

Figure 2.3: The mean absolute error and root mean square error as a function of forecast power classes eighteen hours after the forecast was taken (six hours into the dispatch schedule) for 2008 for the wind farm used in this study.

As discussed later in section 2.5, the expected hourly profits from wind-CAES dispatch are calculated with nonlinear equations. For this reason we replaced point values in the forecast with empirical pdfs of wind generation. Using Monte Carlo simulations, the wind generation pdf was sampled multiple times to determine the expected hourly profit for each hour of the trading day.

2.3.4 Wind with CAES

A wind farm operator selling energy on the day-ahead market will determine optimal hourly energy commitments which maximize the wind farm's hourly profit. Hourly profit is defined as hourly income from energy sold to the market less the cost of using energy from CAES. We assume the marginal cost of energy from the CAES changes with the price of natural gas, but the

heat rate remains constant. Annual profits are the cumulative hourly profits over one year less the annualized capital costs. Additional assumptions used in the model include: (1) the wind farm is a price taker, (2) all electricity offered to the day-ahead market is accepted for dispatch, (3) transmission is not constrained and (4) the wind farm is not participating in non energy markets (e.g. ancillary services).

Dispatch quantities are determined before wind generation and market prices become known. As stated above, wind forecasts are integral in scheduling dispatch. In order to properly use the forecasts, uncertainty associated with the point values of the forecast must be accounted for in dispatch decisions (see previous section). We assume perfect price knowledge each day when dispatch schedules are calculated. In reality, uncertain price forecasts are used to schedule generation. Assuming perfect price knowledge provides an upper bound for the annual profit results from the model.

Parameters used for the wind farm and CAES facility are shown in Table 2.1. Due to the desire of the wind farm that supplied forecast and actual power production data to remain anonymous, we will not mention details about the farm beyond the capacity factor and that it is a large installation in a good onshore wind area of the U.S.

The CAES output power was sized to be 45% of the wind capacity. This size provided the highest annual profit values over eight different price scenarios used in the model (section 2.7). Ramping capabilities of CAES are much faster than the time resolution of our model (Succar, 2011). Therefore, we neglect the ramp up and down times of the CAES output. Sensitivity analysis on the base case parameters is presented in Section 2.4.4.

Table 2.1: Base case parameters for wind-CAES system used in optimization model

Wind Power Capacity Factor	0.31
Wind Generation per Installed MW of Capacity	2716 MWh
CAES Expander Capacity to Wind Farm Capacity Ratio	0.45
Expander to Compressor Power Ratio	1
Energy Output to Energy Input Ratio	1.35
Storage Capacity	15 hrs
Heat Rate	4000 Btu/kWh
Variable Cost of Storage	$3/MWh

The wind farm hourly optimal dispatch quantities for the day-ahead market were computed with a dynamic programming model coded in FORTRAN. For each hour of the day-ahead market there are two decision variables, (1) the amount of energy to place or remove from CAES and (2) the amount of energy to sell. The objective function used in the model is maximizing expected hourly profit from day-ahead market electricity sales. Constraints and exogenous variables are explained in the next sections. Dispatch schedules are based on wind power forecasts and market prices. Uncertainty of wind generation was modeled with empirical probability density functions created from nearly one year of wind forecast and generation data. Resulting dispatch quantities were then used with actual wind generation data to determine actual profits gained from the dispatch quantities.

2.3.5 Optimal dispatch

A wind energy provider will offer energy in the market according to a schedule that maximizes profits. The objective of optimizing hourly profit over *T* periods is given by

$$\max_{\{Q_t, S_t\}} \sum_{t=1}^{T} \pi_t(Q_t, S_t) \tag{2.1}$$

Hourly profit is given by π_t and is a function of the hourly energy storage levels, S_t, and hourly energy dispatch quantities, Q_t. The solution of Equation 2.1 gives the set of dispatch

quantities that will maximize hourly profits over the time horizon *T*. While we are interested in dispatch quantities for a 24 hour period, we solved Equation 2.1 for a time period of 48 hours and used only the first 24 hours. Using a time horizon of 48 hours allows optimal management of the energy storage for the dispatch day while considering trading for the next day.

Energy dispatch from an intermittent source with limited storage ability is similar to the classic inventory problem studied extensively in the operations research community (e.g. Hillier and Lieberman, 2001). Rather than having uncertain demand, as is common in operations research, the supply of wind energy for storage or dispatch is treated as uncontrollable and stochastic. Dynamic programming is well suited to multistage processes such as energy dispatch from storage. Equation 2.1 is solved in a backward recursive manner in which decisions for the last time step are optimized first and the solution progresses to the first time interval. The optimal storage transition at time *t* is computed with regards to all time intervals from *t* to *T*.

In order to apply dynamic programming, the system is described by a state variable. In this case, the energy storage level is the natural choice. As the solution progresses backwards in time, the optimal path for the current time interval is computed for all possible storage levels. Uncertainty of energy supply was modeled with empirical wind generation probability density functions. For more information on dynamic programming models with stochastic supply see Nandalal and Bogardi (2007).

An illustration of the optimal dispatch algorithm is shown in Figure 2.4 with the first and last three hours visible and quantities of energy in storage shown for each hour of the optimization horizon. If energy from wind generation is stored during a particular hour, then the storage level will increase. A decrease in the storage level indicates energy has been removed from storage to be dispatched. Three possible paths for the energy storage level are shown in the diagram. Many potential paths exist from hour 1 to hour 48. Our algorithm to determine optimal dispatch

seeks the path that produces the greatest profit over the time horizon. Initially, there is some given level of energy in the CAES. All paths of energy storage levels through the horizon must start at the initial energy level. In order to optimally use energy storage, no excess energy should remain at the end of the horizon. Therefore, all possible paths end at the minimum storage level.

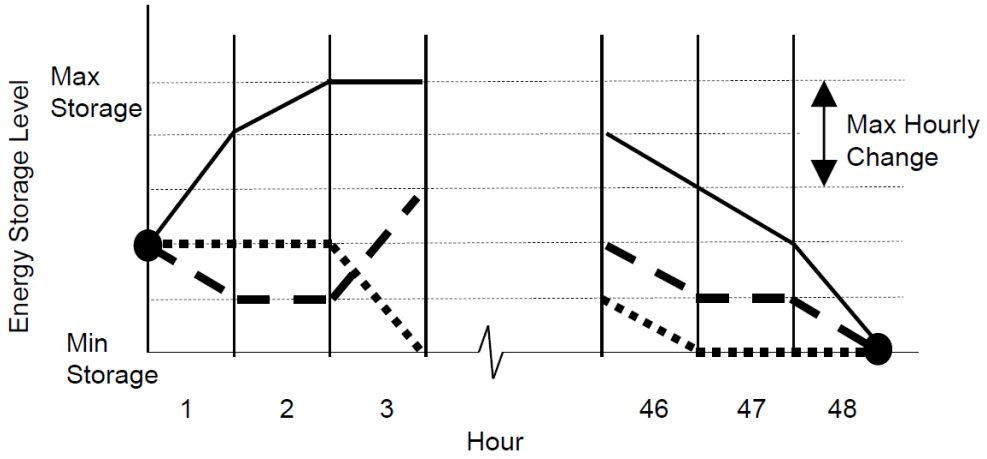

Figure 2.4: Illustration of potential energy storage level changes during the optimization horizon.

Three things should be noted in Figure 2.4. First, this algorithm limits energy storage levels to discrete quantities. This is necessary in order to analyze a finite number of changes in the amount of stored energy during each hour. Second, the change in stored energy during one hour is constrained by the CAES power. If the CAES is fully charged, it cannot use all stored energy in one hour because the expander can only produce electricity at a rate up to its rated power. This is illustrated in Figure 2.5. Finally, all potential paths have the same known initial and final storage levels. It should also be noted that the optimal path is based on wind forecast data and may not be feasible for the actual wind generation during the day of dispatch.

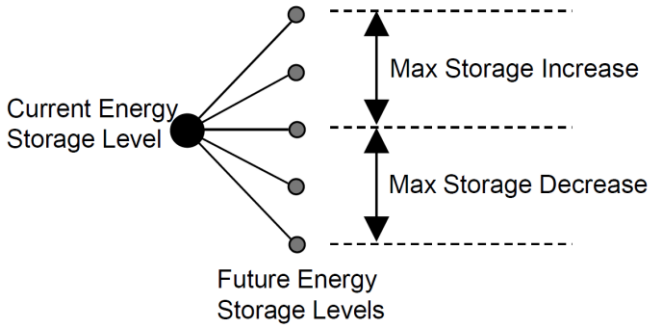

Figure 2.5: Illustration of five possible energy storage level transitions during one hour for a given current energy level. The limitation to transitions of ± 2 levels represents the CAES charge/discharge rate limitation. In the actual model, transitions covered many more than 2 levels.

In our model, the storage level resolution was set so that 300 discrete levels existed between the minimum and maximum storage capacity. During each hour of the optimization horizon, profits for all possible transitions in the energy storage level were calculated. For each storage level transition the dispatch quantity giving the largest expected profit for the current and all remaining hours was determined. Once all transitions have been analyzed, the optimal path for the amount of energy stored in the CAES is determined to give the maximum profit. Expanding the marginal profit function in Equation 2.1 produces Equation 2.2. Marginal profit for each hour denoted by the subscript *t* is calculated as income from electricity sold less the cost of using energy from the CAES.

$$\pi_t = \begin{cases} Q_t p_t - c_S s_t^e & if\ Q_t \le W_t^e + s_t^e \\ (W_t^e + s_t^e)p_t - \left(Q_t - (W_t^e + s_t^e)\right)p_t\alpha - c_S s_t^e & otherwise \end{cases} \quad (2.2)$$

$$\text{where } s_t^e = \begin{cases} (\Delta S_t)\eta, & if\ energy\ is\ dispatched\ from\ CAES \\ 0, & otherwise \end{cases}$$

Total energy available for dispatch during any given hour is the sum of the estimated wind generation (W^e) and the expected amount of energy used from CAES (s^e) during that hour. The superscripts indicate that the variable is estimated. When energy is added to the CAES the value

for s^e is zero. Otherwise s^e is equal to the decrease in stored energy (ΔS) multiplied by the energy output to input ratio (η).

When dispatch does not exceed total available energy, profit is the dispatched quantity (Q) multiplied by the market price (p) minus the cost of using energy from storage ($c_s s^e$) as shown in the first line of Equation 2.2. The parameter c_S represents the marginal cost of using energy from CAES. If energy is overcommitted, profit is calculated from the second line in Equation 2.2 as total available energy (W^e+s^e) multiplied by the market price less the cost of purchasing additional energy in the market to meet the obligation. A weighting factor (α) is also used in the second line of Equation 2.2 to reduce the occurrence of overcommitments. Offering too much energy on the market can be costly and lead to negative profit. Therefore, we increased the effect of overcommitments on the objective function by using a large weighting factor which improved actual profits.

Calculating hourly profits occurs before the wind generation values are known so we must maximize expected profits each hour. Since wind generation is a random variable, we used Monte Carlo sampling to draw a sample of wind generation values from empirical probability density functions. For every storage transition inspected in the model, the expected hourly profit was computed according to Equation 2.2 for all wind generation values drawn in the sample. A sample size of 100 wind generation values provided adequate results. As stated earlier, we created empirical probability density functions using the 2008 wind data. Given the wind forecast value and look-ahead hour, the appropriate pdf was used for the Monte Carlo sampling.

The optimization formulation used in this model is expressed fully as the maximization of expected profits for a 48 hour period as shown below where E [] denotes the expected value.

$$\max_{\{Q_t, s_t\}} \sum_{t=1}^{T=48} E[\pi_t] \quad (2.3)$$

s.t. $$-P^E \leq \Delta S_t \leq P^C \quad (2.4)$$

$$\Delta S_t \leq E[W_t] \quad (2.5)$$

$$0 \leq S_t \leq C \quad (2.6)$$

$$0 \leq Q_t \leq (P^E + W^U) \quad (2.7)$$

Equation 2.3 expresses the objective function as the expected profit over a 48 hour period. Accumulation of stored energy during a one-hour interval cannot exceed P^C, the rated power of the CAES compressor. Alternatively, the maximum drop in stored energy is the negative value of the rated power of the expander in the CAES, P^E. These constraints are expressed in Equations 2.4. Equation 2.5 states that the storage level cannot increase in an hour by more than the expected wind generation for that hour. Equation 2.6 states that the energy storage level cannot exceed the storage capacity, *C*. Equation 2.7 constrains hourly dispatch to a positive value not greater than the sum of the wind farm capacity and the output power of the CAES, P^E.

2.3.6 Realized profit

Hourly dispatch quantities resulting from the optimal dispatch algorithm are used with actual wind generation values to determine hourly profits. In this part of the model, Equation 2.8 is used to calculate realized hourly profits. When scheduled dispatch (*Q*) is less than available energy, hourly profits are calculated with the first line. This is calculated as the revenue from the sale of dispatched energy plus the sale of excess, non-scheduled energy (*q*) minus the cost of energy pulled from the CAES. Excess energy occurs when the CAES system is filled to capacity or the amount of excess wind generation is beyond the CAES compressor charge rate. Revenue from excess energy is reduced by the factor (1- φ) where φ is the market penalty factor and has a value between 0 and 1.

The second line in Equation 2.8 is used when dispatch is greater than available energy. Here profit is calculated as revenue from all available energy minus the cost of purchasing overcommitted energy minus the cost of using stored energy. The market penalty factor is used here to increase the cost of purchasing energy by the amount (1+ φ). In our model we used the value 0.2 for φ which increases the cost of overcommitted energy by twenty percent and decreases revenue from excess energy by twenty percent.

$$\pi_t = \begin{cases} Q_t p_t + q_t p_t (1-\varphi) - c_S s_t & if\ Q_t \le W_t + s_t \\ (W_t + s_t)p_t - \big(Q_t - (W_t + s_t)\big)p_t(1+\varphi) - c_S s_t & otherwise \end{cases} \tag{2.8}$$

$$\text{where } s_t = \begin{cases} (\Delta S_t)\eta, & if\ energy\ is\ dispatched\ from\ CAES \\ 0, & otherwise \end{cases}$$

Equation 2.8 is applied to each hour of the year. After each hour, the amount of stored energy is updated. During hours when dispatch is less than the wind generation, stored energy accumulates up to the storage capacity within the CAES charge limit. When dispatch is greater than wind generation, the storage level is reduced to make up the difference. Each hour that dispatch quantities cannot be met with available wind generation and energy storage (within CAES limitations) energy must be purchased to make up the difference.

In summary, the model maximizes expected hourly profit for a wind-CAES system operating in the day-ahead market. The decision variables are the expected hourly dispatch and storage quantities. Expected hourly profit is calculated via Monte Carlo sampling of wind energy generation from empirical probability density functions. Optimal dispatch quantities are then used to determine profit for one year. Since our data limited us to 325 days, we estimated profits for 365 days by multiplying the summed hourly profits by 365/325.

2.3.7 Market clearing price data

We ran the model with eight different market price scenarios to observe how the results change with market prices. Four separate years, 2006 to 2009, of wholesale market prices from the western zone of the Electricity Reliability Council of Texas (ERCOT) and the Iowa zone of the Midwest ISO (MISO) market were used for the eight price scenarios. Prices for ERCOT came from the balancing market while MISO prices were from the day-ahead market. ERCOT had no central day-ahead market during the years covered with the data. Over 90 percent of energy traded in ERCOT is through bilateral contracts. However, the average bilateral contract price in 2009 was similar to the average balancing market price, although balancing prices were more volatile (Potomac Economics, 2009). We used market prices from these regions because they have high levels of wind capacity compared to the rest of the US. Note that each year of price data was treated as a separate price scenario and run with the same year of wind data. Descriptive statistics for each price scenario are shown in Table 2.2 below.

Table 2.2: Market price statistics for the eight price scenarios used in the model. All values are in U.S. dollars per megawatt-hour.

Scenario	Market	Year	Mean	Std. Dev.	Min	Max
1	ERCOT (West zone)	2006	51.37	31.91	-209.49	823.75
2		2007	52.86	51.67	-407.91	929.99
3		2008	54.64	84.37	-523.20	2076.30
4		2009	25.89	41.53	-235.38	1664.30
5	MISO (Iowa Zone)	2006	42.72	26.30	2.86	329.51
6		2007	49.89	28.52	3.23	236.59
7		2008	47.75	37.94	-233.52	449.72
8		2009	24.09	21.08	-133.27	284.00

MISO prices have been less volatile than ERCOT prices, primarily because balancing market prices were used for the ERCOT scenarios. Due to depressed demand and low natural gas

prices, electricity prices in 2009 were much lower than previous years throughout the U.S. (Wiser and Bolinger, 2011).

2.4 Results

2.4.1 Wind and CAES annual costs

Lawrence Berkley Laboratory found that a sample of 181 wind farms built in 2009 and 2010 had an average capacity-weighted cost of $2.1 million per MW with a range from $1.3 to over $4 million per MW (Wiser and Bolinger, 2011). The Energy Information Administration estimates the average wind installation cost at nearly $2 million per MW (EIA, 2010b). A range of wind farm costs arise due to differences in site requirements and fees. Based on cost numbers in the literature above, we assume a range of installation costs from $1.5 to $2.5 million per MW.

CAES cost estimates also span a large range due to an overall lack of construction experience with CAES plants and differences in sight suitability for a CAES facility. Past cost estimates range from $0.6 - $0.9 million per MW of expander capacity for storage sizes in the 15 to 20 hr range (Denholm and Sioshansi, 2008; Sullivan et al., 2008; EPRI-DOE, 2004). Table 2.3 shows base case cost estimates used in this study for a wind farm and a CAES facility.

Table 2.3: Wind and CAES base case cost estimates

	Wind Farm	CAES Plant
Capital Cost ($/MW)	2 million	0.75 million
Fixed Annual Cost ($/MW)	30 thousand	10 thousand
Economic Life (years)	25	25
Discount Rate	10%	10%
Total Annual Cost ($/MW)	192 thousand	71 thousand

Annual nominal costs for a wind farm and CAES facility were calculated for the full range of cost values assuming a discount rate of 10% representing the blended cost of capital. In our

uncertainty calculations we used a range of discount values from 6 to 14%. Based on our range of estimates discussed above, annual costs for a wind farm range from $145 to $240 thousand per MW of installed capacity. Annual costs for the CAES facility range from $55 to $85 thousand per MW of expander capacity.

To put these costs into perspective, consider a wind farm with a capacity factor of 0.3. For every MW of installed capacity, the farm will generate 2628 MWh of electricity per year. If all of that electricity is sold at a flat rate of $70 per MWh through a PPA, the wind farm will receive $184 thousand per MW of capacity annually. Wiser and Bolinger (2011) calculated the average price of wind power from farms constructed in 2010 to be $73 per MWh (sample size = 26 projects). This does not include the PTC of $22 per MWh or revenue from the sales of renewable energy credits. Once the PTC is included, the revenue increases to $240 thousand per MW of installed capacity.

2.4.2 Annual profits

Results from the model presented in Section 2.2 are shown in Figure 2.6. For all price series used in the model, the annual profits were negative. The model was run multiple times to determine the CAES power output value that produced the highest average annual profit for all price scenarios. Price scenarios 4 and 8 came from 2009 data when electricity prices were quite low resulting in much less profit for the simulated wind farm. Scenarios 1 through 4 came from ERCOT which had higher prices and higher volatility than MISO leading to higher profits. The CAES system was able to take advantage of the greater spikes in the ERCOT prices.

Annual profits for the wind-CAES system are shown alongside annual profits for a standalone wind farm. Scenarios 4, 5 and 8 have low prices and low volatility resulting in more loss when CAES is added. For these scenarios, the CAES value is less than its costs.

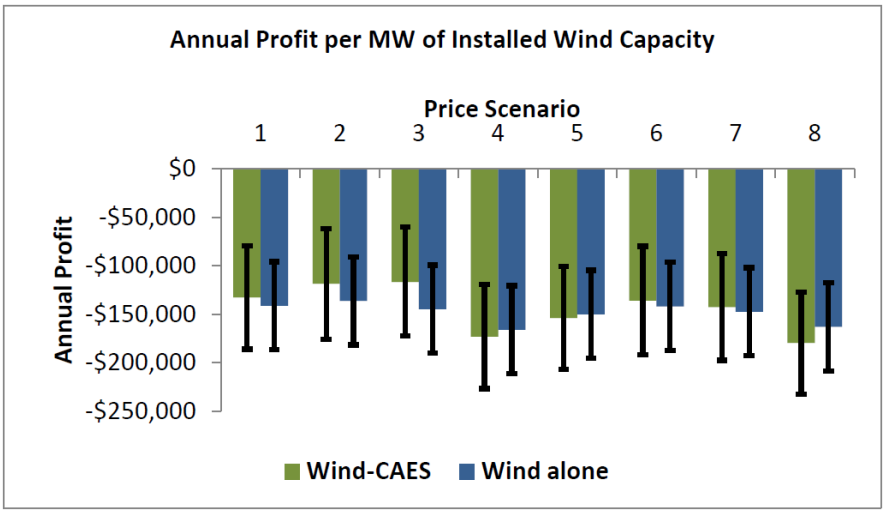

Figure 2.6: Annual profits for the wind farm with CAES using eight different market price input scenarios. All profits are negative (i.e. losses). Uncertainty ranges were created by running the model with different cost assumptions for the wind-CAES system.

The model assumes that overcommitted energy must be purchased from the wholesale market to meet the wind farm's contractual obligation. In using a market penalty factor equal to 0.2, we assumed that overcommitted energy is purchased at twenty percent above the day-ahead price and excess energy is sold at twenty percent below this price. In Section 2.4.4 we show how profits for scenario 6 change with a range of penalty factors.

If the wind farm with CAES is paid a direct subsidy equal to its annual loss, we can determine the cost of carbon emissions avoided in this situation. These costs would be reduced if revenues for ancillary services or price arbitrage were included. The amount of carbon emissions displaced by the wind-CAES system can be estimated by multiplying the amount of electricity dispatched from the wind and CAES with the average carbon emission factor for the U.S. This amount is then reduced by the amount of carbon dioxide emissions from the CAES to give the net emissions displaced. Using these numbers gives the range of values in Table 2.4 expressed in dollars per tonne of carbon dioxide emissions avoided.

Table 2.4: Cost ranges for each tonne of carbon dioxide emissions avoided with the wind farm and CAES.

Price Scenario	Cost range per tonne of carbon dioxide
1	$60 – 150
2	$45 – 130
3	$45 – 130
4	$130 – 240
5	$80 – 180
6	$60 – 150
7	$70 – 160
8	$125 – 220

Due to CAES limitations, a small fraction of the wind energy generated beyond scheduled dispatch cannot be stored. Figure 2.7 shows the excess, non-stored energy for each price scenario for the wind-CAES and standalone wind farm. Figure 2.8 shows the amount of overcommitted energy for each price scenario. Adding a CAES plant to the wind farm reduced overcommitted energy by roughly half.

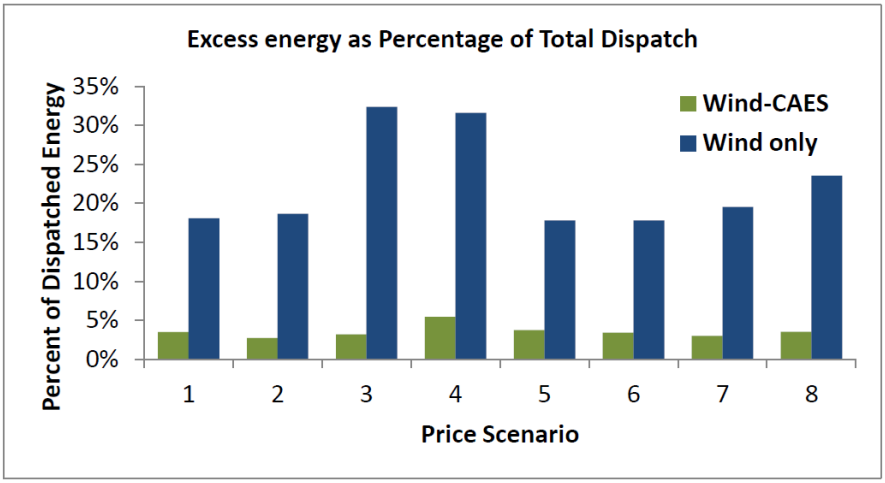

Figure 2.7: Excess, non-stored energy as a percentage of total dispatch for each price scenario used in the model (base case cost assumptions).

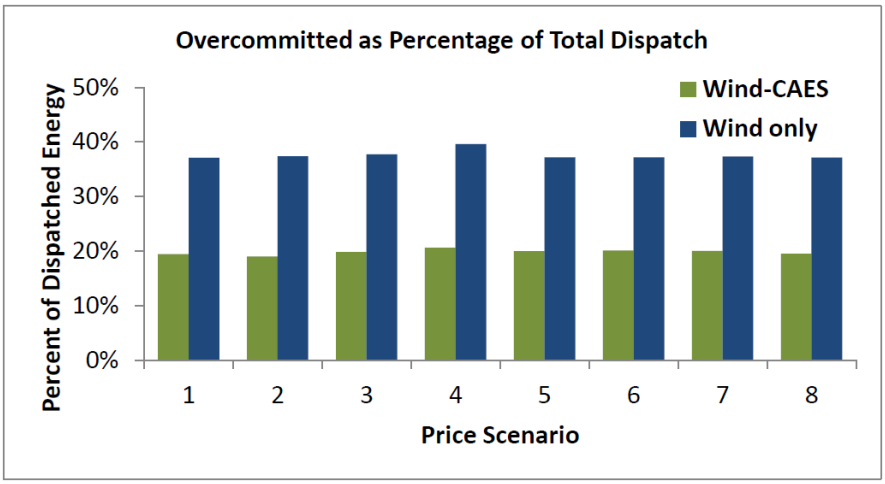

Figure 2.8: Overcommitted energy as a percentage of total dispatch for each of the price scenarios used in the model (base case marginal costs for stored energy).

2.4.3 Annual profit with perfect forecasts

Since the overcommitted energy results in profit loss for the wind farm, wind forecast accuracy directly affects the profits of the wind farm. In order to determine the potential profits from perfect wind forecasts, the model was run using actual wind generation in place of forecasted generation. Results are shown in Figure 2.9, which indicate an upper bound to the profits that can be obtained with improved forecasting techniques. As indicated in the figure, perfect wind knowledge reduces annual losses and also brings profits from the standalone wind farm more in line with the wind-CAES system. All price scenarios still produce annual losses.

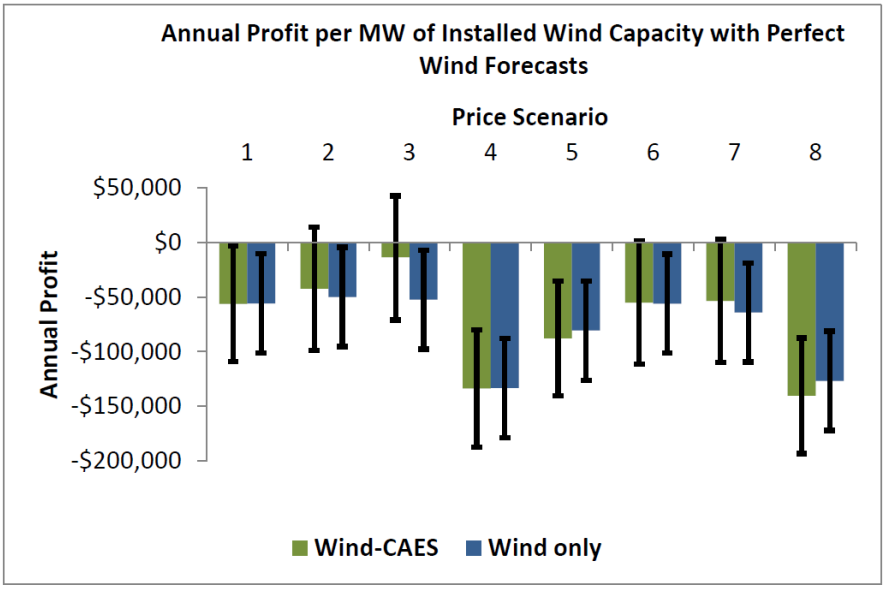

Figure 2.9: Annual profit per installed MW of capacity with perfect wind forecasts for each price scenario. Uncertainty bars were created by running the model with different cost assumptions for the wind-CAES system.

2.4.4 Model sensitivity analysis

Sensitivity analysis for CAES parameters was carried out to determine how the annual profit is affected by the assumptions made. Since price scenario 3 (ERCOT 2008) provided the largest annual profit, we used these prices with base case cost assumptions for stored energy to determine how the profit might change if storage parameters are altered. This was done for storage capacity, ratio of CAES expander to compressor power rating and CAES energy output to energy input ratio. For the base case assumptions with scenario 3, annual profit was calculated to be -$116,000 per MW of installed wind capacity. Sensitivity results are shown in Figure 2.10.

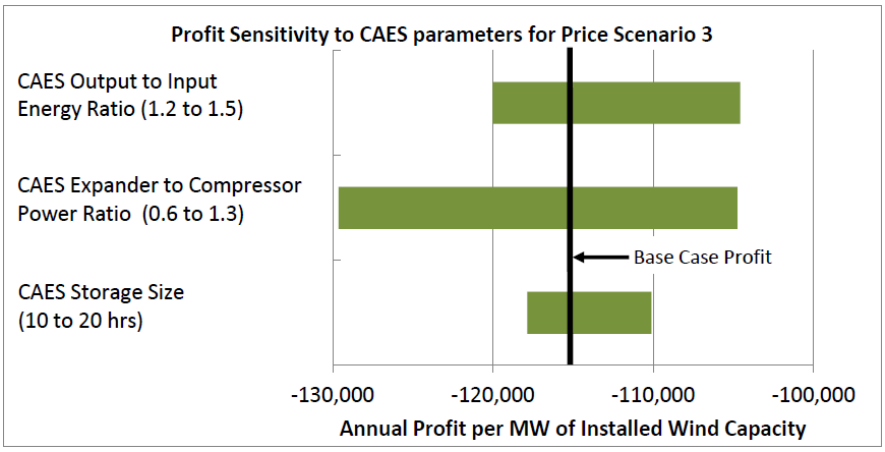

Figure 2.10: Sensitivity of three selected CAES parameters on annual profit results from the model used in this study.

