ISSN 1503-299X

WORKING PAPER SERIES

No. 6/2019

Modelling and forecasting wind drought

Gunnar Bårdsen[#]
Department of Economics
Norwegian University of Science and Technology

Stan Hurn [¤]
School of Economics and Finance, QUT, Australia

Kenneth Lindsay [§]
Department of Mathematics, University of Glasgow, Scotland

Department of Economics

Norwegian University of Science and Technology
N-7491 Trondheim, Norway
http://www.ntnu.edu/econ/working-papers

*This research was undertaken while Gunnar Bårdsen was wisiting QUT and the excellent research facilities provided are gratefully acknowledged. The paper has benefitted from comments from seminar participants at Princeton-QUT-SMU-SHJT Workshop and David F. Hendry. The paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank.

[#] Gunnar Bårdsen, gunnar.bardsen@ntnu.no
[¤] Stan Hurn, s.hurn@qut.edu.au.
[§] Kenneth Lindsay, Kenneth.Lindsay@glasgow.ac.uk.

Modelling and forecasting wind drought

Gunnar Bårdsen[1], Stan Hurn[2], and Kenneth Lindsay[3]

[1]Department of Economics, NTNU, and Norges Bank, Norway.
[2]Corresponding author; School of Economics and Finance, QUT, Australia.
[3]Department of Mathematics, University of Glasgow, Scotland.

November 15, 2019*

Abstract

The paper examines several simple dynamic probit models in terms of their usefulness in forecasting wind drought, defined as 5 or more hours of wind speed less than 3.5 m/sec during the busiest periods of the day for the demand for electricity. Dynamic probit models work well in terms of their ability to forecast and are robust by comparison with an approach based on modelling counts. There seems little advantage to moving to modelling counts unless there is added advantage to market participants in knowing the actual prediction for the number of hours of low wind. Future research should focus on the problem of identifying the first day in a series of days with slow winds, and the first day of reasonable wind after a spell of drought. Both the probit and count models could be improved in this regard.

Keywords

JEL Classification Numbers
C22, G00.

*This research was undertaken while Gunnar Bårdsen was wisiting QUT and the excellent research facilities provided are gratefully acknowledged. The paper has benefitted from comments from seminar participants at Princeton-QUT-SMU-SHJT Workshop and David F. Hendry. The paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank.

1 Introduction

The Australian National Electricity Market (NEM) is one of the world's largest deregulated electricity markets with a transmission network to final users of around 750,000 kms. Based on this infrastructure, the trading between and around 300 registered generators with total installed capacity of approximately 48,000 MW and nine million customers operates as a pooled market under the supervision of the Australian Energy Market Operator (AEMO). Thermal, coal-fired generators have traditionally provided base-load generation for the grid, operating continuously with some ramping of production to meet daily consumer patterns of demand. Collectively, these generators produce approximately 70 terawatt hours (TWh), or 70,000 gigawatt hours (GWh), of energy each year which is close to one-third of total NEM consumption.

Over the 20-year plan period, AEMO anticipates the retirement of a substantial portion of the NEM's conventional generation fleet. A significant number of coal-fired generators in the NEM have either advised that they are closing or will reach the expected end of their technical life during this period. As coal generation retires, however, there is an increasing need for renewable wind and solar resources both in the context of emissions policy and due to their falling costs of production. Almost 80% of all currently announced, proposed, advanced, or committed projects in the NEM are wind or solar generators. Around 5 GW of new solar and wind generation projects are in an advanced stage or are committed to become operational in the next two years, displacing the energy contribution provided by both gas and coal-fired generation. Beyond committed and advanced developments in the next decade, renewable energy targets set at state level such as the Queensland Renewable Energy Target (QRET) and Victorian (VRET), influence the magnitude and location of new renewable energy infrastructure and influence the transmission requirements to enable these developments.

In this new era of reliance on renewable energy, and in particular wind energy, accurate forecasting of the power output from wind turbines is essential. Highly accurate site-specific weather forecasts are a vital component in achieving this objective if Australian wind farms are to maximise their production of renewable energy into the national energy market. It has been estimated that inaccurate short-term forecasts relating to wind and solar generation have cost Australia's renewable energy sector about $5 million in the past decade.[1] The existing forecasting systems for wind and solar are designed for longer-term time horizons and have led to multiple issues over the years. These shortcomings highlight the need for reliable short-term forecasts to provide confidence to both renewable generators and the entire industry.

There are three strands of literature on modelling wind speed. The first of these focusses on the statistical distribution of wind speed in terms of fitting the parameters of the Weibull distribution (Conradsen et al., 1984; Carrillo et al., 2014) or using extreme value analysis to investigate the probability of very high wind speeds (Cook, 1982; Coles and Walshaw, 1994; Simiu and

[1] See
https://reneweconomy.com.au/five-minute-forecasting-to-boost-solar-and-wind-and-battery-investments-35982

Heckert, 1996). A second area of emphasis is the prediction and modelling of wind intensity and its volatility with a view to pricing wind-based options (Benth and Benth, 2009; Campo-rin and Preś, 2012). The third strand of the literature is directly concerned with forecasting wind speed, for surveys see Lei et al. (2009); Chang (2014), which suggest that ARIMA time series modelling, spatial modelling and artificial intelligence-based methods are the methods of choice. There is evidence that the latter class including ANNs, support vector machines and other computer intensive hybrid models are increasingly the methods of choice.

At present in the Australian context, little is known about the actual representativeness of time series data relating to wind, beyond average wind speeds and associated power feed-in. In particular, the frequency and duration of extreme events such as long periods of low wind speeds and low renewable feed-in have not yet been analysed in detail. Consequently, rather than looking at forecasting wind speed directly, this paper is concerned primarily with the absence of wind or "wind drought". There are several reasons why the ability to forecast wind drought is important to energy market participants.

(i) From a market operator's perspective, wind drought implies the need to schedule alternative generation sources to guarantee power supply. Long phases of no or little wind power are a potential threat to future energy systems with a high proportion of renewables. Accurate day-ahead forecasting would facilitate the maintenance of supply.

(ii) Individual market players need to manage generation losses due to turbine stoppage and the cost of storing unnecessary fuel for quick-start gas-fired peaking plants to provide readily accessible capacity in the event of turbine outages.

(iii) From a purely operational perspective, maintenance teams need to perform their activities in conditions of low wind speed for safety reasons. For this purpose wind forecasts are directly useful, not only for isolating seasonal wind drought for routine scheduled maintenance, but also on an ad hoc basis for non-scheduled essential maintenance.

