Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

A non-Gaussian Ornstein-Uhlenbeck model for pricing wind power futures

Benth, Fred Espen; Pircalabu, Anca

Published in:
Applied Mathematical Finance

DOI (link to publication from Publisher):
[10.1080/1350486X.2018.1438904](https://doi.org/10.1080/1350486X.2018.1438904)
[10.2139/ssrn.2979341](https://doi.org/10.2139/ssrn.2979341)

Publication date:
2018

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Benth, F. E., & Pircalabu, A. (2018). A non-Gaussian Ornstein-Uhlenbeck model for pricing wind power futures. *Applied Mathematical Finance*, 25(1), 36-65. <https://doi.org/10.1080/1350486X.2018.1438904>, <https://doi.org/10.2139/ssrn.2979341>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Manuscript

A non-Gaussian Ornstein-Uhlenbeck model for pricing wind power futures

F. E. BENTH* & A. PIRCALABU^{†‡}

*Department of Mathematics, University of Oslo, PO Box 1053 Blindern, N-0316 Oslo, Norway, fredb@math.uio.no, [†]Department of Mathematical Sciences, Aalborg University, Fredrik Bajers Vej 7G, 9220 Aalborg Øst, Denmark, anca@math.aau.dk, [‡]Quantitative Analytics, Neas Energy, Skelagervej 1, 9000 Aalborg, Denmark.

(14 December 2017)

ABSTRACT *The recent introduction of wind power futures written on the German wind power production index has brought with it new interesting challenges in terms of modeling and pricing. Some particularities of this product are the strong seasonal component embedded in the underlying, the fact that the wind index is bounded from both above and below, and also that the futures are settled against a synthetically generated spot index. Here, we consider the non-Gaussian Ornstein-Uhlenbeck type processes proposed by Barndorff-Nielsen and Shephard (2001) in the context of modeling the wind power production index. We discuss the properties of the model and estimation of the model parameters. Further, the model allows for an analytical formula for pricing wind power futures. We provide an empirical study, where the model is calibrated to 37 years of German wind power production index that is synthetically generated assuming a constant level of installed capacity. Also, based on one year of observed prices for wind power futures with different delivery periods, we study the market price of risk. Generally, we find a negative risk premium whose magnitude decreases as the length of the delivery period increases. To further demonstrate the benefits of our proposed model, we address the pricing of European options written on wind power futures, which can be achieved through Fourier techniques.*

KEY WORDS: wind power futures, weather derivatives, Ornstein-Uhlenbeck process, market price of risk

1. Introduction

Following the significant expansion in wind turbine installations that some European countries have experienced over the past years, the demand for financial instruments that can be used to address the problem of volumetric risk in wind power generation has grown. This has led to the launch of a standardized product written on the wind power production index, namely the so-called wind power futures (or wind index futures). Currently, wind power futures can be traded on NASDAQ OMX and the European Energy Exchange (EEX) on the German wind power production index. The index is obtained by measuring the German wind power generation relative to the available installed capacity; hence, the index has a lower bound of 0 and an upper bound of 1, corresponding to a 0% and a 100% wind power utilization,

Correspondence Address: A. Pircalabu, Department of Mathematical Sciences, Aalborg University, Fredrik Bajers Vej 7G, 9220 Aalborg Øst, Denmark. Email: anca@math.aau.dk.

respectively.

To clarify the payoff structure of wind power futures, let us denote by $F(t, T)$ the wind power futures price at time t and delivery during day T , with $0 \leq t < T$ and $0 \leq F(t, T) \leq 1$. Further, let $P(T)$ be the wind index measured at day T . Then, a long position in a wind power futures contract entered at time $t \leq T$ for delivery at T yields the payoff

$$24(P(T) - F(t, T)) \cdot x,$$

where 24 denotes the usual number of hours in a day and x denotes a known fixed tick size. For the wind power futures traded at NASDAQ OMX and EEX, $x = 100$ EUR. Moreover, the futures are settled against an externally provided spot index $P(T)$, which is synthetically generated based on weather data and an individual power curve for every grid point in Germany.

Natural sellers of wind power futures are the wind power producers and companies with considerable wind park portfolios, as they are interested in protection against the low wind scenarios, which are likely to lower revenues. Although one could argue that day-ahead electricity prices tend to increase in times of low wind, wind power generators usually receive a fixed price per generated unit of electricity, and do not participate in the wholesale market themselves. Hence, volumetric risk is the only risk source left to be addressed, and wind power futures can be an obvious tool for stabilizing the revenue of the wind power generators. Typical buyers are conventional power plants acting in e.g. the day-ahead market, whose profitability drops in times of high wind due to the negative relation between wind power production and spot electricity prices.

In this paper, we propose a non-Gaussian Ornstein-Uhlenbeck process in the spirit of Barndorff-Nielsen and Shephard (2001) to model the wind power production index. The model is very straightforward, allowing for an easy estimation of the parameters and analytical pricing of wind power futures, with the latter facilitating the study of the market price of risk. Based on one year of observed German wind power futures curves, we perform an empirical analysis of the risk premia in this newly established market.

Wind power futures are characterized as weather derivatives, and fall in this category together with derivatives written on temperature, rainfall, snowfall, humidity, etc. While the existing literature on temperature derivatives is extensive and broad in terms of modeling approach (see e.g. Davis (2001), Brody *et al.* (2002), Cao and Wei (2004), Campbell and Diebold (2005), Platen and West (2005), Härdle and López Cabrera (2012) and Benth and Šaltytė Benth (2011)), literature related specifically to wind derivatives is very scarce. To the best of our knowledge, the first study concerned with the pricing of wind derivatives is that of Benth and Šaltytė Benth (2009), which was motivated by the introduction of futures and options on wind speed indexes at different wind farm locations in the US back in 2007. However, trade in these products never really picked up, explaining perhaps the scarcity of related studies.

Almost ten years after the first attempt to establish a market for wind derivatives, the introduction of the German wind power futures on NASDAQ OMX and EEX awakens interest again. The study of Gersema and Wozabal (2017) is the first to provide a thorough introduction to the German wind power futures market, the market players and their risks. Further, Gersema and Wozabal (2017) propose an equilibrium pricing model, and based on different case studies they conclude that

a negative risk premium is to be expected in wind power futures markets. A second related study concerning wind power futures is that of Pircalabu and Jung (2017), where the authors focus on the hedging benefits of wind power futures in the context of energy trading companies entering into long-term agreements with wind power generators, where the fluctuating wind power production is bought at a pre-determined fixed price. Here, wind power futures are not the main focus of the analysis, and they are thus treated on a conceptual basis, disregarding some practical aspects concerning the data foundation in their pricing application. In this paper, we shall address this aspect in detail, highlighting its importance.

The paper is structured as follows: In Section 2, we present the data and comment on key features as to motivate the model choice. In Section 3, we introduce the model for the wind power production index and provide an empirical study where the model is applied to German data. Analytical futures prices are derived in Section 4, and based on one year of market prices for wind power futures with different delivery periods, the market price of risk is studied. In Section 5, we elaborate on further applications of the proposed model in derivatives valuation. Section 6 concludes.

2. Data presentation

Since wind power futures are only traded on the German wind power index at the moment, the empirical analysis performed in this paper is based on German data. We consider a time series of daily wind power production indexes for the German market, which was synthetically constructed by MeteoGroup for a period of 37 years (1 January 1979 to 31 December 2015). The synthetic index is displayed in Fig. 1, and measures how the utilization of installed wind power capacity would have looked like in the German market zone in the past, conditional on the *present level* of available capacity and geographical location of wind turbines. Specifically, the *present level* we consider here corresponds to September 2016. To construct such an index, a bottom-up approach was implemented based on historical weather data and power curves. Clearly, since the wind index measures the wind power production relative to the installed capacity, it must be bounded between zero and one. For the data in Fig. 1, the lowest and highest values recorded are 0.35% and 83.05%, respectively.

Figure 1. Index constructed based on the total installed wind power capacity observed in September 2016.

In the context of pricing wind power futures, which is the main focus of the present paper, we argue that fitting a model to the type of data in Fig. 1 seems much more reasonable than considering the historical evolution of the wind index. This is an essential point, since the wind power futures price today is clearly not influenced

by how the available installed capacity evolved over time in Germany, but rather on the present and ideally the future installed capacity level.