Annual profits would improve modestly with better efficiency, more charging power and more storage. This assumes that these improvements could be made without increasing the cost of the CAES. It is apparent that the model is still far from giving positive results for the profits.

We assumed a twenty percent penalty costs occurred when additional energy had to be purchased in order to fulfill a day-ahead contract. We show the sensitivity of our results to the penalty factor in Figure 2.11 using price scenario 6. Profits decrease linearly as the penalty factor is increased. The wind farm without CAES suffers more from higher penalties due to the large amounts of overcommitments.

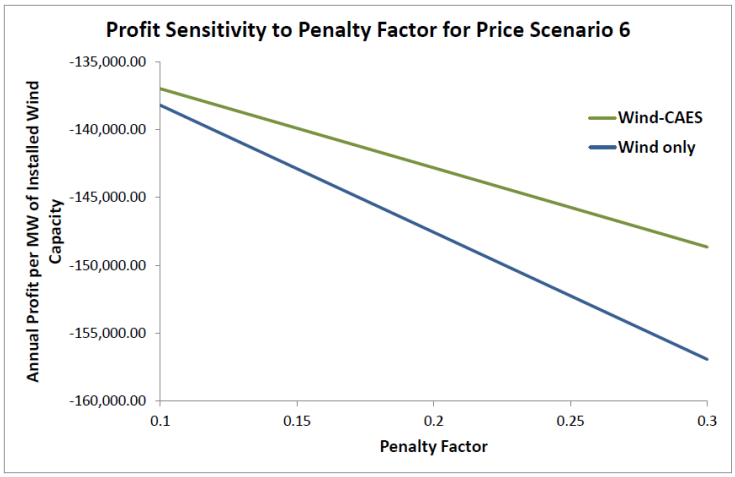

Figure 2.11: Sensitivity of annual profit to the market penalty factor for price scenario 6.

2.4.5 Annual profit with a CO_2 price

If an energy policy is enacted to place a price on carbon dioxide emissions, wind would benefit from higher prices without an increase in generation costs. We are interested in knowing how carbon prices affect the profitability of the wind-CAES system in our model. Since the wind-CAES system in our model is constrained to storing energy only from the wind farm, and we are only considering profits from the day-ahead market, these results do not apply to wind farms in general. We ran the model with market prices adjusted to reflect a carbon dioxide price.

The effect of a carbon dioxide price was estimated for the ERCOT region using the method described by Newcomer et al. (2008). First, we obtained generator data from the Environmental Protection Agency's eGrid database (EPA, 2007) to create a short run marginal cost curve. Next, we used ERCOT load data to estimate market prices. Hourly prices are greatly affected by transmission congestion, generator outages, the volume of electricity sold in the balancing market, and other events. Due to the simplicity of this method, the estimated prices did not match the actual balancing market prices. However, after separating prices by season, we were able to get estimated average hourly prices within 15% of actual average hourly prices. A second marginal cost curve was created with carbon dioxide prices added to the marginal costs for each generator according to its heat rate and fuel type.

Increased electricity prices will reduce demand. To estimate market prices with a carbon dioxide price, we assumed a price elasticity of demand to be -0.1, the reported typical short term value for elasticity by Spees and Lave (2007). As shown in Newcomer et al. (2008), generator merit order in the marginal cost curve will change only slightly for carbon prices up to $50 per tonne. To get price inputs for the model in the hypothetical carbon dioxide pricing situation we first subtracted our estimated prices without carbon dioxide pricing from the actual prices. We then added the residuals to our estimated prices in the carbon dioxide pricing scenario. For price

spikes greater than the largest generator marginal cost we did not alter the price. The goal was to create market prices resembling a short term reaction to a carbon dioxide price.

Figure 2.12 shows our results for carbon prices of $20/tonne and $50/tonne. The EU carbon trading price for 2011 fluctuated between $13 and $26/tonne (Turner, 2011). We feel that a $20/tonne price on carbon dioxide may be possible in U.S. in the future, but higher prices are much less likely. As evidence, the American Clean Energy and Security Act of 2009 was expected to create a price of $15/tonne rising to $26 by 2019 (CBO, 2009), however it failed to become a law. Anything beyond a $50 per tonne price on carbon dioxide seems impossible in the near future in the U.S.

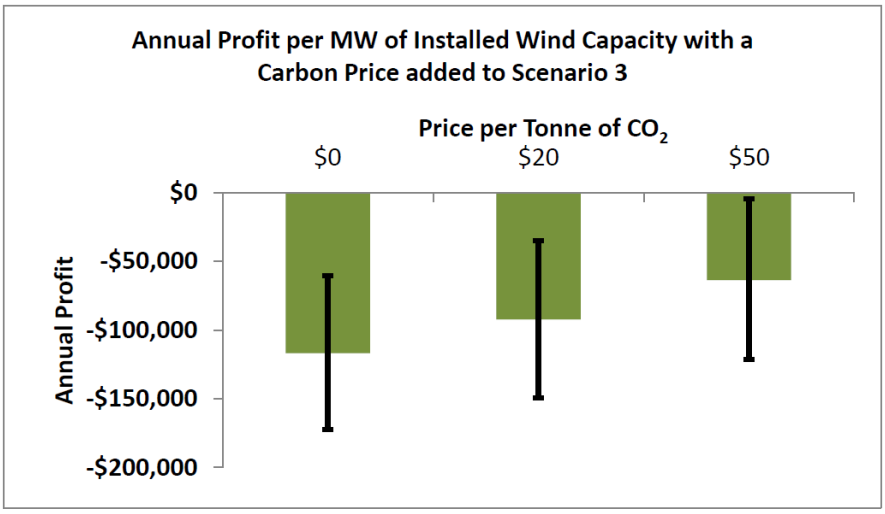

Figure 2.12: Annual profit per installed MW of capacity with carbon prices included in price scenario 3 (ERCOT 2008). Uncertainty bars were created by running the model with different cost assumptions for the wind-CAES system.

A carbon dioxide price of $50 per tonne resulted in annual losses for the wind-CAES system. Price scenario 3 was much more favorable for the model than the other scenarios tested. Revenue from the day-ahead market with a carbon dioxide price would not cover the capital costs. Our analysis with carbon dioxide prices looked only at the short run price change. In the

long run, it is more likely that market prices would drop slightly as generation companies adapted to a carbon dioxide price.

We also ran the model with a constant wholesale electricity price increase. This is essentially what the production tax credit does for a wind farm. It is not clear how the federal production tax credit would apply to a wind farm with CAES since energy from CAES is not considered renewable. What is clear is that our hypothetical wind-CAES system will not be economically feasible unless wholesale market prices are increased. Therefore, we added $60 to every hourly price in the price scenarios considered to observe the results. Figure 13 shows the annual profit for each price scenario along with the estimated annual costs explained earlier. The uncertainty bars span above zero in seven of the price scenarios indicating that a wind-CAES system with low capital costs would be profitable with this subsidy.

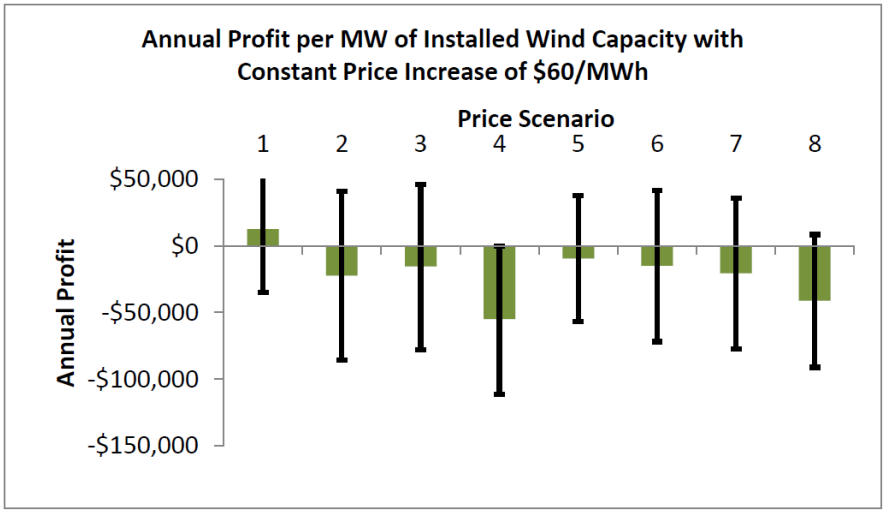

Figure 2.13: Annual profits per installed MW of wind capacity with a constant price increase of $60 per MWh using eight different market price scenarios. Uncertainty bars were created by running the model with different cost assumptions for the wind-CAES system.

2.5 Conclusions and discussion

In time, wind generation may grow to the point where it no longer needs tax support and RPS obligations, and must bid into markets as do other generators. We investigated the option of reducing dispatch uncertainty in systems with very large wind energy by trading wind power in day-ahead markets from a wind-CAES system. In this case, the CAES was sized and used for the sole purpose of reducing the risk associated with dispatch uncertainty. Other potential revenue streams such grid electricity arbitrage, ancillary services, and capacity markets were not considered.

The profitability of collocating a CAES plant with a wind farm to allow participation in the day-ahead energy markets is dependent on several factors including 1) the uncertainty of wind power forecasts, 2) wholesale market prices and 3) wind generation capacity factor. We analyzed one wind-CAES example with eight different market price scenarios using data from ERCOT and MISO to determine its economic viability. In the case studied, it was shown that adding a CAES plant to a wind farm did not provide the ability to make a profit trading energy in the day-ahead market. A wind farm in a more favorable location may have different results since forecast accuracy, markets prices and wind resources vary from one location to another.

Costs and forecast uncertainty are the largest contributors to the unprofitability of the modeled plant. Forecast uncertainty can be reduced significantly by aggregating wind farms (Lang et al, 2006). A CAES plant centrally located among multiple wind farms may allow wind trading on the day-ahead market. In this scenario, the CAES dispatch would be optimized according to price forecasts and aggregated wind forecast accuracy from the corresponding wind farms. One could envision a set of power purchase agreements between the wind farms and the CAES plant.

We find that a wind-CAES system is unlikely to be profitable in the day ahead market without tax incentives or the requirements of an RPS. We examined the effects of a price on carbon (highly uncertain in the USA), finding that a price of $50/tonne CO_2 would not be sufficient for profitability. A large increase in wholesale market prices or a substantial decrease in wind-CAES costs is needed to allow a wind-CAES system used in the manner modeled here to profit from day-ahead energy sales.

2.6 References

1. Bludszuweit, H., Dominguez-Navarro, J. A., Llombart, A., 2008. Statistical analysis of wind power forecast error. IEEE Transactions on Power Systems, 23, 983-991.
2. Castronuovo E.D., Lopes, J.A., 2004. Optimal operation and hydro storage sizing of a wind-hydro power plant. Electrical Power and Energy Systems, 26, 771-778.
3. Congressional Budget Office (CBO), 2009. Analysis of H.R. 2454. Report presented to the House Committee on Energy and Commerce. http://www.cbo.gov/ftpdocs/102xx/doc10262/hr2454.pdf
4. Database of State Incentives for Renewables and Efficiency (DSIRE), http://www.dsireusa.org/
5. Denholm, P., Sioshansi, R., 2009. Value of compressed air energy storage with wind in transmission-constrained electric power systems. Energy Policy, 37(8), 3149-3158.
6. EPRI-DOE, 2004. Handbook supplement of energy storage for grid connected wind generation applications. 1008703
7. Energy Information Administration (EIA), 2010a. Electric net summer capacity: Total (all sectors) 1949-2007, http://www.eia.doe.gov/emeu/aer/txt/stb0811a.xls
8. Energy Information Administration (EIA), 2010b. Electricity market module. DOE/EIA-0554; http://www.eia.gov/forecasts/aeo/assumptions/pdf/electricity.pdf
9. Energy Information Administration (EIA), 2011a. Annual electricity outlook. DOE/EIA-0383; http://www.eia.gov/forecasts/aeo/
10. Energy Information Administration (EIA), 2011b. Annual energy review 2010. DOE/EIA-0384; http://www.eia.gov/totalenergy/data/annual/pdf/aer.pdf
11. Environmental Protection Agency (EPA), 2007. eGrid 2007 version 1.1, http://www.epa.gov/cleanenergy/energy-resources/egrid/index.html
12. Fertig, E., Apt, J., 2011. Economics of compressed air energy storage to integrate wind power: A case study in ERCOT. Energy Policy, 39, pp. 2330-2342
13. Focken, U., Lange, M., Mönnich, K., Waldl, H.-P, Beyer, H. G., and Luig, A., 2002. 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, 90, pp. 231–246

14. Garcia-Gonzalez, J., de la Muela, R. M. R., Santos, L. M., Gonzalez, A. M., 2008. Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. IEEE Transactions on Power Systems, 23, 460-468.
15. Greiner, C., Korpas, M., Gjengedal, T., 2009. Optimal operation of energy storage systems combined with wind Power in short-term power markets. European Wind Energy Conference and Exhibition (EWEC).
16. Harper, J. P., Karcher, M. D., Bolinger, M., 2007. Wind project financing structures: A review and comparative analysis. Lawrence Berkley Laboratory, LBNL-63434, http://eetd.lbl.gov/ea/emp
17. Hillier, F. S., Lieberman, G. J., 2004. Introduction to Operations Research, Eigth ed. McGraw-Hill Professional, New York.
18. Kim, J.H., Powell, W.B., 2009. Optimal energy commitments with storage and intermittent supply. Operations Research, 59 (6), pp. 1347-1360
19. Nandalal, K., Bogardi, J., 2007. Dynamic Programming Based Operation of Reservoirs – Applicability and Limits. Cambridge University Press, New York.
20. Newcomer, A., Blumsack, S. A., Apt, J., Lave, L. B., Morgan, M. G., 2008. Short Run Effects of a Price on Carbon Dioxide Emissions from U.S. Electric Generators. Environmental Science and Technology, 42(9), 3139-3144.
21. Potomac Economics, 2009. ERCOT State of the Market Report 2008.
22. Sullivan, P., Short, W., Blair, N., 2008. Modeling the Benefits of Storage Technologies to Wind Power. National Renewable Energy Laboratory Conference Paper NREL/CP-670-43510.
23. Spees, K., Lave, L. B., 2007. Demand Response and Electricity Market Efficiency, Carnegie Mellon University, Carnegie Mellon Electricity Industry Center Working Paper CEIC-07-01, http:// wpweb2k.gsia.cmu.edu/ceic/papers/ceic-07–01.asp.
24. Succar, S., 2011. Large Energy Storage Systems Handbook, Chapter 5, CRC Press, Boca Raton, Florida.
25. Turner, G., 2011. Emissions are Under-Priced in Europe. Bloomberg New Energy Finance, http://www.bnef.com/PressReleases/view/165
26. Windustry, http://www.windustry.org, last accessed September 2010
27. Wiser, R., Bolinger, M., 2011. 2010 Wind Technologies Market Report. Lawrence Berkley Laboratory.

Chapter 3 - An effective method for modeling wind power forecast uncertainty[2]

3.1 Abstract

Wind forecasts are an important tool for electric system operators. Proper use of wind power forecasts to make operating decisions must account for the uncertainty associated with the forecast. Data from different regions in the USA with forecasts made by different vendors show the forecast error distribution is strongly dependent on the forecast level of wind power. At low wind forecast power, the forecasts tend to under-predict the actual wind power produced, whereas when the forecast is for high power, the forecast tends to over-predict the actual wind power. Most of the work in this field neglects the influence of wind forecast levels on wind forecast uncertainty and analyzes wind forecast errors as a whole. The few papers that account for this dependence, bin wind forecast data and fit parametric distributions to actual wind power in each bin. In the latter case, different parameters and possibly different distributions are estimated for each data bin. We present a method to model wind power forecast uncertainty as a single closed-form solution using a logit transformation of historical wind power forecast and actual wind power data. Once transformed, the data become close to jointly normally distributed. We show the process of calculating confidence intervals of wind power forecast errors using the jointly normally distributed logit transformed data. This method has the advantage of fitting the entire dataset with five parameters while also providing the ability to make calculations conditioned on the value of the wind power forecast.

[2] The work in this paper is currently under review in Renewable Energy

3.2 Introduction

Wind power experienced substantial growth over the past decade in the U.S. and Europe. Installed capacity in the U.S. increased tenfold from 4.2 GW in 2001 to 47 GW in 2011 and now provides nearly 3% of total electrical energy (DOE, 2012). European installed wind capacity increased fivefold in the same time period, from 17.3 GW to 94 GW, and produces over 6% of total electrical energy (Wilkes et al., 2011). Since wind power is not fully dispatchable, wind forecasts are useful for planning and operations in electric power systems.

Electric system operators rely on wind power forecasts for decision making. For system operators, it is important not only that forecasts are accurate, but that the degree of inaccuracy is known. Operating reserves must be procured in advance to cover the uncertainty of wind power forecasts (and load forecasts). This not only applies to short term operation planning, but also to long term resource analysis. Here we present a model for wind power forecast uncertainty dependent on the wind power forecast based on historical wind forecast errors.

Broadly speaking, there are two approaches in the research literature to model wind forecast errors. One method models all wind power forecast errors as having been drawn from a single population with some known parametric probability distribution. Often a normal distribution is used under the assumption that the aggregation of many geographically diverse wind generators in a system justifies the application of the central limit theorem (e.g. Ortega-Vazquez and Kirschen (2009); Bouffard and Galiana (2008) and Doherty and O'Malley (2005)). In other work, forecast errors were fit with a Cauchy distribution (Hodge and Milligan, 2011), a hyperbolic distribution (Hodge et al., 2012), and a doubly truncated normal distribution (Makarov et al., 2009).

Modeling all forecast errors with a single distribution makes the assumption that there is no dependence of forecast error on the forecast wind power level. Lange (2005) transformed wind

speed forecast errors into wind power forecast errors using the nonlinear relation of wind power to wind speed, which produced a distribution of power forecast errors more sharply peaked than the distribution of wind speed forecast errors. Power forecast errors were also shown to be skewed at the extreme values of the forecast range and symmetric near the center of the forecast range.

The other broad approach to modeling wind forecast errors is to condition forecast error distributions on the expected level of wind power. Neilsen et al. (2006) determined wind power forecast confidence levels conditioned on the forecast wind power using quantile regression. Luig et al. (2001) and Bludszuweit et al. (2008) binned wind data by the wind power forecast values and fit Beta distributions to the observed wind power values associated with each bin of power forecasts. Al-Awami and Sharkawi (2009) presented a similar analysis with both beta and extreme value distributions. In each case, separate distribution parameters were required for each bin.

We present data showing that forecast error distributions for the largest USA wind regions are dependent on forecast wind values, and present a method to model wind forecast errors conditioned on the forecast value by applying a logit (or logistic) transformation to the wind forecast and actual wind power data. We find that the logit transformed variables can be reasonably modeled with a bivariate normal distribution, which is considerably easier to analyze than the original data. This method fits a single model to the entire set of data, meaning that only one set of parameters must be estimated, and produces smooth results without discontinuities that arise between bins.

Lau and McSharry (2010) and Pinson (2012) applied logit transformations to wind power time series to make them "more Gaussian" before fitting time series models for wind power forecasting. Both papers found that logit transformations of wind power data can be accurately

modeled with normal distributions. They then fit statistical time series models to the transformed data to generate short term wind power forecasts from 1 to 48 hours. Here we utilize wind power forecasts independently generated by electric power system operators, show that these forecasts are also well described by the logit-normal model, and represent the observed and forecast wind power pairs as a bivariate logit-normal distribution in which the forecast error is implicitly included. This allows a characterization of the forecast error over the range of time scales for which these forecasts are provided.

This chapter is organized as follows. In section 3.2 we briefly describe the data used in this study. Section 3.3 contains some discussion on the dependence of wind forecast uncertainty to wind forecast values. In section 3.4 we outline a method to model this uncertainty by applying a logit transformation to the wind forecast and wind power data. Results are shown for day-ahead wind power forecasts at the ISO level. We also present results using hour-ahead forecasts. Section 3.5 summarizes our conclusions.

3.3 Wind data used

We used wind forecast and actual wind power data from the Electric Reliability Council of Texas (ERCOT) and the Midwest Independent System Operator (MISO). ERCOT's territory covers most of the state of Texas while MISO covers most of the Midwestern portion of the U.S.

The ERCOT data included hourly wind forecast values for 1 to 48 hour look-ahead times covering the years 2009 and 2010. Included with the forecast data were hourly actual wind generation values and hourly estimates of what wind generation would have been if there were no curtailments. When analyzing wind power uncertainty, we used the estimated values of uncurtailed wind power since wind curtailments are not considered in wind power forecasts.

ERCOT curtailed an estimated 17% of wind generation in 2009 (Rogers et al. 2010) and 10% of wind generation in 2010 (Wiser and Bolinger, 2011).

The estimates of uncurtailed hourly wind generation data were made for ERCOT by AWS Truepower based on actual wind generation, meteorological data and curtailment instructions sent from ERCOT to individual wind farms. The AWS analysis assumes that all curtailment instructions were followed and that wind turbine availability was known. In reality, it is not known how well wind farms followed curtailments instructions, and it was apparent that not all wind farms reported wind turbine availability status. Whenever our analysis below refers to "uncurtailed wind power" the AWS estimate is meant.

MISO day-ahead wind forecast and actual wind power data from February 2011 to May 2012 were obtained from the MISO website (MISO, 2012a). MISO wind forecasts are produced by Energy and Meteo GmbH (Kehler et al. 2010). During this period, wind curtailments in MISO were estimated to be 2 to 6% each month (MISO, 2012b) which create minor problems with our analysis as will be shown in the next section. Further information on the data is presented in the appendix.

3.3.1 Wind power forecast error characteristics

Wind power forecasts are provided for system operators of electric power networks in order to assist with decision making. Look-ahead times range from 5 minutes to several days. Wind forecasts for time periods up to 1 hour ahead provide information used in economic dispatch or real-time trading. Longer look-ahead times assist system operators in unit commitment decisions and provide wind farm operators with information for day-ahead market bids.

Decision makers who rely on wind power forecasts must not only prepare for the amount of wind power expected in the grid, they must also prepare for the chance that the forecast is wrong. Wind forecast uncertainty here is characterized by the distribution of wind forecast errors. We

define wind forecast error (e) as the forecast value (F) minus the actual wind power (W) with all variables normalized by the installed wind capacity.

$$e = F - W \tag{3.1}$$

Modeling wind power uncertainty is challenging due both to the highly variable nature of wind speed over different time scales and to the non-linear and variable (for different turbines) transformation relating wind speed to wind power. This causes wind power forecast errors to be non-Gaussian in general and have much different distributions at low wind forecast values than high forecast values (Lange, 2005). Figure 3.1 (a) shows a scatter plot of uncurtailed wind power levels plotted against wind power forecasts for ERCOT during the years 2009 and 2010. As indicated in the plot, points above the unit-slope line are under-forecast while points below the line are over-forecast. Forecast errors are plotted against forecast levels of wind power in Figure 3.1 (b). Since wind power is subtracted from forecasts, positive errors represent over-forecasts and negative errors represent under-forecasts. The three bottom plots in Figure 3.1 (c, d and e) show forecast error distributions at three different wind forecast levels. As evident in the plot, error distributions are skewed left near the low end of the forecast range and skewed right for high wind forecasts.

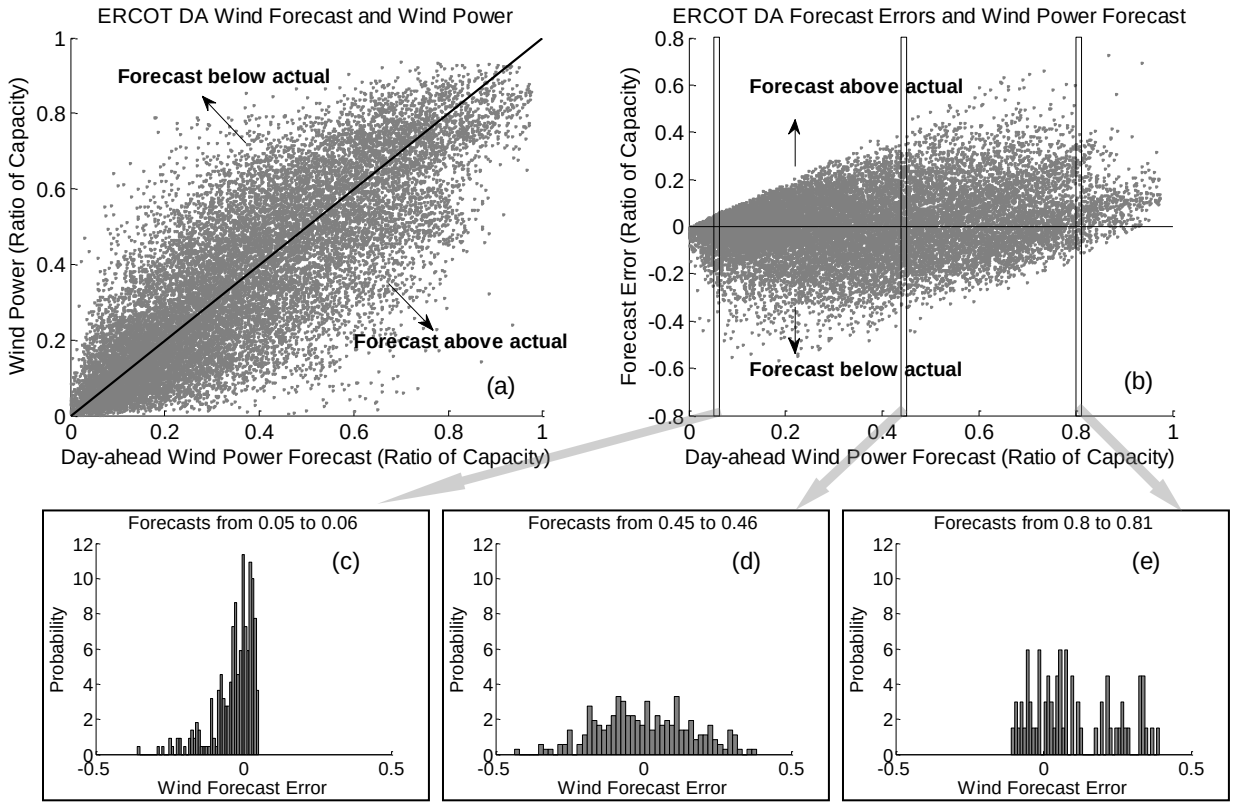

Figure 3.1: ERCOT Estimated uncurtailed wind power (a) and wind forecast errors (b) plotted against forecast wind power. Bottom plots show probability distributions of wind forecast errors corresponding to the three bins highlighted in the (b). Forecast value ranges in the bottom plots are (c) 0.05 to 0.06, (d) 0.45 to 0.46 and (e) 0.8 to 0.81. All values are shown as ratios of installed wind capacity.

Figure 3.2 shows the forecast bias calculated over the range of forecast values with a moving window of size 0.1 of the normalized wind power forecast. Unlike previous work mentioned earlier, we did not attempt to remove the conditional forecast bias in our analysis. Overall, the mean forecast error in the ERCOT data is 0.007 indicating the forecasts are unbiased, but when conditioned on the forecast value a bias emerges that is related to the forecast level of wind. When the predicted wind power is small, the actual power averages higher than the forecast, producing a negative bias. For high wind power predictions the actual wind power averages lower than forecast producing a positive bias. A similar bias pattern occurs in the MISO wind data and has been observed in wind forecasts for California (Makarov et al. 2010) and Germany (Luig et al., 2001). More discussion on the conditional bias is presented in the appendix.

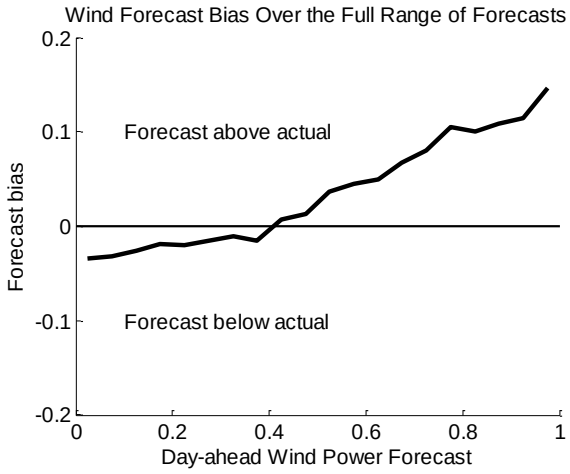

Figure 3.2: ERCOT forecast bias over the range of possible forecasts.

3.4 Logit transformation of wind data

3.4.1 Day-ahead wind power forecasts

We applied logit transforms to data in order to obtain a dataset that can be modeled with a normal distribution. Logit transforms are valid for data constrained between the values 0 and 1. Wind data normalized by installed wind capacity by definition must lie in the range from 0 to 1. Normalized wind data aggregated over an entire electric grid rarely register values equal to 0 or 1 making these data likely candidates for analysis with logit transforms. The underlying assumption in fitting logit transformed data to a normal distribution is that the forecast and wind power data fit a logit-normal distribution (Frederic and Lad, 2008), as discussed in greater detail in the appendix. If we define normalized wind forecast and wind power data as F and W, the transformed variables, F^* and W^* are defined in Equation 3.2. Once transformed, the variables can take values ranging from negative to positive infinity.

$$F^* = \ln\left(\frac{F}{1-F}\right) \qquad W^* = \ln\left(\frac{W}{1-W}\right) \tag{3.2}$$

The transformed variables were fit to a normal distribution with a density function

$$f(X^*) = \frac{1}{\sqrt{2\pi}\sigma_{X^*}} exp\left(-\frac{(X^*-\mu_{X^*})^2}{2\sigma_{X^*}^2}\right) \quad (3.3)$$

The symbols μ_{x*} and σ_{x*} are the mean and standard deviation of the transformed data. Figure 3.3 shows the distributions of F^* and W^* with a fitted normal distribution overlaid. As the graphs show, the normal distribution fits the transformed data well.