The occurrence and duration of phases with low or no wind and their consequences for energy system models has not yet been the subject of much research. Recent work by Plötz and Michaelis (2014) analyzed phases of low wind power feed-in and low residual load in Germany for the period 2006–2012. They used extreme value theory applied to these phases of low power and concluded that periods of wind power feed-in below eight percent of installed power lasting one week occurs once every two years, while a period of more than ten days occurred once every ten years. As interesting as these conclusions are, they do not provide much help from an operational point of view.

This paper looks at wind drought on a daily basis and seeks to provide one-day-ahead forecasts of the probability of wind drought. The fundamental underlying philosophy of the work is to explore whether using only simply dynamic models allows wind drought to be forecast with any degree of accuracy. A simple dynamic Probit model is proposed and wind direction, temperature and air pressure are the only meteorological variables used to provide day-ahead forecasts of wind drought for the Cairns weather station in northern Queensland, Australia. This site is chosen specifically because of it proximity to the Mount Emerald Wind Farm, which is a 180 MW wind power station currently in operation.[2]

2 The Data

The raw data used in this paper were obtained from the Australian Bureau of Meteorology and comprise daily observations on wind speed together with a number of other meteorological variables including various measures of temperature and air pressure for the Cairns weather station in Queensland, Australia (Latitude: -16.8736. Longitude: 145.7458.). This station is chosen primarily because it is located near to proposed major wind farm developments such as Mount Emerald.

To comply with the Systeme Internationale d'Unites (SI) wind speed is expressed in metres per second (m/s). Using this unit of measurement, the commonly used descriptive scale for wind speed is as follows: $< 5\,\mathrm{m/s}$ is *very slow* wind; $5 - 10\,\mathrm{m/s}$ is *rather slow* wind; $10 - 20\,\mathrm{m/s}$ is *considerable* wind; $20 - 30\,\mathrm{m/s}$ is very windy with caution advised; and $> 30\,\mathrm{m/s}$ is *extreme* wind which is dangerous. At very slow wind speeds, there is insufficient torque exerted by the wind on the turbine blades to make them rotate. The speed at which the turbine first starts to rotate and generate power is called the cut-in speed and is typically between $3 - 4\,\mathrm{m/s}$.

The definition of wind drought adopted in this paper is essentially ad hoc. The wind speed of $3.5\,\mathrm{m/s}$ will be adopted as be the wind speed below which power output from wind turbines is negligible. In addition, in the NEM there are three tariff periods in each 24 hour period, namely the overnight off-peak period (22:00 - 07:00), the shoulder period (07:00 - 14:00) and the peak period (14:00 - 22:00). The lack of wind power generation in off-peak periods is likely to be of little consequence to the grid. The distribution of the actual daily counts of hours for which wind speed is less than $3.5\,\mathrm{m/s}$ is shown in Figure 6. After the large peak at 0 hours per day of slow wind, the distribution tails off, but at a slower rate than would be observed in exponential decay. After 5 hours per day the distribution flattens out before tailing off after 10 hours. Consequently, a *wind drought* will be defined to occur if 5 or more hours of the 15 hours during the shoulder and peak periods on any given day experience an hourly average wind speed of less than $3.5\,\mathrm{m/s}$.

[2] The power station is located approximately 49 km south west of Cairns.

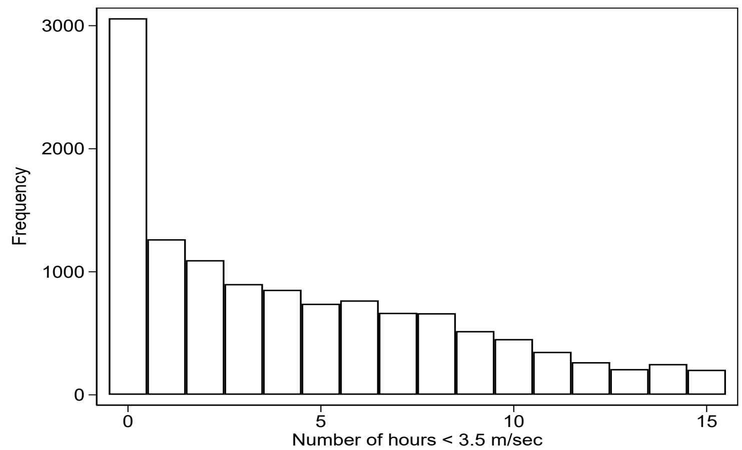

Figure 1: The distribution of the daily counts of the number of hours in the shoulder and peak tariff periods for which wind speed is less than 3.5 m/sec (maximum is 15 hours).

Only three other meteorological variables will be used in this study to model wind speed, namely, air temperature (measured in degrees Celsius), station level air pressure (measured in hecto-Pascals) and wind direction measured in compass degrees, $0^\circ - 360^\circ$. Both temperature and pressure are expected to have a negative effect on wind drought, or in other words, a positive effect on wind speed. In terms of temperature, warm air is less dense than cold air, so warm air rises creating a low pressure area. At ground level wind blows horizontally from high pressure to low pressure areas and so the effects of temperature and pressure are likely to be interdependent. There may also be diurnal and seasonal elements to temperature. Because of the sun's warming effect, winds are usually stronger during daylight hours, and around the globe, winds are usually stronger in the winter.

Wind direction is also an important element of wind speed. The rotation of the Earth on its axis causes winds and creates so-called prevailing winds. Prevailing winds tend to be stronger than wind coming from other directions. Because it is measured in compass degrees, wind direction is an example of a *circular statistic* and must therefore be treated with care. Although it is also a continuous variable, it would make no sense to use its value (in terms of degrees of a circle) without recognizing the innate nature of the measurement.[3] From hourly data on wind direction, mean daily wind direction, $\overline{\theta}$ is computed as

$$\overline{\theta} = \tan^{-1}\left(\frac{\sum_{i=1}^{24}\sin\theta_i}{\sum_{i=1}^{24}\cos\theta_i}\right).$$

The daily time series is then divided into 8 compass directions and the resultant distribution is tabulated in Table 1. It is clear that the prevailing wind at Cairns is South South East (SSE) with the wind coming from that direction 70% of the time.

[3]For instance, the simple average of a wind direction of 10° and 350° is 180°, representing almost precisely

Table 1: Mean daily wind direction for the Cairns weather station in northern Queensland.

