



PERGAMON

Renewable Energy 28 (2003) 585–602

**RENEWABLE
ENERGY**

www.elsevier.com/locate/renene

Using medium-range weather forecasts to improve the value of wind energy production

M.S. Roulston ^{a,b,*}, D.T. Kaplan ^c, J. Hardenberg ^{a,b},
L.A. Smith ^{a,b}

^a *Pembroke College, Oxford University, Oxford, UK*

^b *Centre for the Analysis of Time Series, London School of Economics, London, UK*

^c *Macalester College, Minnesota, USA*

Received 12 February 2002; accepted 30 April 2002

Abstract

The value of different strategies for consolidating the information in European Centre for Medium Range Weather Forecasting (ECMWF) forecasts to wind energy generators is investigated. Simulating the performance of generators using the different strategies in the context of a simplified electricity market revealed that ECMWF forecasts in production decisions improved the performance of generators at lead times of up to 6 days. Basing half-hourly production decisions on a production forecast generated by conditioning the climate on the ECMWF operational ensemble forecast yields the best results of all the strategies tested.

© 2002 Elsevier Science Ltd. All rights reserved.

1. Introduction

One of the challenges faced by wind energy producers is the variation in power output caused by unpredictable fluctuations in wind velocity over a time scale of hours or days. As wind energy increases its penetration of the overall electricity market, it comes into competition with other forms of electricity generation that are less susceptible to output variations. To maintain the total power supply equal to total demand, it is necessary to compensate for fluctuations in wind energy pro-

* Corresponding author.

E-mail address: roulston@maths.ox.ac.uk (M.S. Roulston).

duction with other forms of generation; this uncertainty in output reduces the value of wind generation.

Weather forecasts can, in principle, be used by wind energy producers to make better decisions. These decisions can have genuine value in the context of an electricity market.

The fact that weather forecasts are imperfect affects how they can be used in making wind prediction decisions. For very short lead times (≈ 1 hour) time series forecasting methods can be used to forecast power production [18]. In this paper, we focus on forecasts for lead times from 1 to 10 days, so-called “medium-range” forecasts. We consider the forecasts from the numerical weather prediction (NWP) computer model run by the European Centre for Medium Range Weather Forecasting (ECMWF). This model has a spatial resolution of about 60 km and forecasts are issued every 12 h. These forecasts must be converted into the corresponding power output of a wind farm. The forecast end product should be a probabilistic description of wind energy production for each period for which a production decision must be made (e.g. 30 minutes in the UK market). Given that producers must specify production levels 3.5 h before delivery in the UK market and 12 to 36 h before delivery in Nordpool, readers might question the value of medium-range forecasts, even very accurate ones, to the wind energy sector. The purpose of this paper is to explore the *potential* value of wind forecasts in a fully commoditized energy market, with active future and forward contract trading.

In the meteorological literature there are several methods for assessing the value of probabilistic forecasts [10]. These methods include the Brier Score [2], Ranked Probability Score [4,11], Relative Operating Characteristics (ROC) [19,8] and Rank Histograms [1,5,20,6] and information content [16]. The above methods are general methods for assessing forecast skill and are largely user-independent.

The value of a forecast, however, does not stem solely from the forecast’s accuracy, but from how the forecast can be used to improve decision making. Previous studies of the impact on decision making of ECMWF forecasts have focused on binary decision scenarios [15], although the framework for how they might be used in continuous decision-making situations has been outlined [17]. In this report, we consider several strategies for consolidating the information in an ECMWF medium-range forecast with other information available to the decision maker. These strategies are evaluated in the context of a specific decision-making environment, namely a simple model of an electricity market. While the market model we used is very artificial, the synthetic power production data were generated from measured wind speed data from the CLRC Rutherford Appleton Laboratory, Energy Research Unit (ERU) test site in Oxfordshire, UK. We show that the information consolidation strategy can have a substantial impact on the value of a wind generation production facility.

While the specific decision model we consider here may not be appropriate to all wind energy producers, the general framework we provide for consolidating forecast information with other sources of information and the method of evaluation the strategies should prove useful in a variety of decision-making situations. The framework

should also be of interest to meteorologists, as it is within decision-making contexts that their forecasts must perform.

2. A model of wind energy decision making

In deregulated markets, it is crucial for electricity generators to be able to make reliable promises of energy for delivery to the grid at a future time. In this respect wind generators are at a disadvantage compared with other types of generators because of the inherent unpredictability of their production levels. Partly for this reason some deregulated markets exempt wind energy producers from penalty payments for failure to satisfy supply obligations. The exact rules differ between markets and such exemptions may not be sustainable as wind energy penetration increases. In a fully liberalized market the ability to forecast future production levels could yield substantial financial advantage.