2.1 Seasonality

Aside from the wind index being bounded on $[0,1]$, another key feature is the yearly seasonality we observe in the data illustrated in Fig. 1. To emphasize the annual pattern, we complement the time series plot in Fig. 1 with the empirical autocorrelation function of the index in Fig. 2(a). Following the related literature (see e.g. Benth *et al.* (2008), Härdle and López Cabrera (2012), and Benth and Šaltytė Benth (2011)), the yearly seasonality can be addressed by the following seasonality function:

$$\Lambda(t) = a_1 + a_2 \sin(2\pi t/365) + a_3 \cos(2\pi t/365). \quad (1)$$

Fitting this function to the wind index by ordinary least squares yields the parameter estimates reported in Table 1. In Fig. 2(b), the wind power production index is plotted together with the fitted seasonal function. For better clarity, we display a snapshot of the last 10 years, i.e., from year 2006 to 2015.

Table 1. OLS estimates for the parameters of the seasonal function.

	Estimate	Standard error
$\hat{a}_1$	0.2164	0.0014
$\hat{a}_2$	0.0102	0.0020
$\hat{a}_3$	0.0839	0.0020

Figure 2. Empirical autocorrelation function of the synthetic wind power production index and fitted seasonal function.

3. A model for the wind power production index

Motivated by the two key features of the wind index enhanced in Section 2, i.e., boundedness on $[0,1]$ and yearly seasonality, we specify a model for the wind power production index as follows. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space with a

filtration $\{\mathcal{F}_t\}_{t \geq 0}$ satisfying the usual conditions. We denote by P the wind power production index obtained by measuring the wind power production relative to the available installed capacity, implying that $P(t) \in [0, 1]$ for all t .

We define $P(t)$ as

$$P(t) = \Lambda(t) \exp(-X(t)), \quad (2)$$

where $\Lambda(t)$ describes the deterministic seasonal component of the wind power production index and $X(t)$ follows a non-Gaussian Ornstein-Uhlenbeck process as in the stochastic volatility model proposed by Barndorff-Nielsen and Shephard (2001). Specifically,

$$dX(t) = \alpha(\mu - X(t))dt + dL(t), \quad (3)$$

with L being a driftless subordinator, and $\mu > 0$, $\alpha > 0$ denoting two constants. From standard theory, the solution of the Ornstein-Uhlenbeck process is

$$X(t) = X(0)e^{-\alpha t} + \mu(1 - e^{-\alpha t}) + \int_0^t e^{-\alpha(t-s)} dL(s),$$

where $X(0) = \ln(\Lambda(0)/P(0))$. The constant μ is connected to $\Lambda(t)$, and its purpose is to ensure that $P(t)$ never exceeds 1. In order to elaborate on this, we include the following Proposition regarding the stationarity of $X(t)$.

Proposition 3.1. Let $\ell(dz)$ denote the Lévy measure corresponding to the Lévy process $L(t)$. If

$$\int_{|z|>2} \ln |z| \ell(dz) < \infty,$$

then $X(t)$ given by the Ornstein-Uhlenbeck process in Eq. (3) has a limiting distribution. The stationary solution of $X(t)$ is

$$X(t) = \mu + \int_{-\infty}^t e^{-\alpha(t-s)} dL(s),$$

where L here is a two-sided Lévy process.

We refer to Sato (1999), Thm. 17.5, for a proof and more details on this result. Regarding the stationary solution X , we refer to an extensive discussion in Basse-O'Connor *et. al* (2014).

Returning to the connection between μ and $\Lambda(t)$, let $M = \max(\Lambda(t))$. Then, we must have that

$$\begin{aligned} \max(P(t)) &= \max(\Lambda(t) \exp(-X(t))) \\ &\leq M \exp(-\min(X(t))). \end{aligned}$$

Owing to L being a subordinator, it follows from the stationary solution in Proposition 3.1 that $X(t) \geq \mu$. Further, since we also have that $0 \leq P(t) \leq 1$, we choose

μ such that $M \exp(-\mu) = 1$. Thus, we let

$$\mu = \ln M, \quad (4)$$

and obtain an exact upper bound of 1 as a possible case. This way of introducing seasonality in the model has its advantages and disadvantages, and we refer to Appendix A for a detailed discussion on the subject.

Next, we state the limiting distribution of $X(t) - \mu$ for a specific case, since this will be used in our empirical study.

Proposition 3.2. If $L(t)$ is a compound Poisson process with exponentially distributed jumps,

$$L(t) = \sum_{k=1}^{N(t)} J_k, \quad (5)$$

where $N(t)$ is a Poisson process with frequency λ and J_k are independent identically distributed exponential random variables with density function

$$f_J(x) = \kappa e^{-\kappa x}, \quad (6)$$

then the limiting distribution of $X(t) - \mu$, where $X(t)$ evolves according to Eq. (3), is the Gamma distribution with density function given by

$$f_\Gamma(x) = \frac{\kappa^{\lambda/\alpha} x^{\lambda/\alpha-1} e^{-\kappa x}}{\Gamma(\frac{\lambda}{\alpha})}. \quad (7)$$

Proof. See Appendix B.1.

While the model proposed in this section captures key features of the wind index, there are other alternatives when it comes to modeling data with range $[0,1]$. In particular, we mention the Jacobi processes. A Jacobi process, which is in fact an extension of the Heston model, will have values in any desirable positive interval [Ackerer *et. al* (2017)]. In our case, we could consider a process of the type

$$dP(t) = -a(P(t) - b)dt + \sqrt{cP(t)(1 - P(t))}dW(t),$$

where $a > 0$, $c > 0$, $0 < b < 1$ and W denotes a Brownian motion.

On one hand, the Jacobi approach is simpler compared to our proposed model in that the wind index is modeled directly, and the Lévy process is replaced by a Brownian motion. On the other hand, since we let $\ln(\Lambda(t)/P(t))$ be an Ornstein-Uhlenbeck, our approach is advantageous from a calibration perspective. In fact, estimation of model parameters for the Jacobi process is not straightforward cf. Gouriéroux and Valéry (2002). Furthermore, the marginal distribution of the Jacobi process is fixed to the beta distribution, whereas our model allows for great flexibility in choosing marginal distributions. In terms of derivatives pricing, both models have their advantages when it comes to the pricing of wind power futures. For the Jacobi process, making use of its polynomial property could result in simple (possibly explicit) expressions for the futures price. Regarding our proposed process, explicit

pricing formulas for wind power futures are attainable, as we shall illustrate later in the paper. Unlike our model however, it is unclear how a measure change is included in the Jacobi model as to preserve the Jacobi-structure. In light of the discussion above, we favor the model in Eqs. (2)-(3), and shall not pursue the Jacobi processes in the present paper. Nevertheless, we stress that the Jacobi approach is an interesting and unexplored alternative for modeling the wind power production index.

3.1 An empirical analysis on German wind index data

In this section, we turn to the empirical study of the German wind index time series in Fig. 1. Recalling that the seasonal function entering Eq. (2) has already been estimated in Section 2.1, an estimate for μ immediately follows from Eq. (4). We obtain $\hat{M} = 0.3009$, implying that

$$\hat{\mu} = -1.2010. \quad (8)$$

Using the expression for $P(t)$ in Eq. (2), the variable $X(t) - \mu$ is then constructed by

$$X(t) - \mu = - \left(\ln \left(\frac{P(t)}{\Lambda(t)} \right) + \mu \right). \quad (9)$$

According to Appendix A, our way of incorporating seasonality in the model introduces the potential of having

$$X(t) - \mu < 0. \quad (10)$$

When considering the time series $X(t) - \mu$, we do indeed observe negative values; however, the percentage of negative data points is very low, corresponding to 0.95%, which we find acceptable.

Next, we consider the parameter α entering the dynamics of the $X(t)$ process cf. Eq. (3), and note that

$$X(t+1) - \mu = e^{-\alpha}(X(t) - \mu) + \int_t^{t+1} e^{-\alpha(t+1-s)} dL(s). \quad (11)$$

Clearly, it follows from Eq. (11) that α can be obtained by fitting an AR(1) model to $X(t) - \mu$, and by using the relation $\phi = \exp(-\alpha)$, where ϕ denotes the slope coefficient in the AR(1). However, this procedure requires residuals to be normally distributed, which is not the case here. Recalling that the lag s correlation between observations s periods apart can be expressed as

$$\text{Corr}(X(t+s), X(t)) = \phi^s = e^{-\alpha s},$$

we fit instead the function $\exp(-\alpha t)$ to the sample autocorrelation of $X(t) - \mu$ using

nonlinear least squares¹, and obtain

$$\hat{\alpha} = 0.5455 \quad (12)$$

with a standard error of 0.0011. For comparison purposes, we provide in Appendix C detailed results from implementing the AR(1)-estimation approach, including a residual analysis.