Direction	Frequency	Percent	Cumulative
NNE	575	4.69	4.69
ENE	554	4.52	9.20
ESE	1689	13.77	22.97
SSE	8295	67.60	90.57
SSW	197	1.61	92.18
WSW	67	0.55	92.72
WNW	233	1.90	94.62
NNW	660	5.38	100.00

The months from May to October are dominated by the sub-tropical ridge, with Cairns under the influence of the south east trade stream. The prevailing winds are East to Southeasterly with strongest winds (cyclones excluded) usually occurring during April and August. During the summer months, North to Northeasterly sea breezes dominate the winds along the coast.

3 Dynamic Probit Models

Let y_t be an index defined as 1 if there is a wind drought, as defined previously, and zero otherwise:

$$y_t = \begin{cases} 1, & \text{if wind drought} \\ 0, & \text{otherwise.} \end{cases}$$

The question arises as to why the recurrent events, in this case instances of wind drought, are being summarized in this way. In other words, why not use the data on wind speed directly rather than working with some summary of its behaviour? Harding and Pagan (2016) provide a number of reasons for this choice, some of which are appropriate in the current context.

1. The choice of a binary indicator emphasizes a feature of the original data that is not immediately obvious and also focusses attention on the frequency and length of time spent in the wind drought state.

2. This choice reduces the dimensionality of the problem and isolates the characteristics that a model seeking to interpret the data needs to incorporate. In this case the added complexity of the behaviour of wind speed in off-peak periods is avoided.

3. Large and short-lived movements in the original data have no effect upon the constructed pattern of wind drought and irrelevant features, such as sudden and extreme gusts of wind, do not confound the models purpose of predicting the absence of wind.

the opposite direction.

4. This use of the data is meaningful to decision makers in the form of an operator in the energy market as well as individual generators and retailers. It is clear that there is a general interest in whether or not it is possible to predict y_{t+1} given information at time t. Predicting regressions have long been of interest and the setup here mimics the proposals for early warning systems of financial crisis, see for example, Candelon et al. (2009).

The Probit model assumes that the probability of $y_t = 1$ conditional on an information set I_t follows a cumulative normal distribution:

$$P(y_t = 1 \mid I_t) = \Phi(\pi_t), \qquad \pi_t = x_t' \beta$$

where π_t is a latent index, x_t is a vector of explanatory variables and Φ is the cumulative standard normal distribution function.

In a much of the literature the y_t are taken to be independently distributed. The autocorrelation and partial autocorrelation functions of the constructed wind drought variable are show in Figure 2. The pattern of autocorrelations is strongly suggestive of a low-order autoregressive behaviour. This kind of pattern is often observed in constructed binary variables which take their value from an underlying time series (Pagan and Harding,**Reference missing**). An important implications of this observation is that simple statical models for binary data, which rely on the assumption that the observations are independently and identically distributed, are likely to be misspecified. Care must therefore be exercised in the modelling to ensure that this dependence in the constructed indicator is adequately accounted for.

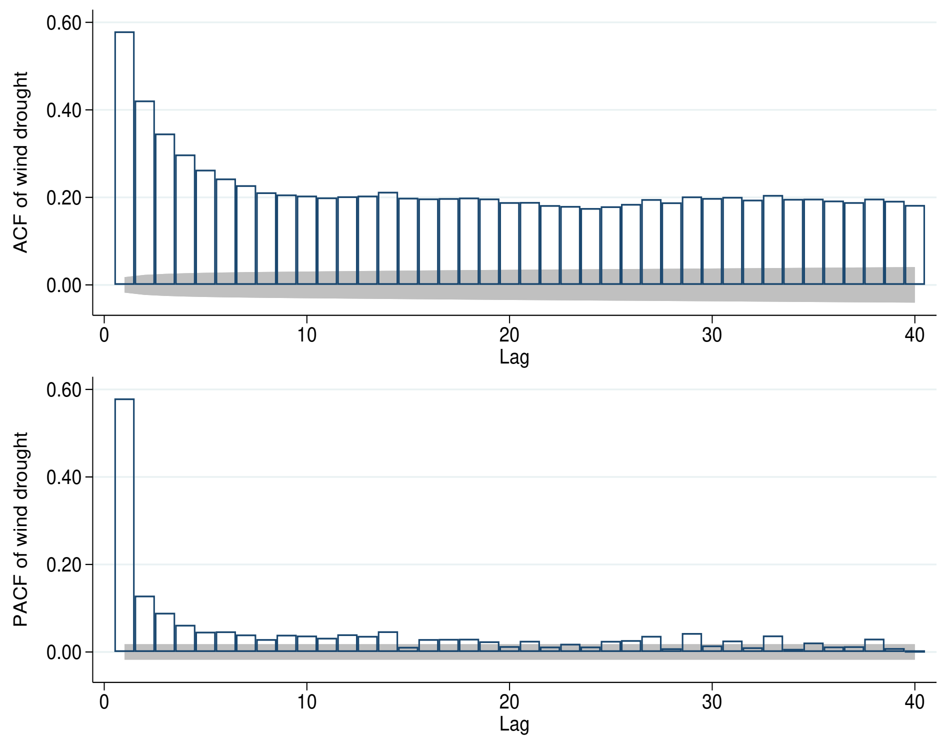

Figure 2: The autocorrelation and partial autocorrelation functions of wind drought.

To capture the dynamics, there has been a tendency to include y_{t-1} to produce extra dynamics, see for example, Dueker (1997) and Candelon et al. (2009) in the business cycle literature. Harding and Pagan (2011) point out that this approach is flawed in the case of business cycles because y_{t-1} is often a function of both y_t and y_{t-2} and therefore cannot be treated as pre-determined.[4] These considerations present no difficulty in the current context, and so y_{t-1} can be legitimately included in a dynamic Probit approach to give a model of the form

$$P(y_t = 1 \mid I_t) = \Phi(\pi_t), \qquad \pi_t = \sum_{j=1}^{q} \delta_j y_{t-j} + x_t'\beta \, .$$

Kauppi and Saikkonen (2008) extend this specification to a general autoregressive dynamic Probit model of form

$$\begin{aligned} P(y_t = 1 \mid I_t) &= \Phi(\pi_t), \\ \pi_t &= \sum_{j=1}^{p} \alpha_j \pi_{t-j} + \sum_{j=1}^{q} \delta_j y_{t-j} + x_t'\beta + y_{t-d} x_t'\gamma, \end{aligned} \tag{1}$$

which allows for dynamics in both the generating mechanism, the index, the explanatory variables, and interaction terms. The inclusion of interaction terms goes some way to inducing

[4]This relationship is due to the requirement that a recession be defined as two quarters of negative growth.

the extra dependence by allowing the coefficients of the Probit model to be stochastic and to evolve. This proposal is due to Bellégo and Ferrara (2009) and may be expected to introduce some extra non-linear dependence.