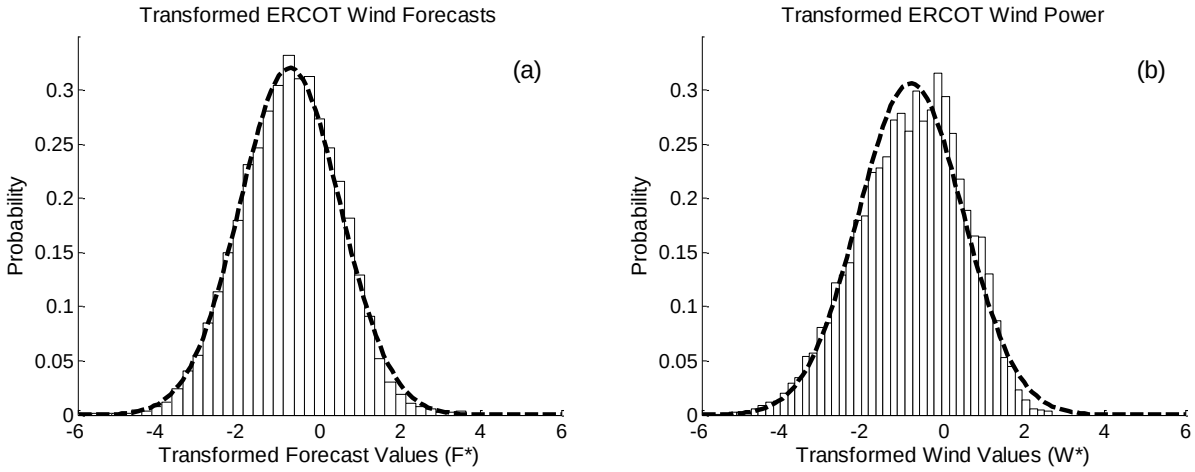

Figure 3.3: Relative frequency distributions (histogram bars) for the transformed wind power forecast (a) and wind power (b) in ERCOT with a fitted normal distribution (dashed line) overlaid. Analogous MISO data are shown in Figure 10.

The transformed variables are plotted against each other in Figure 3.4. The solid lines represent the contour of a bivariate normal distribution fit to the data.

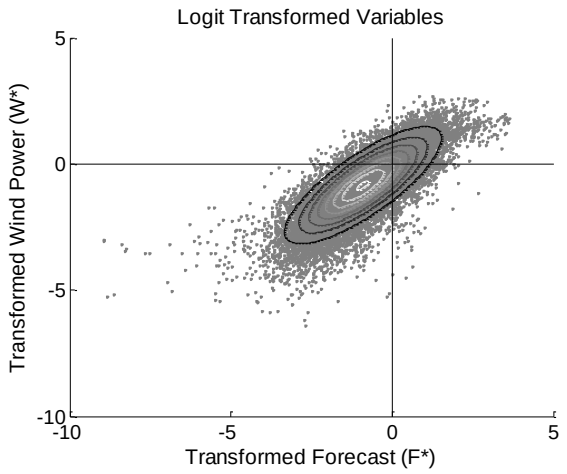

Figure 3.4: Logit transformed wind power data plotted against transformed wind forecast data. The data shown here are transformations of the data displayed in Figure 1 (a). Solid lines are contours of a fitted bivariate normal distribution with μ_{F^*} = -0.74, μ_{W^*} = -0.81, σ_{F^*} = 1.55, σ_{W^*} = 1.70 and ρ = 0.80

We modeled the transformed variables as jointly normally distributed with the bivariate normal distribution:

$$f(F^*, W^*) = \frac{1}{2\pi\sigma_{F^*}\sigma_{W^*}\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(F^*-\mu_{F^*})^2}{\sigma_{F^*}^2}+\frac{(W^*-\mu_{W^*})^2}{\sigma_{W^*}^2}-\frac{2\rho(F^*-\mu_{F^*})(W^*-\mu_{W^*})}{\sigma_{F^*}\sigma_{W^*}}\right]\right)} \quad (3.4)$$

where $$\rho = \frac{cov(W^*,F^*)}{\sigma_{W^*}\sigma_{F^*}} \quad (3.5)$$

and cov(,) is the covariance.

Using the fitted normal distribution provides a convenient way to model wind forecast errors. For a particular forecast value, the transformed wind power is modeled with a normal distribution with the conditional probability density defined by.

$$f(W^*|F^*) = \frac{1}{\sqrt{2\pi}\sigma_{W^*|F^*}} exp\left(-\frac{(W^*-\mu_{W^*|F^*})^2}{2\sigma_{W^*|F^*}^2}\right) \quad (3.6)$$

where $$\mu_{W^*|F^*} = \mu_{W^*} + \frac{\rho\sigma_{W^*}}{\sigma_{F^*}}(F^* - \mu_{F^*}) \quad (3.7)$$

$$\sigma_{W^*|F^*} = \sigma_{W^*}\sqrt{1-\rho^2} \quad (3.8)$$

Given a wind power forecast value, a decision maker is interested in knowing a particular confidence interval for the expected wind power. A confidence interval of 95% ranges from the 2.5th percentile to the 97.5th percentile calculated from the inverse of the cumulative distribution function (*CDF*). In the case of a normal distribution, this calculation simplifies to a function of the conditional mean and standard deviation. The end values of a given confidence interval are calculated by

$$CI_{\alpha}\{W^{*}|F^{*}\} = \left[\mathbf{CDF}^{-1}_{W^{*}|F^{*}}\left(\frac{1-\alpha}{2}\right), \quad \mathbf{CDF}^{-1}_{W^{*}|F^{*}}\left(\frac{1+\alpha}{2}\right)\right] \tag{3.9}$$

$$= [\mu_{W^{*}|F^{*}} - Z_{\alpha}\,\sigma_{W^{*}|F^{*}}, \quad \mu_{W^{*}|F^{*}} + Z_{\alpha}\,\sigma_{W^{*}|F^{*}}] \tag{3.10}$$

The symbol α ranges between 0 and 1 and indicates the desired confidence interval level (i.e. 0.95 for a 95% confidence interval). The value of Z_{α} in Equation 10 is selected from standard normal distribution tables to produce a particular confidence interval, CI_{α}. Table 3.2 shows selected values of Z_{α} for desired confidence intervals.

Table 3.1: Z values for selected confidence intervals (CI).

α	70%	80%	85%	90%	95%
Z_{α}	1.047	1.282	1.440	1.645	1.960

Figure 3.5 shows the 95% normal distribution confidence interval of actual wind power in the transformed space conditioned on the forecast wind power. The solid line is the mean of W^{*} as a function of F^{*}, and the region within the dashed lines represents 1.96 standard deviations from the mean which contains 95% of the W^{*} values assuming a bivariate normal distribution.

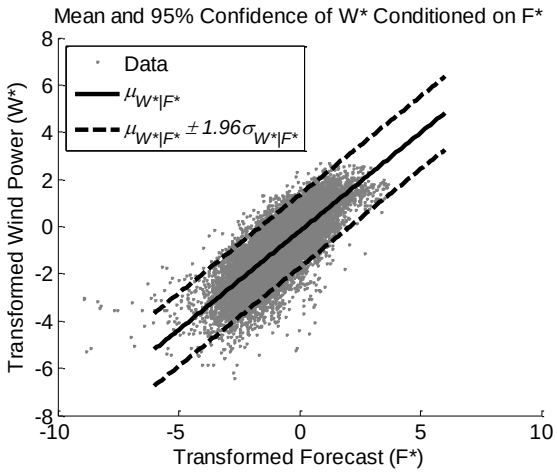

Figure 3.5: Transformed ERCOT data with conditional mean for W* and 95% confidence interval of W* as a function of F*.

A confidence interval in the transformed space is converted to the confidence interval in the original data space using Equation 11 for each variable. The result is displayed as solid lines in Figure 3.6.

$$F = \frac{1}{1+e^{-F^*}} \qquad W = \frac{1}{1+e^{-W^*}} \tag{3.11}$$

Once a confidence interval is determined for the observed wind power given a wind forecast, we can determine a confidence level for wind forecast errors by subtracting the confidence interval values of wind power from the wind forecast values. Figure 3.7 shows a range of confidence intervals for actual wind power based on the day-ahead wind forecast and the confidence levels of wind forecast errors. As shown in Figure 3.7 (b), the confidence intervals are generally asymmetric and not centered on zero.

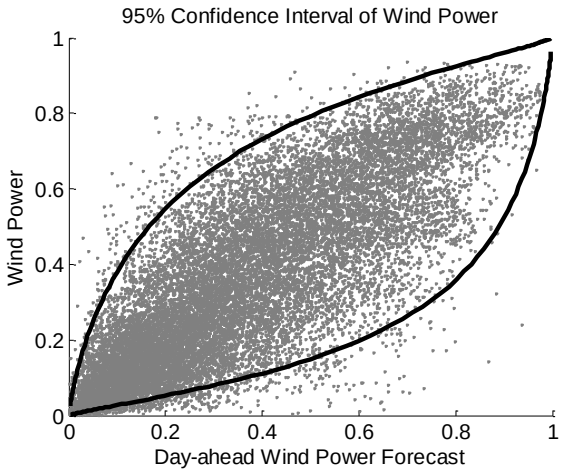

Figure 3.6: Ninety-five percent confidence interval for estimated uncurtailed ERCOT wind power in the original space plotted as a function of the day-ahead forecast level of wind power.

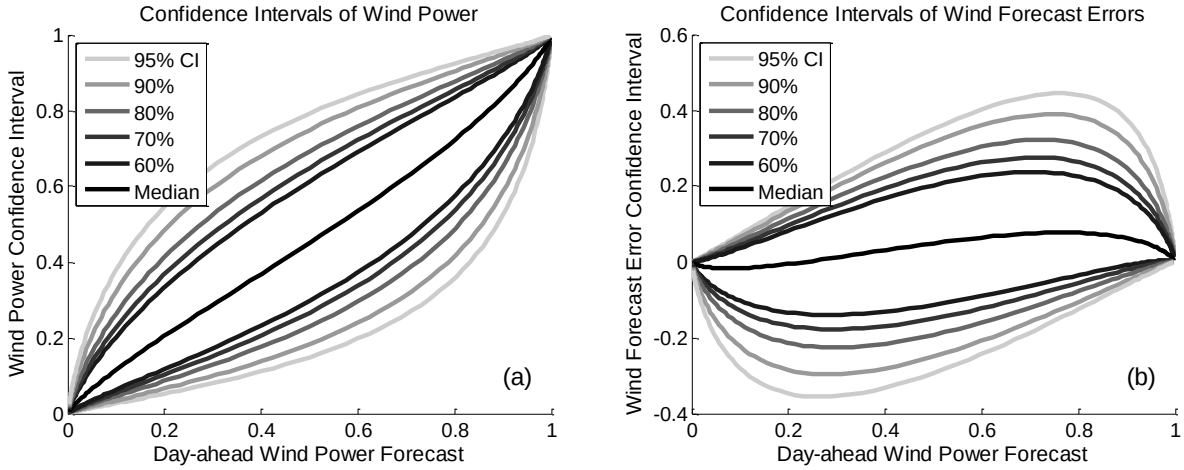

Figure 3.7: ERCOT wind power confidence levels (left plot) based on wind forecast level and wind error confidence levels (right plot) based on wind forecast level.

We plotted 95, 90 and 80% confidence intervals of wind forecast errors using the logit transform and compared these to the same confidence intervals calculated with a moving window of width 0.1 in Figure 3.8. Fitting a bivariate normal distribution to the logit transform provides a good closed-form representation of the wind forecast errors requiring a fit of only 5 parameters; μ_{F^*}, μ_{W^*}, σ_{F^*}, σ_{W^*} and ρ. In comparison, fitting a Beta distribution to 50 bins of wind

data requires the estimation of 100 parameters (2 parameters for each bin) as was done in Bludszuweit et al. (2008).

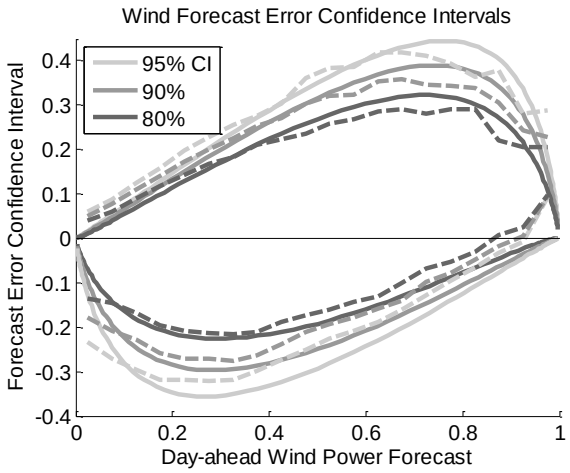

Figure 3.8: Confidence intervals conditioned on ERCOT wind forecast level. Solid lines are calculated using the logit transformation, and dashed lines are calculated by binning data according to the forecasts.

Figures 3.2-3.8 were created using day-ahead wind forecasts and estimated uncurtailed wind power data from ERCOT. We also modeled wind forecast errors with MISO wind data. The MISO data did not contain wind curtailment estimates so we used observed wind power with wind forecasts. As mentioned earlier, curtailments in MISO ranged from 2 to 6% of total wind energy generated, depending on the month. We feel that these levels of wind curtailments are low enough to provide meaningful results. A system operator analyzing wind forecast errors would likely have uncurtailed wind generation estimates available, and would be able to model the forecast errors more accurately.

Figure 3.9 (a) shows the actual wind power against the day-ahead forecast wind power for each hour of the sample period with all data normalized by the wind capacity. Note that wind power levels in MISO rarely reach 80% of installed capacity. This may be due to the fact that wind farms in MISO are distributed over a much wider and more geographically diverse region

than ERCOT wind farms. Figure 3.9 (b) displays the same data after applying the logit transformation.

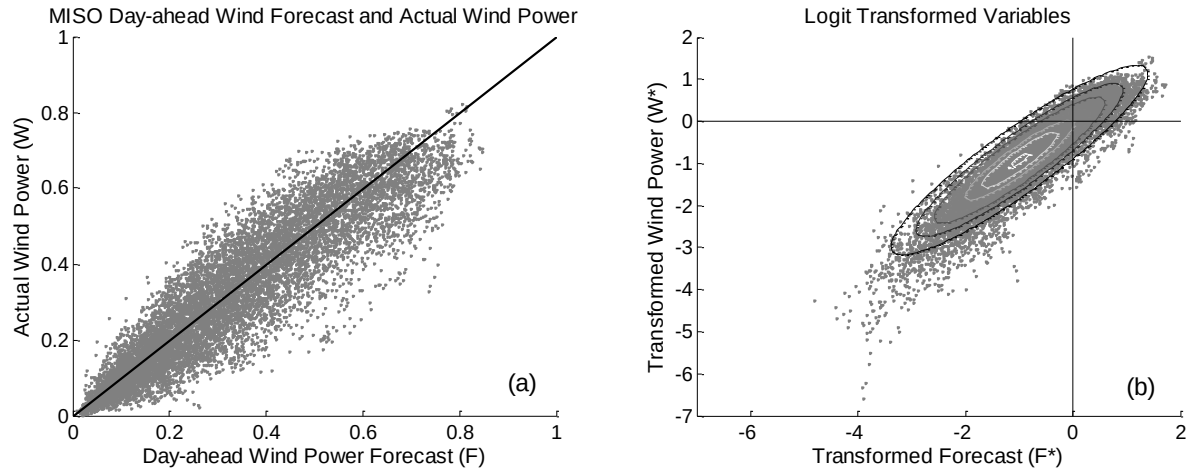

Figure 3.9: (a) Actual wind power in MISO plotted against the day-ahead forecast values and (b) the logit transformation of the same data. A contour (solid lines) of the fitted bivariate normal distribution is plotted over the data in (b) with μ_{F^*} = -0.82, μ_{W^*} = -0.89, σ_{F^*} = 1.05, σ_{W^*} = 1.18 and ρ = 0.92

In order to evaluate how normally distributed the transformed data are, the relative frequency distributions of the transformed data are displayed in Figure 3.10.

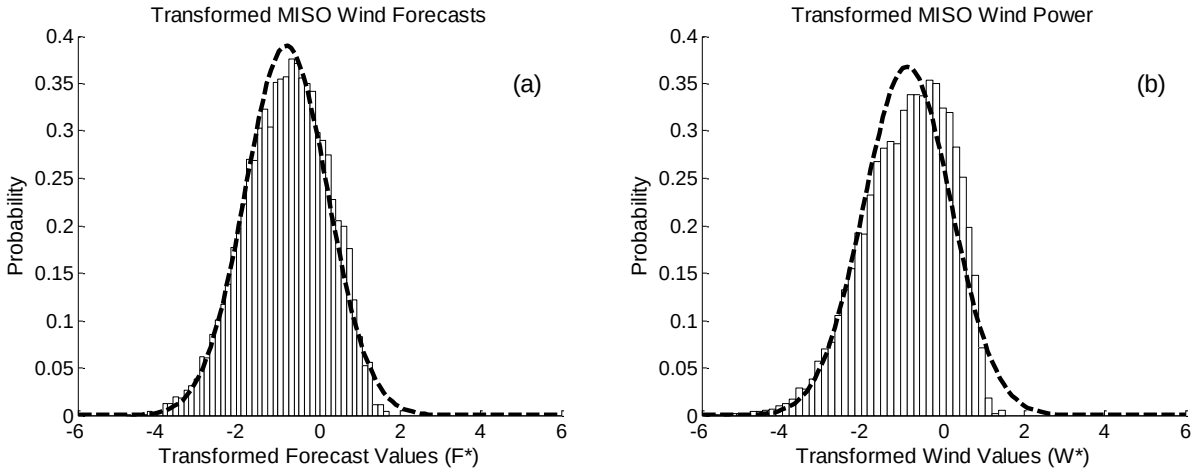

Figure 3.10: Relative frequency distributions (histogram bars) for the transformed wind power forecast (a) and wind power (b) in ERCOT with a fitted normal distribution (dashed line) overlaid.

The transformed actual wind power (W^*) distribution is skewed left. One likely reason for this skewness is the effect of wind curtailments on the data. Since wind power in MISO rarely exceeds 80% of installed capacity, verification of confidence intervals was impossible for forecasts near a normalized value of 1. Figure 3.11 shows the 95% confidence envelope for MISO wind errors calculated with the logit transformation (solid lines) and a moving window of width 0.1 (dashed lines).

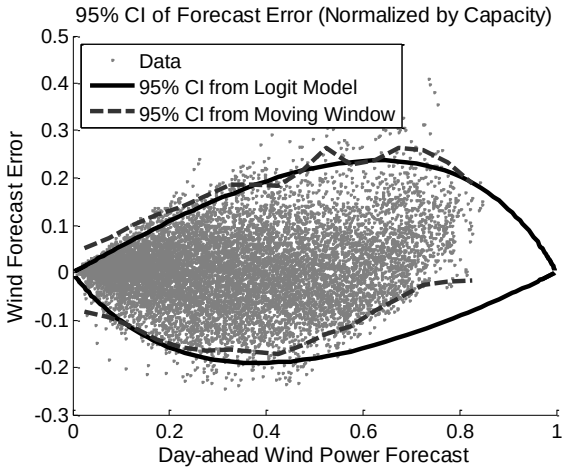

Figure 3.11: Wind forecast errors in MISO plotted against wind forecast values with 95% confidence intervals. Solid lines are calculated using the logit transformation, and dashed lines are calculated with a moving window of width 0.1.

3.4.2 Hour-ahead wind power forecasts

We applied the same analysis to ERCOT and MISO wind data using hour-ahead forecasts. Forecasts were created using the persistence method where the wind power forecast for one hour is equal to the actual wind power from the previous hour (Equation 3.12).

$$F(t) = W(t-1) \tag{3.12}$$

As in the previous section, we used the estimated uncurtailed values for wind power in ERCOT. Figure 3.12 (a) shows the actual wind power values plotted against the hour-ahead forecast wind power values. Hour-ahead wind forecast error confidence intervals are shown in Figure 3.12 (b) where dashed lines show intervals calculated with a moving window of width 0.1 and solid lines are derived from a logit transformation. Figure 3.13 shows the same plots for the MISO data.

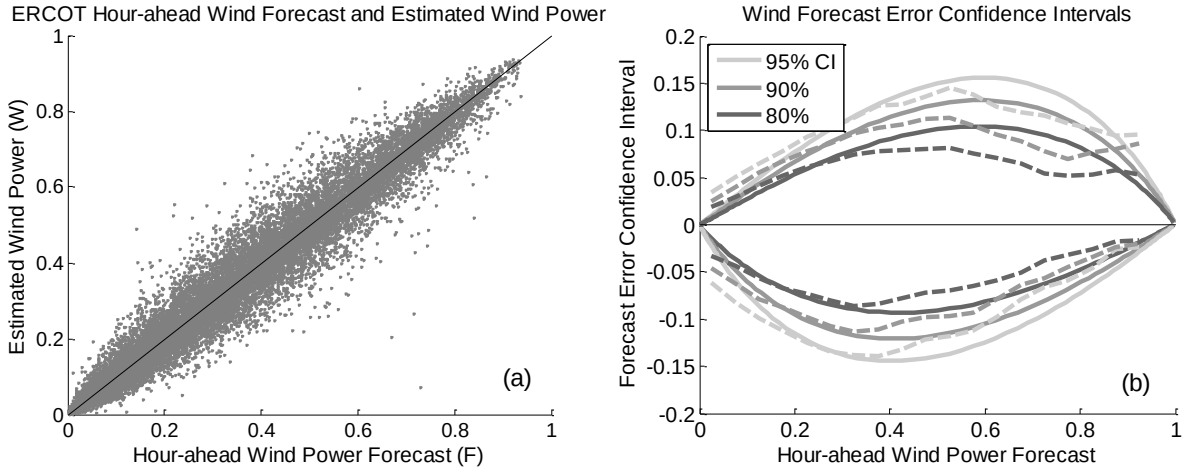

Figure 3.12: ERCOT estimated uncurtailed wind power plotted against the hour-ahead persistence forecasts (a), and confidence intervals for wind forecast errors (b). Solid lines are calculated using the logit transformation, and dashed lines are calculated with a moving window.

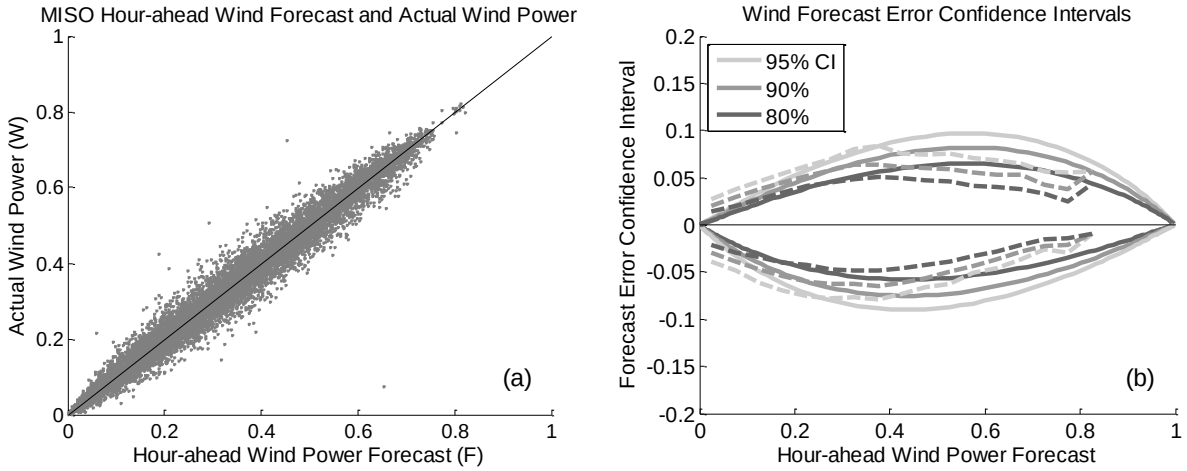

Figure 3.13: MISO wind power plotted against the hour-ahead persistence forecasts (a), and confidence intervals for wind forecast errors (b). Solid lines are calculated using the logit transformation, and dashed lines are calculated with a moving window.

The agreement between confidence intervals calculated with logit transforms and the moving window are similar to that obtained with the day-ahead forecast data. The scale of the vertical axis is smaller in Figures 3.12 (b) and 3.13 (b) than in the figures for the day-ahead results. Once again, verification of confidence intervals was impossible for forecasts near a normalized value of 1 with the MISO data since the maximum wind power observed was around 0.8. While the agreement between confidence intervals calculated with the logit transformation model and the moving window is not perfect, it is reasonably close. A logit transformation model is shown to be an effective tool to estimate uncertainty in wind power forecasts.

3.5 Conclusion

Observed wind power forecast error distributions are highly dependent on the forecast level of wind power. At low wind forecast power, the forecasts under-predict the actual wind power produced, whereas when the forecast is for high power, the forecast tends to over-predict the actual wind power.

Thus, forecast errors modeled with a single distribution do not provide adequate information for system operators. We presented a method to analyze wind forecast errors with a logit transformation. Transforming wind data with a logit transform in this manner is a straightforward method to determine the amount of uncertainty associated with wind forecasts using historical data. The advantage of this method is that transformed data can be accurately modeled with a bivariate normal distribution. This greatly simplifies the analysis since one set of parameters is estimated instead of multiple parameters for different forecast levels. Calculations of confidence intervals with this method use a model fit to the entire dataset while providing the ability to condition wind uncertainty on the wind forecast value for a given time period. We

applied the logit transform to hourly data, but there is no reason that this method should not work well with different time scales.

Electric grid system operators can use this model in their respective decision making analysis. Proper confidence intervals of wind forecast errors for a given forecast level of wind power is essential for decision making. Uncertainty associated with wind power and load forecasts determine the amount of reserve requirements for reliable grid operation.

3.6 References:

1. Al-Awami, A.T., El-Sharkawi, M.A., 2009. Statistical characterization of wind power output for a given wind power forecast, North American Power Symposium , pp. 1-4.
2. Bludszuweit, H., Dominguez-Navarro, J.A., Llombart, A., 2008. Statistical analysis of wind power forecast error, IEEE Transactions on Power Systems, 23 (3), pp. 983-991.
3. Bouffard, F., Galiana, F.D., 2008. Stochastic security for operations planning with significant wind power generation, 2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, pp. 1-11.
4. DOE, 2012. Wind powering america: Installed U.S. wind capacity and wind project locations, U.S. Department of Energy, http://www.windpoweringamerica.gov/wind_installed_capacity.asp .
5. Doherty, R., O'Malley, M., 2005. A new approach to quantify reserve demand in systems with significant installed wind capacity, IEEE Transactions on Power Systems, 20 (2), pp. 587- 595.
6. Focken, U., Lange, M., Mönnich, K., Waldl, H.P., Beyer, H.G., Luig, A., 2002. 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, 90, pp. 231–246.
7. Frederic, P., Lad, F., 2008. Two moments of the logitnormal distribution, Communications in Statistics - Simulation and Computation , 37 (7),pp. 1263-1269.
8. Giebel, G., Sørensen, P., Holttinen, H., 2007. Forecast error of aggregated wind power, TradeWind Deliverable Report, Risø-I-2567(EN).
9. Hodge, B., Milligan, M., 2011. Wind power forecasting error distributions over multiple timescales, 2011 IEEE Power and Energy Society General Meeting, pp. 1-8.
10. Hodge, B., Florita, A., Orwig, K., Lew, D., Milligan, M., 2012. A Comparison of wind power and load forecasting distributions, 2012 World Renewable Energy Forum, NREL/CP-5500-54384, http://www.nrel.gov/docs/fy12osti/54384.pdf

11. Kehler, J., Ming, H., McMullen, M., Blatchford, J., 2010. ISO perspective and experience with integrating wind power forecasts into operations, 2010 IEEE Power and Energy Society General Meeting, pp. 1-5.

12. Lange, M., On the uncertainty of wind power predictions—Analysis of the forecast accuracy and statistical distribution of errors, Journal of Solar Energy Engineering, Vol. 127, (2005) 177–184.

13. Lau, A., McSharry, P., 2010. Approaches for multi-step density forecasts with applications to aggregated wind power, The Annals of Applied Statistics, 4 (3), pp. 1311–1341.

14. Luig, A., Bofinger, S., Beyer, H.G., 2001. Analysis of confidence intervals for the prediction of regional wind power output, Proceedings of the European Wind Energy Conference, Copenhagen, Denmark, pp. 725-728.

15. Makarov, Y.V., Loutan, C., Ma, J., de Mello, P., 2009. Operational impacts of wind generation on California power systems, IEEE Transactions on Power Systems, 24 (2), pp. 1039-1050.

16. Makarov, Y.V., Guttromson , R.T., Huang, Z., Subbarao, K., Etingov, P.V., Chakrabarti, B.B., Ma, J., 2010. Incorporating wind generation and load forecast uncertainties into power grid operations, Report PNNL-19189, PNNL.

17. MISO, 2012a. Day Ahead Wind Forecast, https://www.midwestiso.org/MarketsOperations/ RealTimeMarketData/ Pages/DayAheadWindForecast.aspx

18. MISO, 2012b. Reliability Subcommittee Report June 27, 2012. https://www.misoenergy.org/Library/Repository/Meeting Material/Stakeholder/RSC/2012/20120626/20120626 RSC Item 08 Wind Curtailment Data.pdf

19. Nielsen, H.A., Madsen, H., Nielsen, T.S., 2006, Using quantile regression to extend an existing wind power forecasting system with probabilistic forecasts, Wind Energy, 9, pp. 95–108.

20. Ortega-Vazquez, M.A., Kirschen, D.S., 2009. Estimating the spinning reserve requirements in systems with significant wind power generation penetration, IEEE Transactions on Power Systems, 24 (1), pp. 114-124.

21. Pinson, P., 2012. Very-short-term probabilistic forecasting of wind power with generalized logit–normal distributions, Journal of the Royal Statistical Society: Series C (Applied Statistics), 61 (4), pp. 555-576.

22. Rogers, J., Fink, S., Porter, K., 2010. Examples of wind energy curtailment practices, NREL Subcontract Report, NREL/SR-550-48737.