To get an idea of the goodness-of-fit of the proposed exponential function, we plot in Fig. 3 the empirical autocorrelation function together with the fitted exponential. The fit is satisfactory, capturing rather well the sample autocorrelations at the first lags, which are also the most significant. We do however note that the fitted exponential drops to zero slightly quicker than the sample autocorrelation does.

Figure 3. Fitted exponential to the empirical autocorrelation function of $X(t) - \mu$.

The remaining part of the fitting procedure relates to estimating the parameters of the stationary distribution of $X(t) - \mu$, i.e. the Gamma distribution cf. Proposition 3.2. The choice of a Gamma distribution is motivated by its correspondence to an $L(t)$ being a compound Poisson process with exponential jumps, as well as its reasonable description of the data which we shall illustrate shortly.

Due to the presence of dependence, fitting the Gamma distribution to the actual data (the positive part) would not necessarily yield accurate estimates. Consequently, we wish to fit the Gamma distribution to an *iid* sample generated from the actual data. To achieve this, we consider the ‘opposite’ of a block bootstrap, in the sense that we do not wish to generate a bootstrapped sample that preserves the autocorrelation structure the we observe in the data; on the contrary, we wish to ensure independence. Specifically, we follow the procedure described below to obtain a sample with the desired properties:

- (1) Estimate an optimal block-length l by following the procedure in Politis and White (2004) and Patton *et. al* (2009).²
- (2) Draw a number $x_1 \geq 0$ from the empirical distribution of $X(t) - \mu$ and let $B_{1,l} = \{\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_l\}$ denote the block consisting of l consecutive indexes, with $\tilde{x}_1$ corresponding to the position of x_1 in $X(t) - \mu$.
- (3) Let $\tilde{T}$ equal the length of the original time series ($\tilde{T} = 13,514$ cf. Fig. 1) and repeat the following for $j = 2, \dots, \tilde{T}$.

¹We applied the `nlinfit` function in Matlab.

²The procedure is intended for e.g. carrying out the so-called stationary block bootstrap introduced in Politis and Romano (1994), which is generally applicable for stationary weakly dependent time series.

- (a) Draw (with replacement) a new number $z_j \geq 0$ from the empirical distribution of $X(t) - \mu$, and let $\tilde{z}_j$ be the corresponding index.
- (b) If $\tilde{z}_j \subseteq B_{j-1,l}$, discard the draw and repeat step (a). Otherwise, set $x_j = z_j$, $B_{j,l} = \{\tilde{x}_j, \dots, \tilde{x}_{j+l-1}\}$ and proceed.

We implement the above procedure with $\hat{l} = 45$, and fit a Gamma distribution to the generated bootstrap sample of the data of length $\tilde{T}$. Stressing that the parameter α in Eq. (3) coincides with α in Eq. (7), we retrieve $\hat{\lambda}$ and $\hat{\kappa}$ conditional on $\hat{\alpha}$ cf. Eq. (12). By repeating this $N = 10,000$ times, a bootstrapped distribution of $\{(\hat{\lambda}_i, \hat{\kappa}_i)\}_{i=1}^N$ is produced. Based on these bootstrapped distributions, we then obtain the estimates reported in the first column block of Table 2. For comparison, we also fit a Gamma distribution to the actual data (the positive part). The results are displayed in the second column block of Table 2, and we find that they are very similar to the ones obtained with the bootstrap method. To provide some evidence for the goodness-of-fit of the Gamma distribution, we plot in Fig. 4 an example of a bootstrapped sample of the data and the empirical distribution of $X(t) - \mu$, together with corresponding fitted Gamma distributions. Disregarding the few negative values in the empirical distribution of $X(t) - \mu$, the results show that the Gamma distribution provides an acceptable fit to the data.

Table 2. Parameter estimates for the Gamma distribution, conditional on $\hat{\alpha}$.

	Bootstrap procedure		Empirical distribution of $X(t) - \mu$	
	Estimate	Standard error	Estimate	Standard error
$\hat{\lambda}$	1.3649	0.0183	1.3645	0.0157
$\hat{\kappa}$	1.6201	0.0207	1.6187	0.0206

Figure 4. Bootstrapped and empirical distribution of $X(t) - \mu$ together with the corresponding fitted Gamma distribution.

We remark in passing that other stationary distributions could be chosen as long as they are within the class of self-decomposable distributions. However, a different choice of a stationary distribution does not always result in the Lévy process being easily characterisable, as is the case with the Gamma distribution. For detailed discussions on self-decomposability, we refer the interested reader to Barndorff-Nielsen and Shephard (2001) and Halgreen (1979).

Based on the empirical results obtained in this section, we conclude that the model proposed in Eqs. (2)-(3) provides a good overall fit and is thus a reasonable model for the German wind power production index.

4. Pricing of wind power futures

Motivated by the recent introduction of futures written on the German wind power production index, we derive in this section futures prices based on our proposed model. We denote by $F(t, T)$ the wind power futures price at day $t \geq 0$, with delivery at day $T \geq t$. As it is usual in these types of markets, if we want to consider martingale pricing, we must *define* the futures price as the conditional expectation of the wind index at delivery, since the buy-and-hold argument does not hold. The expectation is not to be taken under the objective measure $\mathbb{P}$, but under a pricing measure $\mathbb{Q}$ that is equivalent to $\mathbb{P}$, and hence, $\mathbb{Q}$ -dynamics for the wind power production index must be established.

Since wind is naturally not a tradable asset, there are many potential martingale measures $\mathbb{Q}$. In order to choose such one, we consider here the class of parametrized equivalent measures that can be obtained from the Esscher transform [Esscher (1932)]. Restricting our discussion to a constant market price of risk which we shall denote by θ , and following Benth *et al.* (2008), we define the probability $\mathbb{Q}$ through

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \exp(\theta L(t) - \psi_{L(1)}(-i\theta)t), \quad (13)$$

with $\psi_{L(1)}$ being the cumulant function of $L(1)$ defined as

$$\psi_{L(1)}(x) = \ln \mathbb{E}[e^{ixL(1)}]. \quad (14)$$

Furthermore, to ensure that the Esscher transform is well-defined, we assume that there exists a non-negative constant c such that

$$\mathbb{E}[e^{cL(1)}] < \infty. \quad (15)$$

Hence, the Esscher transform is well-defined for all $\theta \leq c$.

Narrowing the discussion down to our context, where we let $L(t)$ be a compound Poisson process with exponentially distributed jumps, the cumulant function of $L(1)$ becomes

$$\psi_{L(1)}(x) = \lambda \frac{ix}{\kappa - ix}. \quad (16)$$

For a detailed derivation of this result we refer to the proof in Appendix B.1. Also, since we have established that the limiting distribution is the Gamma distribution, we get explicit conditions for the non-existence of the cumulant, as we shall illustrate shortly.

In the following proposition, we derive an explicit expression for the futures price.

Proposition 4.1. Let $0 \leq t \leq T$ and assume that $P(t)$ and $X(t)$ evolve according to the model in Eq. (2) and Eq. (3), respectively. Further, let $L(t)$ be a compound Poisson process as specified in Proposition 3.2. Assuming that the exponential moment condition in Eq. (15) holds for a $c \geq 0$, we have that the wind power futures

price $F(t, T)$ is given by

$$F(t, T) = \Lambda(T)H_\theta(t, T) \left(\frac{P(t)}{\Lambda(t)} \right)^{\exp(-\alpha(T-t))}, \quad (17)$$

where

$$H_\theta(t, T) = \exp \left(-\mu(1 - e^{-\alpha(T-t)}) \right) \left(\frac{\kappa_\theta + e^{-\alpha(T-t)}}{\kappa_\theta + 1} \right)^{\lambda_\theta/\alpha},$$

and

$$\begin{aligned} \kappa_\theta &= \kappa - \theta, \\ \lambda_\theta &= \frac{\lambda\kappa}{\kappa - \theta}. \end{aligned}$$

Proof. See Appendix B.2.