We model the mechanism by which producers promise to deliver energy to the grid in terms of a contract to provide a target amount of electricity, E_c , for a particular half-hour in the future. The producer is paid a fixed unit price, P_c , for this promised electricity. If, when the half-hour arrives, wind production at the generator's wind farm is not sufficient to satisfy the contract, then the generator must satisfy the contract by purchasing the shortfall on the spot market at a cost of P_s per unit. If E_a is the actual amount of electricity produced by the wind farm for the half-hour in question, then the net income, I , of the wind farm for the period is given by

$$I = \begin{cases} E_c \cdot P_c & \text{when } E_a \geq E_c \\ E_c \cdot P_c - (E_c - E_a) \cdot P_s & \text{when } E_a < E_c \end{cases} \quad (1)$$

The actual production level, E_a , and the spot price, P_s , are both unknown when the contract is written.¹ We assume that the user is *risk-neutral*, with the sole aim of maximizing their expected income.² In making its decision, the producer needs to balance two factors: the income $E_c \cdot P_c$ that comes from promising to deliver energy E_c versus the risk of a penalty if E_a falls short of E_c . Given the values of E_a and P_s , the income, I , of the wind farm for the contract period is given by Eq. (1).

Consider first a simplified situation where the spot price, P_s , is known and only the actual production E_a is uncertain. Presumably, $P_s > P_c$, otherwise there would be no incentive for the wind energy producer to keep E_c at realizable levels.

From the wind energy producer's perspective, the uncertainty in E_a at the time the contract is agreed to can be summarized by a probability distribution. This probability

¹ Note that overproduction, i.e. $E_a > E_c$, has no value in this model, but the model could be changed to incorporate a value or a penalty for production beyond E_c .

² Risk-averse users may build a risk penalty into their utility function. In this case, access to alternative generation can be used to reduce this risk penalty, but it will not affect the decision made by a risk-neutral user—it merely changes the actual cost of having to buy any shortfall on the spot market into the opportunity cost of having excess capacity. Aggregation of wind energy production will also reduce the risk penalty.

distribution can be represented as a density $p(E_a)$, or, equivalently, in terms of a

cumulative distribution $P(E_a) = \int_0^{E_a} p(x)dx$, that gives the probability of E_a being less

than any indicated value. For reasons that will become clear, we prefer the cumulative format for describing probabilities.

The general form of $P(E_a)$ is shown in Fig. 1. From fundamental considerations, we know that $P(0) = 0$ and that $P(E_{a,\max}) = 1$, where $E_{a,\max}$ is the maximum capacity of the wind energy generating facility. Between these two extremes we know that $P(E_a)$ is a non-decreasing function, but the details of its shape depend on the specific information \mathfrak{I} available to the wind energy producer, and we prefer to write $P(E_a)$ as a conditional probability, conditioned on the information \mathfrak{I} and denoted $P(E_a|\mathfrak{I})$. We shall consider how weather forecast information can be consolidated with other information to produce an overall \mathfrak{I} and the consequent $P(E_a|\mathfrak{I})$ in Section 3, but for now we merely stipulate that the generator has some $P(E_a|\mathfrak{I})$.

To maximize total expected income the producer will set E_c at the level that maximizes I . But, since E_a and P_s are uncertain when the contract is signed, the producer sets E_c to maximize the expected value of net income, I .

The optimal value of E_c is such that the marginal income from an additional unit of promised electricity is exactly balanced by the expected penalty from failing to meet E_c for that unit. This balance occurs when

$$P_c - P(E_c) \cdot P_s = 0 \text{ or, equivalently, } P(E_c) = \frac{P_c}{P_s}. \quad (2)$$

Since $P_s > P_c$, the ratio P_c/P_s is always between 0 and 1. The optimal level, E_c , can be calculated from inverting $P(E_a|\mathfrak{I})$ at the level P_c/P_s , i.e. setting E_c at the $100 \times P_c/P_s$ percentile of $P(E_a|\mathfrak{I})$, as shown in Fig. 1. This is actually the familiar




Fig. 1. The general form of $P(E_a|\mathfrak{I})$, the cumulative distribution of the probability of E_a conditioned on information \mathfrak{I} .

binary cost–loss scenario [7]. In this case the decision that must be made is whether to promise to supply the next unit of electricity.