4 Estimation Results

Four variations of the general dynamic probit specification in (1) are estimated over the period 1 January 1987 to 31 December 2017. The meteorological regressors are limited to a very small set, so the results are completely transparent in terms of interpretability. These regressors are

$$x_{it} = \begin{bmatrix} x_{1t} \\ x_{2t} \\ x_{3t} \\ x_{4t} \end{bmatrix} = \begin{pmatrix} \text{temperature}_{t-1} \\ \text{pressure}_{t-1} \\ \text{difference}_{t-1} \\ \text{wind speed}_{t-1} \end{pmatrix}$$

where x_{3t} is the difference between the daily maximum and daily minimum wind speeds and x_{4t} is the lagged value of the actual count variable from which the drought indicator is constructed. There may be an argument to the effect that all the meteorological variables are jointly determined. Consequently only lagged variables are used to guard against an anticipative regression where variables are used before they are observed. This convention also means that testing for weak exogeneity is not required. In addition, as the interest here is primarily in terms of one-step-ahead forecasts, there is no scope for feedback effects to distort the results and therefore the question of strong exogeneity is not addressed. Further define two sets of dummy variables, namely, D^d_{it} being a set of 8 dummy variables for wind direction as in Table 1, and D^m_{it} a set of 12 monthly dummy variables.

The models estimated here are as follows, where the symbol # is shorthand for the interaction between two variables.

$$\pi_t = \alpha_0 + \delta_1 y_{t-1} + \sum_{i=2}^{7} \theta_i D^d_{it} + \sum_{i=2}^{12} \omega_i D^m_{it} \qquad \textbf{(Model 1)}$$

$$\pi_t = \alpha_0 + \delta_1 y_{t-1} + \sum_{i=2}^{7} \theta_i D^d_{it} + \sum_{i=2}^{12} \omega_i D^m_{it} + \sum_{i=1}^{4} \beta_i x_{it} \qquad \textbf{(Model 2)}$$

$$\pi_t = \alpha_0 + \delta_1 y_{t-1} + \sum_{i=2}^{7} \theta_i D^d_{it} \# x_{4t} + \sum_{i=2}^{12} \omega_i D^m_{it} + \sum_{i=1}^{4} \beta_i x_{it} \qquad \textbf{(Model 3)}$$

$$\pi_t = \alpha_0 + \delta_1 h_{t-1} + \sum_{i=2}^{7} \theta_i D^d_{it} \# x_{4t} + \sum_{i=2}^{12} \omega_i D^m_{it} + \sum_{i=1}^{4} \beta_i x_{it}. \qquad \textbf{(Model 4)}$$

Model 1 is the simplest possible dynamic Probit model which includes only the lagged dependent variable and the deterministic dummy variables for wind direction and month of the year. Model

2 extends Model 1 by including meteorological regressors. Model 3 deals with the issue of the prevailing wind direction by interacting lagged wind speed with wind direction. Define h_t as the time series of counts with distribution illustrated in Figure 6. Model 4 uses the lagged counts to capture dynamic effects rather than simply including the lagged binary dependent variable as in Models 1 to 3.

A number of features of the results reported in Table 2 are worth emphasizing. The most of important of these is the statistical significance of the term capturing the dynamics. In Models 1–3 the lagged binary dependent variable, y_{t-1}, is always significant and positive reflecting the positive autocorrelation illustrated in Figure 2. Moreover, broadly speaking the coefficient estimate on y_{t-1} is the same as the first order partial autocorrelation coefficient indicating that the regressors are roughly orthogonal. In Model 4 the lagged count variable, h_{t-1}, which takes the place of y_{t-1} is also significant. The summary statistics indicate that Model 4 is to be preferred suggesting the lagged count variable, with its slight increase in variability, is a useful addition to the model.

Table 2: Estimating the dynamic Probit models over the period 1 January 1987 to 31 December 2017. Coefficient estimates alone are reported. Statistical significance based on robust standard errors are denoted: * (p<0.05), ** (p<0.01) and *** (p<0.001).

Variable	Model 1	Model 2	Model 3	Model 4
Constant	1.131***	49.577***	40.520***	32.722***
y_{t-1}	1.329***	0.253***	0.270***	
h_{t-1}				0.107***
temp_{t-1}		−0.064***	−0.063***	−0.055***
press_{t-1}		−0.044***	−0.036***	−0.030***
wv_{t-1}		−0.530***	0.031	0.256***
diff_{t-1}		−0.057***	−0.068***	−0.065***
NNE				
ENE	−0.874***	−0.830***		
ESE	−1.631***	−1.530***		
SSE	−2.159***	−1.891***		
SSW	−1.150***	−1.005***		
WSW	0.293	0.483		
WNW	−0.732***	−0.612***		
NNW	−0.803***	−0.750***		
Jan				
Feb	0.111	0.096	0.101	0.104
Mar	0.014	0.134	0.123	0.111
Apr	−0.160*	0.045	0.028	0.003
May	−0.157*	−0.025	−0.050	−0.068
Jun	−0.135	−0.071	−0.103	−0.104
Jul	−0.160*	−0.071	−0.105	−0.104
Aug	−0.149*	−0.093	−0.141	−0.137
Sep	−0.229**	−0.142	−0.193*	−0.170
Oct	−0.340***	−0.197*	−0.242**	−0.223*
Nov	−0.336***	−0.272***	−0.304***	−0.280***
Dec	−0.298***	−0.251***	−0.262***	−0.232**
NNE# wv_{t-1}			0.000	0.000
ENE# wv_{t-1}			−0.285***	−0.292***
ESE# wv_{t-1}			−0.510***	−0.530***
SSE# wv_{t-1}			−0.613***	−0.637***
SSW# wv_{t-1}			−0.392***	−0.424***
WSW# wv_{t-1}			0.053	0.031
WNW# wv_{t-1}			−0.245***	−0.257***
NNW# wv_{t-1}			−0.257***	−0.267***
R^2_m	0.350	0.416	0.421	0.430
R^2_c	0.518	0.554	0.560	0.563
aic	9884.035	8871.496	8801.552	8667.024
bic	10030.727	9047.526	8977.582	8843.054