23. Wilkes, J., Moccia, J., Drangan, M., 2012. Wind in power: 2011 European wind statistics, European Wind Energy Association, http://www.ewea.org/ fileadmin/ewea_documents/documents/publications/statistics/Stats_2011.pdf

24. Wiser, R., Bolinger, M., 2011. DOE 2010 Wind technologies market report, Department of Energy, DOE/GO-102011-3322.

3.7 Appendix

Table 3.2 provides details on the ERCOT and MISO data used in this study. Figure 3.14 shows the frequency distributions of the estimated uncurtailed wind power in ERCOT, actual wind power in ERCOT and actual wind power in MISO over the respective time periods below.

Table 3.2: Summary statistics for the ERCOT and MISO wind data used in this study.

	ERCOT	MISO
Time Period	Dec 2008 – Dec 2010	Feb 2011 – May 2012
Average Load	35 GW	61 GW
Maximum Load	66 GW	104 GW
Installed Wind Capacity	8.325 – 9.53 GW	9.125 – 10.79 GW
Actual Wind Capacity Factor	0.28	0.33
Uncurtailed Wind Capacity Factor	0.35	N/A
Ratio of actual wind energy to load	0.069	0.055
Ratio of uncurtailed wind energy to load	0.087	N/A

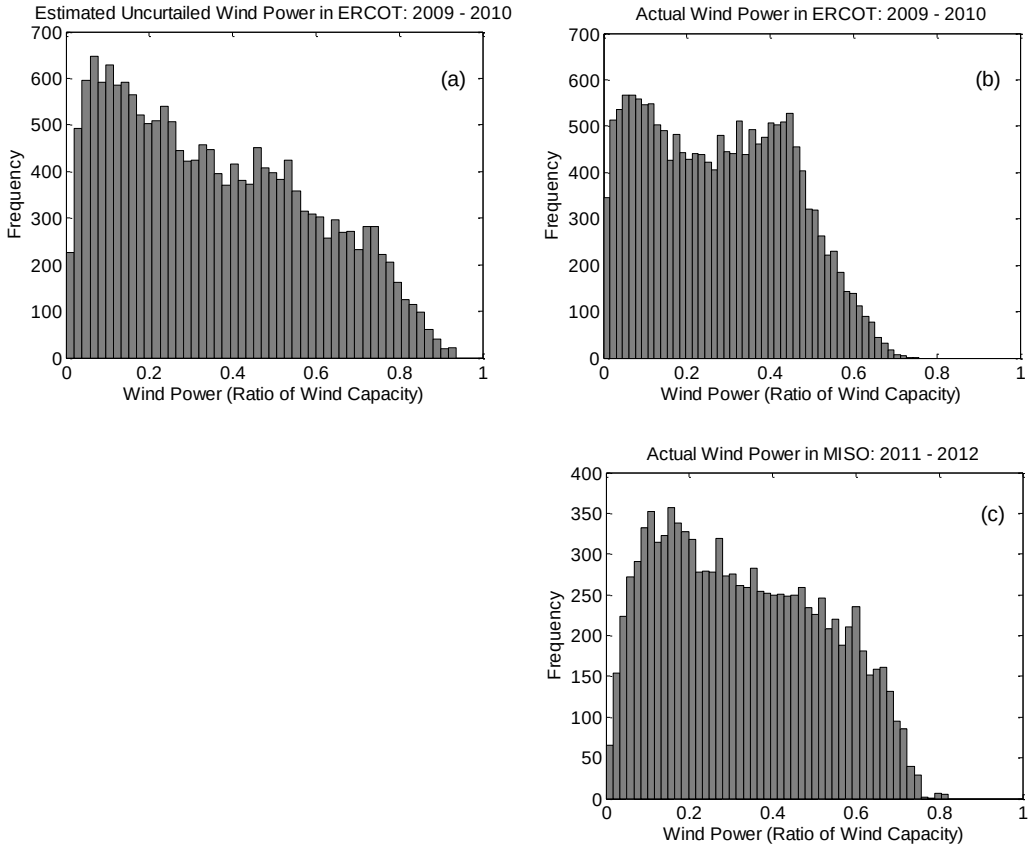

Figure 3.14: Frequency distributions of (a) uncurtailed hourly wind power in ERCOT (b) actual hourly wind power in ERCOT and (c) actual hourly wind power in MISO.

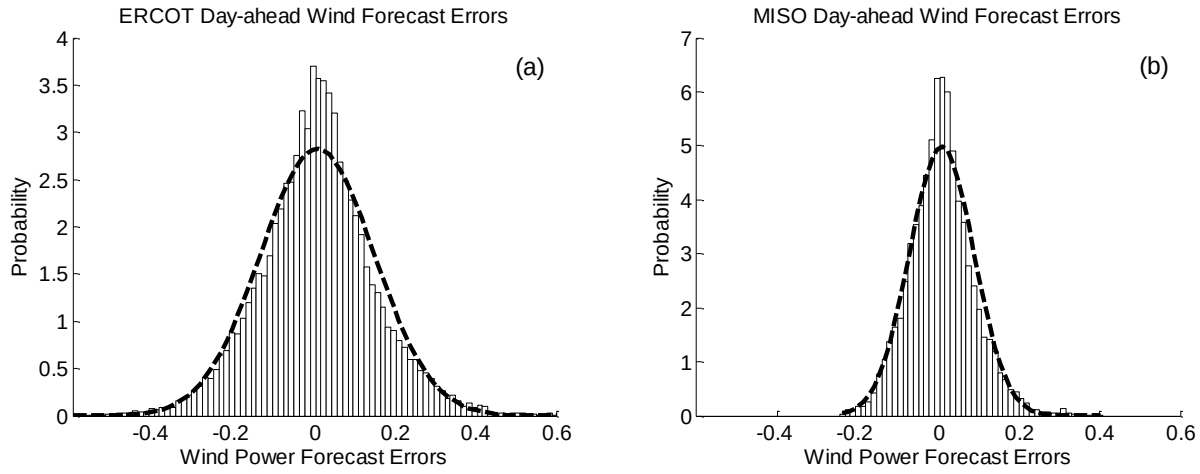

Figure 3.15: Frequency distribution (histogram bars) of wind power forecast errors in ERCOT (a) and MISO (b) with a fitted normal distribution (dashed line) overlaid.

Figure 3.15 shows the distribution of wind forecast errors in ERCOT and MISO for the respective time periods in our data. The figures above indicate that the forecasts are unbiased on the whole for both regions. As is well-known, normal distributions are not well suited for modeling wind forecast errors (Hodge and Milligan, 2011). MISO day-ahead wind forecasts tend to be more accurate than ERCOT (although the wind capacity is similar in both regions). This may be due to the much larger territory covered by MISO. Large geographic diversity tends to reduce total forecast errors if forecast errors from sub-regions are weakly correlated (Focken et al., 2002; Giebel et al., 2007).

3.8.1 Logit-normal distribution for wind data

Previously, we showed that the logit transformations of the wind power forecast and actual wind power data fit a normal distribution reasonably well. If the logit transformation of a variable is normally distributed, $N(\mu, \sigma^2)$, then the variable itself is distributed logit-normal, $LN(\mu, \sigma^2)$. For the wind forecast (F) and wind power (W), the logit transformations are

$$F^* = \ln\left(\frac{F}{1-F}\right) \qquad W^* = \ln\left(\frac{W}{1-W}\right) \tag{3.13}$$

Each of the transformed variables is normally distributed with a probability density

$$f(X^*) = \frac{1}{\sqrt{2\pi}\sigma_{X^*}} exp\left(-\frac{(X^*-\mu_{X^*})^2}{2\sigma_{X^*}^2}\right) \tag{3.14}$$

The symbols μ_{x^*} and σ_{x^*} represent the mean and standard deviation of the transformed variables. The original variables are logit-normal distributed with a probability density

$$f(X) = \frac{1}{\sqrt{2\pi}\sigma_X X(1-X)} exp\left(-\frac{(X^*-\mu_X)^2}{2\sigma_X^2}\right) \tag{3.15}$$

The symbols have the same definitions as in Equation 3.14. Given that the logit transformed variables are jointly normal, the variables *F* and *W* are jointly logit-normal distributed with probability density defined as

$$f(F,W) = \frac{1}{2\pi\sigma_F\sigma_W\sqrt{1-\rho^2}F(1-F)W(1-W)} e^{\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(F^*-\mu_F)^2}{\sigma_F^2}+\frac{(W^*-\mu_W)^2}{\sigma_W^2}-\frac{2\rho(F^*-\mu_F)(W^*-\mu_W)}{\sigma_F\sigma_W}\right]\right)} \tag{3.16}$$

Figure 3.16 shows the distribution of *F* and *W* from the ERCOT data with a fitted logit-normal distribution overlaid.

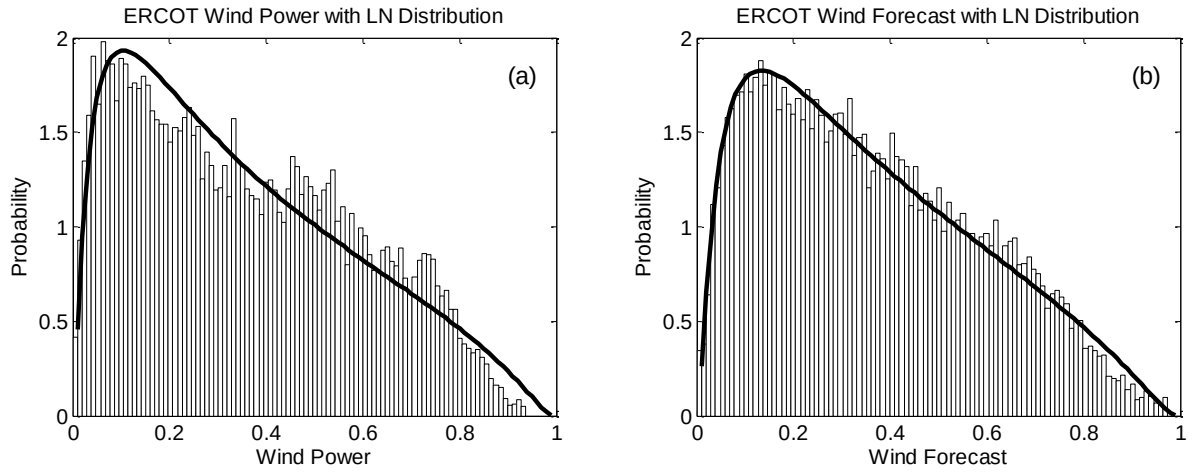

Figure 3.16: Distributions of wind power and wind power forecasts from ERCOT with logit-normal distributions overlaid.

3.7.2 Conditional wind forecast bias

As shown in section 3.4, wind power forecasts exhibit a bias that varies thorough the range of forecast values. At low wind power forecasts the bias is negative while high wind power forecasts have a positive bias. We calculated the mean value of the forecast error conditioned on the forecast value for the ERCOT day-ahead wind power forecasts in the following manner with E[·] denoting the expectation operation.

$$\mathrm{E}[e|F] = \mathrm{E}[(F - W)|F] = F - \mathrm{E}[W|F] \tag{3.17}$$

For a given forecast, the bias is the forecast minus the expectation of W conditioned on the forecast value. The expectation is determined with the logit-normal distribution as

$$\mathrm{E}[W|F] = \int_0^1 \frac{1}{\sqrt{2\pi}\sigma_{W|F}W(1-W)} exp\left(-\frac{(W^*-\mu_{W|F})^2}{2\sigma_{W|F}^2}\right) W dW \tag{3.18}$$

Since there are no closed form solutions for the moments of the logit-normal distribution the expected value must be obtained with numerical integration. We used the quad function in Matlab, which implements Simpson's rule, to evaluate the integral in Equation 3.18 over the range of possible values for *F*. Figure 3.17 shows the expected values *W* conditioned on *F* and the expected error conditioned on *F*. In each plot the solid lines were calculated with the logit-normal model, and the dashed lines were calculated using a moving window of width 0.1.

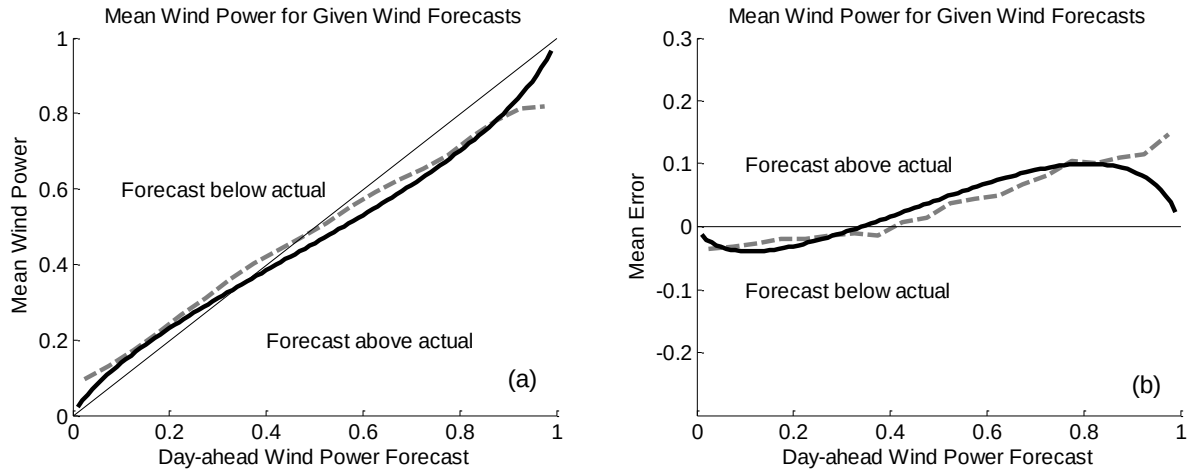

Figure 3.17: Mean wind power conditioned on the wind power forecast (a) and mean forecast error conditioned on the wind power forecast (b). Solid lines were calculated directly with binned data, and dashed lines were calculated using the logit-normal distribution.

In order to examine the effect of a forecast algorithm that accounts for the observed forecast bias, we subtracted the bias values calculated with the logit-normal model in Figure 3.17 (b) from the ERCOT day-ahead forecasts to remove the conditional bias in the data. Figure 3.18 shows the day-ahead wind power forecast distributions in ERCOT with the original data and the unbiased forecasts. The plot of original error distribution is the same as in Figure 3.2 in the main text. Errors from the original data have a mean value of 0.007 and variance of 0.02. Errors calculated from the unbiased forecasts have a mean value of -0.005 and variance of 0.019.

Removing the conditional bias in the forecasts mean that low forecasts are increased and high forecasts are decreased. As previously shown in Figure 3.16 (b), most forecasts are near 0.2 as a ratio of the wind capacity. Therefore, most errors resulting from the unbiased forecasts are slightly increased compared with the errors from the biased forecasts. The resulting unbiased forecast error distribution in Figure 3.18 (b) is skewed, but retains an overall mean near zero and a variance comparable to the biased errors.

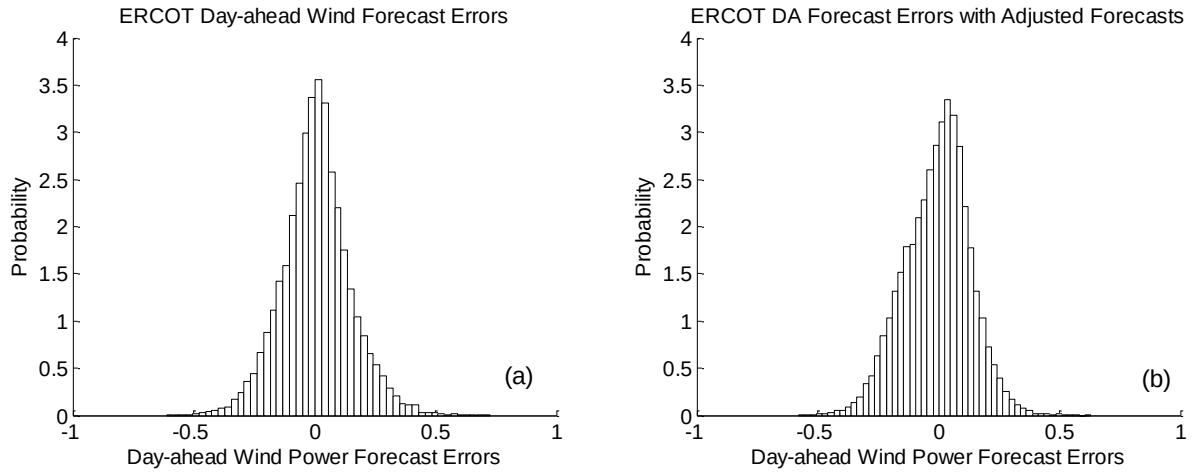

Figure 3.18: Wind forecast error distributions from ERCOT with original data (a) and with the conditional bias removed from the forecasts.

Figure 3.19 shows the forecast errors plotted against the forecast values with the original ERCOT data and with the data adjusted to remove the bias. In the adjusted case the upward trend in the forecast errors is reduced.

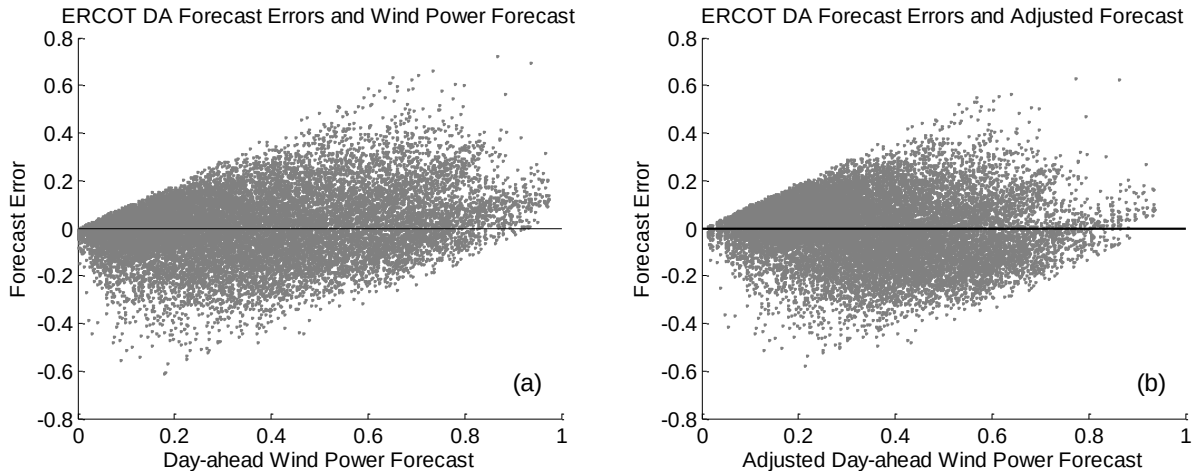

Figure 3.19: Wind forecast error distributions from ERCOT with original data (a) and with the conditional bias removed from the forecasts.

As evident in Figure 3.17, the data do not fit a logit-normal distribution perfectly. Since the bias was calculated from the fitted logit-normal model, it was also not calculated perfectly. Figure 3.20 shows the conditional bias remaining in the forecast errors after adjusting the forecast data to remove the modeled bias. Residual bias was calculated with a moving window of width 0.1. When compared to the bias calculated for forecast bins in Figure 3.17 (b), one can see that the conditional bias is mostly eliminated.

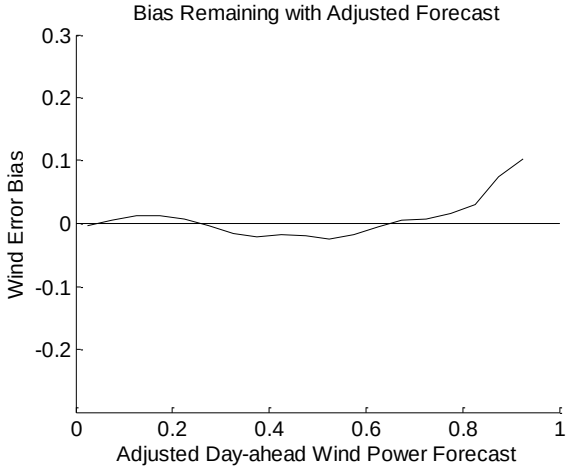

Figure 3.20: Bias calculated with binned forecast data after removing the conditional bias calculated with the logit-normal distribution; compare with Figure 3.17 (b).

Chapter 4 - What day-ahead reserves are needed in electric grids with high levels of wind power?

4.1 Abstract

Wind power forecasts provide useful information for operational decision making, but they are imperfect and the level of uncertainty must be known for proper use. In order to ensure grid reliability, dispatchable generation must offset errors in the load and wind power forecasts. We have previously shown the dependence of wind forecast uncertainty on the wind forecast level. Here we calculate increases in net load uncertainty introduced from day-ahead wind power forecasts. Our analysis uses data from two different electric grids in the U.S. with similar levels of installed wind capacity and large differences in wind and load forecast accuracy due to geographic characteristics. We present a probabilistic method to calculate the dispatchable generation capacity required to balance day-ahead wind and load forecast errors for a given level of reliability. Using empirical data, we show that the capacity requirements for 95% day-ahead reliability range from 2100 MW to 5700 MW for ERCOT and 1900 MW to 4500 MW for MISO, depending on the amount of wind and load forecast for the next day. We briefly discuss the additional requirements for higher reliability levels and the effect of correlated wind and load forecast errors. Additionally, we show that each MW of additional wind power capacity in ERCOT must be matched with 0.30 MW day-ahead dispatchable generation capacity. For MISO, the value drops to 0.13 MW of dispatchable capacity for each MW of additional wind capacity.

4.2 Introduction

Wind power has experienced substantial growth over the past decade in the U.S. and in Europe. Installed capacity in the U.S. increased more than tenfold from 4.2 GW in 2001 to 47 GW in 2011. Electricity generated from wind in the U.S. accounted for under 0.2% of the total electricity production in 2001, and now provides nearly 3% of total electrical energy (DOE, 2012). European installed wind capacity increased fivefold in the same time period, from 17.3 GW to 94 GW, and produces over 6% of the total electrical energy produced (Wilkes et al., 2012).

Current levels of wind capacity do not seem to pose much of a problem for grid operators, but future growth will require operational changes in order to ensure a reliable grid (Xie et al, 2011). Since wind power is not dispatchable, wind forecasts are used along with load forecasts to make operational decisions ranging from minutes to days in advance (Botterud et al., 2010). Forecasts are not perfect and, therefore, back-up generation capacity is needed to ensure reliability in the operation of the grid. Increases in wind power capacity raise the level of uncertainty in operational decisions requiring additional back-up generation.

Here, we present a model to combine the uncertainty of load and wind forecasts, and show how this uncertainty changes for different levels of wind forecasts. Since the majority of operational decisions are made a day in advance, we focus our analysis on the uncertainty associated with day-ahead forecasts. Using data from two different U.S. independent system operators (ISOs), we show that increases in net load day-ahead uncertainty is on the order of 15%-30% of the installed wind forecast capacity. In other words, for every MW of expected wind power, an ISO should have 0.13- 0.3 MW of additional dispatchable capacity to cover the uncertainty. The exact amount depends on several factors including forecast accuracy, load level and desired reliability level.

4.2.1 Wind power integration in electric grids

When wind power accounts for a small fraction of the energy generation in electric grids, the uncertainty added to grid operations is insignificant compared to the uncertainty in load levels. As a result, generation from wind has been taken on an as-available basis. In recent years, however, wind power has reached levels in some parts of the U.S. and Europe where the uncertainty of wind affects the efficiency and reliability of grid operations (Xie et al, 2011). With policies in place to encourage electricity generation from renewable resources, such as wind, this trend is expected to continue (EIA, 2012). For this reason, many ISOs around the world implemented centralized wind forecasting systems for use as a decision support tool (DOE, 2011; Botterud et al., 2010). In grids where forecasts are used for unit commitment and reliability analysis, the accuracy of wind power forecasts is a primary concern (DOE, 2011). Forecasts have improved in recent years, but uncertainty inherent in wind forecasts must be quantified in order to ensure that adequate generation is available to meet demand at all times.

The term "net load" describes the amount of load remaining after wind power is subtracted. Conventional generation is scheduled according to forecasts for net load. Net load forecast errors are compensated with adjustments in generator output to match observed conditions. In order to ensure that a given level (e.g. 99%) of errors is covered, a grid must have generation capacity available beyond the amount expected to be dispatched. This extra capacity is procured in multiple markets; balancing energy (real-time market), 10 minute reserves, 30 minute reserves and regulation.

In recent years, several wind integration studies have been performed to determine operational effects of high levels of wind energy (Enernex, 2006; GE, 2008; GE, 2010). In all cases, they used simulated data of wind power and wind power forecasts to model electric grid operations with specified levels of wind. Results from the studies indicate that load and wind

forecast errors display little to no systematic correlation; meaning the probability of large wind and load forecasts occurring simultaneously is quite low. They also found that operating reserves would need to be increased in order to cope with wind forecast uncertainty once the wind reached levels significant levels.

Most peer-reviewed research addresses the role of wind uncertainty in unit commitment algorithms (Wang et al., 2008) and operating reserve procurements (Doherty and O'Malley, 2005). Other work considers the tradeoffs between grid reliability and the social cost of achieving levels of reliability in electric grids with significant levels of wind power (Ortega-Vazquez and Kirschen, 2009; Matos and Bessa, 2009). Finally, the role of wind forecast uncertainty has been considered in the day-ahead energy market clearing process (Bouffard and Galiana, 2008). In most work analyzing the effect of wind uncertainty on grid operations, load and wind forecast errors are modeled as normally distributed random variables. Applying the central limit theorem to the aggregation of many uncorrelated or weakly correlated wind farms seems to justify this assumption, but analysis of real data indicate that wind forecast errors cannot be adequately modeled using normal distributions (Lange, 2005; Bludszuweit, 2008). Moreover, data show that load forecast errors are not normally distributed (Hodge et al. 2012). Realistic wind forecast error distributions are skewed towards one direction, depending on the amount of wind forecast. Better representation of wind forecast uncertainty seems to have a significant effect on unit commitment results (Lowery and O'Malley, 2011).

Here we focus on net load forecast uncertainty as a function of the amount of wind forecast, as it has been shown that these two values are highly related (Lange, 2005; Bludszuweit, 2008). At wind levels where net load uncertainty is dominated by wind forecast uncertainty, operational decisions should be based on a systematic assessment of net load uncertainty reflecting changes in wind forecast levels. Using data from two different ISOs, we show that both wind and load

forecast errors are not normally distributed random variables, and we also show the influence of wind forecast levels on net load uncertainty. We then scale-up the wind data by multiplying all values by a constant to simulate a future grid with much higher levels of wind capacity and determine how much dispatchable capacity must be retained to balance the wind uncertainty. Accurate quantification of net load uncertainty is essential for optimal grid management. It is also important in long term grid planning to ensure that enough reserve capacity is available to compensate for increased in uncertainty due to wind capacity growth.

4.3 Wind and Load Data

We used data on wind and load forecasts along with actual values of wind and load from Electric Reliability Council of Texas (ERCOT) and the Midwest Independent System Operator (MISO). The ERCOT data consist of hourly wind forecast values for 1 to 48 hour look-ahead times covering the years 2009 and 2010. Actual wind generation values along with estimated uncurtailed wind generation potential were provided in the data. In order to analyze wind power uncertainty, we used estimated values of uncurtailed wind power to remove the effect of wind curtailments which are not considered in wind power forecasts. AWS Truepower estimated hourly uncurtailed wind power values using measurements taken from wind farms in ERCOT and curtailment instructions issued by ERCOT. This is important considering that ERCOT curtailed an estimated 17% of wind generation in 2009 and 10% of wind generation in 2010 (Rogers et al., 2010). Day-ahead load forecasts were obtained from the Ventyx database.

MISO wind forecast data came from their website, which provides day-ahead wind forecasts and actual wind generation. Our data span the time period of February 2011 to May 2012. Day-ahead load forecasts and actual loads were also obtained from the Ventyx database. Wind

curtailments in MISO were estimated to range from 2% to 6% each month (MISO, 2012). We neglect these curtailments in our analysis.

Table 4.1 shows summary statistics from ERCOT and MISO during our data periods. ERCOT's territory covers most of the state of Texas, while MISO covers most of the Midwestern portion of the U.S. Summary information on day-ahead wind and load forecast accuracy in both grids is also included

Table 4.1: Summary of wind and load in ERCOT and MISO during the time periods of the data we used in this study.

	ERCOT	MISO
Date Range	Jan 2009 – Dec 2010	Feb 2011 – May 2012
Average Load	36 GW	61 GW
Maximum Load	66 GW	104 GW
Average Nameplate Wind Capacity	8.8 GW	10 GW
Average Wind Capacity Factor	0.35	0.33
Percentage of load served by wind	8.7%	5.5%
Mean absolute wind forecast error (percentage of installed wind capacity)	11%	6.1%
Root mean squared wind forecast error (percentage of installed wind capacity)	14%	8.1%
Mean absolute forecast error (percentage of mean load)	3.1%	1.7%
Root mean squared load forecast error (percentage of mean load)	4.2%	2.2%

As shown in Table 4.1, aggregate wind power in MISO has a much lower forecast error than does ERCOT. During the time periods covered in our dataset, both ERCOT and MISO had similar levels of installed wind capacity (~10 GW). However, the wind farms in MISO's territory span a much larger area than ERCOT covers (Figure 4.1). Large wind power footprints reduce the correlation of wind generation between wind farms, and lower variability in aggregate wind generation (Katzenstein et al., 2010). The accuracy of wind forecasts at the ISO level improves as the variability decreases (Focken et al., 2002).

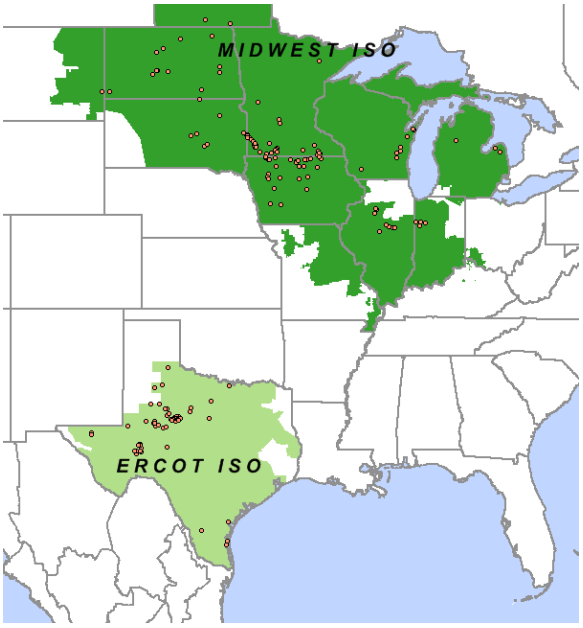

Figure 4.1: Map of wind farms in ERCOT and MISO. Data source: Ventyx.