An alternative way of deriving the same result is to consider the total expected income. This is

$$\langle I \rangle = E_c \cdot P_c - \int_0^{E_c} P_s \cdot (E_c - E_a) p(E_a | \mathfrak{I}) dE_a. \quad (3)$$

The generator seeks to maximize $\langle I \rangle$ by adjusting E_c . Taking the derivative of $\langle I \rangle$ with respect to E_c gives

$$\frac{d\langle I \rangle}{dE_c} = P_c - P_s \int_0^{E_c} p(E_a | \mathfrak{I}) dE_a = P_c - P_s P(E_a | \mathfrak{I}). \quad (4)$$

Setting $d\langle I \rangle / dE_c = 0$ yields the maximum value of $\langle I \rangle$ in agreement with Eq. (2).

In the more general case, P_s is also unknown at the time the contract is signed. The income maximization problem then involves P_s in the integral; we need to consider the joint distribution $p(E_a, P_s | \mathfrak{I})$ rather than $p(E_a | \mathfrak{I})$, and integrate over the possible values of P_s :

$$\langle I \rangle = E_c \cdot P_c - \int_0^{E_c} \int_0^\infty P_s \cdot (E_c - E_a) p(E_a, P_s | \mathfrak{I}) dE_a dP_s. \quad (5)$$

In the special case when P_s is independent of E_a , this reduces to

$$\frac{d\langle I \rangle}{dE_c} = P_c - \int_0^\infty P_s p(P_s | \mathfrak{I}) dP_s \int_0^{E_c} p(E_a | \mathfrak{I}) dE_a = P_c - \langle P_s \rangle P(E_a | \mathfrak{I}). \quad (6)$$

Only the expectation value of P_s enters into the decision. It is this special case that we consider here, since we have no data indicating a relationship between E_a and P_s . However, any hypothesized dependence of P_s on E_a could easily be incorporated into our method. Note that while the E_c decision depends only on $\langle P_s \rangle$, the actual income depends on P_s itself.

3. Constructing $P(E_a | \mathfrak{I})$

In the conditional probability $P(E_a | \mathfrak{I})$ that is used to set the production level E_c , the conditioning variable \mathfrak{I} is the totality of information available to the decision maker. Information \mathfrak{I} might include specialized knowledge of local conditions and past performance of the wind energy generator, information from local monitoring

of wind and atmospheric conditions; knowledge of the season and seasonal trends, weather forecasts from numerical weather prediction (NWP) models, and so on.

An important component of \mathfrak{I} is information about how to combine these various sources of information: we call these “information consolidation strategies”. In this paper, we compare six different strategies which involve information that we expect would be generally available to wind energy decision makers. These different strategies lead to different probability distributions and therefore potentially different production decisions. In Section 4 we use a simulation method to evaluate the consequences of these decisions in terms of the total income produced by the wind energy generation facility.

The strategies are described below.

3.1. Climatology

We assume that the wind generator has a record of past energy output organized by time. Since energy output is highly correlated with wind velocity (cubed), a finely sampled time series of wind velocity might provide the requisite information. Such a historical record is considered climatological information by meteorologists since it is conditioned on time and not on the measured state of the atmosphere. We therefore adopt the term “climatology” to describe this strategy. Typically there is a regular diurnal variation in the strength of the wind, so we normalized the power production with respect to this diurnal cycle. The normalized power production is

$$W_*(\text{day}, \text{hour}) = \frac{W(\text{day}, \text{hour})}{\langle W(\text{hour}) \rangle}, \quad (7)$$

where $W(\text{day}, \text{hour})$ is the power production on a given day at a given hour and $\langle W(\text{hour}) \rangle$ is the power production at the given hour averaged over the data set.

For this report, we generate $P(E_a | \mathfrak{I}(\text{climate}))$ by fitting a parametric model to the historical data of W_* , where the parameters vary according to calendar date. This model assumes that, for a given day of the year, the normalized, average half-hourly production is a random variable with a gamma distribution given by

$$\rho(W_*) = \frac{\lambda^r}{\Gamma(r)} W_*^{r-1} e^{-\lambda W_*}, W_* > 0. \quad (8)$$

For the CLRC ERU site used in our study, the parameter values used were

$$r = 0.6 \quad (9)$$

and

$$\lambda = -0.095 \sin\left(\frac{2\pi d}{365.25}\right) - 0.51 \cos\left(\frac{2\pi d}{365.25}\right) + 0.73, \quad (10)$$

where d is the day of the year. Fig. 2 compares the distribution of half-hourly production given by Eqs. (8) and (10) with the histogram of half-hourly productions for six days of the year. In each case the histogram was constructed using a two-




Fig. 2. A comparison of the distribution of mean half-hourly power production given by the climatological model (thick line) [Eqs. (8) and (10)] and the actual distribution of synthetic half-hourly production constructed using the high-frequency ERU wind data (thin line). The actual distributions were constructed using the two-week period centred on the date in question.

week period centered on the date in question. The value of W_* was then converted to a value of E_a by multiplying by $\langle W(\text{hour}) \rangle$ corresponding to the time of day being predicted.