Turning now to the effects of the vector of explanatory variables, the coefficient estimates on temperature and station level pressure are negative and significant. The negative coefficient on air pressure is expected as wind flows from areas of high pressure to areas of low pressure implying that high pressure is associated with low wind speeds. The temperature effect is less obvious as higher temperatures usually mean increasing wind flow in order to fill the void as the hot air rises. The negative coefficient estimate may however be due to the fact that temperature is playing the role of the month dummy variables and capturing the generally known fact that wind speeds are usually higher in the winter months. This intuition is supported by the fact that the monthly dummy variables, with a few exceptions, become largely irrelevant once temperature is introduced into the model. The difference between the maximum and minimum recorded speed on a given day has a significant negative effect indicating that the lower the maximum wind speed on the previous day the higher the probability of a wind drought. Interestingly, the effect of interacting the lagged wind speed variable with direction provides a significant step up in the explanatory power of the model. The behaviour of lagged wind speed while significant in many cases is slightly perverse when the interaction terms are included and so should be excluded.

While it may be that the effects of temperature and pressure on wind speed may be highly interactive and nonlinear, these relationships are not pursued here. Although the linear model may be overly simplistic it has the advantage of being transparent and easily interpretable. It addition it performs relatively well for such a simple specification as shown in Table 2 which reports the estimated coefficients of Models 1 - 4 and some goodness of fit statistics. The McFadden R^2 is the ratio of the unrestricted log-likelihood function to the log-likelihood function of a restricted model in which the latter contains only an intercept and the adjusted version includes a correction for the number of estimated parameters. The count R^2 represents the proportion of correct predictions (classifications into drought or non-drought days) from the model and the adjusted count statistic is the proportion of correct guesses beyond the number that would be assigned by simply choosing the largest class.

5 Sensitivity and Forecasting

An important advantages of an estimated Probit model is that it provides predicted probabilities associated with the state of the binary dependent variable. As the point of the current exercise is to make predictions about wind drought using the model, it is therefore important that at the very least the model does relatively well in classifying the binary outcomes correctly. The standard procedure is to classify predictions using a cutoff probability of 0.5, so that $\widehat{y}_t = 1$, when predicted probability of a drought obtained from the model is greater than 0.5. The probability of correctly predicting a wind drought ($y_t = 1$) using the model is known as the model's *sensitivity*. On the other hand, the probability of correctly predicting the alternative ($y_t = 0$) is known as the *specificity* of the model.

Ideally, we would like to maximize both sensitivity and specificity. For a given model, lowering the probability cutoff point is one way of improving the sensitivity of the model, but this improvement comes with a cost because specificity is necessarily reduced. Consequently, there is a trade-off between sensitivity and specificity when manipulating the probability cutoff. Figure 6 illustrates the trade-off between sensitivity and specificity based on the parameter estimates of Model 4. The plot suggests that using a probability cutoff of 0.5 gives increased specificity at the cost of sensitivity. The plots of specificity and sensitivity intersect at 0.38 indicating that the optimal balance between these factors in this particular model may require the use of a cutoff probability significantly less than 0.5.

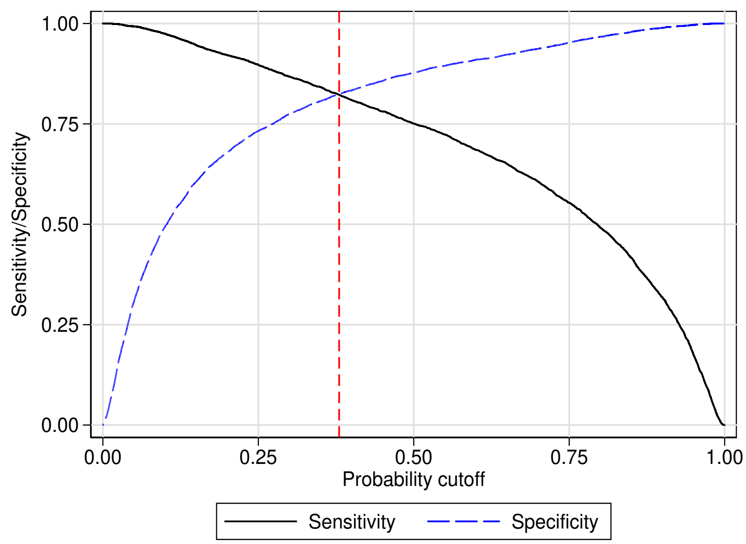

Figure 3: Sensitivity versus specificity.

Table 3 provides the formal results for the classification of events based on cutoffs of 0.5 (top panel) and 0.38 (bottom panel). From the top panel of Table 3 it may be deduced that $100 \times (3369 + 5992)/11323 = 82.678\%$ of the outcomes are correctly classified by Model 4. The sensitivity of the estimated model is $100 \times 3369/4489 = 75.05\%$ and the specificity is $100 \times 5592/6834 = 87.68\%$. The quantitative results of changing the cutoff probability to 38% are reported in the bottom panel of Table 3. The result is an increase in sensitivity from 75.05% to 82.22%, while the specificity is reduced from 87.68% to 82.48%, indicating a much better balance between these factors. Overall it seems that a good balance between sensitivity and specificity is important as missing a drought or falsely predicting a drought both have consequences in terms of costs.

Table 3: The classifications of wind drought ($y_t = 1$) and no wind drought ($y_t = 0$) obtained from Model 4 with cutoff of 0.5 (top panel) and 0.38 (bottom panel), respectively.

		True		
		Drought	Non-drought	Total
Categorized	Drought	3369	842	4211
	Non-drought	1120	5992	7112
	Total	4489	6834	11323

		True		
		Drought	Non-drought	Total
Categorized	Drought	3691	1197	4888
	Non-drought	798	5637	6435
	Total	4489	6834	11323

The real indication of the quality of the simple dynamic Probit models estimated here relates to their ability to predict wind droughts. Consequently, the models are estimated over the period 1 January 1987 to 31 December 2017, and these estimates are then used to provide a dynamic one-step-ahead forecast for the period 1 January 2018 to the end of the sample on 18 August 2018 as plotted in Figure 4. The grey shaded areas are actual wind droughts ($y_t = 1$), the solid lines are the forecast probabilities of wind drought and the dashed line is the cutoff probability of 0.38 used in Model 4 and which is superimposed on all the graphs for consistency. Note, however, that it is probably correct to expect the optimal cutoff probability for each model to be different.