Since $P(t) \leq 1$, it follows that $\mathbb{E}^\mathbb{Q}[P(T)|\mathcal{F}_t] \leq 1$, and so $F(t, T) \leq 1$; moreover, if $P(t) \geq 0$, we also have that $F(t, T) \geq 0$.

Considering the expression for $F(t, T)$ in Proposition 4.1, we note that the condition $\theta < \kappa$ must be imposed to ensure exponential integrability of L and thus the existence of an Esscher transform. Since the estimated κ is positive cf. Table 2, we have no sign restriction on the market price of risk θ . Also note that while the distributional properties of the jump process remain unchanged, the jump intensity and jump size are impacted by the Esscher transform: $L(t)$ is still a compound Poisson process with exponentially distributed jumps, but now with intensity λ_θ and mean jump size $1/\kappa_\theta$. A positive θ will emphasize the jump intensity and the mean jump size, while a negative θ will have the opposite effect.

According to Eq. (17), the shape of the futures curve $T \rightarrow F(t, T)$ depends explicitly on the seasonal function Λ , a function H_θ that incorporates the market price of risk θ and a term that includes today's spot wind power index $P(t)$. The seasonal component gives a contribution to the futures curve corresponding to the fitted seasonal function plotted in Fig. 2(b).

To illustrate the contribution from the second term entering the expression for $F(t, T)$, we plot in Fig. 5 the evolution of $H_\theta(t, T)$ as a function of θ for five different maturities, and using the parameter estimates for α , λ and κ obtained in Sec. 3.1 for the German data. When considering the different maturities T , we observe that $H_\theta(t, T)$ converges very fast to a fixed shape as T increases; with no market price of risk, that is $\theta = 0$, the contribution from the second term is very close to 1 meaning that the futures price is almost unaffected by this term. Generally, we observe that a negative θ implies a value of $H_\theta(t, T) > 1$ and hence an increase in the futures price. Equivalently, a positive θ implies a decrease in the futures price.

Unlike the first two terms, the third term in Eq. (17) gives rise to a stochastically varying shape for the futures curve in the short end. As the time to maturity increases, this term will either decrease or increase to 1 depending on whether $P(t) > \Lambda(t)$ or $P(t) < \Lambda(t)$. To depict this behavior, we set $t = 0$ and α to its estimated value from Sec. 3.1, and plot in Fig. 6 two situations: First, we let $P(0) = 0.40$ and $\Lambda(0) = 0.30$ and second, we let $P(0) = 0.20$ and $\Lambda(0) = 0.30$.

Figure 5. Values of $H_\theta(t, T)$ for $\theta \in [-1.5, 1.5]$ and different maturity periods. α , λ and κ are fixed to the estimated values obtained in Sec. 3.1.

Figure 6. The shape of the third term entering the expression for the futures price in Eq. (17). We consider the case where $\Lambda(0) = 0.30$, and $P(0) = 0.20$ and $P(0) = 0.40$, respectively.

Combining the three terms discussed above, we obtain two futures curves which we plot in Fig. 7. The same values for $\Lambda(0)$ and $P(0)$ as in Fig. 6 are employed, with $\theta = 0$ and the parameter estimates obtained for the German data. On the short end, the shape of the futures curve is highly influenced by the behavior of the contributing term from Fig. 6. On the long end, the shape is mostly influenced by the seasonal function, as $F(t, T) \sim v\Lambda(T)$ for a constant v and $T \gg t$. The decreasing pattern of both curves in Fig. 7 in the long end is due to the yearly seasonal cycle and the fact that the initial value of the seasonality curve, i.e. $\Lambda(0)$, corresponds to data as of 1 January. We stress that the annual pattern of the term structure is not clear in Fig. 7 as we restrict our attention to 100 days.

Since the market price of risk is rarely zero in reality, we also investigate the contribution to the futures curve implied by $\theta = \pm 0.1$; that is, we again compute $H_\theta(t, T)$, but now for two fixed values of θ and maturities $T \in [1, 50]$. The results are displayed in Fig. 8, showing that the contribution of a constant market price of risk $\theta \neq 0$ corresponds to a function that is either decreasing or increasing exponentially.

4.1 An empirical study of the market price of risk

Since wind power futures on the German wind power index have been traded for a while now, historical futures prices quoted in the market are available, allowing us to perform an empirical study of the market price of risk. Like with commodity

Figure 7. Theoretical futures curves implied by the proposed model fitted to the German data. The market price of risk θ is set to zero, and the start values for $\Lambda(0)$ and $P(0)$ set to illustrate the same instances as in Fig. 6.

Figure 8. Contribution of a constant market price of risk to the futures curve. Here, we consider the cases $\theta = \pm 0.1$.

futures such as power or gas futures, delivery periods for wind power futures are usually an entire week, month, year, etc. This is in contrast with the type of curve implied by our proposed model, which is smooth and made up of daily futures prices (contracts with non-overlapping delivery periods). To convert the single-day delivery prices $F(t, T)$ obtained with our model to prices of contract types quoted in the market, we assume that

$$f(t, T_1, T_2) = \frac{1}{T_2 - T_1 + 1} \sum_{\tau=T_1}^{T_2} F(t, \tau),$$

where T_1 and T_2 denote start and end delivery dates, respectively.

As we shall illustrate shortly, wind power futures prices are given in EUR/wph (wind production hour) with a tick size fixed to EUR 100, which is used to convert differences between futures and spot values into a monetary measure (see NASDAQ OMX (2017)). Given a wind power futures contract with delivery during e.g. a week, a difference of 0.01 (1%) between the value of the futures contract at time t and the average realized index for the same delivery will yield a profit or loss of 1 (EUR/wph) $\times$ 24 (hours) $\times$ 7 (days) = 168 (EUR). With our model, we established that $0 \leq F(t, T) \leq 1$, and we will simply multiply this value by 100 such that

theoretical and quoted prices are comparable.

To provide an example of a wind power futures curve quoted in the market, we illustrate in Fig. 9 the observed curve on $t = 1$ September 2016. The observations correspond to NASDAQ OMX closing prices and are plotted using horizontal lines from start to end delivery, where time is measured in days. We note that we make up the observed curve using 13 contracts, namely 3 front weeks, 5 front months, 4 front quarters and 1 front year, relative to the valuation date t . For comparison purposes, we also add in Fig. 9 prices implied by our model with $\theta = 0$.

Figure 9. Observed wind power futures curve on 1 September 2016 together with the corresponding theoretical curve implied by our proposed model with zero market price of risk. All prices are given in EUR/wph, and a tick size equal to EUR 100, i.e. we use the conversion 1% = 1 EUR/wph.

Generally, our model produces prices that are above the quoted prices in the market, translating to the fact that $\theta > 0$ according to our discussion earlier. Further, note the strong seasonality pattern in both the theoretical and the market quoted futures curves, with winter contracts being much more expensive than summer contracts.

Having an explicit futures price formula facilitates the calibration of θ , which can be achieved through a minimization of the distance between theoretical and observed prices. To distinguish between theoretical prices implied by our model and market prices, let $f_\theta(t, T_1, T_2)$ denote the theoretical price, emphasizing its dependence on θ . Further, let $f^{Obs}(t, T_1, T_2)$ denote the corresponding closing price observed in the market. To extract the implied market price of risk associated with the contracts in Fig. 9, we consider the following:

$$\hat{\theta}(t, T_1, T_2) = \min_{\theta} |f^{Obs}(t, T_1, T_2) - f_\theta(t, T_1, T_2)|.$$

Implementing this procedure³ yields a market price of risk per contract and valuation date. The obtained values are tabulated in Table 3, confirming that the implied values for $\hat{\theta}$ are generally positive.

Next, we briefly turn our attention to the risk premium, defined as

$$RP(t, T_1, T_2) = f^{Obs}(t, T_1, T_2) - f_{\theta=0}(t, T_1, T_2).$$

Owing to our model construction, notice that we will generally have an alternating sign between the implied market price of risk and the risk premium, i.e. $\theta > 0$

³We applied the `fmincon` function in Matlab.

Table 3. The implied market price of risk on 1 September 2016.

Contract type	Delivery period	Implied $\hat{\theta}$ 1 September 2016
1	Week 1	0.2411
2	Week 2	0.0980
3	Week 3	0.2385
4	Month 1	0.1213
5	Month 2	0.1140
6	Month 3	0.0437
7	Month 4	-0.0005
8	Month 5	0.0382
9	Quarter 1	0.0684
10	Quarter 2	0.0252
11	Quarter 3	0.0268
12	Quarter 4	0.0634
13	Year 1	0.0376

implies $RP < 0$ and vice versa.