The simple conditioning on calendar date reflects the relative paucity of historical detailed wind measurements available to us. With a richer database, the conditioning might be extended to include the prevailing wind direction, and so on.

3.2. Persistence

As is well known, tomorrow's weather tends to be like today's. Extending this principle to wind generation, we can set $P(E_a|\mathfrak{F}(\text{persistence}))$ to be based on the recent few days' past history of power production. Time series analysis techniques for this are described in Ref. [18].

In this report, we construct $P(E_a|\mathfrak{F}(\text{persistence}))$ as the empirical distribution of diurnally normalized half-hourly energy production, W_* , for the 24 h prior to the production decision multiplied by the value of $\langle W(\text{hour}) \rangle$ corresponding to the time of day being predicted.

3.3. ECMWF best-guess forecast

The ECMWF (T319) NWP model forecasts atmospheric conditions at a set of grid points spaced horizontally by approximately 60 km. We selected the surface wind speed at the nearest station for which the ECMWF separately archives its forecasts. Systematic biases between the wind at the simulated wind farm and at the station were corrected and the corrected wind was converted to a power production using a turbine output profile described in Section 4. The ECMWF forecasts are issued every 12 h although for this study we only used the forecasts issued for midday; for each half-hour interval the forecast was taken to be that of the closest-in-time ECMWF forecast.

The high-resolution ECMWF forecast is a single point, not a probability distribution. For the purposes of decision making, we set the target production level E_c to be the forecast production level. This simple scheme corresponds to a step-shaped $P(E_a|\mathfrak{F}(\text{ECMWF hires}))$. It is easy to imagine other ways of translating the best-guess forecast to a decision, for instance adjusting it based on past local experience. One such method will be described below that conditions climate data on the best-guess forecast.

3.4. ECMWF ensemble prediction system

In addition to generating a traditional single, high-resolution, best-guess forecast, ECMWF also run an *ensemble prediction system* (EPS). This consists of an ensemble of 51 forecasts made at a lower resolution. The leading ensemble member is initialized with an initial condition corresponding to the one used to initialize the high-resolution forecast, while the other 50 members have initial conditions obtained by perturbing this condition. The ensemble prediction system is an attempt to quantify the state-dependent predictability of the atmosphere [9]. This state-dependent predictability can manifest itself as a sensitivity to initial conditions popularly known as the "butterfly effect".

As with the best-guess forecast, the wind velocity ensemble from the ECMWF model was translated to power output using the function described in Section 4.

3.5. Climate conditioned on the ECMWF high-resolution forecast

Rather than treating the ECMWF output as an actual forecast of physically observable variables, we treat it as a concise description of the state of the atmosphere and we use it as a conditioning variable into the climate data.

For this paper, using a database of historical NWP forecasts, we searched the historical forecasts for the 10 days whose historical forecasts were closest to the current ECMWF high-resolution forecast (smallest absolute difference). The 480 points of diurnally normalized half-hourly wind production for these 10 days were extracted from the climate record and multiplied by $\langle W(\text{hour}) \rangle$ to produce $P(E_a | \mathfrak{I}(\text{ECMWFhires}_{\text{climate}}))$.

3.6. Climate conditioned on the ECMWF ensemble forecast

This is quite similar to climate conditioned on the best-guess forecast. Each ensemble of windspeed forecasts was converted into a 3-vector, \mathbf{F} , by ordering the ensemble members, V_i , and then using the 5th, 25th and 46th members of the ensemble as coordinates

$$\mathbf{F} = [V_5, V_{25}, V_{46}]. \quad (11)$$

The coordinates were thus the 10th percentile, median and 90th percentile of the 51 member ensemble. The historical forecasts were then searched to find the 10 days that had the smallest Euclidean distance between their own \mathbf{F} -vector and the \mathbf{F} -vector of the current forecast. These 480 diurnally normalized half-hour points were multiplied by $\langle W(\text{hour}) \rangle$ to generate a 480-step form for $P(E_a | \mathfrak{I}(\text{ECMWFeps}_{\text{climate}}))$.