There are a number of interesting observations. The first is that the predicted probabilities from the simple benchmark model are not as dynamic as those from Models 2 and 4. For the most part the predictions represent the unconditional probabilities of a wind drought. The fact that the probabilities do not fall very far in non-drought periods is problematic as there are a number of instances when false positives are found, but these are not close to a period of drought. This feature of the model is particularly troubling. The predicted probabilities from Models 2 and 4 are far more dynamic. These richer model specifications seem to pay dividends in late January, mid April, mid June and late July, where periods of wind drought that are completely missed by the benchmark model are identified. Interestingly, there does not seem to be much to choose between Models 2 and 4 despite the fact that the goodness of fit statistics indicating a strong preference for Model 4.

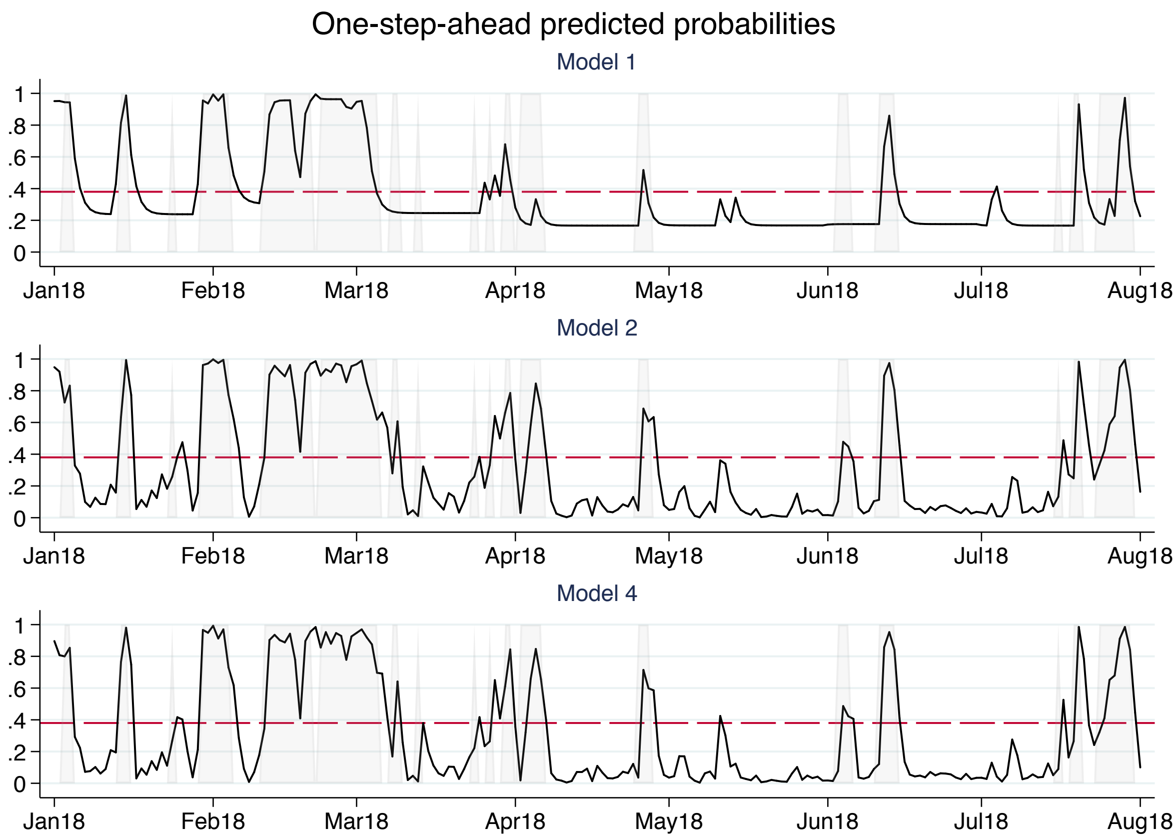

Figure 4: Comparison of one-step-ahead forecasts for Models 1, 2 and 4. The models are estimated over the period 1 January 1987 to 31 December 2017 and these estimates are then used to provide dynamic one-step-ahead forecasts for 1 January 2018 to the end of the sample on 18 August 2018. Grey shaded areas denote actual wind droughts ($y_t = 1$), the solid lines are the forecast probabilities of wind drought and the dashed line is the cutoff probability of 0.38 used in Model 4.

Finally, and most importantly, the results indicate a common problem with predictions obtained from dynamic Probit models, namely, that the predicted probabilities result in failure to predict the first occurrence of an event and tend to miss the transition out of a period of repeated events. Consider for example the 24 January 2018; a wind drought occurs but Model 4's predicted probability is 0.21. Given the memory in the model, the predicted probability on 25 January is 0.41 indicating a drought at the cutoff of 0.38. This tendency to miss the first occurrence is endemic to models of this kind (see for example Christensen et al. (2012) in the context of forecasting spikes in electricity prices.) Similarly, dynamic Probit models find it difficult to accurately date the exit point from a string of events. This is particularly evident in March 2018 where two occurrences of non-drought are missed. Major improvements to these simple models will follow if a first trigger for an event can be found. Similarly, a way to override the memory in the process and turn off the probability of an event in a timely fashion will also substantially improve the overall performance of the model.

6 Robustness

For all the reasons given previously, the discretised random variable, wind drought, is used as the major object of the econometric investigation in this paper. Despite this focus, the original time series of counts of the number of hours per day of wind speed less than 3.5 m/s is still available. A useful robustness check of the estimated models is therefore to model these counts directly. Regression models for counts, like other limited or discrete independent variable models, are nonlinear and so dealing with the time series memory in the counts is a difficult problem. Consequently, the approach adopted here is straightforward and is designed simply as a robustness check and makes no claim to be a comprehensive treatment of time series models for count data.

Given a set of regressors, x_t, for a dependent variable of counts y_t, the natural stochastic models for counts are either the standard Poisson regression model given by

$$\Pr[Y = y_t] = \frac{e^{-x_t'\beta}(x_t'\beta)^{y_t}}{y_t!},$$

or the negative binomial regression model

$$\Pr[Y = y_t] = \frac{e^{-v_t x_t'\beta}(v_t x_t'\beta)^{y_t}}{\Gamma(y_t + 1)},$$

where $\Gamma(\cdot)$ is the Gamma function and v_t is an unobserved disturbance with a $\Gamma(\alpha^{-1}, \alpha)$ distribution with $\alpha > 0$. The Poisson model is a special case of the negative binomial model corresponding to the restriction $\alpha = 0$. The negative binomial regression model is appropriate if the time series of counts has a variance which is too large to be consistent with a Poisson distribution, a case known as over-dispersion. Figure 5, which shows the unconditional Poisson and negative binomial distributions with parameters estimated from the actual count data, indicates that the negative binomial is likely to be the better choice.