So far in our analysis, we have restricted our attention to a single observed futures curve. Based on this, it is of course difficult to comment on general tendencies regarding the market price of risk in the wind power futures market. In a stylized situation, to have a time series for a given contract could be very interesting, since this would render the time series properties of the market price of risk visible. However, it may be problematic that the contracts move in time to maturity (time to start of delivery), suggesting that the various market prices of risk are not directly comparable. An alternative approach would be to find one market price of risk every day, given by a θ that minimizes the distance of the theoretical curve to all the available futures contracts that day. Then, we would get a series of market prices of risk for the whole market. This hides potential dependencies on time to delivery and length of delivery, but will nonetheless reveal the risk premium sign, and potentially if there are any interesting time series properties for the market price of risk. To gain more insight, we address next both types of investigations mentioned here.

The data we consider are observed wind power futures curves for the period from 1 February 2016 to 31 January 2017, amounting to a total of 257 curves. Each curve consists of 13 observed prices corresponding to the contract types specified in Table 3. Further, we consider static parameter estimates, that is, the ones obtained in Sec. 3.1. Ideally, the model should be recalibrated each day in the interval from 1 February 2016 to 31 January 2017, but lack of a synthetic index time series constructed for each of the valuation dates impedes such analysis. Nevertheless, we do not believe that the market has undergone significant changes in the period 1 February 2016 to 31 January 2017 relative to September 2016, thus justifying our study.

Performing the same analysis as the one given in Table 3 on all wind power futures curves yields a time series of implied $\hat{\theta}$ s for each one of the 13 contract types. In Fig. 10, we plot some examples. Despite the lack of direct comparability caused by the presence of a strong seasonal effect, Fig. 10 highlights some interesting features. First, notice that the implied market price of risk corresponding to shorter deliveries is more volatile. Especially the front week contract series exhibits a singular behavior, with one possible explanation being the valuable information encompassed in short-term weather forecasts. Since this forward-looking information is not included in our model, the computed values for $\hat{\theta}$ for the front week contracts contain both

Figure 10. Implied $\hat{\theta}$ s obtained by minimizing the distance between the observed futures prices and the theoretical prices for a given day and contract type. The time series stretches from 1 February 2016 to 31 January 2017.

a market price of risk as well as a sort of information premium. For longer delivery periods or start deliveries that lie further away from the valuation date, the information from weather forecasts becomes less reliable and hence its effect diminishes. Second, we mention that the market for wind power futures is still very illiquid, and especially Figs. 10(e) and 10(f) illustrate this through the long periods with an unchanged implied $\hat{\theta}$. Illiquidity can also explain the increased volatility for contracts

with shorter delivery lengths as these contracts roll more often, thus ‘forcing’ the price to change regardless of the trading activity.

Averaging across the implied $\hat{\theta}$ s for each of the contract types produces the values displayed in Fig. 11. We observe that all mean values are positive, consolidating our earlier findings relating to a positive θ (and hence a negative risk premium). Also notable is the decay in mean values with the length of delivery period. Possible explanations for this behavior can be different actors operating in different segments of the market, viable weather forecasts for very near and short delivery periods, illiquidity and seasonality.

Figure 11. Mean implied $\hat{\theta}$ per contract type, obtained by averaging across the daily implied θ s corresponding to each of the contract types. The numbering of contract types coincides with that of Table 3.

Next, we compute a single implied θ at each t , based on all contracts making up the futures curve. That is, we consider the following minimization problem:

$$\hat{\theta}(t) = \min_{\theta} \sum_{k=1}^K \left| f^{Obs} \left(t, T_1^{(k)}, T_2^{(k)} \right) - f_{\theta} \left(t, T_1^{(k)}, T_2^{(k)} \right) \right|,$$

where $K = 13$ in our case, since each curve consists of 13 wind power futures contracts. As mentioned previously, an investigation of this type would produce more comparable values for the market price of risk. The results are presented in Fig. 12, yielding that the implied $\hat{\theta}$ for the whole market is positive, which is not surprising considering our previous empirical findings.

4.1.1 Correlations and sign of risk premia. All the empirical analyses performed in this section point convincingly towards a negative risk premium, implying a wind power futures market that is in backwardation. Generally in typical commodity markets, the normal backwardation case is an expected situation, since the hedgers are usually the producers who are willing to accept a lower price (e.g. the futures price) than what is predicted in the spot. This seems to be the case for the German wind power futures market as well, and a possible explanation, as also stated in Gersema and Wozabal (2017), goes as follows: The production of single (or a collection of) wind parks is generally much more correlated to the average German wind power production than e.g. the production of a gas-fired power plant. Thus, wind power futures are a more powerful hedging tool for wind power generators than for conventional generators. As a result, the former group exhibits a higher demand and

Figure 12. Daily implied $\hat{\theta}$ s obtained by minimizing the distance between all market prices and theoretical prices for a given valuation date. We consider historical wind power futures curves for the period 1 February 2016 to 31 January 2017.

is willing to accept a lower price when selling wind power futures – and hence the negative risk premium in the German wind power futures market.

To substantiate the claims stated above, we perform a concise empirical investigation: On one hand, we compute the correlation between the German synthetic wind index (cf. Fig. 1) and the historical wind power production index of 26 different German wind parks. This data consists of daily measurements from 1 January 2012 to 31 December 2015, and is provided to us by Neas Energy. Further, the 26 wind parks we consider differ in e.g. total installed capacity, number of wind turbines in the park and geographical location.

On the other hand, we compute the correlation between the German synthetic wind index and the historical day-ahead spark spread. Note that the decision to run/not run of gas-fired power plants depends on whether or not the spark spread is positive, and hence the spark spread is a measure for the profitability of such plants. We compute the spark spread as the difference between the day-ahead electricity price in Germany and the day-ahead gas price in the NetConnect Germany hub scaled by a heat rate h , with $h \in [1.9, 2.4]$. This interval corresponds to an efficiency between approximately 42% and 53%, which reflects a realistic level according to e.g. figure 14 in the report by Ecofys (2014).

Since the interval from 1 January 2012 to 31 December 2015 is the ‘common denominator’ for the many different time series we consider here, all correlations are computed based on this time interval. The correlations between the generation of the 26 wind parks and the German synthetic wind index are illustrated in Fig. 13, hereby also showing the approximate geographical location of each individual wind park we consider. Fig. 13 shows a very strong positive relation (generally) between the production index of one specific wind park and the German index. When computing the linear correlation between the day-ahead spark spread and the German index, we get a value of -0.47 for the lowest h , i.e. $h = 1.9$. Increasing h , that is assuming a less efficient gas-fired power plant, weakens the negative relation between the spark spread and the German index, rendering wind power futures less attractive as hedging instruments (for $h = 2.4$, we get a value of -0.44).

Figure 13. Linear correlations between the German synthetic index and the actual wind power production index of 26 distinct wind parks.

4.2 Synthetic vs. historical wind index

An essential point in the context of pricing wind power future is the distinction between the synthetic index (illustrated in Fig. 1) and the historical evolution of the actual German wind power production index. To highlight the importance of this distinction, we perform a comparison study in what follows.

Let us start with introducing the historical data we shall use for comparison: Because data for the actual historical index is not directly available to us, we construct the index using its two underlying data components. First, we consider the total wind power production on a daily basis corresponding to the period from 1 January 2012 to 31 December 2015 (a total of 1461 observations). Second, we consider monthly observations for the total installed wind power capacity for the same period, with monthly measurements having the 1st of each month as time stamps (a total of 49 observations, counting the measurement corresponding to 1 January 2016). The two data components are illustrated in Fig. 14(a) and Fig. 14(b), respectively, revealing the impressive growth that Germany has experienced over the considered period.

A proxy for the actual daily index $\tilde{P}$ is then obtained as

$$\tilde{P}(t) = \frac{W(t)}{24C(t)}, \quad (18)$$

where $W(t)$ denotes the total wind power production in Germany at day t , and $C(t)$ denotes the total installed capacity in Germany at day t . Since the installed capacity data is measured at monthly intervals, intermediate daily values are obtained by linear interpolation. The evolution of the index $\tilde{P}$ is displayed in Fig. 14(c), confirming that all measurements lie above 0 and below 1, as expected. The mini-

Figure 14. Historical data for Germany.

imum measurement corresponds to a value of approx. 0.01 (a 1% utilization of the installed capacity), and the maximum measurement reaches approx. 0.76 (a 76% utilization of the installed capacity).