4. Evaluating the strategies

Each information consolidation strategy leads to a probability distribution $P(E_a | \mathfrak{I})$. The six strategies we consider here have \mathfrak{I} corresponding to one of $\mathfrak{I}(\text{climate})$, $\mathfrak{I}(\text{persistence})$, $\mathfrak{I}(\text{ECMWFhires})$, $\mathfrak{I}(\text{ECMWFeps})$, $\mathfrak{I}(\text{ECMWFhires}_{\text{climate}})$ or $\mathfrak{I}(\text{ECMWFeps}_{\text{climate}})$. The probability distribution, which reflects the generator's knowledge, determines the optimal level of E_c for each half-hour contract by setting $P(E_c | \mathfrak{I}) = P_c / \langle P_s \rangle$. The income or loss that results from the chosen E_c depends on the actual energy production E_a and the spot price P_s for the half-hour.

If we know the true probability distribution for each half-hour, $P(E_a | \mathfrak{I}(\text{Omniscience}))$, we could readily compute the difference between the decision made based on each strategy's $P(E_a | \mathfrak{I})$ and the omniscient-optimal decision based on $P(E_a | \mathfrak{I}(\text{Omniscience}))$. Pricing out the cost of the strategy's decision using P_c and P_s would give the relative value of each strategy.

Since we lack omniscience, we take a different approach: a simulation calculation where we use actual measurements of wind velocity at the wind energy production site and actual spot prices to calculate the income generated by each strategy. This can be thought of, somewhat abstractly, as a kind of Monte Carlo simulation using




Fig. 3. A one-week sample of the high-frequency wind data from the CLRC Rutherford Appleton Laboratory Energy Research Unit in Oxfordshire. The data were taken at a height of 10 m and sampled at 1-m intervals.

$P(E_a|\mathfrak{I}(\text{Omniscience}))$, where we use the actual record during the simulation period in place of truly random samples from the unknown $P(E_a|\mathfrak{I}(\text{Omniscience}))$.

We simulated a wind energy production site using high sampling frequency wind speed data from the CLRC Rutherford Appleton Laboratory, Energy Research Unit (ERU) test site in Oxfordshire, UK ($51^{\circ}34'9''\text{N}$, $1^{\circ}19'16''\text{W}$). This data was measured at a height of 7 m and at intervals of 1 minute. Fig. 3 shows a sample of these data. The wind speed at the site showed a clear diurnal cycle, as revealed in Fig. 4. To convert the high-frequency wind speed data into a wind turbine output level we used




Fig. 4. The average diurnal cycle in the windspeed at the ERU test site.

the idealized power production profile shown in Fig. 5. We thus generated a synthetic time series of half-hourly mean wind power production at the ERU site. This time series was normalized by the diurnal cycle in production.

We used the ECMWF forecasts made for Heathrow Airport (LHR) (51°29'47"N, 0°26'10"W). Heathrow Airport is approximately 100 km from the ERU site.

As a proxy for the electricity spot price, P_s , we used the half-hourly System Marginal Price (SMP) for the UK, which we obtained from ESIS Ltd. This is not a spot price but the price charged by the most expensive generator for production during that half-hour period. The period of simulation was 1999–2000. Fig. 6 shows samples of UK system demand and the corresponding System Marginal Price during winter and summer.

For each half-hour period, a sample of 480 half-hour productions was generated. For the climatology this sample was generated by picking 480 productions distributed according to Eq. (8). For persistence, the 480 samples were picked with replacement from the most recent 48 half-h. In the case of the ECMWF high-resolution forecast all of the 480 samples were identical to this forecast, while when the ensemble forecast was used the 480 samples were picked with replacement from the 51 member ensemble. For the climate conditioned on the ECMWF forecasts, the 10 days that had forecasts closest to the current forecast were used as the sample of 480 half-hour productions. Each of the 480 samples of power production was then paired with a spot price generated by the stochastic spot price model described in the Appendix. For a given value of E_c , the income averaged over the 480 pairs of (E_a, P_s) could then be calculated. The value of E_c that maximized this average when $P_c = 10$ GBP/MWh was then found using a brent optimization algorithm [13].³ Once




Fig. 5. The profile of wind turbine output as a function of wind speed that was used for the study.

³ We note that since in this case predictions of P_s were independent of predictions of E_a , only the value of $\langle P_s \rangle$ was actually required. However, the method we used of sampling the joint probability distribution $p(E_a, P_s)$ can be used when predictions of E_a and P_s are not independent. Since the weather




Fig. 6. Samples of UK system demand and the system marginal price for January and July 1999. In the demand panels the thick line is the actual demand and the thin line is the demand predicted using Eq. (A2) described in the Appendix.

this value of E_c had been selected, the actual income that would have been made for this half-hour was calculated using the synthetic production data from the ERU site for E_a , the UK SMP for P_s and $P_c = 10$ GBP/MWh. Thus for each strategy a time series of half-hourly net income was generated. For the strategy based on persistence and the strategies based on ECMWF forecast products the simulations were

can influence the spot price of electricity as well as wind power production, some dependence between E_a and P_s would not unreasonable.

repeated for lead times of 1 to 10 days.⁴ The time series of half-hourly income were then converted to time series of net weekly income. A bootstrap technique was used to estimate the mean and variance of the annual income [3]. Fifty estimates of the total annual income were made by picking 52 weeks with replacement from the time series and summing them. This sample of 50 total annual incomes was then used to estimate the mean and variance of this quantity.