4.3.1 Day-ahead load uncertainty

Load forecasts have been used for planning and operations by electric grids for many years. Due to the repeatability of daily patterns in electricity demand, load forecasts tend to be more accurate than wind power forecasts (Table 1). Electric grids compute load forecasts at many intervals for daily operations. Here we are concerned with the forecasts made one day in advance, which are used to make decisions in unit commitment and operating reserve levels. Figure 4.2 shows actual hourly load versus forecasted load in ERCOT for one week in January. The bottom line in Figure 4.2 illustrates the forecast error. We define the absolute forecast error as the forecast value minus the actual value.

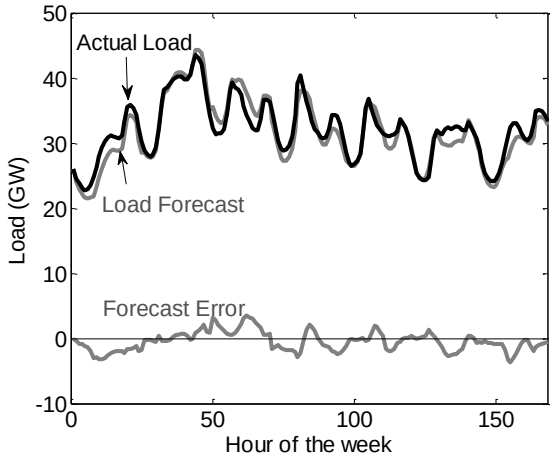

Figure 4.2: Forecasted load, actual load and load forecast errors in ERCOT from the week of January 4 – 10, 2009.

We used the @Risk software package to fit parametric distributions to load forecast errors in ERCOT and MISO. In both sets of load forecast errors the logistic distribution provided a much better fit to the normal distribution as measured by likelihood values (Figure 4.3). While no distribution passed a Pearson's chi-squared test with the data, we found the logistic distribution to be adequate for our analysis. Similar to the work by Hodge et al. (2012), we found that the actual load forecast error distribution has more mass around zero than what a normal probability density function (PDF) would predict. Therefore, we chose the logistic distribution, as it seemed to fit the data better. The actual distribution used to model load forecast uncertainty is not crucial as long as it accurately represents the data. A logistic distribution is defined by a location and scale parameter (α and β) with the PDF given in Equation 4.1. The mean is equal to α and the variance is equal to $(\pi^2/3)\ \beta^2$.

$$f_x(x) = \frac{e^{-(x-\alpha)/\beta}}{\beta\left(1+e^{-(x-\alpha)/\beta}\right)^2} \tag{4.1}$$

Figure 4.3 shows histograms of the load forecast errors with fitted logistic and normal distributions for ERCOT and MISO. Load forecast errors are shown plotted against load forecast levels in Figure 4.4. In our analysis we divided the load forecast data into three forecast bins: low load (forecasts less than 90% of mean load), medium load (between 90% and 120% of mean load) and high load (forecasts greater than 120% of mean load). We fit a separate logistic distribution to each load forecast bin.

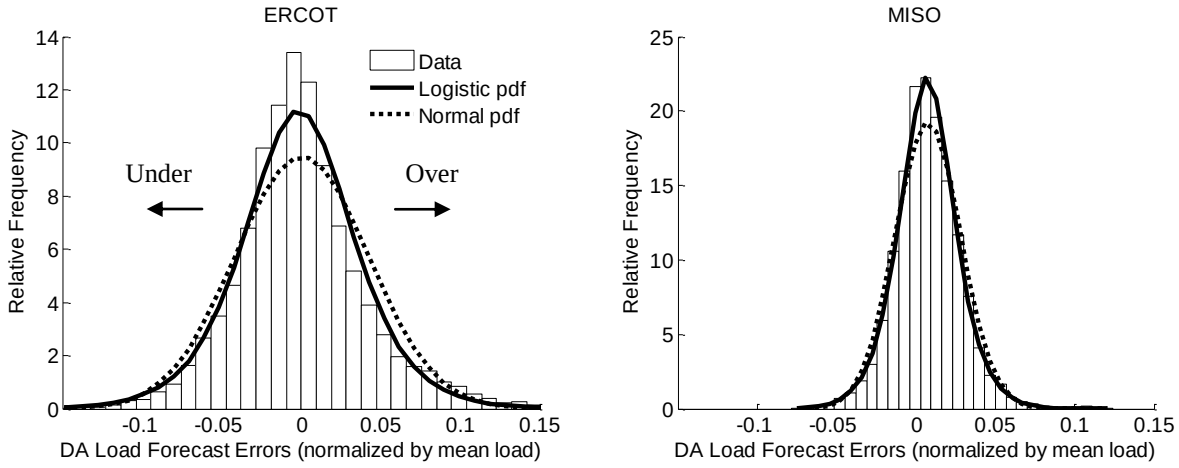

Figure 4.3: Load forecast error histograms for ERCOT and MISO using data from the time periods given in Table 1. Included in the figures are logistic (solid line) and normal (dotted line) density functions fit to the data. All values are normalized by the mean load for each region. Logistic distribution parameters are α = -0.0017 and β = 0.022 for ERCOT and α = 0.0070 and β = 0.0112 for MISO. Normal distribution parameters are μ = 0 and σ = 0.0421 for ERCOT and μ = 0.0075 and σ = 0.0207 for MISO.

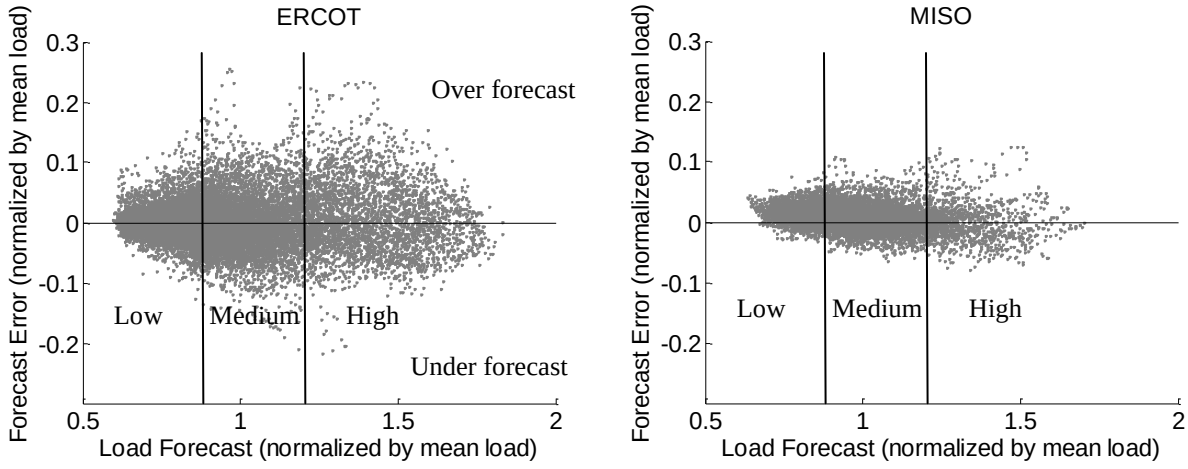

Figure 4.4: Day-ahead load forecast errors plotted against the load forecast values for ERCOT and MISO.

4.3.2 Day-ahead wind uncertainty

In comparison to day-ahead load forecasts, day-ahead wind forecasts have much greater relative uncertainty. Figure 4.5 shows a histogram of wind forecast errors determined with hourly day-ahead forecast data from ERCOT and MISO. All values are shown as a fraction of installed wind capacity. A normal probability density function fit to the error values is included in the plots. As before, the forecast errors are defined as the forecast value minus the observed value. In both datasets, a normal distribution under-predicts the probability of wind forecast errors near zero.

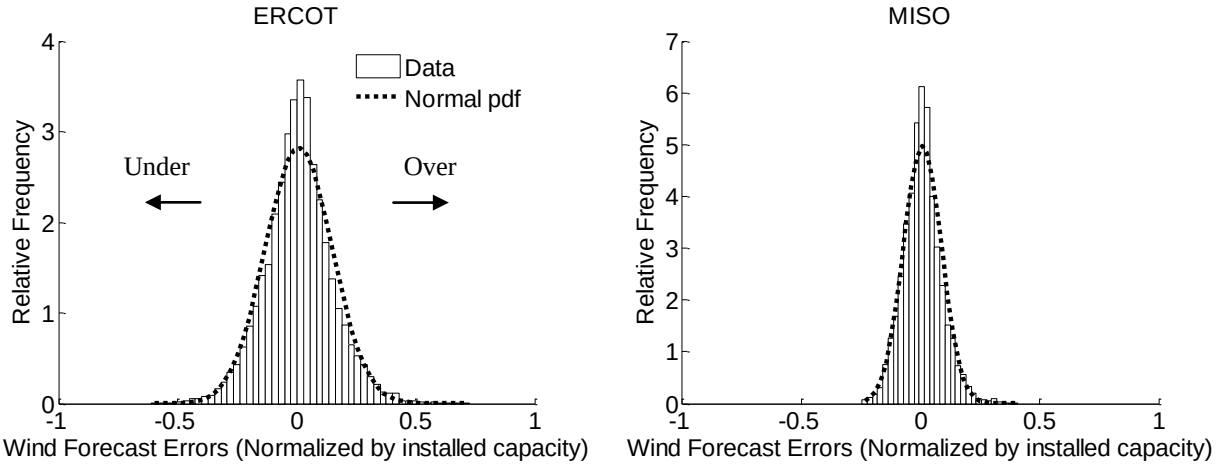

Figure 4.5: ERCOT and MISO day-ahead wind forecast errors for the time periods in Table 1. ERCOT errors are calculated using estimated uncurtailed wind power data while MISO errors use actual wind power data. Histogram bars represent the forecast errors and the dotted line shows the fitted normal density function.

Analyzing wind forecast errors from the perspective of Figure 4.5 misses a key element of wind forecast uncertainty, i.e. that the distribution of wind forecast errors is dependent on the wind forecast level. As in many bounded distributions, this behavior seems perfectly reasonable when one considers that for wind forecasts near zero, the observed wind values are likely to be greater than the forecast values. The converse is true for wind forecast levels near the maximum

value. Figure 4.6 shows the wind forecast errors plotted against the wind forecast values. Here we see that there is indeed a strong dependence of forecast errors on the forecast level.

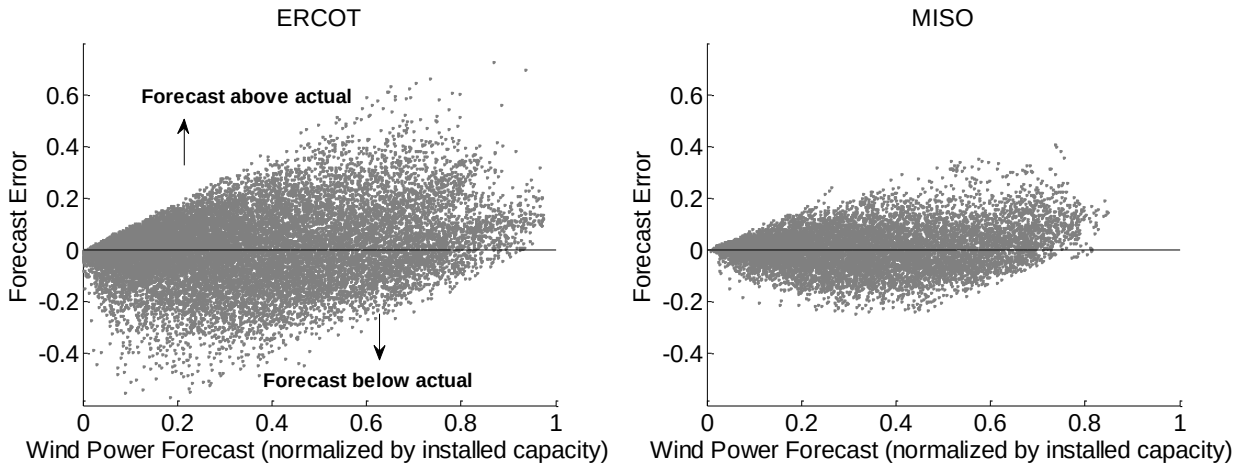

Figure 4.6: Day-ahead wind forecast errors (vertical axis) plotted against the day-ahead wind forecast values (horizontal axis) in ERCOT and MISO. All values are expressed as a ratio of installed wind capacity.

In order to account for the dependence of wind forecast error distributions on the wind forecast level, we used the logit-normal model to represent wind forecast errors as described in Chapter 3. Using this method, we fit normal distributions to logit transformations of hourly, day-ahead wind forecasts (F) and observed wind power (W) values as per Equation 4.2. Both variables are normalized by the installed wind capacity in order to ensure that F and W have ranges from 0 to 1. The resulting transformed variables range from negative to positive infinity

$$F^* = \ln\left(\frac{F}{1-F}\right) \tag{4.2}$$

$$W^* = \ln\left(\frac{W}{1-W}\right).$$

Once the transformed variables are fit to a normal distribution, confidence intervals are straightforward to calculate and transform back into the original variable space. Different confidence levels are displayed as a function of the wind forecast level in Figure 4.7 using logit-normal parameters derived from fitting the data from Figure 4.6.

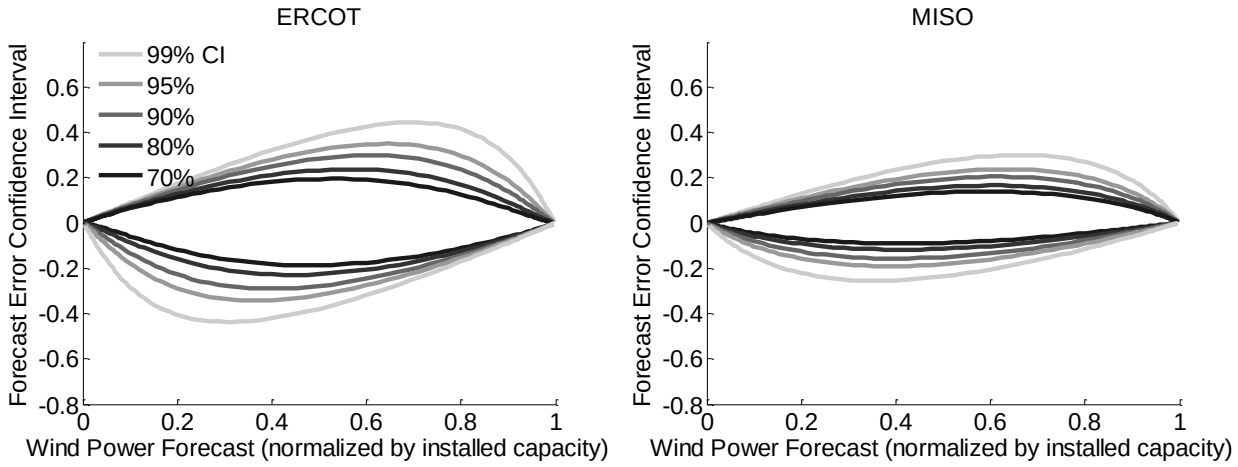

Figure 4.7: Confidence levels for the data in Figure 5 calculated with the logit-normal model.

If the logit transformation of a variable is normally distributed, $N(m, s^2)$, then the variable itself is distributed logit-normal, $LN(m, s^2)$. In this model the wind forecast values and observed wind values are assumed to be jointly distributed logit-normal with the joint probability density function

$$f_{F,W}(F,W) = \left(\frac{1}{2\pi\sigma_F\sigma_W\sqrt{1-\rho^2}F(1-F)W(1-W)}\right) e^{\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(F^*-\mu_F)^2}{\sigma_F^2}+\frac{(W^*-\mu_W)^2}{\sigma_W^2}-\frac{2\rho(F^*-\mu_F)(W^*-\mu_W)}{\sigma_F\sigma_W}\right]\right)}. \quad (4.3)$$

The conditional distribution of observed wind power given a wind forecast level (W|F) then becomes

$$f_{W|F}(W|F) = \left(\frac{1}{\sqrt{2\pi}\sigma_{W|F}W(1-W)}\right) e^{\left(-\frac{(W^*-\mu_{W|F})^2}{2\sigma_{W|F}^2}\right)}. \quad (4.4)$$

Distributions of observed wind power values are displayed for three different wind forecast levels in Figure 4.8. As expected from the data displayed in Figure 4.6, the distribution of the observed wind variable is skewed to the right at low forecast levels and skewed left when the forecast level is close to the installed capacity. Also, we made no attempt to remove the forecast bias which is apparent in the figure. See the Chapter 3 appendix for more discussion on the forecast bias.

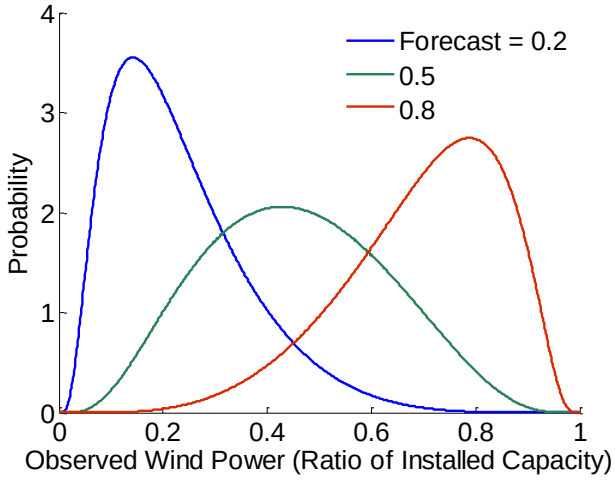

Figure 4.8: Observed wind distributions for 3 different wind forecast levels (as a fraction of installed capacity) derived from the logit-normal model fit to ERCOT day-ahead wind data for the years 2009 and 2010.

Since wind forecast errors are defined as the forecast level minus the observed wind value, forecast error distribution are constructed by shifting and translating the PDFs of the observed wind values. The resulting distributions for wind errors conditioned on the same three forecast levels are shown in Figure 4.9.

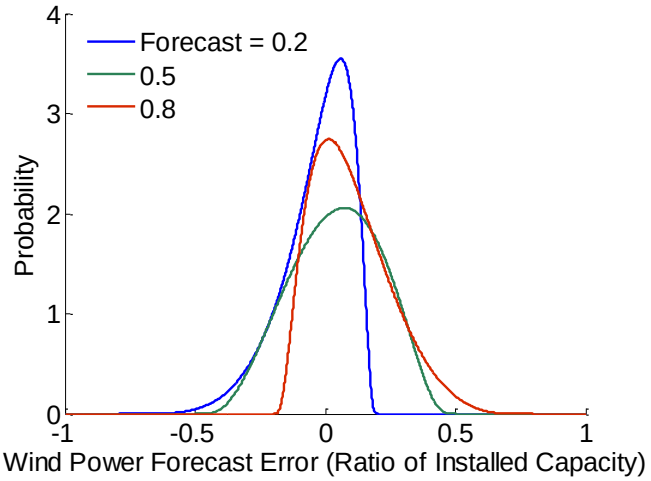

Figure 4.9: Wind forecast error distributions constructed with the logit-normal model fit to the ERCOT wind data for the years 2009 and 2010. Three different forecast levels are shown. All data are normalized by the installed wind capacity.

4.3.3 Net Load Uncertainty

Net load is defined as load (L) minus wind (W) power

$$N = L - W. \tag{4.6}$$

The actual (subscript: A) levels of load and wind subtracted from the forecasted values (subscript: F) give the net load forecast errors

$$e_N = (L_F - W_F) - (L_A - W_A). \tag{4.7}$$

Rearranging terms gives

$$e_N = (L_F - L_A) - (W_F - W_A), \tag{4.8}$$

$$e_N = e_L - e_W. \tag{4.9}$$

Load forecast errors and wind forecast errors are commonly assumed to be independent variables (Matos and Bessa, 2009; Ortega-Vazquez and Kirschen, 2009. The correlation coefficient of wind and load forecast errors in our data was 0.09 for ERCOT and 0.05 for MISO (both significant at the 95% level). These values are small, and we will assume no correlation between wind and load forecast errors as a base case. We briefly show the effect of correlated load and wind forecast errors on net load uncertainty. The appendix contains more discussion on the correlation between wind and load forecast errors. Since the actual correlation between wind and load errors is slightly positive, our assumption of no correlation will lead to slightly higher estimates of uncertainty for net load than actual. In this case, the probability density function of the net load errors is derived by taking the convolution[3] of the load errors and the negative of the wind errors. The integral was solved numerically.

$$f_{e_N}(e_N) = \int_{-\infty}^{\infty} f_{e_L}(x) f_{e_W}(e_N - (-x))dx \quad (4.10)$$

The cumulative distribution function (CDF) is then the integral of the PDF, which allows one to determine confidence intervals for a desired reliability level.

$$F_{e_N}(e_N) = \int_{-\infty}^{e_N} f_{e_N}(z)dz \quad (4.11)$$

Figure 4.10 shows CDFs for the load and net load errors calculated using data from ERCOT for a medium load forecast (90% to 120% of the mean load) and wind forecast level of 5000 MW with 10,000 MW of installed wind capacity. The dashed lines show the 2.5th and 97.5th

[3] Technically, we are calculating the cross-correlation since this is a difference of two random variables.

percentiles. As expected, the addition of 5 GW of wind energy in the system moves the endpoints of the 95% confidence interval away from zero. In this example, the net load uncertainty at the 95% confidence level ranges from -4650 MW to 3850 MW. In order to cover this level of uncertainty, a grid must have 4650 MW of generation capacity available beyond the net load forecast in the case of under forecasts and 3850 MW of down generation capacity in the case of over forecasts. Since under forecasts are more severe from a grid reliability perspective, we focus on the portion of the net load error confidence interval that occur from under forecasting.

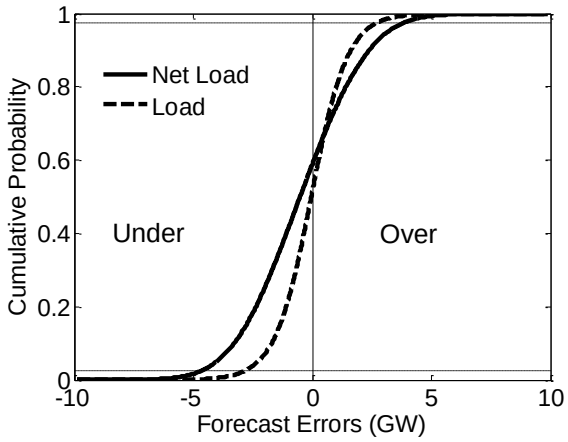

Figure 4.10: Cumulative distribution functions of net load errors (solid line) and load errors (dashed line) produced with the statistical model of net load uncertainty fit to data from ERCOT. The two dashed horizontal lines represent the 2.5th and 97.5th percentiles spanning 95% of the forecast errors.

Considering only the errors that results from an under forecast of net load, we can look at the probability of the forecast error being less than some given value conditioned on the event that the forecast error is negative.

$$P(e_N \leq z | e_N \leq 0) = \frac{F_{e_N}(e_N)}{F_{e_N}(e_N=0)} \tag{4.12}$$

Using Equation 4.12 we calculate the value for *z* that gives a probability of 0.05 to get the generation capacity required to cover 95% of potential under forecast errors.

4.4 Results

At the time the data were collected, ERCOT installed wind capacity averaged 8.8 GW and MISO wind capacity averaged 10 GW. We scaled up the ERCOT wind data to simulate an installed wind capacity of 10 GW in both grids and compared dispatchable generation capacity requirements due to wind uncertainty for both grids. This provides insight into reserve requirements for two grids with large differences in wind power forecast accuracy.

We assume a system operator procures reserves in the form of dispatchable capacity during day-ahead operation planning. The amount of reserves procured at this stage will be complemented with energy traded on the spot market during the dispatch day. For this reason, it is not necessary to procure enough reserves to cover all of the forecast uncertainty a day in advance.

Figure 4.11 shows the amount of generation capacity required to cover 95% of the day-ahead net load under-forecast errors for a range of wind power forecasts. We choose a 95% uncertainty level since that is the stated ERCOT metric in operational reserve procurements (ERCOT, 2012). Each plot displays separate curves for three load forecast levels. In ERCOT, the three curves are evenly spaced as load uncertainty increased gradually with the load forecast value. The MISO load uncertainty increased more abruptly at high load forecasts. Each curve shows the dispatchable capacity requirement as a function of the wind forecast. As the wind forecast increases, uncertainty also increases, resulting in greater dispatchable generation requirements. Each grid showed a particular wind forecast value where uncertainty reached a maximum and then decreased. This occurs because the wind forecast error confidence interval converges to

zero at the maximum wind forecast level (Figure 4.6). Uncertainty in ERCOT peaked at a wind forecast of 7500 MW and in MISO this occurred at 6000 MW. Back-up requirements for load with no wind are shown at the y-intercepts in the graphs.

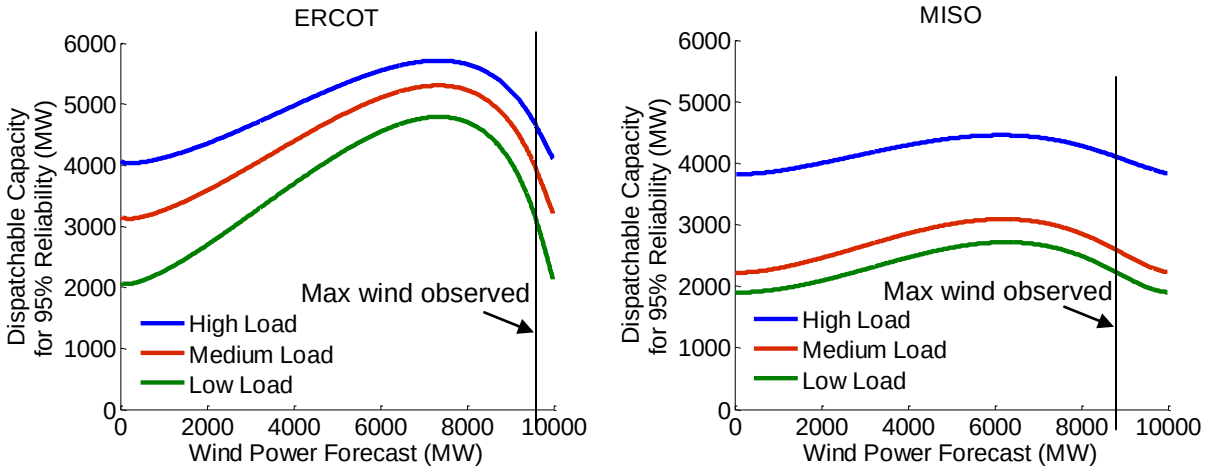

Figure 4.11: Amount of dispatchable generation capacity required to cover 95% of net load under forecast errors as a function of wind forecast level for 10 GW of installed wind capacity. Back-up requirements peak in ERCOT at a wind forecast of 7500 MW and in MISO at 6000 MW.

Our treatment of load forecast uncertainty in Figure 4.11 indicates that ERCOT requires between 2100 and 4100 MW of dispatchable generation to cover load forecast uncertainty (y-intercepts of curves). In MISO the range is between 1900 and 3800 MW. When the uncertainty of wind forecasts is included, the peak levels of dispatchable generation reach 5700 MW for ERCOT and 4500 MW for MISO for 10 GW of wind capacity. If operational reserves levels are not based on forecast values, there will be hours when reserves are over procured and hours when they are under procured.

Figure 4.11 is based on a day-ahead reliability level of 95%. In the following plots we show the amount of dispatchable reserves required for a high load forecast in each grid to cover 95 to 99% of day-ahead net load errors. In each grid, 2000 MW of additional capacity is required to move from covering 95% of net load errors to 99% of net load errors.

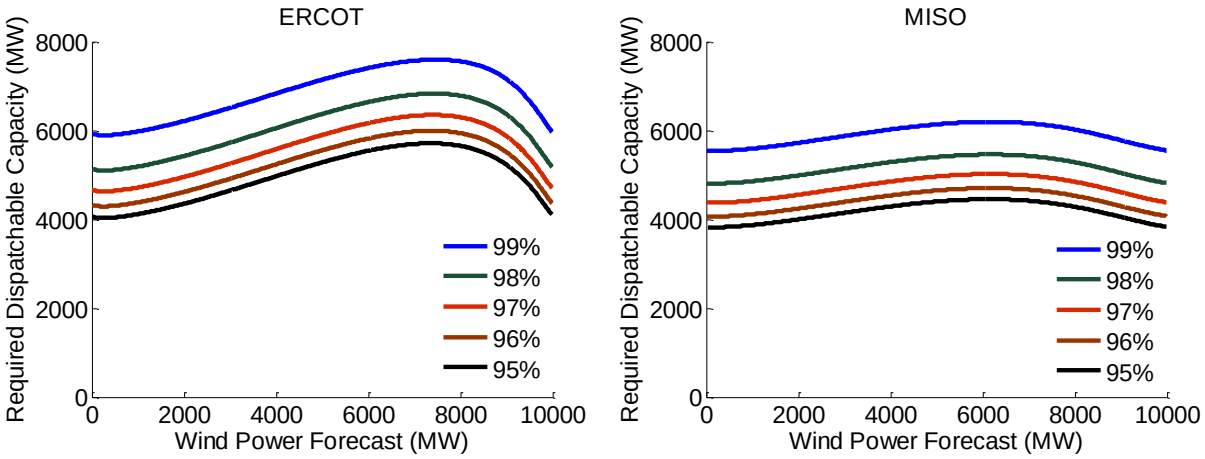

Figure 4.12: Dispatchable generation capacity required to cover 95% to 99% of net load forecast errors in ERCOT and MISO for a high load forecast with 10 GW of installed wind capacity.