Next, we perform a linear regression of the actual index against the synthetic index, based on data in the interval 1 January 2012 to 31 December 2015 (corresponding to the four years that the two indexes have in common). We obtain an estimate for the intercept of 0.0055 (with standard error 0.0011) and an estimate for the slope of 0.8303 (with standard error 0.0039); a scatter plot of the actual index against the synthetic index is displayed in Fig. 15.

Figure 15. Scatter plot of the actual index against the synthetic index.

The regression results translate to the fact that using the actual instead of the synthetic index data for model calibration would lead to an underestimation of the wind power futures price, which is indeed not surprising. One factor that helps explain this finding is that the synthetic index does not include information concerning intentional temporary switch off of turbines to reduce output, whereas the actual data does. With everything else being equal, this entails that the actual historical index must generally yield lower values than the synthetic index. Another contributing factor - and likely the most important - is the fixed installed capacity used to compute the synthetic index as opposed to the varying installed capacity used to compute the actual historical index, cf. Fig. 14(b). In continuation hereof, the expansion in wind turbine installations is centered in the wind-rich northern part of Germany, which is expected to have pulled the German index upwards. Hence, the synthetic index which is based on the newer installed capacity numbers for September 2016 is expected to have a higher mean than the actual index which is based

on a varying installed capacity. In the context of pricing wind power futures, we are only interested in the available installed capacity on the valuation date, and not on its historical evolution. Lastly, we mention that advancements in wind turbine technology can also be of relevance in this context.

To investigate how the use of the actual historical index influences the conclusions drawn in Section 4.1 regarding the risk premium, we have rerun all the computations performed in Sections 2.1, 3.1 and 4.1, based on this data. Surprisingly, not only does the model produce lower futures prices (with $\theta = 0$) as argued above, but the conclusions on the risk premium change drastically. Based on the newly calibrated model, we obtain a risk premium that is generally positive, implying a wind power futures market that is in contango, which is in contrast with our earlier findings. In a nutshell, using the ‘wrong’ data, i.e. the actual historical index, for the calibration of the model parameters has a significant impact, leading to very misleading conclusions.

5. Pricing options on wind power futures contracts

As a further demonstration of the advantages of our proposed model and the measure change using the Esscher transform, we consider here the pricing of European options written on wind power futures. While we acknowledge that these options are not traded on an exchange at the current time, they are potentially interesting, and hence this section is intended to provide an outlook.

Let us consider a call option on a wind power futures contract, where the exercise time of the option is T , the strike price is K , and r denotes a constant risk-free rate. To simplify calculations in what follows, we further assume that the maturity of the futures contract coincides with the exercise of the option, i.e., the call option is written on the actual wind power production. The call option price $C(t; T, K, T)$ can be expressed as the discounted conditional expectation of the future payoff under $\mathbb{Q}$, which is the pricing measure under the Esscher transform cf. Section 4. Hence,

$$\begin{aligned} C(t; T, K, T) &= e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[\max(F(T, T) - K, 0) | \mathcal{F}_t] \\ &= e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[\max(P(T) - K, 0) | \mathcal{F}_t] \\ &= e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[\max(A(T)e^{Z(T)} - K, 0) | \mathcal{F}_t], \end{aligned}$$

where

$$\begin{aligned} A(T) &= \Lambda(T) \exp(-X(t)e^{-\alpha(T-t)} - \mu(1 - e^{-\alpha(T-t)})), \\ Z(T) &= - \int_t^T e^{-\alpha(T-s)} dL(s). \end{aligned}$$

Note that $A(T)$ can easily be computed given $P(t)$, the estimated seasonality function $\hat{\Lambda}$, and the speed of mean reversion $\hat{\alpha}$. To compute the call option price based on our model framework, it is convenient to employ Fourier techniques, as suggested in Benth *et al.* (2008). Following Benth *et al.* (2008), we define the Fourier transform

of a function $g \in L^1(\mathbb{R})$ as

$$\widehat{g}(y) = \int_{\mathbb{R}} g(x) e^{-iyx} dx. \quad (19)$$

If $\widehat{g} \in L^1(\mathbb{R})$, the inverse Fourier transform can be expressed as

$$g(x) = \frac{1}{2\pi} \int_{\mathbb{R}} \widehat{g}(y) e^{iyx} dy.$$

Before proceeding to computing the price $C(t; T, K, T)$, we state the payoff function in terms of the Fourier transform.

Lemma 5.1. For $a > 1$, we define

$$g_T(x) = e^{-ax} \max(A(T)e^x - K, 0).$$

Then, we have that

$$\widehat{g}_T(y) = \frac{K}{(a-1+iy)(a+iy)} \left(\frac{K}{A(T)} \right)^{-(a+iy)},$$

where $\widehat{g}_T$ is the Fourier transform of g_T .

The result in Lemma 5.1 follows from employing the definition in Eq. (19). We note that the factor $\exp(-ax)$ in the definition of g_T is introduced due to the call option payoff not being a square-integrable function. For more details, we refer to Benth *et al.* (2008), Lemma 9.1, and Carr and Madan (1999). In the next Proposition, we derive the price $C(t; T, K, T)$.

Proposition 5.2. Let $C(t; T, K, T)$ denote the price of a call option written on a wind power futures contract with strike K , exercise T , and delivery period of the futures contract T . The price $C(t; T, K, T)$ at time $t \leq T$ is given as

$$C(t; T, K, T) = e^{-r(T-t)} \frac{1}{2\pi} \int_{\mathbb{R}} \widehat{g}_T(y) \Xi(t, T) dy, \quad (20)$$

where

$$\Xi(t, T) = \left(\frac{\kappa_\theta + (a+iy)e^{-\alpha(T-t)}}{\kappa_\theta + a + iy} \right)^{\lambda_\theta/\alpha}.$$

Proof. See Appendix B.3.

We note that by having an analytical expression for the cumulant $\psi_{L(1)}^{\mathbb{Q}}$, the call option price $C(t; T, K, T)$ can easily be determined by solving the integral in Eq. (20) numerically. Concerning the estimation of put option prices, these follow from the put-call parity.

We conclude this section by illustrating in Fig. 16 call option prices obtained by applying the formula in Proposition 5.2 for a series of strike indexes K . The valuation date t equals 31 December 2015, and we consider two different maturities, 1 July

2016 and 1 December 2016, as to emphasize the seasonal effects. Not surprisingly, the yearly seasonality in the wind index translates to the call options being cheaper for delivery during summer than during winter. Lastly, we note that the option prices could be multiplied by a tick size of EUR 100 in order to achieve comparability with the forward prices quoted in the market, see e.g. Fig 9.

Figure 16. Estimated call option prices as functions of the strike K at $t = 31$ December 2015. The computations are performed with $r = 0$, $a = 1.1$ and the parameter estimates obtained in Sections 2.1 and 3.1. We assume $\theta = 0$, and thus $\kappa_\theta = \kappa$, $\lambda_\theta = \lambda$.

6. Conclusion

In this paper we propose a non-Gaussian Ornstein-Uhlenbeck model for the wind power production index. The model has appealing characteristics, among others straightforward estimation of model parameters and analytical tractability. Motivated by the recent introduction of the German wind power futures on NASDAQ OMX and EEX, we employ the proposed model to conduct an empirical study on German data. First, the model is fitted to a synthetically generated time series of German wind power production indexes, revealing a good overall fit. Then, explicit prices for wind power futures are derived in the framework of no-arbitrage pricing. This facilitates the study of the market price of risk, which can be obtained by the usual practice of minimizing the distance between theoretical prices produced with our model and actual prices observed in the market. Based on historical wind power futures curves made up of closing prices from NASDAQ OMX, we perform different studies of the market price of risk.

Generally, we find evidence of a negative risk premium, whose magnitude decreases as the length of the delivery period increases. The negative risk premium suggests that wind power producers are willing to accept a lower price when selling wind power futures. As also argued in Gersema and Wozabal (2017), this behavior is due to wind power futures being a more powerful hedging tool for wind power generators than for conventional generators. This argument is enhanced by a brief empirical study, which demonstrates that the production of individual wind parks at different locations in Germany is more correlated to the German index than the production of conventional generators (here gas-fired power plants). Also, we find that the market price of risk is more volatile for shorter delivery periods, and argue that this behavior might be related to liquidity aspects and the information contained in short-term weather forecasts, which our model does not incorporate.