5. Results

Fig. 7 compares the mean weekly net income obtained by generators using strategies based on the different wind forecasting methods. Persistence forecasting underperforms relative to climatology even at a lead time of one day (Fig. 7a).

All of the forecasts based on ECMWF forecast products outperform climatology at lead times of up to 5 days (Fig. 7b–e). Using the climate conditioned on the ECMWF forecasts to estimate the distribution of half-hourly production yields better results than treating the ECMWF forecasts in a deterministic manner. Fig. 7f compares the mean weekly income of a generator using the climate conditioned on the ECMWF ensembles and a generator using the climate conditioned on the ECMWF best-guess forecast. The ensemble user can expect a higher mean income at all lead times beyond 1 day and obtains the largest relative advantage at a lead time of 4 days. This is due to the performance of the deterministic user decaying faster with lead time than the ensemble user. The ensemble conditioned climate shows value at a lead time of 6 days. Fig. 8 compares the net daily income of a generator using the climatological forecasts with that of a generator using the forecast generated by conditioning the climate on the 4-day ensemble forecast.

We repeated the simulations using a simpler spot price model in the decision-making process. The simpler model assumed a static log normal (demand-invariant) distribution for spot price. We also repeated the simulations using a fixed contract price of $P_c = 20$ GBP/MWh and also a variable contract price of $P_c = 0.5\langle P_s \rangle$. In all cases the relative performance of the strategies remained unchanged.

6. Discussion

Our results indicate that, in the context of the market model we used, the ECMWF forecasts could be very valuable to wind energy producers. The single best-guess forecasts provide most of the improvement over climatology. These forecasts double net income at short lead times and even at 3 days income was boosted by 75%. Use of the ensemble forecasts raised net income by up to 20% more and extended the useful range of the forecasts by an extra day.

⁴ For all lead times, a single contract price of $P_c = 10$ GBP/MWh was used. In a real market one would expect P_c to move closer to P_s as the lead time decreases.