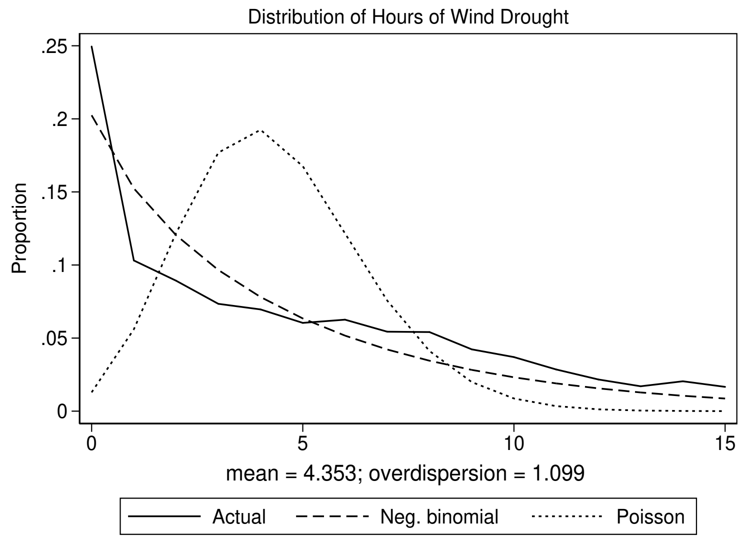

Figure 5: The unconditional distribution of the daily hourly counts of wind drought during the shoulder and peak periods. Superimposed are the best fitting unconditional negative binomial and Poisson distributions.

This simple model is, of course, inappropriate because as has been established previously that the time series of counts has memory (see Figure 2). The simple expedient of adding a lagged dependent variable, y_{t-1}, as an additional regressor is problematic. Note that

$$\begin{aligned} \mu_t &= \exp(x_t'\beta - \phi y_{t-1}) \\ \Rightarrow \log \mu_t - \log \mu_{t-1} &= x_t'\beta - x_{t-1}'\beta + \phi\, y_{t-1} - \phi\, y_{t-2}. \end{aligned}$$

Taking expectations of this expression for the growth rate of the mean gives

$$\mathbb{E}[\log \mu_t - \log \mu_{t-1}] = \phi(y_t - y_{t-1}),$$

which implies a non-zero growth rate for the conditional mean. Rather than explore more complex time series models for counts, the robustness checks here will adopt the proposal of Zeger and Qaqish (1988) and use

$$\mu_t = \exp\left(x_t'\beta - \phi\, \log(0.5 + y_{t-1})\right),$$

as the specification of the conditional mean. The small correction to y_{t-1} is necessary to deal with zero counts. An additional robustness check will use the negative binomial distribution instead of the Poisson distribution in specifying the count model.

The results of the estimation of the Poisson and negative binomial dynamic counts models over the period 1 January 1987 to 31 December 1017 are reported in Table **??**. The general specification implied by Model 4 of Sections 4, 5 looks to be appropriate, and there is little to choose between both models in terms of coefficient estimates. It is apparent, however, that the negative binomial specification is to be preferred. The over-dispersion parameter α is estimated as $\log \alpha$ to enforce the restriction that $\alpha > 0$. The test of the hypothesis that $\alpha = 0$ is rejected at the 1% level, indicating that the negative binomial specification is to be preferred, and this conclusion is also supported by the information criteria. Given the plots in Figure 5, this result is not unexpected.

Table 4: Poisson and negative binomial count regression based on the Zeger and Qaqish (1988) specification for the conditional mean and using the Model 4 specification. Only coefficient estimates are reported. Statistical significance based on robust standard errors are indicated as follows: * denotes (p<0.05), ** denotes (p<0.01) and *** denotes (p<0.001).

Variable	Poisson	Negative Binomial
Constant	19.157***	27.192***
$\log(\max(0.5, h_{t-1}))$	0.387***	0.336***
temp_{t-1}	−0.014***	−0.021***
press_{t-1}	−0.017***	−0.024***
diff_{t-1}	−0.036***	−0.034***
$\text{NNE\#wv}_{t-1}$	−0.048***	−0.092***
$\text{ENE\#wv}_{t-1}$	−0.077***	−0.118***
$\text{ESE\#wv}_{t-1}$	−0.133***	−0.175***
$\text{SSE\#wv}_{t-1}$	−0.217***	−0.270***
$\text{SSW\#wv}_{t-1}$	−0.060***	−0.102***
$\text{WSW\#wv}_{t-1}$	−0.006	−0.046
$\text{WNW\#wv}_{t-1}$	−0.079***	−0.119***
$\text{NNW\#wv}_{t-1}$	−0.086***	−0.129***
Jan		
Feb	0.033	0.037
Mar	0.048*	0.052
Apr	0.034	0.022
May	0.048	0.027
Jun	0.036	0.025
Jul	0.012	0.029
Aug	0.035	0.049
Sep	−0.022	−0.005
Oct	−0.060*	−0.060
Nov	−0.065**	−0.074*
Dec	−0.087***	−0.097**
$\log\alpha$		−1.452***
AIC	50996.255	48824.743
BIC	51172.285	49008.108

The real test of the dynamic count model, however, is in the forecasting performance of the model. The estimated negative binomial model is used to produce one-step-ahead forecasts of counts for the period 1 January 2018 to 20 August 2018. The actual and predicted values together with a shaded area to indicate the definition of wind drought are shown in Figure

6. Also shown is a horizontal line at 5 hours indicating the cutoff for the definition of a wind drought as adopted previously.