In this paper, we have restricted our attention to a constant market price of risk θ ; admittedly, it is possible to allow for e.g. a seasonally varying θ in the Esscher transform. While it remains unclear whether this is backed by the data, one could potentially imagine a seasonality in the market. For a more general measure change, stochastic θ 's (even being state dependent) could be considered as well, however this aspect is left for future research.

To highlight the importance of fitting our proposed model to a wind index that is generated assuming a constant as opposed to a varying level of installed capacity, we show through an empirical example that building on the ‘wrong’ data foundation can lead to the opposite conclusion regarding the sign of the risk premium. Finally, we address the pricing of European options written on wind power futures contracts, as to elaborate further on the benefits of the proposed modeling approach. Since an analytical expression for the cumulant is readily available, we show that the pricing of calls and puts can be achieved without difficulty.

Funding

Fred Espen Benth acknowledges support from FINEWSTOCH, funded by the Norwegian Research Council. Anca Pircalabu is supported by the Innovation Fund Denmark under Grant 4135-00082B.

Appendix A. Seasonality in the non-Gaussian Ornstein-Uhlenbeck model

The purpose of this appendix is to elaborate on issues related to seasonality in the non-Gaussian Ornstein-Uhlenbeck model for the wind power production index. Let us start by considering a simplified version of our model proposed in Eqs. (2)–(3), where $\Lambda(t) = 1$ and $\mu = 0$:

$$\begin{aligned} P(t) &= \exp(-X(t)), \\ dX(t) &= -\alpha X(t)dt + dL(t). \end{aligned}$$

Since L is a subordinator, it follows that $X(t)$ is non-negative, and thus $P(t) \in [0, 1]$ is not violated with this model specification. While it is highly important to comply with the bound restrictions for $P(t)$, we cannot ignore the shortcoming of the above model regarding seasonality: The wind power production index has a strong seasonal component embedded in its dynamics, causing $\Lambda(t) \neq 1$ and $\mu \neq 0$ in reality.

To include seasonality, one possibility is to relax the assumptions imposed on $\Lambda(t)$ and μ above. This has led to our model proposed in Eqs. (2)–(3), and based on the variable $P(t)$, we have argued that $\mu = \ln M$, where $M = \max(\Lambda(t))$. However, if we let $m = \min(\Lambda(t))$ and instead regard the variable $P(t)/\Lambda(t)$, a renewed analysis yields that $P(t)/\Lambda(t) \in [0, 1/m]$, while $\exp(-X(t)) \in [0, \exp(-\mu)] = [0, 1/M]$. Since $m < M$, we have that $1/m > 1/M$, implying that the span of the data $P(t)/\Lambda(t)$ will be bigger than what our model, that is $\exp(-X(t))$, can capture. Hence, there is a potential of having

$$X(t) - \mu < 0,$$

which is somehow in contradiction with Proposition 3.2, where we obtained a Gamma distribution as the limiting distribution of $X(t) - \mu$.

Despite the slight inconsistency produced with our model specification in Eqs. (2)–(3), there are strong arguments in favor of this model. First, the above mentioned issue proves to have very limited impact in practice and second, the task of parameter estimation becomes more straightforward, since $\Lambda(t)$ can be fitted by ordinary least squares, and an estimate for μ immediately follows from the established relation $\mu = \ln M$.

Lastly, we mention that there are of course other alternatives concerning the inclusion of a seasonal component in the model. In our case for example, where we assume $L(t)$ to be a compound Poisson process with exponentially distributed jumps, such an alternative could be to consider seasonal instead of constant intensity. Such an approach will, in mean, provide us with a seasonally varying $P(t)$, while satisfying $P(t) \in [0, 1]$. However, this will come at the cost of having a complex empirical analysis, for which reason we have not pursued this approach in the present paper.

Appendix B. Proofs

B.1 Proof of Proposition 3.2

Because

$$X(t) = X(0)e^{-\alpha t} + \mu(1 - e^{-\alpha t}) + \int_0^t e^{-\alpha(t-s)} dL(s),$$

the characteristic function of $X(t)$ becomes

$$\begin{aligned} \mathbb{E}[e^{ixX(t)}] &= \exp((X(0)e^{-\alpha t} + \mu(1 - e^{-\alpha t}))ix) \mathbb{E}[e^{ix \int_0^t e^{-\alpha(t-s)} dL(s)}] \\ &= \exp\left(ixX(0)e^{-\alpha t} + ix\mu(1 - e^{-\alpha t}) + \int_0^t \psi_{L(1)}(xe^{-\alpha(t-s)})ds\right). \end{aligned}$$

Further, the cumulant function of $L(1)$ denoted by $\psi_{L(1)}$ above is defined as

$$\begin{aligned} \psi_{L(1)}(x) &= \ln \mathbb{E}[e^{ixL(1)}] \\ &= \ln \mathbb{E}\left[\mathbb{E}\left[\exp\left(ix \sum_{k=1}^{N(1)} J_k\right) \middle| N(1)\right]\right] \\ &= \ln \sum_{n=0}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} (\mathbb{E}[e^{ixJ}])^n \\ &= \ln\left(e^{-\lambda} e^{\lambda \mathbb{E}[e^{ixJ}]}\right) \\ &= \lambda(\mathbb{E}[e^{ixJ}] - 1), \end{aligned}$$

with the second equality following from the definition of $L(t)$. Since J is an exponentially distributed random variable with density function given in Eq. (6), its

characteristic function entering the expression of $\psi_{L(1)}$ is simply

$$\mathbb{E}[e^{ixJ}] = \int_0^\infty e^{ixy} \kappa e^{-\kappa y} dy = \frac{\kappa}{\kappa - ix}.$$

Hence,

$$\psi_{L(1)}(x) = \lambda \frac{ix}{\kappa - ix}.$$

In the limit as $t \rightarrow \infty$, we obtain

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{E}[e^{ixX(t)}] &= \exp \left(ix\mu + \int_0^\infty \psi_{L(1)}(xe^{-\alpha s}) ds \right) \\ &= \exp \left(ix\mu + \lambda \int_0^\infty \frac{ixe^{-\alpha s}}{\kappa - ixe^{-\alpha s}} ds \right) \\ &= \exp \left(ix\mu + \frac{\lambda}{\alpha} \ln \frac{\kappa}{\kappa - ix} \right) \\ &= e^{ix\mu} \left(1 - i \frac{x}{\kappa} \right)^{-\frac{\lambda}{\alpha}}, \end{aligned}$$

where we recognize the second factor of the above product as the characteristic function of the Gamma distribution. The result follows.

B.2 Proof of Proposition 4.1

Appealing to the adaptedness of $X(t)$ and recalling Eq. (2) and Eq. (3), we have that

$$\begin{aligned} F(t, T) &= \mathbb{E}^\mathbb{Q}[P(T)|\mathcal{F}_t] \\ &= \Lambda(T) \exp \left(-X(t)e^{-\alpha(T-t)} - \mu(1 - e^{-\alpha(T-t)}) \right) \\ &\quad \times \mathbb{E}^\mathbb{Q} \left[\exp \left(- \int_t^T e^{-\alpha(T-s)} dL(s) \right) \middle| \mathcal{F}_t \right], \end{aligned}$$

where $\mathbb{Q}$ is the pricing measure obtained from the Esscher transform. Further, since L is a $\mathbb{Q}$ -Lévy process and thus characterized by independent increments, we find

$$\begin{aligned} \mathbb{E}^\mathbb{Q} \left[\exp \left(- \int_t^T e^{-\alpha(T-s)} dL(s) \right) \middle| \mathcal{F}_t \right] &= \mathbb{E}^\mathbb{Q} \left[\exp \left(- \theta \int_t^T e^{-\alpha(T-s)} dL(s) \right) \right] \\ &= \exp \left(\int_t^T \psi_{L(1)}^\mathbb{Q}(ie^{-\alpha(T-s)}) ds \right) \\ &= \exp \left(\int_0^{T-t} \psi_{L(1)}^\mathbb{Q}(ie^{-\alpha s}) ds \right), \end{aligned}$$

where $\psi_{L(1)}^\mathbb{Q}$ denotes the cumulant function of $L(1)$ under the Esscher transformed measure $\mathbb{Q}$. Using the Radon-Nikodym derivative in Eq. (13) and the definition in