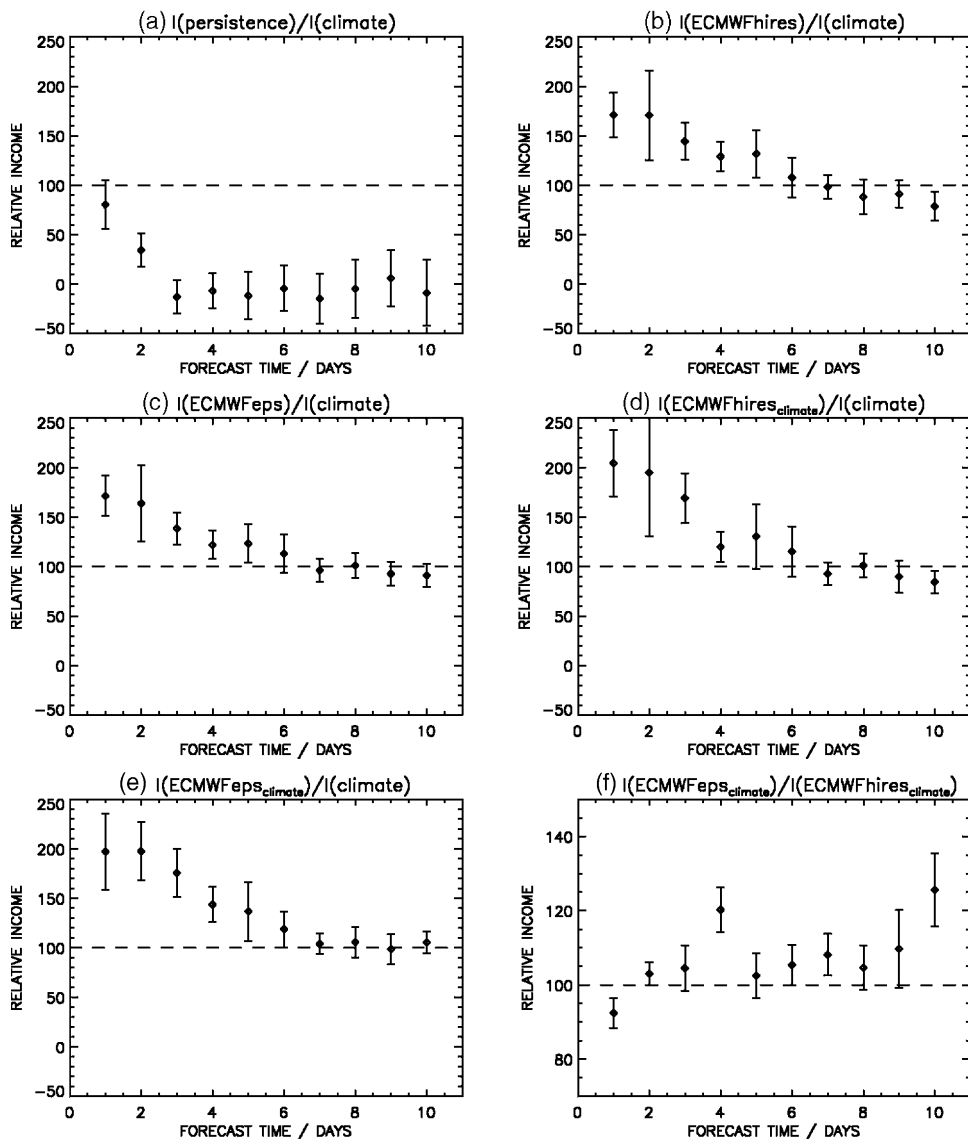


Fig. 7. Relative mean weekly net income as a function of forecast lead time for five different wind forecasting methods described in the text. Panels (a)–(e) are normalized with respect to the profits of a generator using the climatological wind model. Panel (f) compares the ensemble forecasts with the climate conditioned on the best-guess forecast. The error bars are the standard errors obtained by resampling with replacement [3].




Fig. 8. Net daily income over the simulation period. Panel (a) is for a generator using the climatological wind forecasts and (b) is for a generator using the climate conditioned on the 4-day ECMWF ensemble forecast.

Fig. 8 indicates that the unseasonably calm period during late January 2000 caused the climatology user to suffer losses; however, even the ensemble user generally outperforms the climatology user, and not just during this calm period. The 4-day ensemble forecast user has a higher daily income on 60% of the days and a higher net weekly income on 80% of the weeks. The medium-range forecast users raise their income by both reducing losses due to underproduction and reducing the opportunity cost associated with underpromising when production is unseasonably high.

Our results assume that, without the forecasts, the producer would use the climatological model. This climatological model was actually constructed by fitting to both the historical data and the data used for the simulation period. This use of in-sample data to construct the climatology could have biased the results in favour of climatology. Therefore, the value of the medium range forecasts may be even greater than our study suggests.

In this paper we made several simplifying assumptions. The assumptions were not necessary but including additional complexities would make the study less general.

1. The relationship between wind speed and power output was idealized. In practice this would be more complicated, particularly in the case of multi-turbine wind farms where wake effects would vary depending on the wind direction.

2. It was assumed that producers can write their contracts for specific half-hours at arbitrary times before delivery. This is not the case in all contemporary electricity markets. For example, Nordpool allows only weekly averaged “contracts-for-differences” at lead times longer than the day before delivery.
3. The convergence between P_c and P_s with shortening lead time was completely ignored. This means that while direct comparisons can be made between different strategies at the same lead time, comparisons involving different lead times are not meaningful.
4. The impact of weather on electricity demand could be incorporated into the spot price model.
5. The risk tolerance of the producer could be allowed for when optimizing E_c .

All of these effects would have to be investigated to ascertain the value of forecasts to a specific site in a specific market. In addition, it should be remembered that the forecasts for Heathrow Airport were used. Using the forecasts corresponding to the ERU site may improve the value of the forecasts further.

Acknowledgements

The authors would like to thank the CLRC Rutherford Appleton Laboratory's Energy Research Unit for providing the data for this study. They would also like to thank ECMWF for providing historical forecasts, and Tim Palmer, Tom Hamill and Jim Halliday for useful discussions. L.A.S. and M.S.R. were supported by ONR DRI grant N00014-99-1-0056.

Appendix. The spot price model

In order to capture the seasonal dependence of spot price and other time correlations in spot price, we constructed models of the seasonal dependence of demand and of the dependence of spot price on demand. These models are described here. However, quite similar overall results of the strategy values were obtained using a simple model where spot price was random but independent of time.