The difference between generation capacity requirements with and without wind show the amount of additional dispatchable generation needed to cover wind uncertainty. In ERCOT, the additional requirement reaches between 1500 and 2600 MW for 10 GW of installed wind capacity depending on the hourly load forecast; MISO requires a maximum of 900 MW (Figure 4.13). The large difference in the amount of additional generation requirements is due to the size of the forecast confidence intervals and the ratio of wind-to-load uncertainty in each region (the 95^{th} percentile positive wind error / 95^{th} percentile negative load error). As shown in Figure 4.14, the ERCOT wind-to-load uncertainty ratios range from 1 to 2, while in MISO it fluctuates between 0.6 and 1.2. In ERCOT wind power would have provided 8.7 percent of the total electricity generated without curtailments (Table 4.1). Wind power only produced 5.5 percent of the electricity in MISO. The different ratios of electricity generated from wind power in each grid contribute to the different wind-to-load uncertainty ratios.

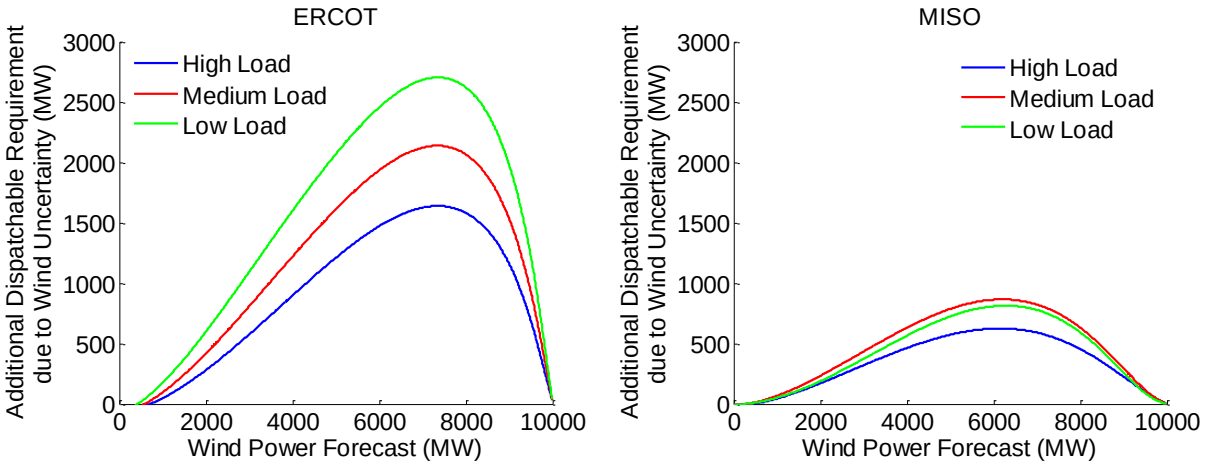

Figure 4.13: Additional dispatchable generation capacity required to cover 95% of wind forecast errors in ERCOT and MISO with 10 GW of installed wind capacity. This is the amount of capacity required for net load uncertainty less the amount of load uncertainty.

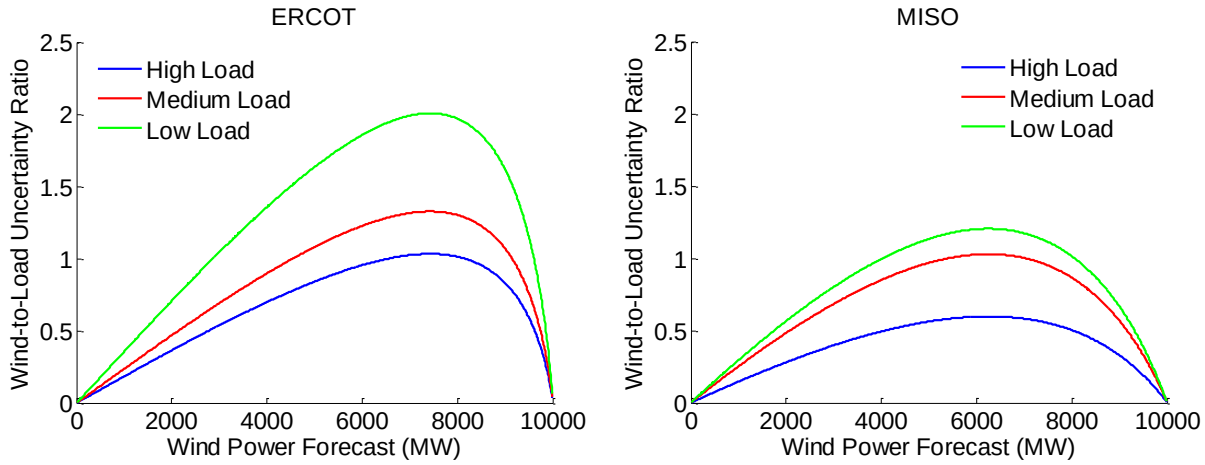

Figure 4.14: Ratio of the 95th percentile of positive wind forecast errors to the 95th percentile of negative load forecast errors in ERCOT and MISO with 10 GW of installed wind capacity. Wind is relatively more uncertain in ERCOT than in MISO.

Figure 4.13 indicates that the optimal levels of reserves to compensate wind power forecast uncertainty must be determined from the wind forecasts. Simply using a percentage of wind capacity or a percentage of the wind forecast can lead to suboptimal procurement levels. Figure 4.15 shows the additional reserves from wind forecast uncertainty as a percentage of the wind power forecast. Dispatchable capacity requirements as a percentage of the wind forecast peak near the middle of the forecast range.

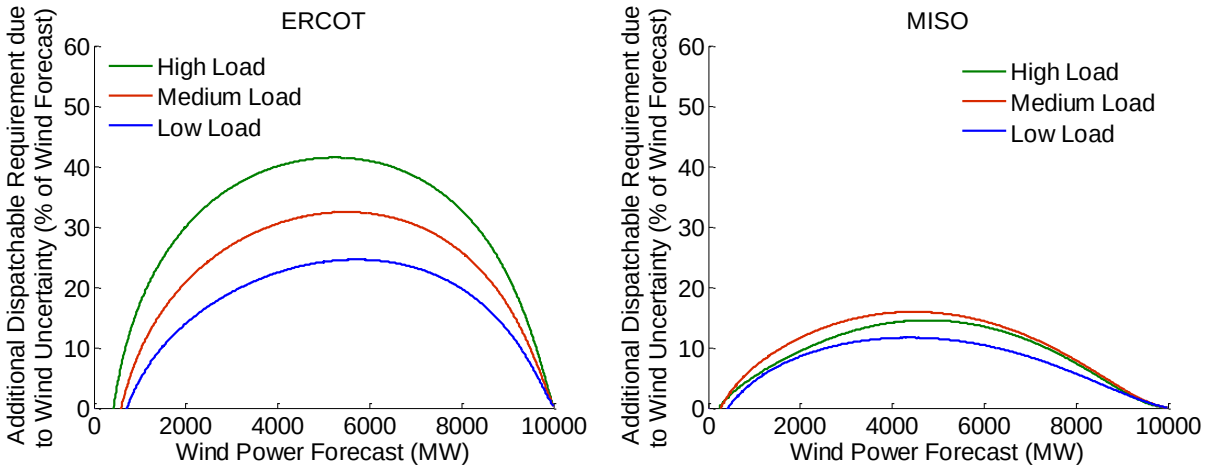

Figure 4.15: Additional dispatchable generation capacity required to cover 95% of wind forecast errors shown as a percentage of the wind power forecast in ERCOT and MISO with 10 GW of installed wind capacity.

In the results presented above, wind and load forecast errors are assumed to be uncorrelated; however, in the appendix we show that this is not always true. The correlation coefficient varied from -0.4 to +0.4 when the data were divided by the time of day or the forecast levels. Using a Monte Carlo simulation, we created simulated wind and load forecast errors with a predefined correlation coefficient which were used to generate net load errors. The simulated net load errors were used to calculate dispatchable generation capacity requirements for 95, 97.5 and 99% day-ahead reliability. Figure 4.15 shows the dispatchable capacity required for a range of correlation coefficients when the load forecast is in the medium range and forecasted wind is at the point where uncertainty is the maximum: 7500 MW for ERCOT and 6000 MW for MISO.

Positive wind forecast errors and load forecast errors have the opposite effect on net load errors. For this reason, positively correlated wind and load forecast errors reduce net load uncertainty while negatively correlated errors increase the uncertainty. As seen in Figure 4.16, when wind and load forecast errors are correlated with medium load forecasts, reserve requirements vary by up to 900 MW in ERCOT and 600 MW in MISO to cover 95% of day-ahead net load uncertainty.

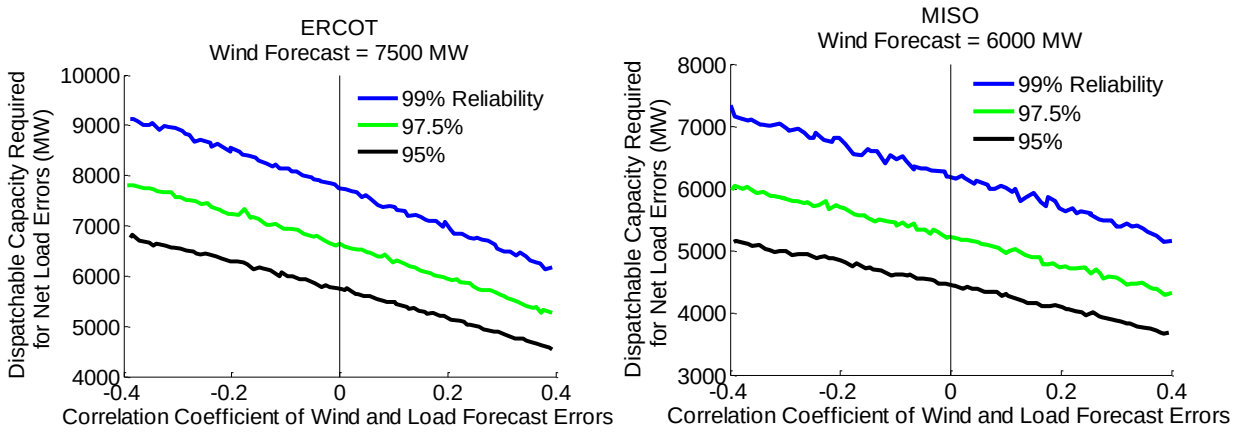

Figure 4.16: Dispatchable generation capacity required to cover net load forecast errors if wind and load forecasts are correlated. Three different reliability levels are shown. Load forecasts are between 90 and 120% of mean load and wind forecast is set such that wind uncertainty is at a maximum.

Future levels of installed wind capacity will require greater levels of operating reserves to balance the additional forecast uncertainty. Applying the analysis shown above we proceed to estimate the amount of dispatchable generation required to balance wind forecast errors if installed wind capacity is increased to 30 GW. Adding wind to a system will change the aggregate wind forecast accuracy. Whether wind forecasts will become more accurate is unclear. On one hand, the geographic diversity of wind farms may increase, which would reduce aggregate forecast errors. On the other hand, if future wind farms cluster in the areas with the best wind resources, forecast accuracy will change very little. MISO wind farms seem to already exhibit close to the maximum geographical dispersal. The ERCOT situation is much different. Wind is highly concentrated in the western section and future wind developments may reduce the correlation of output between wind farms (see Figure 4.1). For this reason, we include a scenario in which wind forecasts improve such that the forecast RMSE normalized by wind capacity decreases by 10%.

We stepped through a range of installed wind capacities up to 30 GW and determined the additional dispatchable generation required to cover 95% of net load errors due to wind uncertainty. For each wind capacity value, we calculated the maximum net load uncertainty and subtracted the load uncertainty. Load uncertainty corresponding to high load forecasts in both ERCOT and MISO was used as this allows us to determine the amount of additional capacity needed beyond the highest level required for load uncertainty.

Figure 4.17 shows the results of our calculations. In ERCOT, 9000 MW of dispatchable generation is required to cover the additional uncertainty of wind forecasts when wind capacity reaches 30 GW. If wind forecast errors are reduced by 10% as measured by the RMSE, the level of dispatchable generation required drops to 7500 MW. This is in addition to the 4000 MW required to cover forecast uncertainty from the load. The MISO grid requires 4000 MW of dispatchable generation to cover forecast errors from 30 GW of installed wind capacity. Improved forecasts reduce this amount to 3300 MW.

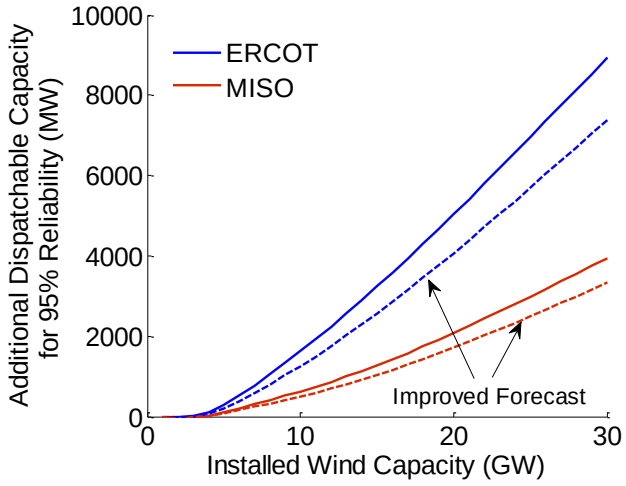

Figure 4.17: Additional dispatchable generation capacity requirements for installed wind capacity up to 30 GW in ERCOT and MISO. This is the capacity required beyond amounts for load forecast errors. Solid lines assume no change in wind forecast accuracy. Dashed lines show the effect of reducing RMSE of wind forecasts by 10%.

Figure 4.18 shows additional dispatchable capacity due to wind forecast uncertainty as a percentage of the installed wind capacity for capacity values up to 30 GW.

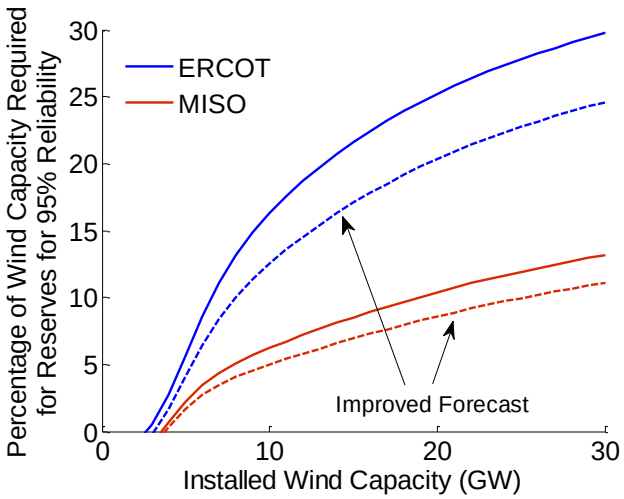

Figure 4.18: Additional reserve capacity as a percentage of installed wind capacity for installed wind capacity up to 30 GW in ERCOT and MISO. Solid lines assume no change in wind forecast accuracy. Dashed lines show the effect of reducing RMSE of wind forecasts by 10%.

The amount of additional dispatchable generation required to cover 95% of net load forecast errors is between 25% and 30% of installed wind capacity for ERCOT and between 11% and 13% for MISO. Adding wind capacity to a grid will displace energy generated by conventional power plants. At higher levels of wind power, firm capacity requirements increase meaning that conventional generators will provide more ancillary services while selling less energy in the markets.

Figure 4.19 shows the additional dispatchable capacity due to wind forecast uncertainty in ERCOT and MISO for wind penetration levels up to 30% assuming the capacity factor does not change. MISO still requires less reserve capacity than ERCOT due to the more accurate wind forecasts in MISO. The gap between the curves is less due to the fact that MISO requires much more installed capacity than ERCOT to reach a penetration level of 30%.

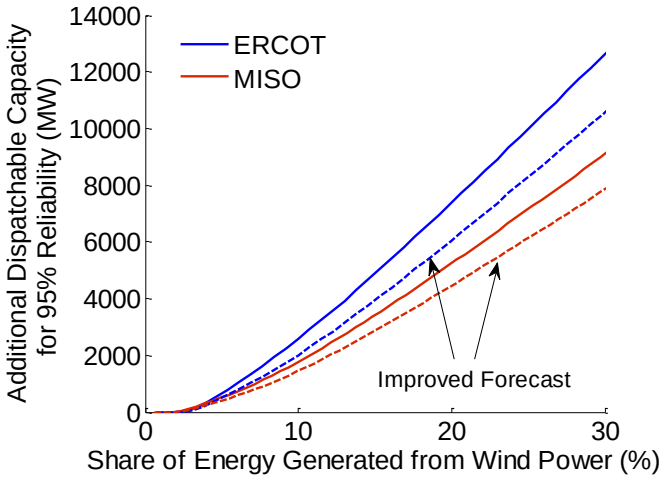

Figure 4.19: Additional reserve capacity for wind penetration up to 30% in ERCOT and MISO. Solid lines assume no change in wind forecast accuracy. Dashed lines show the effect of reducing RMSE of wind forecasts by 10%.

4.5 Conclusion

Proper quantification of wind forecast uncertainty and how it affects net load uncertainty is essential for optimal reserve procurements. Currently, most grids do not have high enough levels of wind power to make this a big concern, but the use of wind is growing. Procuring the economically optimal amount of dispatchable generation in the day-ahead market will become important. We used an empirical analysis of wind forecast uncertainty and load forecast uncertainty to develop a method that can be used in making day-ahead procurement decisions.

Using data from ERCOT and MISO we calculated the amount of dispatchable generation capacity required to cover the day-ahead wind and load forecast errors at different desired confidence levels. The data show that wind forecast error distributions are strongly dependent on the amount of wind and load forecasted for a particular hour. In order to cover 95% of the day-ahead net load under-forecast errors, reserve levels ranged from 2100 to 5700 MW in ERCOT and 1900 to 4500 MW for MISO depending on the load and wind forecasts. Hence, decisions on generation capacity procurement to cover the day-ahead forecast uncertainty must

consider the systematic nature of wind and load forecast uncertainty. Optimal decision making requires careful consideration of the forecast uncertainty each hour every day. Procurement decisions made for an entire month do not fulfill this requirement. As wind power uncertainty begins to dominate load uncertainty, system operators will have to adjust decision making to account for the wind uncertainty.

Ramp rate requirements are a factor when planning for hourly changes in net load uncertainty. We did not consider this aspect of reserve planning. However, it is an important issue that requires further investigation. Nor did we consider the economics of ensuring enough dispatchable generation for wind power uncertainty. This work shows that ensuring a particular level of reliability in an electric system with significant levels of wind power capacity requires hourly levels of reserve capacity that depend on the hourly predictions for load and wind power. We quantified such reserve level requirements as a function of the load and wind power prediction. Finally, our work did not take into account grid contingencies leading to forced generator outages. This risk is often covered with reserves equal to the amount of generation loss from the largest potential contingency in the grid.

Wind power is likely to play a larger role in electricity production in future grids. As wind power accounts for a greater percentage of electricity generation, conventional generators will be displaced. While wind displaces energy generation from conventional power plants, it only displaces a small amount of capacity resources. We show that in ERCOT, as installed wind capacity reaches 30 GW, up to 30% of this capacity will need to be matched with dispatchable capacity to cover wind forecast uncertainty. In MISO, up to 13% of the installed wind capacity must be matched with dispatchable capacity. This will have important policy implications for future grids. Sufficient levels of firm, dispatchable capacity must be procured to ensure resource adequacy in girds with high levels of wind power.

4.6 References

1. Bludszuweit, H., Dominguez-Navarro, J.A., Llombart, A., 2008. Statistical analysis of wind power forecast error, IEEE Transactions on Power Systems, 23 (3), pp. 983-991.
2. Botterud, A., Wang, J., Miranda, V., Bessa, R.J., (2010). Wind power forecasting in U.S. electricity markets, The Electricity Journal, Vol. 23, 3, 71–82
3. Bouffard, F., Galiana, F.D., 2008. Stochastic security for operations planning with significant wind power generation, 2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, pp. 1-11.
4. DOE, 2012. Wind powering america: Installed U.S. wind capacity and wind project locations, U.S. Department of Energy, http://www.windpoweringamerica.gov/wind_installed_capacity.asp .
5. Doherty, R., O'Malley, M., 2005. A new approach to quantify reserve demand in systems with significant installed wind capacity, IEEE Transactions on Power Systems, 20 (2), pp. 587- 595.
6. EIA, 2012. Annual energy outlook 2012, Energy Information Administration, DOE/EIA-0383(2012), http://www.eia.gov/forecasts/aeo/
7. Enernex, 2006. Final report – 2006 Minnesota wind integration study, Enernex Corporation, http://www.uwig.org/windrpt_vol%201.pdf
8. ERCOT, (2012). ERCOT Methodologies for Determining Ancillary Service Requirements, April, 2012.
9. Focken, U., Lange, M., Mönnich, K., Waldl, H.P., Beyer, H.G., Luig, A., 2002. 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, 90, pp. 231–246.
10. GE Energy, 2008. Final report – analysis of wind generation impact on ERCOT ancillary services requirements, GE Energy, http://www.uwig.org/AttchB-ERCOT_A-S_Study_Final_Report.pdf
11. GE Energy, 2010 Western wind and solar integration study, NREL/SR-550-47781, http://www.osti.gov/bridge
12. Hodge, B., Florita, A., Orwig, K., Lew, D., Milligan, M., 2012. A comparison of wind power and load forecasting distributions, 2012 World Renewable Energy Forum, NREL/CP-5500-54384, http://www.nrel.gov/docs/fy12osti/54384.pdf
13. Holttinen, H., Milligan, M., Kirby, B., Acker, T., Neimane, T., Molinski, T., 2008. Using standard deviation as a measure of increased operational reserve requirement for wind power, Wind Engineering, 32 (4), pp. 355-378.
14. Katzenstein, W., Fertig, E., Apt, J., 2010. The variability of interconnected wind plants, Energy Policy, 38 (8), pp. 4400–4410
15. Kehler, J., Ming, H., McMullen, M., Blatchford, J., 2010. ISO perspective and experience with integrating wind power forecasts into operations, 2010 IEEE Power and Energy Society General Meeting, pp. 1-5.

16. Lange, M., On the uncertainty of wind power predictions—Analysis of the forecast accuracy and statistical distribution of errors, Journal of Solar Energy Engineering, Vol. 127, (2005) 177–184.

17. Lowery, C., O'Malley, M., 2012. Impact of wind forecast error statistics upon unit commitment, Sustainable Energy, IEEE Transactions on, 3 (4), pp.760-768,

18. Makarov, Y.V., Etingov, P.V., Jian M., Zhenyu H., Subbarao, K., 2011. Incorporating uncertainty of wind power generation forecast into power system operation, dispatch, and unit commitment procedures, Sustainable Energy, IEEE Transactions on, 2 (4), pp.433-442

19. Matos, M.A., Bessa, R., 2009. Operating reserve adequacy evaluation using uncertainties of wind power forecast, PowerTech, 2009 IEEE Bucharest, pp.1-8

20. MISO Reliability Subcommittee Report June 27, 2012. https://www.misoenergy.org/Library/Repository/Meeting Material/Stakeholder/RSC/2012/20120626/20120626 RSC Item 08 Wind Curtailment Data.pdf

21. Ortega-Vazquez, M.A., Kirschen, D.S., 2009. Estimating the spinning reserve requirements in systems with significant wind power generation penetration, IEEE Transactions on Power Systems, 24 (1), pp. 114-124.

22. Rogers, J., Fink, S., Porter, K., 2010. Examples of wind energy curtailment practices, NREL Subcontract Report, NREL/SR-550-48737.

23. U.S. Department of Energy, Strategies and decision support systems for integrating variable energy resources, Technical Report, 2011, online: http://www1.eere.energy.gov/wind/pdfs/doe_wind_integration_report.pdf

24. Wang, J., Shahidehpour, M., Li, Z., 2008. Security-constrained unit commitment with volatile wind power generation, Power Systems, IEEE Transactions on, 23 (3), pp. 1319-1327

25. Wilkes, J., Moccia, J., Drangan, M., 2012. Wind in power: 2011 European wind statistics, European Wind Energy Association Technical Report, http://www.ewea.org/ fileadmin/ewea_documents/documents/publications/statistics/Stats_2011.pdf

26. Wiser, R., Bolinger, M., 2011. DOE 2010 Wind technologies market report, Department of Energy, DOE/GO-102011-3322.

27. Xie, L., Carvalho, P.M.S., Ferreira, L.A.F.M., Juhua L., Krogh, B.H., Popli, N., Ilić, M.D., 2011. Wind integration in power systems: operational challenges and possible solutions, Proceedings of the IEEE, 99 (1), pp.214-232

4.7 Appendix

4.7.1 Correlation of wind and load forecast errors

Wind and load forecast errors are often assumed to be independent. However, this is rarely examined with data. Correlation coefficients for the ERCOT and MISO data were found to be 0.09 and 0.05 respectively. This is near zero, but there are times when the wind and load forecast errors are much higher. Here we show correlation coefficients for the wind and load forecasts errors in ERCOT and MISO with the data separated by time of day, size of forecasts and time of year. We divided the data by hour and calculated the correlation coefficient for the wind and load errors each hour in Figure 4.20. In both ERCOT and MISO the errors are generally positively correlated. In ERCOT the correlation is significant for all but the last three hours while MISO has values that are generally not statistically different than zero.

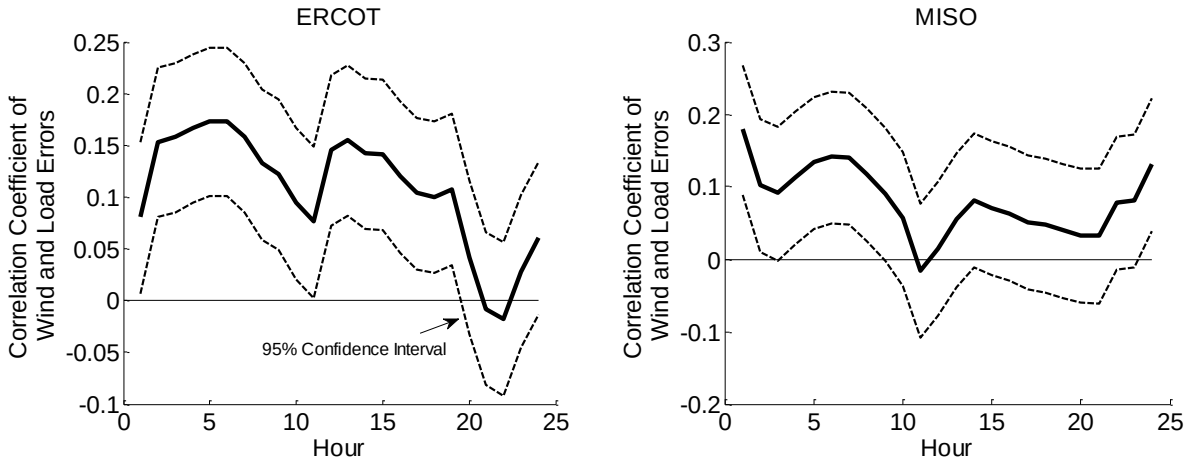

Figure 4.20: Correlation coefficient of wind and load forecast errors by hour of the day in ERCOT and MISO.

In ERCOT the correlation coefficient is high during the early morning and at midday. In MISO the coefficient is high at midnight and in the early morning. To examine this further we separated the data by month and by hour. Data from each month were sorted into 24 bins

according to the hour. For example, we grouped all of the forecast errors that occurred in January 2009 at 10 a.m. into one bin. This produced many values for correlation coefficients. Some may be spurious due to small sample sizes. Figure 4.21 shows the distribution of the correlation coefficients for both ERCOT and MISO. ERCOT data covered 24 months while MISO data covered 15 month.

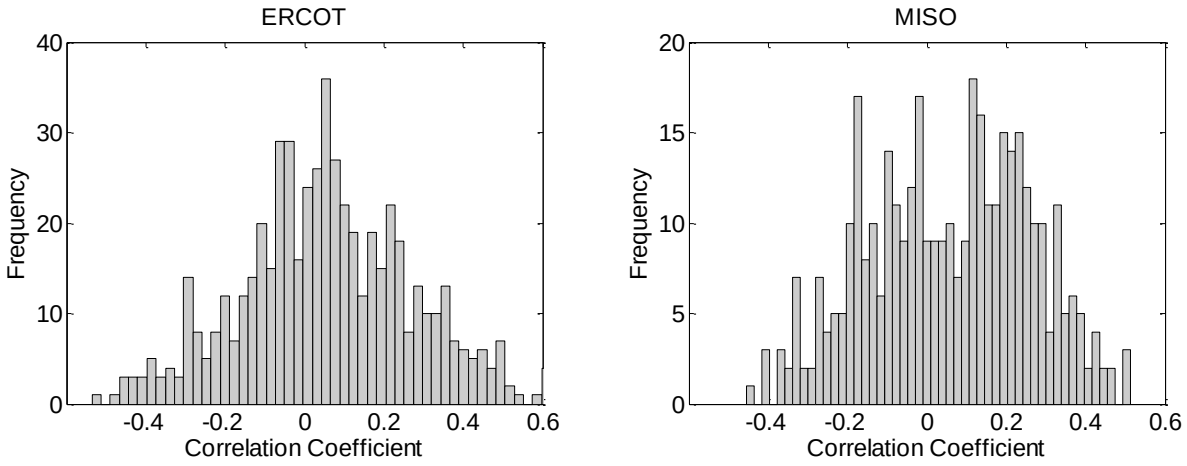

Figure 4.21: Frequency distribution of correlation coefficients calculated for each hour within the same month.

The data show little pattern in the correlation between wind and load errors. From the figure above, we conclude that it is possible to have significant positive and negative correlation between wind and load forecast errors.

Finally, we grouped the data according to forecast values to observe correlation when wind and load forecasts are high, medium and low. Tables 4.2 and 4.2 show the correlation coefficient calculated for 4 different levels of wind and load forecasts in ERCOT and MISO.

Table 4.2: Correlation coefficients calculated in ERCOT for four different forecast levels of wind and load.