Eq. (14), it follows that the characteristic function of $L(1)$ under $\mathbb{Q}$ can be expressed as

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}} \left[e^{ixL(1)} \right] &= \mathbb{E} \left[e^{ixL(1)+\theta L(1)} \right] e^{-\psi_{L(1)}(-i\theta)} \\ &= \exp(\psi_{L(1)}(x - i\theta) - \psi_{L(1)}(-i\theta)).\end{aligned}$$

Thus,

$$\psi_{L(1)}^{\mathbb{Q}}(x) = \psi_{L(1)}(x - i\theta) - \psi_{L(1)}(-i\theta).$$

Recalling the expression for $\psi_{L(1)}$ stated in Eq. (16), we obtain

$$\begin{aligned}\psi_{L(1)}^{\mathbb{Q}}(x) &= \lambda \left(\frac{i(x - i\theta)}{\kappa - i(x - i\theta)} - \frac{i(-i\theta)}{\kappa - i(-i\theta)} \right) \\ &= \lambda_{\theta} \left(\frac{ix}{\kappa_{\theta} - ix} \right),\end{aligned}$$

where

$$\begin{aligned}\kappa_{\theta} &= \kappa - \theta, \\ \lambda_{\theta} &= \frac{\lambda\kappa}{\kappa - \theta}.\end{aligned}$$

Inserting all the information obtained above in the expression for $F(t, T)$ yields

$$\begin{aligned}F(t, T) &= \Lambda(T) \exp \left(-X(t)e^{-\alpha(T-t)} - \mu(1 - e^{-\alpha(T-t)}) + \int_0^{T-t} \lambda_{\theta} \frac{-e^{-\alpha s}}{\kappa_{\theta} + e^{-\alpha s}} ds \right) \\ &= \Lambda(T) \exp(-\mu(1 - e^{-\alpha(T-t)})) \left(\frac{\kappa_{\theta} + e^{-\alpha(T-t)}}{\kappa_{\theta} + 1} \right)^{\lambda_{\theta}/\alpha} \left(\frac{P(t)}{\Lambda(t)} \right)^{\exp(-\alpha(T-t))}.\end{aligned}$$

The proposition follows.

B.3 Proof of Proposition 5.2

From Lemma 5.1, it follows that the call option price can be expressed as

$$C(t; T, K, T) = e^{-r(T-t)} \frac{1}{2\pi} \int_{\mathbb{R}} \widehat{g}(y) \mathbb{E}^{\mathbb{Q}}[e^{(a+iy)Z(T)} | \mathcal{F}_t] dy.$$

Recall from the derivation of the wind power futures price in Appendix B.2 that the conditional expectation above is given by

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}[e^{(a+iy)Z(T)} | \mathcal{F}_t] &= \mathbb{E}^{\mathbb{Q}}[e^{-(a+iy) \int_0^{T-t} e^{-\alpha s} dL(s)} | \mathcal{F}_t] \\ &= \exp \left(\int_0^{T-t} \psi_{L(1)}^{\mathbb{Q}}((ai - y)e^{-\alpha s}) ds \right).\end{aligned}$$

The result follows from a straightforward calculation, where the expression for $\psi_{L(1)}^{\mathbb{Q}}(x)$ derived in Appendix B.2 must be employed.

Appendix C. Further details on the estimation of α

In this appendix, we consider the estimation of α cf. Eq. (3) by performing an AR(1)-estimation. Recalling the discussion and notation in Section 3.1, fitting an AR(1) to $X(t) - \mu$ yields $\hat{\phi} = 0.6067$, with a standard error of 0.0070. This implies that

$$\hat{\alpha} = 0.4997,$$

which is different (but not too far) from the estimate obtained in Section 3.1.

Considering the sample autocorrelation of the resulting residuals (Fig. 1(a)) and residuals squared (Fig. 1(b)), we note that an autoregressive model of higher order could be beneficial. By computing the partial autocorrelation function of $X(t) - \mu$, this is indeed confirmed. Specifically, the pacf cuts off at lag 3, indicating that an AR(3) is preferred. With its continuous-time analogous, the CAR(3) model, we are however not ensured positivity, which is clearly an essential point in our modeling of the wind power production index. The proposed Ornstein-Uhlenbeck process in Eq. (3) is positive by design, while a similar CAR(3) process is not necessarily so. One could check case by case, but this is somehow cumbersome, and therefore not pursued further in the present study.

Figure C1. The acf of residuals and residuals squared after having fitted an AR(1) to $X(t) - \mu$.

References

- Ackerer, D., Filipović, D. and Pulido, S. (2017). The Jacobi stochastic volatility model. Swiss Finance Institute Research Paper No. 16-35.
- Barndorff-Nielsen, O.E. and Shephard, N. (2001). Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, **63**(2), 167–241.
- Basse-O'Connor, A., Graversen, S.-E. and Pedersen, J. (2014). Stochastic integration on the real line. *Theory of Probability and Its Applications*, **58**(2), 193–215.

- Benth, F.E., Kallsen, J. and Meyer-Brandis, T. (2007). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, **14**(2), 153–169.
- Benth, F.E., Šaltytė Benth, J. and Koekebakker, S. (2008). *Stochastic Modeling of Electricity and Related Markets*. World Scientific.
- Benth, F.E. and Šaltytė Benth, J. (2009). Dynamic pricing of wind futures. *Energy Economics*, **31**(1), 16–24.
- Benth, F.E. (2011). The stochastic volatility model of Barndorff-Nielsen and Shephard in commodity markets. *Mathematical Finance*, **21**(4), 595–625.
- Benth, F.E. and Šaltytė Benth, J. (2011). Weather derivatives and stochastic modelling of temperature. *International Journal of Stochastic Analysis*, Vol. 2011, Article ID 576791, 21 pages.
- Brody, D., Syroka, J. and Zervos, M. (2002). Dynamical pricing of weather derivatives. *Quantitative Finance*, **2**(3), 189–198.
- Campbell, S.D. and Diebold, F.X. (2005). Weather forecasting for weather derivatives. *Journal of the American Statistical Association*, **100**(469), 6–16.
- Cao, M. and Wei, J. (2004). Weather derivatives valuation and market price of weather risk. *Journal of Futures Markets*, **24**(11), 1065–1089.
- Carr, P. and Madan, D.B. (1999). Option valuation using the Fast Fourier Transform. *Journal of Computational Finance*, **2**, 61–73.
- Davis, M. (2001). Pricing weather derivatives by marginal value. *Quantitative Finance*, **1**(3), 305–308.
- Ecofys (2014). International comparison of fossil power efficiency and CO_2 intensity - Update 2014 [online]. Available online at: <http://www.ecofys.com/files/files/ecofys-2014-international-comparison-fossil-power-efficiency.pdf> (accessed 27 May 2017).
- Esscher, F. (1932). On the probability function in the collective theory of risk. *Skandinavisk Aktuarietidskrift*, **15**, 175–195.
- Gersema, G. and Wozabal, D. (2017). An equilibrium pricing model for wind power futures. *Energy Economics*, **65**, 64–74.
- Gouriéroux, C. and Valéry, P. (2002). Estimation of a Jacobi process. Working paper, version of November 2002.
- Halgreen, C. (1979). Self-decomposability of the generalized inverse Gaussian and hyperbolic distributions. *Probability Theory and Related Fields*, **47**(1), 13–17.
- Härdle, W.K. and López Cabrera, B. (2012). Implied market price of weather risk. *Applied Mathematical Finance*, **19**(1), 59–95.
- NASDAQ OMX (2017). Nasdaq Renewables Wind Index Germany (DE) [online]. Available online at: <http://www.nasdaqomx.com/transactions/markets/commodities/markets/renewables> (accessed 10 March 2017).
- Pircalabu, A. and Jung, J. (2017). A mixed C-vine copula model for hedging price and volumetric risk in wind power trading. *Quantitative Finance*, **17**(10), 1583–1600.
- Platen, E. and West, J.M. (2005). A fair pricing approach to weather derivatives. *Asia-Pacific Financial Markets*, **11**(1), 23–53.
- Politis, D.N. and Romano, J.P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, **89**(428), 1303–1313.
- Politis, D.N. and White, H. (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews*, **23**(1), 53–70.
- Patton, A., Politis, D.N. and White, H. (2009). Correction to “Automatic block-length selection for the dependent bootstrap”. *Econometric Reviews*, **28**(4), 372–375.
- Sato, K.-I. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press.