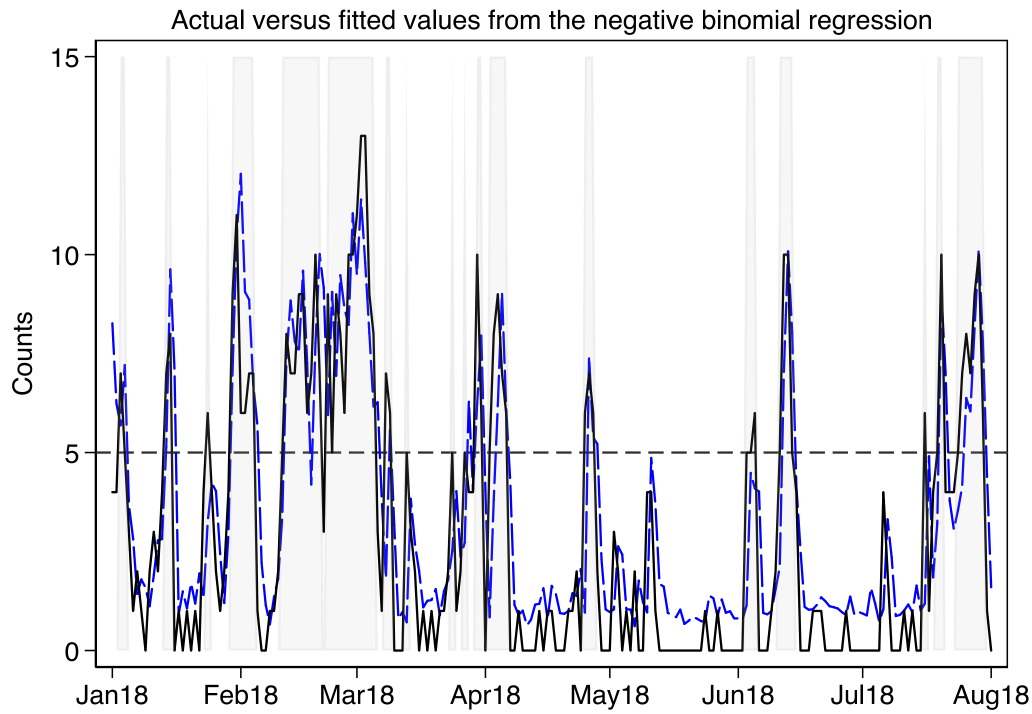

Figure 6: Actual and predicted values obtained from the negative binomial regression model using the specification in Model 4. The predictions are for the period 1 January 2018 to 18 August 2018. The areas shaded grey represent periods of wind drought as defined in the paper.

If a classification exercise based on the definition of a wind drought as defined previously (number of hours ≥ 5 in the shoulder and peak periods), then the negative binomial model returns a *sensitivity* of 71.06 and a specificity of 88.80. The sensitivity of this approach is well below that of the dynamic Probit model, but this loss of predictive accuracy in classifying periods of wind drought should be balanced against the additional information provided by a count model. For example, predicting 4 hours of low wind if 6 actually eventuate is a misspecifiction in terms of the definition, but there is valuable information in the forecast nonetheless. From a purely metric driven perspective, however, the usefulness of the binary dependent variable models are confirmed and additional work to sharpen their prediction is a potentially useful avenue for future research.

7 Conclusion

This paper has examined a series of simple dynamic Probit models in terms of their usefulness in forecasting wind drought. Wind drought for the purpose of this paper is defined as 5 or more hours of wind speed less than 3.5 m/s during the busiest periods of the day in terms of the demand for electricity. The dynamic Probit models work well in terms of their ability to forecast and are robust by comparison with an approach based on modelling counts. There seems little advantage in moving to modelling counts unless there is an added advantage to market participants in knowing the actual prediction for the number of hours of low wind. Future research should focus on the problem of identifying the first day in a series of days with slow winds and the first day of reasonable wind after a spell of drought. Both the Probit and count models could be improved in this regard.

References

Bellégo, C., and Ferrara, L. 2009. *Forecasting Euro-area recessions using time-varying binary response models for financial variables.* Working Paper 259. Banque de France.

Benth, F., and Benth, J.Š. 2009. Dynamic pricing of wind futures. *Energy Economics*, **31**, 16–24.

Camporin, M., and Preś, J. 2012. Modelling and forecasting wind speed intensity for weather risk management. *Computational statistics and data analysis*, **56**, 3459–3476.

Candelon, B., Dumitrescu, E-I., and Hurlin, C. 2009. *How to evaluate an early warning system? Towards a united statistical framework of assessing financial crisis forecasting methods.* Research Memorandum 46. Maastricht University.

Carrillo, C., Cidrás, J., Díaz-Dorado, E., and Obanado-Montaño, A.F. 2014. An approach to determine the Weibull parameters for wind energy analysis: the case of Galicia (Spain). *Energies*, **7**, 2676–2700.

Chang, W-Y. 2014. A literature review of wind forecasting methods. *Journal of Power and Energy Engineering*, **2**, 161–168.

Christensen, T., Hurn, A.S., and Lindsay, K.A. 2012. Forecasting spikes in electricity prices. *International Journal of Forecasting*, **28**, 400–411.

Coles, S.G., and Walshaw, D. 1994. Directional modelling of extreme winds. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, **43**, 139–157.

Conradsen, K., Nielsen, L.B., and Prahm, L.P. 1984. Review of Weibull statistics for estimation of wind speed distributions. *Journal of Cllimate and Applied Meteorology*, **23**, 593–608.

Cook, N.J. 1982. Towards better estimation of extreme winds. *Journal of Wind Engineering and Industrial Aerodynamics*, **9**, 295–323.

Dueker, M. 1997. Strengthening the case for the yield curve as a predictor of U.S. recessions. *Federal Reserve Bank of St. Louis Review*, **79**, 41–51.

Harding, D., and Pagan, A.R. 2011. An econometric analysis of some models of constructed binary random variables. *Journal of Business and Economic Statistics*, **29**, 86–95.

Harding, D., and Pagan, A.R. 2016. *The Econometric Analysis of Recurrent Events in Macroeconomics and Finance*. Princeton, New Jersey: Princeton University Press.

Kauppi, H., and Saikkonen, P. 2008. Predicting U.S. Recessions with Dynamic Binary Response Models. *The Review of Economics and Statistics*, **90**, 777–791.

Lei, M., Shiyan, L., Chuanwen, J., Hongling, L., and Yan, Z. 2009. A review on the forecasting of wind speed and generated power. *Renewable and Sustainable Energy Reviews*, **13**, 915–920.

Plötz, P., and Michaelis, J. 2014. *The probability of long phases without wind power and their impact on an energy system with high share of renewable energies*. Tech. rept. Fraunhofer Institute for Systems and Innovation Research ISI, Karlsruhe, Germany.

Simiu, E., and Heckert, N.A. 1996. Extreme wind distribution tails: a 'peaks over threshold' approach. *Journal of Structural Engineering*, **122**, 539–547.

Zeger, S.L., and Qaqish, B. 1988. Markov regression models for time series: a quasi-likelihood approach. *Biometrics*, **44**, 1019–1031.