ERCOT Wind and Load Forecast Correlation				
Wind Forecast as a Percentage of Installed Capacity	Load Forecast as a Percentage of Mean Load Forecast			
	LF ≤ 80%	80%≤ LF<100%	100%≤ LF<120%	LF>120%
WF ≤ 20%	0.04	0.04	0.07	0.04
20% < WF ≤ 40%	0.26	0.19	0.17	0.19
40% < WF ≤ 60%	0.26	0.09	0.05	0.12
WF > 60%	0.19	0.07	0.14	0.39

Table 4.2: Correlation coefficients calculated in MISO for four different forecast levels of wind and load.

MISO Wind and Load Correlation				
Wind Forecast as a Percentage of Installed Capacity	Load Forecast as a Percentage of Mean Load Forecast			
	LF ≤ 80%	80%≤ LF<100%	100%≤ LF<120%	LF>120%
WF ≤ 20%	0.17	0.03	0.05	0.07
20% < WF ≤ 40%	-0.07	0.06	0.07	0.23
40% < WF ≤ 60%	-0.02	0.14	0.07	0.26
WF > 60%	0.33	0.18	0.09	-0.01

Once again there are occasions where the errors have significant levels of correlation. The highest values occur with high levels of forecasted load and also very low levels of forecasted load. Given the tables and figures above we developed a model of the load and wind forecast errors with correlation levels ranging from -0.4 to 0.4.

We used a Monte-Carlo simulation to draw wind and load forecast errors from the same probability distributions presented in the main text. In order to create correlated errors we followed the procedure outlined below.

1. Generated correlated values (U) from a multivariate standard normal distribution.
2. Calculate the standard normal cumulative distribution values, $\Phi(U)$. These values are uniformly distributed and correlated.
3. Use the inverse CDF of the wind and load probability models to generate correlated wind and load forecast errors.
4. Calculate the net load errors by subtracting the wind errors from the load errors.
5. Determine an empirical CDF of the net load errors and determine the 5^{th} percentile of the negative net load errors.

We used an iterative process to determine the correct initial correlation of the variables in step one to create the proper level of correlation in the simulate error values.

4.7.2 Look-ahead time effect on forecasts

Our treatment of load forecast errors combines load forecasts made for each hour of the day. The forecast is made once per day meaning that the look-ahead time is different for each hour. If the load forecast accuracy is greatly affected by the look-ahead time, then a better method would be model load forecast errors by hour. Since we only have one forecast per day, it is not possible to observe the effects of look-ahead time on the load forecasts. Here we consider the effect of the look-ahead time by comparing forecast errors from the first (midnight to 1 a.m.) and last (11 p.m. to midnight) hours of each day on consecutive days. The errors come from two different forecasts. These errors are plotted against each other in Figure 4.22 for ERCOT and MISO. We assume that load is very similar during these two hours. While this is not definitive, it shows evidence that the load level is more important than the look-ahead time.

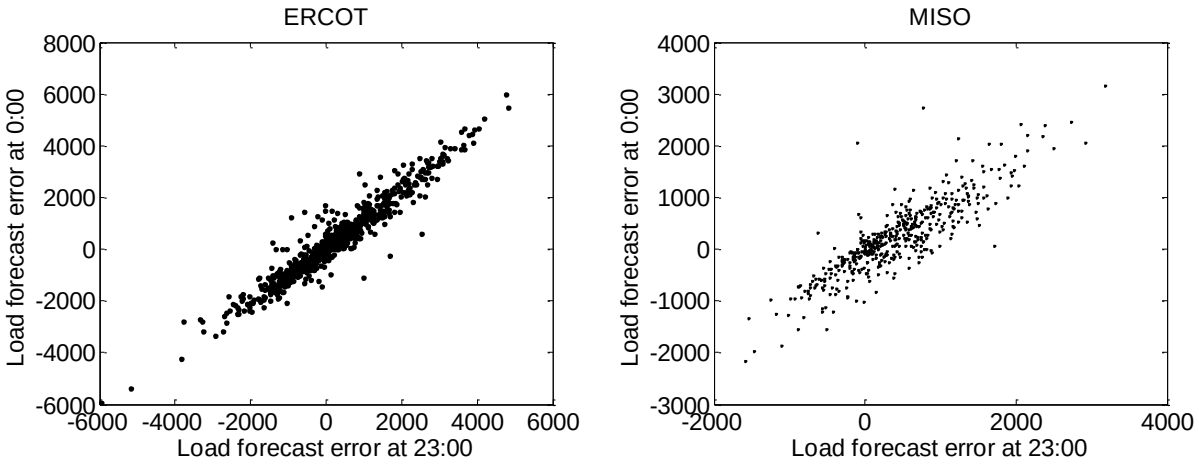

Figure 4.22: Load forecast errors for hour 1 and hour 24 on consecutive days made with different forecasts (two consecutive hours). The forecasts have very different look-ahead times but the errors are highly correlated.

Wind forecast uncertainty is obviously dependent on look-ahead time. However, for the time period used in day-ahead forecasts the uncertainty increases only a modest amount with look-ahead time. Figure 4.23 shows the root mean squared error for our data set. Day-ahead forecasts used look-ahead times from 19 to 36 hours. The range of RMSE values as a fraction of installed wind capacity over these look-ahead times is 0.13 to 0.15.

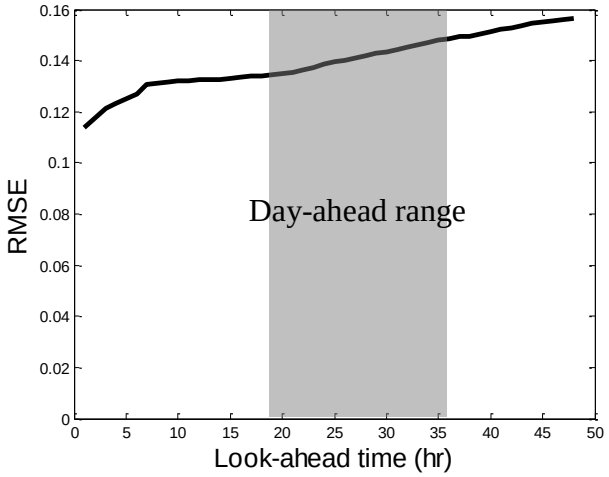

Figure 4.23: Wind forecast root mean squared error values in ERCOT for look-ahead times 1 to 48 hours. The day-ahead time frame is shown in the figure. RMSE is displayed as a fraction of installed wind capacity.

When looking at the RMSE value for different wind forecast levels it becomes apparent that the wind forecast level has a much greater effect on uncertainty than the look-ahead time. Figure 4.24 shows the same data separated according the forecast level. Wind forecast errors were placed into four different bins; forecasts less than 20% of wind capacity, forecasts between 20% and 40% of wind capacity, forecasts between 40% and 60% capacity and forecasts greater than 60% of wind capacity. RMSE values increase to a greater degree with forecast values compared with look-ahead time.

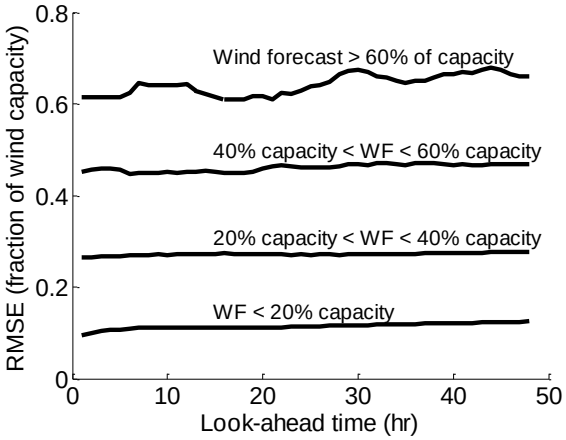

Figure 4.24: Wind forecast root mean squared error values in ERCOT for look-ahead times 1 to 48 hours with forecast errors are separated into four classes according to the forecast level. High forecasts are much more uncertain than low forecasts. RMSE is displayed as a fraction of installed wind capacity.

5 Forecasting, measurement and verification for direct load control of residential air-conditioners in energy markets[4]

5.1 Abstract

Direct load control (DLC) of residential air conditioners is a common practice among electric utilities. Historically, the objective of DLC was to reduce peak demand. Today, utilities in some parts of the U.S. can bid the load reduction from DLC into energy markets. We present a censored regression model to forecast load from residential air conditioners using historical load data, daily hour and ambient temperature. Load forecasts assist in grid management by reducing the uncertainty of load reduction during a DLC event. Our model also provides a method for verification of load reduction after an event. We tested the forecast model with hourly data from 467 air conditioners located in three different utilities. We use two months of data to train the model and then run day-ahead forecasts over a six week period. The model produced mean square errors ranging from 4% to 8% of the mean air conditioning load. Accurate forecasts were produced with much lower data requirements than other methods presented in literature.

5.2 Introduction

Independent system operators (ISOs) and electrical utilities implement demand response (DR) to alter electricity consumption in order to maintain grid reliability or provide electrical service at a lower cost. Some state legislatures require utilities to use DR programs. One method of DR is direct load control (DLC) where electrical appliances are remotely powered off. Air conditioners are often used in DLC where a load aggregator[5] controls many ACs. In deregulated electricity markets, the load aggregator can then bid this load reduction into many of the same

[4] Collaborative research in the Department of Engineering & Public Policy is the norm. This chapter is based on work done in close collaboration with Shira Horowitz.

[5] A load-aggregator can be a load serving entity such as a utility or it can be an independent aggregator.

markets that suppliers can bid conventional generation into, such as the energy and capacity markets.

Unlike traditional generation where the supply is deterministic (barring extreme events that lead to a forced outage), the available DLC resource is uncertain and must be forecasted. While generators are paid according to the quantity of energy supplied, DLC participants are paid based on the amount of load reduction. Load reductions cannot be directly measured; they are estimated by subtracting actual load during a DLC event from the amount of load a customer was assumed to have without the DLC event. In this paper we propose a new method for forecasting and measuring DLC of residential ACs using a tobit model with both upper and lower censoring.

5.2.1 Residential direct load control

Electric utilities began implementing demand response and energy efficiency programs in the 1970s in response to increased fuel prices and growing demand for electricity. Electric demand grew 35% from 1991 to 2011 (EIA, 2012a). To satisfy demand growth, electric power providers must build more power plants and transmission lines which are costly and take years to build. Effective DLC programs allow system planners to delay construction of new power plants or transmission lines by shifting peak demand to other times. DLC is also used in some regions to provide reserve capacity for contingencies in the grid. This allows grid operators to schedule less conventional generation for reserves.

In residential DLC programs, load aggregators remotely turn off appliances. Several appliances are used for DLC such as ACs, water heaters and pool pumps. We focus on DLC applied to ACs. Direct control of residential ACs began in the 1970s to reduce peak demand (Flanigan and Hadley, 1994). ACs are "cycled" by cutting power to the compressor for short periods of time during a DLC event. ACs are well suited for DLC since they comprise a large portion of residential loads (roughly 20% of electricity consumption) (EIA, 2012b), and are often

at their peak use during afternoons when electricity prices are high (Reddy et al., 1992; Sastry et al., 2010). Also, unlike many appliances such as lights and computers, air conditioners can be powered off for a brief time without much discomfort felt by customers. An early investigation of comfort levels during DLC events showed that only 15% of participants reported discomfort during the events (Kempton et al., 1992). A California utility surveyed customers during a pilot study and found the majority did not notice DLC events lasting 15 minutes or less (Sullivan et al., 2012). Other studies of AC load control programs indicate 10% to 30% of customers override the control signal after 2 hours of control depending on the ambient temperature (Kema, 2006; Kirby, 2003).

Advanced electric meters (i.e. smart meters) enable greater use of DLC in electric grids (Strbac, 2008; Hamilton and Gulhar, 2010). DLC for residential ACs is currently accomplished via wireless signals sent to appliances. Communication occurs in one direction, so there is no ability to verify if an appliance responded to the signal. Advanced meters will alleviate this issue by providing two-way communication. They will also allow the collection of load data at time intervals of one hour or less, greatly increasing the ability to forecast loads. The Federal Energy Regulatory Commission's (FERC) 2011 report on demand response and advanced metering showed the penetration of advanced meters increased from 8.7% to 13.4% of all electricity customers from 2009 to 2011 (FERC, 2011a). Federal spending on advanced metering initiatives from the American Recovery and Reinvestment Act is driving this growth which the Edison Foundation's Institute for Energy Efficiency projects to reach over 50% by 2015 (DOE, 2012).

Recent changes in wholesale electric markets are also likely to increase the use of DLC. In 2011, FERC issued order number 745 which directs wholesale energy market operators to compensate demand side resources the full energy market price as long as dispatching the DR resource is cost-effective (FERC, 2011b). Each market operator sets a threshold price based on

historical data which is used as the minimum price at which DR resources are compensated for load curtailments.

DLC provides flexibility in the grid enabling greater use of wind and solar power (Callaway, 2009; Koch et al., 2010). These renewable sources of energy have environmental benefits, but their variable and unpredictable generation present challenges to grid operators. DLC resources can quickly respond to drops in wind or solar output in a manner that is potentially more cost effective and reliable than dispatching an additional generator (Newell and Felder, 2007). Increased reliance on electricity generation from wind and solar power will likely drive DR programs (DOE, 2008). This will require more accurate load forecasting techniques that are easy to implement like the method we develop in this work.

5.2.2 Load forecasting

Accurate load forecasts for DLC events are essential for many reasons. Payments made to DR resources for curtailing loads are based upon the customer baseline (CBL), an estimate of the counterfactual event of energy use without a load reduction. Inaccuracies in the CBL lead to incorrect and unfair payments. Underpayments for DR resources discourage further participation while overpayments lead to excessive charges towards load serving entities who must pay for the reductions. System planners need to accurately know how much load reduction to expect during a system emergency. As DLC programs grow, uncertainty in the load forecasts will become a bigger issue, especially as DLC resources provide more ancillary services.

A review of literature from independent system operators (ISOs) and regional transmission operators (RTOs) shows that default CBLs differ greatly across markets (Grimm, 2008; Kema, 2011). Most of the CBLs are simple moving averages. All ISO/RTOs accept alternative methods for CBL determination as long as it is approved. In the PJM RTO, the default CBL is the average hourly load profile from the 4 highest load days of the previous 5 similar day types (weekdays,

Saturdays, Sundays/holidays) (PJM, 2012). The California ISO calculates CBLs by averaging loads from the previous 10 similar days (CASIO, 2012). The New York ISO has a similar method using an average of 10 similar day types (NYISO, 2010). New England's ISO also uses ten similar day types for the CBL (ISONE, 2012). The Electric Reliability Council of Texas publishes 3 different default CBL calculations: a linear regression of energy consumption on covariates representing weather conditions, daylight hours, season and day of the week; a moving average of 8 of the previous 10 similar days; or a model that averages days with load profiles similar to the event day (ERCOT, 2012). The Midwest ISO does not implement a default CBL and asks participants to submit their own forecasts for approval (Newell and Hajos, 2010). All ISO/RTOs allow for intercept adjustments to the CBL to better align it with load on the day of the event. This improves the accuracy of verifications (Kema, 2011; Coughlin et al., 2009; Goldberg, 2010).

Moving averages do not produce a good forecast for highly variable loads like residential ACs, so there is substantial work on air conditioner load forecast models. Broadly speaking, the models can be classified into two distinct categories: physical models and statistical models. Both model types attempt to forecast AC load as a function of several variables, primarily: temperature, time of day and relative humidity.

Most of the work in this area is directed at developing physical models of houses by employing an energy balance on a residence to estimate heat flow from the ambient air into the living space. These models consist of a system of differential equations that describe the evolution of indoor temperature and the on/off cycles of the air conditioner compressor given weather variables such as temperature, relative humidity, solar radiation, etc. Implementation of these models requires knowledge of thermal characteristics and thermostat settings of each house for use as parameters (Bargiotas and Birdwell, 1988; Molina et al., 2003; Gustafson et al., 1993).

Parameters can be measured at each house, but a more common approach is to use maximum likelihood algorithms to estimate the parameters from historical data (Pahwa and Brice, 1985; El-Ferik et al., 2006; Kamoun and Malhamé, 1992). The latter method still requires knowledge of the thermostat set points. Once a single residential AC is adequately described, it is used to produce a forecast of aggregate demand from many ACs. Several methods to aggregate individual AC loads have been proposed in the research community (Molina-García et al., 2011; Malhamé and Chong, 1985; Callaway, 2009). Usually this is carried out by expressing one or more variables with a probability density (i.e. indoor temperature or on/off state of compressors). This class of models is very expensive to implement because they require large quantities of data. They may also be sensitive to changes in the physical properties of the residence. Finally, they do not capture the behavioral aspects of AC use such as work schedules.

Statistical models on the other hand do not directly model the dynamics of energy flows. Instead they capture trends in historical AC load data to predict future loads. There is comparatively less work on statistical models applied to AC load forecasts, especially residential AC loads. One proposal to forecast load reduction from AC DLC relied on fitting a model to load measurements at a feeder circuit level (Eto et al., 2012). This method cannot forecast load for individual households. It also requires that a large fraction of ACs on each feeder participate in DLC so that it can distinguish the signal from the noise.

Parametric models of AC duty cycles have also been used to estimate load reductions by comparing controlled and non-controlled AC data (Ryan et al., 1989). Autoregressive models have also been used in AC forecasts for non-residential buildings (Penya et al., 2011), but they would not likely fit highly variable residential data well. Finally, more advanced models have been proposed to forecast building energy consumption using support vector regression (Xuemei et al., 2010) and artificial neural networks (Beccali et al., 2004). These types of models capture

the non-linearities in energy demand, but are very data intensive for each household.

Given that all ISO/RTOs currently implement simple statistical models to forecast loads, a straight forward econometric method to forecast AC load for DLC applications seems likely to gain traction. We apply a doubly censored tobit model to forecast hourly individual air conditioner loads. This accounts for the non-linearities inherent in AC energy consumption while not requiring extreme amounts of data. Our model uses ambient temperature and time of day as covariates to estimate air conditioner use. The individual loads are aggregated via a bootstrap method to create an aggregate load forecast with confidence intervals. Using this approach we calculate day-ahead hourly load forecasts over a six week period for a group of 467 air conditioners. Temperature values for the following day are assumed to be known ahead of time. Forecast models are recalibrated each day with the data available at the time of the forecast.

The remainder of this paper is organized as follows: in section 5.3 we describe the dataset. Section 5.4 describes the tobit model and the theoretical framework of the model. The results are in section 5.5 and section 5.6 covers the policy implications of this work.

5.3 Data

We obtained, under a confidentiality agreement, a dataset from Pepco Holdings, Inc. The dataset contains AC energy consumption data, weather data and meta-data for the ACs for 536 residential ACs from July - September 2010. Due to various issues with data quality we discarded data from 69 units and analyzed data from the remaining 467 units (details on data quality and cleaning protocol are in the appendix). Data loggers were installed on the air conditioners during the month of July so the initial date of data collection varies. Nearly all loggers provided data for the entire months of August and September. The data loggers recorded current measurements for the compressor circuits. During installation, technicians took spot

measurements of voltage and power which were used to convert the current measurements to power measurements.

The raw data were instantaneous power values recorded at three minute intervals. We assumed a constant power level during each three minute period to estimate energy consumption at the hourly time scale. In other words, the compressor was assumed to either be on or off for the entirety of each 3 minute period. In adding up the energy consumed during each of the three minute periods during an hour, the hourly estimates take on 20 discrete values. These errors will not have a significant effect on the final results which are aggregated over all units. We also simplified the data by using the rated power level of the compressor during intervals it was running and zero values when it was off. The raw data showed power values fluctuating mildly around the rated power level of the compressor while it was running.

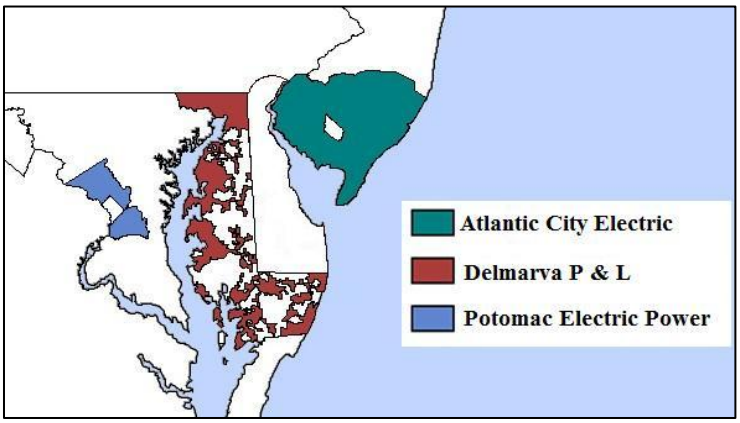

Figure 5.1: Map of regions served by the three utilities where the data were collected.

Air conditioners in this study were located in service territories for three different utilities: Potomac Electric Power Company (PEPCO); Delmarva Power and Light (Delmarva) and Atlantic City Electric (ACE). Our data came from Maryland and New Jersey customers (Figure 5.1). PEPCO's Washington D.C. customers and Delmarva's Delaware customers were not included. Hourly temperature and humidity data were collected from weather stations located

near each utility's territory and were assumed to be uniform throughout each region. Temperature statistics for each region during the time period covered in the data are in Table 5.1. Temperature data for ACE had three missing values out of 2100 hours of observations, two of which were during consecutive hours. We interpolated them from the adjacent hours.

Table 5.1: Temperature data statistics from each region where the air conditioners are located.

	Utility		
	PEPCO	ACE	Delmarva
Minimum Temperature (^{o}F)	57	49	54
Maximum Temperature	98	99	96
Mean Temperature	76	73	74
Standard Deviation	7	9	8

Many ACs included in the dataset belong to customers who participate in a DLC program that cycled air conditioners during periods of extremely high demand. In order to participate in the DLC program, customers agreed to either have switches capable of remote operation installed on the AC compressor circuit or smart thermostats that could be adjusted remotely. Signals indicating a DLC event were sent through a pager network. Customers received notice 24 hours prior to a DLC event. During the time period covered in the data, 8 DLC events occurred ranging in duration from one to four hours. Customers had the option of overriding the signal if they wanted. However, this occurred with only one customer in the dataset.

Customers in the DLC program were on one of three cycling levels: 50%, 75% or 100%. Air conditioner control during a DLC event used a smart algorithm to limit the time a compressor could run. Summary statistics for the dataset are in Table 5.2.

Table 5.2: Summary statistics from AC data set

	Utility			
	PEPCO	ACE	Delmarva	Total
Number of total air conditioners	181	72	214	467
Air conditioners cycling at 50%	58	72	88	218
Air conditioners cycling at 75%	68	0	68	136
Air conditioners cycling at 100%	55	0	58	113
Air conditioner size ≤ 2 kW	50	10	49	109
Air conditioner size > 2 kW and ≤ 3 kW	83	42	110	235
Air conditioner size > 3 kW and ≤ 4 kW	36	17	50	103
Air conditioner size > 4 kW	12	3	5	20
Average air conditioner size	2.6	2.7	2.5	2.5
Air conditioner age ≤ 5 yrs	56	30	70	156
Air conditioner age > 5 and ≤ 10	61	12	79	152
Air conditioner age > 10 and ≤ 15	39	17	46	102
Air conditioner age > 15	25	13	19	57
Average air conditioner age	9.1	9.5	8.6	8.9

5.4 Framework

5.4.1 Tobit model

Preference for AC usage is positively related to temperature. At higher temperatures, consumers want more cooling, even if their AC has reached is maximum capacity, while at cooler temperatures, consumers want less AC, and if it is cold enough, they may even want a negative amount of AC (i.e. heat). Actual AC load however, is constrained by the capacity of the AC: it can never be less than zero and cannot exceed the maximum AC capacity. We therefore model observed AC energy consumption using a doubly censored regression model, also known as a tobit model (Tobin, 1958).

We model each AC individually and add the results for an aggregate forecast. Preference for AC *i* at time *t* (incremented hourly) is modeled as a latent variable y_i^*:

$$y_{i,t}^* = \sum_{h=1}^{24}\left(\beta_{D_h,i} D_{h,t} + \beta_{D_h,i} D_{h,t} T_t\right) + \beta_{T^2,i} T_t^2 + \beta_{T1,i} T_{t-1}^2 + \beta_{E,i} E_t + \beta_{P,i} P_t + \varepsilon_{i,t} \qquad (5.1)$$

Where T_t is $max(0, R_t - 65)$, R_t is the temperature in degrees Fahrenheit during hour t, $D_{h,t}$ is an indicator variable for hour of the day, E_t, an indicator variable for a DLC event during hour t, P_t is an indicator for the three hours immediately after an event and $\beta_{\chi,i}$ is a parameter for AC i for covariate χ. The error $\varepsilon_{i,t} \sim N(0, \sigma_i)$ accounts for unobservables and random shock.

The shifted temperature term and temperature squared account for non-linearities of the observed temperature range. We use indicator variables for the 24 hours of the day and have an interaction variable with the indicators and temperature to account for consumers' diurnal activity cycle which affects preference for AC. The lagged temperature term is included to account for thermal inertia in homes. Only a lag of 1 was included because the partial-autocorrelation function was insignificant beyond a lag of 1.

Certain ACs had one or more hours of the day that had very few uncensored $y_{i,t}$ values (i.e. all the values at 3 a.m. for a particular AC were either 0 or λ_i, the capacity of the AC). This makes it difficult for the optimization routine to converge and results in insignificant estimates. We therefore combined any $D_{h,t}$ with 3 or fewer uncensored values with $D_{h-1,t}$ or $D_{h+1,t}$ (which ever had fewer uncensored values) until each $D_{h,t}$ contains greater than 3 uncensored values.

Consumers do not have a preference for DLC events, however there were DLC events in the dataset so they are accounted for with E_t. A DLC event may increase a customer's preference for AC immediately after the event, since her AC may have cycled when she would have preferred it to be on, so P_t is included to account for this. Data for event and post-event hours had to be included in the analysis because the heteroscedasticity and autocorrelation consistent standard errors require regularly spaced data (see appendix).

We explored additional variables such as lagged AC load, humidity and cooling degree hours, however these variables were all strongly collinear with other covariates and were rejected for this model.

To simplify the notation, we combine all covariates into vector $\boldsymbol{X_t}$ and all $\beta_{\chi,i}$ in vector $\boldsymbol{\beta_i}$ and rewrite Equation 5.1 as

$$y_{i,t}^{*} = \boldsymbol{X_t'\beta_i} + \varepsilon_{i,t} \tag{5.2}$$

We censor $y_{i,t}^{*}$ between 0 and λ_i, the capacity of the AC, to obtain $y_{i,t}$, the actual energy consumption of AC *i* during hour *t*:

$$y_{i,t} = \begin{cases} 0 & y_{i,t}^{*} \leq 0 \\ y_{i,t}^{*} & 0 < y_{i,t}^{*} < \lambda_i \\ \lambda_i & y_{i,t}^{*} \geq \lambda_i \end{cases} \tag{5.3}$$

We estimate $\boldsymbol{\beta_i}$ using maximum-likelihood estimation. The likelihood function is in the appendix.

Serial correlation of the errors is accounted for with heteroscedasticity and autocorrelation consistent (HAC) standard errors. HAC standard errors using White standard errors and Newey-West covariance weights are derived in appendix B using Bernard and Busse (2003). We do not use generalized least-squares to obtain consistent estimates for $\boldsymbol{\beta_i}$ because of difficulties analytically specifying the likelihood function.

5.4.2 Forecasting and confidence intervals

The model described in 5.3.1 was used to forecast AC load for each customer. Individual forecasts were summed to forecast an aggregate load. A load-aggregator would typically bid DLC into the forward (day-ahead) energy market the day before the event is to occur. For example, a load aggregator would place a bid in the forward market on August 14 for a DLC event that is to occur on August 15. The aggregator would have data up to and including August

13 to forecast load for the August 15 event. Our forecasts are computed the same way. For an August 15 forecast, we compute the parameters with all the data up to and including August 13. For an August 16 forecast we use all data up to and including August 14. Each AC, *i* therefore has different parameters $\boldsymbol{\beta_{i,d}}$ for each day *d*. In general, $\widehat{\boldsymbol{\beta}}_{i,d}$ and $\hat{\sigma}_{i,d}$ are estimated using all available $\boldsymbol{X}_{i,t}$ where $t < d - 2$.

The starting values used for the parameters maximum likelihood estimate for the first day are the ordinary least squares estimates. The starting values for each successive day $d > 1$ are the estimates from the previous day: $\widehat{\boldsymbol{\beta}}_{i,d-1}$ and $\hat{\sigma}_{i,d-1}$. Results were not sensitive to changes in the starting values.

We include uncertainty associated with $\sigma_{i,d}$, the variance of the error term, in our forecast and confidence intervals. The bootstrap method is necessary to compute aggregate forecasts and confidence intervals because of the non-linearities created by censoring. The latent residuals are Gaussian, however the censored residual are not. There is a high degree of asymmetry in the residuals since there are many more observed low temperatures than high temperatures. We used $M = 1000$ iterations of the bootstrap. Results were stable beyond 1000 iterations.

The forecasted latent variable estimate for customer *i* at time *t* on day *d* is

$$\hat{y}^{*}_{i,t} = \boldsymbol{X}'_{\boldsymbol{t}}\widehat{\boldsymbol{\beta}}_{\boldsymbol{i}} \; \forall t \in d \tag{5.4}$$

We drew a random error, $\varepsilon_{i,t,m} \; \forall t \; \varepsilon \, d$ for the mth iteration of the bootstrap from $N(0,\Sigma_d)$. For N customers in a utility, the covariance matrix is

$$\Sigma_{\mathrm{d}} = \begin{bmatrix} \sigma_{1,\mathrm{d}}^2 & \rho\sigma_{1,\mathrm{d}}\sigma_{2,\mathrm{d}} & \cdots & \rho\sigma_{1,\mathrm{d}}\sigma_{\mathrm{N,d}} \\ \rho\sigma_{2,\mathrm{d}}\sigma_{1,\mathrm{d}} & \sigma_{2,\mathrm{d}}^2 & \cdots & \rho\sigma_{2,\mathrm{d}}\sigma_{\mathrm{N,d}} \\ \vdots & \vdots & \ddots & \vdots \\ \rho\sigma_{\mathrm{N,d}}\sigma_{1,\mathrm{d}} & \rho\sigma_{\mathrm{N,d}}\sigma_{2,\mathrm{d}} & \cdots & \sigma_{\mathrm{N,d}}^2 \end{bmatrix} \tag{5.5}$$

where ρ is the correlation of errors between customers. We assume the errors have correlation ρ across customers for each t in each utility since customers with geographical proximity are likely exposed to similar shocks.