For the purposes of this study it was assumed that, for a given day, the electricity spot price and the wind energy production are independent. That is, if $p(E_a, P_s)$ is the joint probability conditioned on the day, then $p(E_a, P_s) = p(E_a)p(P_s)$. This is not necessarily true since unseasonally high winds would be associated with high values of E_a and they may also increase P_s by increasing the demand for electricity. A log-normal model for P_s was used

$$P_s = e^X, X: \mathbb{N}(\mu, \sigma^2), \mu = 1.35 + (9.2 \times 10^{-5})D, \sigma = 0.3 \quad (A1)$$

where D is the system demand in megawatts and P_s is in GBP/MWh. Fig. A1 shows the actual and model distribution of system marginal prices associated with four different demands. To condition the model for P_s , an estimate of system demand is




Fig. A1 A comparison of the actual distribution of UK system marginal prices (thin line) and the demand conditioned probability distribution given by Eq. (A1) (thick line) for four different values of the system demand.

needed. The demand model we used estimated demand based solely on the day of the year and the time of day. It did not include demand from previous days or weather information, both of which can be included in medium-range demand forecasting models [14,21]

$$D = a_0 \sin(2\pi T_1) + b_0 \cos(2\pi T_1) + \sum_{n=1}^3 [a_n \sin(2\pi T_2) + b_n \cos(2\pi T_2)] + c, T_1 = d/365.25, T_2 = h/24, \quad (\text{A2})$$

where d is the number of days after 1 January and h is the number of hours since midnight. Two sets of parameters were used; one set for weekdays and one set for weekends and public holidays. Fig. 6 compares the predictions made using Eq. (A1) with the actual system demand.

References

- [1] Anderson JL. A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate* 1996;9:1518–30.
- [2] Brier GW. Verification of forecasts expressed in terms of probabilities. *Mon Weath Rev* 1950;78:1–3.
- [3] Efron B, Tibshirani R. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statist. Sci.* 1986;1:54–77.
- [4] Epstein E. A scoring system for probability forecasts of ranked categories. *J Appl Meteorol* 1969;8:985–7.

- [5] Hamill TM, Colucci SJ. Random and systematic error in NMC's short-range eta ensembles. In: Preprints of 13th Conference on Probability and Statistics in the Atmospheric Sciences (San Francisco) American Meteorological Society. 1996. p. 51–6.
- [6] Hamill TM. Interpretation of rank histograms for verifying ensemble forecasts. *Mon Weath Rev* 2001;129:550–60.
- [7] Katz RW, Murphy AH. Forecast value: prototype decision-making models. In: Katz RW, Murphy AH, editors. *Economic value of weather and climate forecasts*. Cambridge: Cambridge University Press; 1997. p. 183–217.
- [8] Mason IB. A model for assessment of weather forecasts. *Aust Meteorol Mag* 1982;30:291–303.
- [9] Molteni F, Buizza R, Palmer TN, Petroliagis T. The ECMWF ensemble prediction system: methodology and validation. *Quart J Roy Meteorol Soc* 1996;122:73–119.
- [10] Murphy AH, Daan H. Forecast evaluation. In: Murphy AH, Katz RW, editors. *Probability, statistics and decision making in the atmospheric sciences*. Boulder, CO: Westview Press; 1985. p. 379–437.
- [11] Murphy AH. A note on the ranked probability score. *J Appl Meteorol* 1971;10:155–6.
- [13] Press WH, Teukolsky SA, Vetterling WT, Flannery BP. *Numerical recipes in FORTRAN: the art of scientific computing.*, 2nd ed. Cambridge: Cambridge University Press, 1992.
- [14] Ramanathan R, Engle R, Granger CWJ, Vahid-Araghi F, Brace C. Short-run forecasts of electricity loads and peaks. *Int J Forecast* 1997;13:161–74.
- [15] Richardson DS. Skill and relative economic value of the ECMWF ensemble prediction system. *Quart J Royal Meteorol Soc* 2000;126:649–67.
- [16] Roulston MS, Smith LA. Evaluating probabilistic measures using information theory. *Mon. Weath. Rev.* 2001;130:1653–60.
- [17] Smith LA, Roulston MS, Hardenberg J. End to end ensemble forecasting: towards evaluating the economic value of the ensemble prediction system. In: *ECMWF Tech. Memo 336*. 2001.
- [18] Sfetsos A. A comparison of various forecasting techniques applied to mean hourly wind speed time series. *Renewable Energy* 2000;21:23–35.
- [19] Swets JA. The relative operating characteristic in psychology. *Science* 1973;182:990–9.
- [20] Talagrand O, Vautard R, Strauss B. Evaluation of probabilistic prediction systems. In: *Proceedings of the ECMWF Workshop on Predictability*, 20–22 October. Reading, UK: ECMWF; 1997. p. 1–25.
- [21] Taylor JW, Buizza R. Using weather ensemble predictions in electricity demand forecasting. *Int J Forecasting* 2002;in press.