We add the random error to our latent estimate and then censor it to obtain M censored load estimates $v_{i,t,m}$ for each customer i at each time $t \in d$:

$$v_{i,t,m}^* = \hat{y}_{i,t}^* + \varepsilon_{i,t,m} \tag{5.6}$$

$$v_{i,t,m} = \begin{cases} 0 & v_{i,t,m}^* \leq 0 \\ v_{i,t,m}^* & 0 < v_{i,t,m}^* < \lambda_i \\ \lambda_i & v_{i,t,m}^* \geq \lambda_i \end{cases} \tag{5.7}$$

We average across customers to obtain M average aggregated loads at each time, t

$$\Upsilon_{\mathrm{t,m}} = \frac{1}{N}\sum_{i=1}^{N} v_{i,t,m} \tag{5.8}$$

We report the forecasted mean load at each time period, $\overline{\Upsilon}_{\mathrm{t}}$, as the average across all bootstrap iterations:

$$\overline{\Upsilon}_{\mathrm{t}} = \frac{1}{M}\sum_{i=1}^{M} \Upsilon_{\mathrm{t,m}} \tag{5.9}$$

We specify the alpha level CI by ordering the $\Upsilon_{t,m}$ and extracting the $\frac{1+\alpha}{2}$ observations as the CI.

We calculate ρ for each utility by doing a grid search over $0 \leq \rho \leq 1$. We perform this calculation separately for each utility since they would bid in separately. We calculate the percent of observations where the population average fits within the confidence intervals for the 50%, 90% and 95% confidence intervals for the within sample estimate. ρ is chosen to give the closest fit for the percentage of population values to the confidence interval.

5.5 Results

We fit the tobit model described in section 5.3 to data from each data logger in our sample of customers. We used the models to forecast AC load each hour from August 15, 2010 to September 30, 2010. We use the following steps to produce hourly aggregate AC load forecasts each day during the simulation period.

(1) Fit tobit model to logger data collected up to day $d - 2$.

(2) Calculate an hourly forecast for each AC for day d.

(3) Aggregate individual forecasts using the bootstrap method to get expected load and confidence intervals

(4) Repeat 1 - 3 for $d + 1$, $d + 2$....

Figures 5.2, 5.3 and 5.4 shows the median and upper and lower quartiles of the t-statistics for the $\widehat{\boldsymbol{\beta}}_{x,i}$ for all 470 individual AC models.

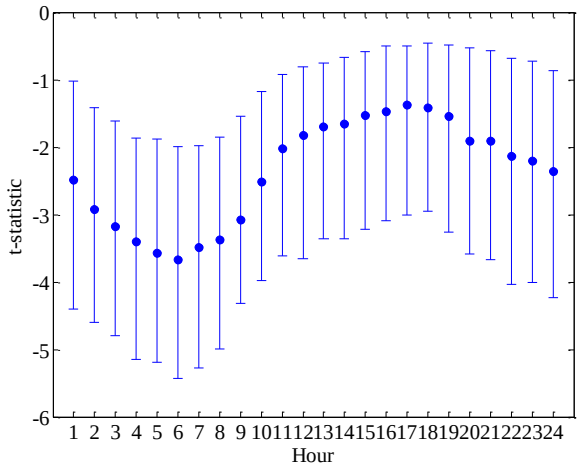

Figure 5.2: Median, upper quartile and lower quartile t-statistics for $\widehat{\boldsymbol{\beta}}_{\boldsymbol{D_h,i}}$

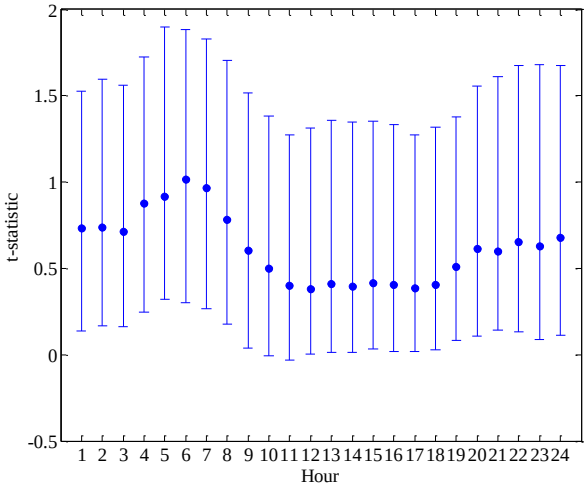

Figure 5.3: Median, upper quartile and lower quartile t-statistics for $\widehat{\boldsymbol{\beta}}_{\boldsymbol{TD_h,i}}$

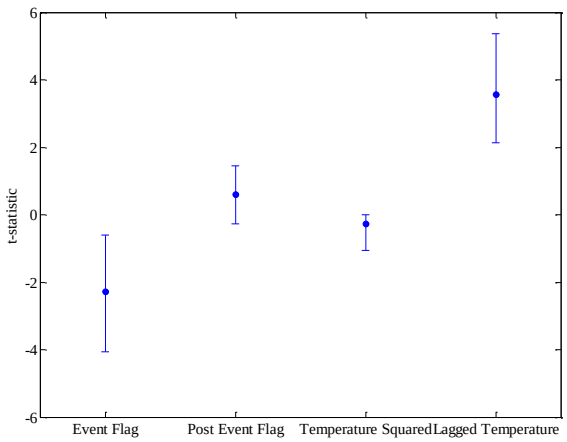

Figure 5.4: Median, upper quartile and lower quartile t-statistics for $\widehat{\boldsymbol{\beta}}_{E,i}$, $\widehat{\boldsymbol{\beta}}_{P,i}$, $\widehat{\boldsymbol{\beta}}_{T^2,i}$ and $\widehat{\boldsymbol{\beta}}_{T1,i}$

Aggregate forecasts were produced using the bootstrap method described in section 5.3. For each utility correlation coefficient ρ was estimated. Table 5.3 shows the correlation coefficients for each utility.

Table 5.3: Correlation coefficients for each utility

	Utility		
	PEPCO	ACE	Delmarva
Correlation Coefficient ρ	0.12	0.27	0.15

Figure 5.5 shows the actual load (purple), tobit forecast (dark blue) and the confidence intervals at the 50%, 90% and 95% level for the forecast for 5 days in August in the PEPCO utility. On the fourth day, August 26, an event occurred in the afternoon which can be seen on the plot.

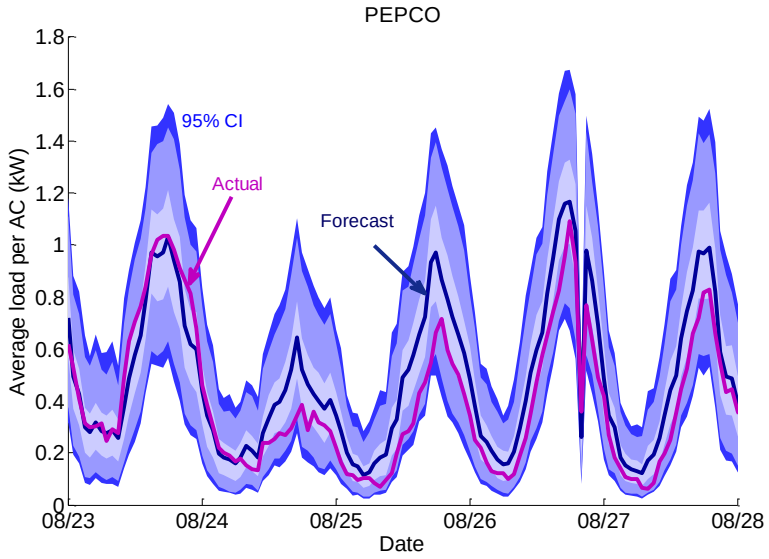

Figure 5.5: Median actual load (purple), tobit forecast (dark blue) and the confidence intervals at the 50%, 90% and 95% level for the forecast for 5 days in August in PEPCO

We compare the tobit estimate to the default CBL used in the PJM RTO since PEPCO is in PJM territory. The default CBL in PJM is the average hourly load profile from the 4 highest load days of the previous 5 similar day types (weekdays, Saturdays, Sundays/holidays) (PJM, 2012). We did not do an intercept adjustment in the manner that PJM does for verification since this is a day-ahead forecast. Figure 5.6 compares the CBL forecast to the tobit forecast. Table 5.4 shows the mean squared error (MSE) for the CBL and tobit models. The tobit model has an MSE that is much lower than the CBL.

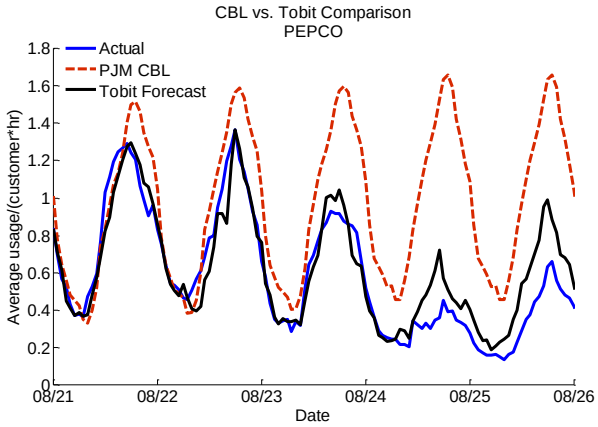

Figure 5.6: Median actual load (purple), tobit forecast (dark blue) and the confidence intervals at the 50%, 90% and 95% level for the forecast for 5 days in August in PEPCO

Table 5.4: Mean squared errors of the tobit forecast and the CBL implemented in PJM

	Mean Squared Error		Ave Load per Customer (kW)
	Tobit	CBL	
PEPCO	0.034	0.26	0.55
ACE	0.041	0.35	0.49
Delmarva	0.027	0.30	0.62

Since the tobit model presented here uses hourly temperatures, it is important to see how well it performs over a range of temperatures. Figure 5.7 shows the forecast errors plotted against the ambient temperature. At high temperatures there is a tendency to over-forecast. One possible explanation for the bias at high temperatures is that vacations occur more frequently at the end of summer when we made our forecasts based on data collected earlier in the summer. A full summer of training data would likely improve the forecasts by allowing monthly indicators in the model.

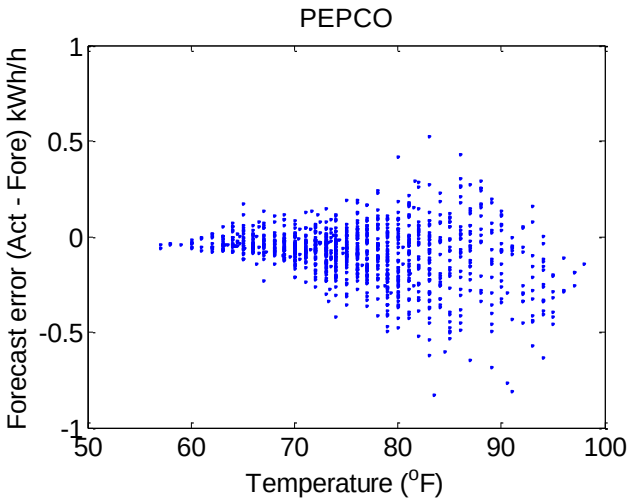

Figure 5.7: Forecast errors plotted against the ambient temperature in PEPCO

5.6 Policy Implications and Discussion

Demand response is increasing in the US as a way to make the electric grid more reliable and provide services at a lower cost. Forecasting, measurement and verification of direct load control are becoming increasingly important, as penetration levels of demand response increase. Forecasting is important for system planning and measurement and verification are necessary to ensure that payments are fair. Forecasting, measurement and verification are difficult because we are measuring the quantity of power that was *not* used, and we must reconstruct a counterfactual situation.

We have developed a new, censored regression based model for forecasting the available direct load control resource. This forecast can be used for measurement and verification to determine AC load in the counterfactual where DLC is not applied. This method is more accurate than the typical moving averages used by most ISO's, and is simple, easy and cheap to implement. This method can be further refined in future work, but introduces censored regression to load forecasting as an improvement on current forecasting methods.

5.7 References

1. Bargiotas D, Birdwell J. 1988. Residential air conditioner dynamic model for direct load control. Power Delivery, IEEE Transactions on. 3, pp. 2119-2125.
2. Beccali M, Cellura M, Lo Brano V, Marvuglia A. 2004. Forecasting daily urban electric load profiles using artificial neural networks. Energy Conversion and Management. 45, pp. 2879–2900.
3. Bernard AB, Busse MR. 2003. Consistent standard errors in panel tobit with autocorrelation. Working paper 03-25, Tuck Business School. http://ssrn.com/abstract=439061orhttp://dx.doi.org/10.2139/ssrn. 439061
4. Callaway D. 2009. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, energy conversion and management. Energy Conversion and Management. 50, pp. 1389–1400.
5. CASIO. 2012. Business practice manual for metering. Manual Version 6, California Independent System Operator. https://bpm.caiso.com/bpm/bpm/version/000000000000168
6. Coughlin K, Piette MA, Goldman C, Kiliccote S. 2009. Statistical analysis of baseline load models for non-residential buildings. Energy and Buildings. 41, pp. 374 – 381.
7. DOE. 2008. 20% wind power by 2030. Technical report, Department of Energy. http://www.20percentwind.org/20p.aspx?page=Report
8. DOE. 2012. Smart grid investment grant program progress report. Technical report, Department of Energy. http://energy.gov/sites/prod/files/Smart Grid Investment Grant Program – Progress Report July 2012.pdf
9. EIA. 2012a. Annual energy review 2011. Technical report, Energy Information Administration. http://www.eia.gov/totalenergy/data/annual/pdf/aer.pdf
10. EIA. 2012b. Annual energy outlook 2012. Technical report, Energy Information Administration. http://www.eia.gov/forecasts/aeo/pdf/0383(2012).pdf
11. El-Ferik S, Hussain SA, Al-Sunni FM. 2006. Identification and weather sensitivity of physically based model of residential air-conditioners for direct load control: A case study. Energy and Buildings. 38, pp. 997 – 1005.
12. ERCOT. 2012., Emergency interruptible load service default baseline methodologies, Technical report, Electric Reliability Council of Texas.
13. Eto J, Nelson-Hoffman J, Parker E, Bernier C, Young P, Sheehan D, Kueck J, Kirby B. 2012. The demand response spinning reserve demonstration–measuring the speed and magnitude of aggregated demand response. In System Science (HICSS), 2012 45th Hawaii International Conference on. pp. 2012 –2019.
14. FERC. 2011a. Assessment of demand response and advanced metering 2011. Technical report, Federal Energy Regulatory Commission. http://www.ferc.gov/legal/staff-reports/11-07-11-demand-response. pdf
15. FERC. 2011b. Demand response compensation in organized wholesale energy markets. Docket No. RM10-17-000 Order No. 745, Federal Energy Regulatory Commission.

http://www.ferc.gov/whats-new/comm-meet/2011/121511/E-4.pdf

16. Flanigan T, Hadley S. 1994. Analysis of successful demand side management at publicly owned utilities. Technical report, Oak Ridge National Laboratory. http://www.ornl.gov/sci/ees/etsd/pes/pubs/3445603869133.pdf
17. Goldberg M. 2010. Measure twice, cut once. Power and Energy Magazine, IEEE. 8, pp. 46 –54.
18. Grimm C. 2008. Evaluating baselines for demand response programs. In 2008 AEIC Load Research Workshop. San Antonio, TX.
19. Gustafson M, Baylor J, Epstein G. 1993. Estimating air conditioning load control effectiveness using an engineering model. Power Systems, IEEE Transactions on. 8, pp. 972 –978.
20. Hamilton K, Gulhar N. 2010. Taking demand response to the next level. IEEE Power & Energy Magazine. 8, pp. 60–65.
21. ISONE 2012. Market rule 1 section III.8. Independent System Operator of New England.
22. Kamoun S, Malhamé R. 1992. Convergence characteristics of a maximum likelihood load model identification scheme. Automatica. 28, pp. 885 – 896.
23. Kema. 2006. 2006 smart thermostat program impact evaluation. Report for San Diego Gas and Electric Co. Technical report, Kema. http://sites.energetics.com/madri/toolbox/pdfs/pricing/kema_2006_sdge_smart_thermostat.pdf
24. Kema. 2011. Pjm empirical analysis of demand response baseline methods. Technical report, Kema. http://www.pjm.com/~/media/committees-groups/subcommittees/drs/20110613/20110613-item-03b-cbl-analysis-report.ashx
25. Kempton W, Reynolds C, Fels M, Hull D. 1992. Utility control of residential cooling: resident-perceived effects and potential program improvements. Energy and Buildings. 18, pp. 201 – 219.
26. Kirby B. 2003. Spinning reserves from responsive loads. Technical report, Oak Ridge National Laboratory. http://certs.lbl.gov/certs-load-pubs.html
27. Koch S, Barcenas FS, Andersson G. 2010. Using controllable thermal household appliances for wind forecast error reduction. In IFAC Conference on Control Methodologies and Technology for Energy Efficiency. Portugal.
28. Malhamé R, Chong CY. 1985. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. Automatic Control, IEEE Transactions on. 30, pp. 854 – 860.
29. Molina A, Gabaldon A, Fuentes J, Alvarez C. 2003. Implementation and assessment of physically based electrical load models: Application to direct load control residential programmes. Generation, Transmission and Distribution, IEE Proceedings-. 150, pp. 61 – 66.
30. Molina-García ,A., Kessler M., Fuentes, J., Gómez-Lázaro E. 2011. Probabilistic characterization of thermostatically controlled loads to model the impact of demand response programs. Power Systems, IEEE Transactions on. 26 pp. 241 –251.

31. Newell S, Felder F. 2007. Quantifying demand response benefits in PJM. Technical report, The Brattle Group. http://conserveland.org/libraries/3/library_items/195

32. Newell S, Hajos A. 2010. Demand response in the Midwest ISO: An evaluation of wholesale market design. Technical report. The Brattle Group.

33. NYISO. 2010. Emergency demand response program manual. New York Independent System Operator. http://www.nyiso.com/public/webdocs/products/demand_response/emergency_demand_response/edrp_mnl.pdf

34. Pahwa A, Brice C. 1985. modeling and system identification of residential air conditioning load. Power Apparatus and Systems, IEEE Transactions on. 104, pp. 1418 – 1425.

35. Penya Y, Borges C, Agote D, Fernandez I. 2011. Short-term load forecasting in air-conditioned non-residential buildings. Industrial Electronics (ISIE), 2011 IEEE International Symposium on. pp. 1359 –1364.

36. PJM. 2012. 2011 DSR annual report. Technical report, PJM.

37. Reddy T, Vaidya S, Griffith L, Bhattacharyya S, Claridge D. 1992. A field study on air conditioning peak loads during summer in college station, texas. Technical report, Energy Systems Laboratory; Texas A & M University. http://hdl.handle.net/1969.1/2107

38. Ryan N, Powers J, Braithwait S, Smith B. 1989. Generalizing direct load control program analysis: implementation of the duty cycle approach. Power Systems, IEEE Transactions on. 4, pp. 293 –299.

39. Sastry C, Srivastava V, Pratt R, Li S. 2010. Use of residential smart appliances for peak-load shifting and spinning reserves: cost/benefit analysis. Technical report, Pacific Northwest National Laboratory. http://www.aham.org/ht/a/GetDocumentAction/i/51596

40. Strbac G. 2008. Demand side management: Benefits and challenges. Energy Policy 36: 4419 – 4426. ISSN 0301-4215. Foresight Sustainable Energy Management and the Built Environment Project.

41. Sullivan M, Bode J, Kellow B, Woehleke S, Eto J. 2012. Using residential AC load control in grid operations: PG&E's ancillary service pilot. Technical report.

42. Tobin J. 1958. Estimation of relationships for limited dependent variables. Econometrica 26: pp. 24–36.

43. Xuemei L, Lixing D, Yan L, Gang X, Jibin L. 2010. Hybrid genetic algorithm and support vector regression in cooling load prediction. In Knowledge Discovery and Data Mining, 2010. WKDD '10. Third International Conference on. pp. 527 –531.

5.8 Appendix

5.8.1 Data Cleaning Protocol

We removed ACs with poor data quality from the analysis. We describe here what metrics we used to determine the data quality.

Three loggers had too few observations (fewer than 9 days; the remaining loggers all had mroe than 52 days) and were removed. The data from 2 loggers were logged at the wrong frequency (1 minute and 5 minute instead of three minute) and were removed. Some of the loggers showed that the AC was rarely used (less than 3 hours during the entire summer). Obviously, a model will predict zero load for a household that never uses the AC. Data from these loggers were not analyzed since they provided no information for our forecasts. Several ACs had 20A loggers even though the AC capacity was greater than 20A. If a logger logged data at 20A more than 10% of the time, we assume that it required a higher amperage logger, and discard the data. There were also several loggers that became stuck on a particular value. We removed any logger from the dataset that switched state (from off to on or vice versa) in fewer than 2% of its observations. Finally, we removed loggers that had unrealistically low readings (all observations below 3 amps). A summary of the number of logger data discarded is in Table 5.5.

Table 5.5: Number of loggers discarded from the dataset.

Data Problem	Number of Loggers
Length of time less than 9 days	6
Incorrect time intervals	2
AC nearly always off	10
Logger maxed out	11
Stuck logger	13
Value unrealistically low	27
Total discarded	69

5.8.2 Tobit derivation

We derive the likelihood function and heteroscedasticity and autocorrelation consistent (HAC) standard errors for the tobit model in this section. Our derivation for the HAC standard errors is based on Bernard and Busse (2003).

This is a general derivation for a doubly-censored tobit model. Lower and upper bounds for censoring are represented here as *a* and *b*. For the model presented in this paper, $a = 0$ and $b = \lambda$. To simplify the notation, we drop *i*, the index for AC from the derivations, since the maximum likelihood estimate for each AC is done separately.

The latent variable is:

$$y_t^* = \boldsymbol{X_t'\beta} + \varepsilon_t$$

The censored variable is:

$$y_t = \begin{cases} a & y_t^* \le a \\ y_t^* & a < y_t^* < b \\ b & y_t^* \ge b \end{cases} \tag{5.10}$$

We define indicator variables:

$$I_t(a) = \begin{cases} 1 & y_t^* \le a \\ 0 & y_t^* > a \end{cases} \tag{5.11}$$

$$I_t(ab) = \begin{cases} 1 & a < y_t^* < b \\ 0 & y_t^* \le a \cup y_t^* \ge b \end{cases} \tag{5.12}$$

$$I_t(b) = \begin{cases} 1 & y_t^* \ge b \\ 0 & y_t^* < b \end{cases} \tag{5.13}$$

We assume the latent variable, y_t^*, has distribution $N(\mu,\sigma^2)$. The entire probability density of the lower censored region is applied at *a*, and the same for the upper censored region at *b*. The probability density function for the censored variable is

$$f(y_t) = \begin{cases} \Phi\left(\frac{a-\mu}{\sigma}\right) & y_t^* \le a \\ \frac{1}{\sigma}\phi\left(\frac{y_t-\mu}{\sigma}\right) & a < y_t^* < b \\ 1 - \Phi\left(\frac{b-\mu}{\sigma}\right) & y_t^* \ge b \end{cases} \tag{5.14}$$

$$f(y_t) = \begin{cases} \Phi\left(\frac{a-\mu}{\sigma}\right) & y_t^* \le a \\ \frac{1}{\sigma}\phi\left(\frac{y_t-\mu}{\sigma}\right) & a < y_t^* < b \\ \Phi\left(\frac{\mu-b}{\sigma}\right) & y_t^* \ge b \end{cases} \tag{5.15}$$

where Φ(z) is the cumulative density function (CDF) and $\phi(z)$ is the probability density function (PDF) for the standard normal distribution $N(0, 1)$.

The likelihood function $\mathcal{L}(\boldsymbol{X_t}, \theta)$ and log-likelihood function $\ell(\boldsymbol{X_t}, \theta)$ are expressed in terms of the vector of parameters $\boldsymbol{\theta} = [\boldsymbol{\beta}' \ \sigma]'$:

$$\mathcal{L}(\boldsymbol{X_t}, \theta) = \prod_{t=1}^{N} f(y_t) \tag{5.16}$$

and the gradient with respect to σ is:

$$\begin{aligned} \ell(\boldsymbol{X_t}, \theta) &= ln(\mathcal{L}) \\ &= \sum_{t=1}^{N} ln\big(f(y_t)\big) \\ &= \sum_{t=1}^{N} \left[I_t(a) ln\left(\Phi\left(\frac{a-\mu}{\sigma}\right)\right) + I_t(ab) ln\left(\frac{1}{\sigma}\phi\left(\frac{y_t-\mu}{\sigma}\right)\right) + I_t(b) ln\left(\Phi\left(\frac{\mu-b}{\sigma}\right)\right)\right] \end{aligned} \tag{5.17}$$

where

$$\mu = \boldsymbol{X}_t'\boldsymbol{\beta} \tag{5.18}$$

$$\frac{1}{\sigma}\phi\left(\frac{y_t-\mu}{\sigma}\right) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\left(\frac{(y_t-\mu)^2}{2\sigma^2}\right)} \tag{5.19}$$

We insert (5.18) and (5.19) into (5.17) to obtain:

$$\ell(\boldsymbol{X_t}, \theta) =$$

$$\sum_{t=1}^{N}\left[I_t(a)ln\left(\Phi\left(\frac{a-\boldsymbol{X_t'\beta}}{\sigma}\right)\right) - I_t(ab)\frac{1}{2}ln(2\pi\sigma^2)\left(\frac{1}{\sigma}\phi\left(\frac{y_t-\boldsymbol{X_t'\beta}}{\sigma}\right)\right)I_t(ab)\left(\frac{\left(y_t-\boldsymbol{X_t'\beta}\right)^2}{2\sigma^2}\right) + I_t(b)ln\left(\Phi\left(\frac{\boldsymbol{X_t'\beta}-b}{\sigma}\right)\right)\right] \tag{5.20}$$

The gradient of the likelihood function is

$$\nabla\ell(\boldsymbol{X_t}, \theta) = \begin{bmatrix}\frac{\partial\ell}{\partial\boldsymbol{\beta}} \\ \frac{\partial\ell}{\partial\sigma}\end{bmatrix} \tag{5.21}$$

where the gradient with respect to $\boldsymbol{\beta}$ is:

$$\frac{\partial\ell}{\partial\boldsymbol{\beta}} = \sum_{t=1}^{N}\left[I_t(a)\frac{\phi\left(\frac{a-\boldsymbol{X_t'\beta}}{\sigma}\right)}{\boldsymbol{\Phi}\left(\frac{a-\boldsymbol{X_t'\beta}}{\sigma}\right)}\left(-\frac{\boldsymbol{X_t}}{\sigma}\right) + I_t(ab)\left(\frac{\left(y_t-\boldsymbol{X_t'\beta}\right)\boldsymbol{X_t}}{\sigma^2}\right) + I_t(b)\frac{\phi\left(\frac{\boldsymbol{X_t'\beta}-b}{\sigma}\right)}{\boldsymbol{\Phi}\left(\frac{\boldsymbol{X_t'\beta}-b}{\sigma}\right)}\left(\frac{\boldsymbol{X_t}}{\sigma}\right)\right] \tag{5.22}$$

and the gradient with respect to $\boldsymbol{\sigma}$ is:

$$\frac{\partial\ell}{\partial\sigma} = \sum_{t=1}^{N}\left[I_t(a)\frac{\phi\left(\frac{a-\boldsymbol{X_t'\beta}}{\sigma}\right)}{\boldsymbol{\Phi}\left(\frac{a-\boldsymbol{X_t'\beta}}{\sigma}\right)}\left(\frac{\boldsymbol{X_t'\beta}-a}{\sigma^2}\right) - I_t(ab)\frac{1}{\sigma} + I_t(ab)\frac{\left(y_t-\boldsymbol{X_t'\beta}\right)^2}{\sigma^3} + I_t(b)\frac{\phi\left(\frac{\boldsymbol{X_t'\beta}-b}{\sigma}\right)}{\boldsymbol{\Phi}\left(\frac{\boldsymbol{X_t'\beta}-b}{\sigma}\right)}\left(\frac{b-\boldsymbol{X_t'\beta}}{\sigma^2}\right)\right] \tag{5.23}$$

A term in the gradient is in indeterminate form for $z < -38$

$$\lim_{z\to-\infty}\frac{\phi(z)}{\boldsymbol{\Phi}(z)} = \frac{0}{0} \tag{5.24}$$

so we apply L'Hôpital's rule:

$$\begin{aligned}\lim_{z\to-\infty}\frac{\phi(z)}{\boldsymbol{\Phi}(z)} &= \lim_{z\to-\infty}\frac{\partial\phi(z)/\partial z}{\partial\boldsymbol{\Phi}(z)/\partial z} \\ &= \lim_{z\to-\infty}\frac{-z\phi(z)}{\phi(z)} \\ &= -z\end{aligned} \tag{5.25}$$

The auto-covariance is:

$$\hat{\gamma}(\delta) = \frac{1}{\tau-\delta}\sum_{i=1}^{\tau-\delta}\nabla\ell(\boldsymbol{X_t}, \theta)\nabla\ell(\boldsymbol{X_{t+\delta}}, \theta)' \tag{5.26}$$

The Newey-West weights are expressed as:

$$\omega(\delta) = 1 - \frac{\delta}{\Delta+1} \tag{5.27}$$

where $\Delta \leq \sqrt{\tau}$, we use $\Delta = \tau^{0.4}$. The variance of the likelihood estimate with the HAC correction is expressed in terms of the auto-covariance and Newey-West weights:

$$Var\ \nabla\ell = \omega(0)\hat{\gamma}(0) + \sum_{\delta=1}^{\Delta}(\omega(\delta)(\hat{\gamma}(\delta) + \hat{\gamma}(-\delta)')) \tag{5.28}$$

The variance with HAC correction are expressed in terms of (5.28) and the Hessian (*H*):

$$Var(\theta) = (-H)^{-1}\ Var\ \nabla\ell\ (-H/\tau)^{-1} \tag{5.29}$$