



UNIVERSITÀ DEGLI STUDI DI PADOVA

Dipartimento di Scienze Economiche “Marco Fanno”

MODELLING AND FORECASTING WIND SPEED INTENSITY
FOR WEATHER RISK MANAGEMENT

MASSIMILIANO CAPORIN
University of Padova

JULIUSZ PREŚ
West Pomeranian University of Technology

December 2009

“MARCO FANNO” WORKING PAPER N.106

Modelling and forecasting wind speed intensity for weather risk management^{*}

Massimiliano Caporin[†]
Department of Economics and Management “Marco Fanno”
University of Padova, Italy

Juliusz Preś
Department of Engineering and Managerial Economics
West Pomeranian University of Technology, Poland
and
Consus S.A., Torun, Poland

First draft: February 2009
This version: December 2009

Abstract. The modelling of wind speed is a traditional topic in meteorological research, where the main interest is on the short-term forecast of wind speed intensity and direction. More recently, this theme has received some interest in the quantitative finance literature for its relationship with electricity production by wind farms. In fact, electricity producers are interested in long-range forecasts and simulation of wind speed for two main reasons: to evaluate the profitability of building a wind farm in a given location and to offset the risks associated with the variability of wind speed for an already operating wind farm. In this paper, we contribute to the increasing literature regarding environmental finance by comparing three approaches that are capable of forecasting and simulating the long run evolution of wind speed intensity (direction is not a concern, given that the recent turbines can rotate to follow wind direction): the Auto Regressive Gamma process, the Gamma Auto Regressive process, and the ARFIMA-FIGARCH model. We provide both in-sample and out-of-sample comparisons of the models, as well as some examples for the pricing of wind speed derivatives using a model-based Monte Carlo simulation approach.

Keywords: Gamma Auto Regressive, Auto Regressive Gamma, ARFIMA-FIGARCH, wind speed modelling, wind speed simulation

JEL codes: C22, C53, G17, G13, G22

^{*} We thank Dominique Guégan, Marius Frunza, Neil Shephard and all participants to the CFE09 workshop held in Cyprus for their helpful comments.

[†] Corresponding author: Massimiliano Caporin, University of Padova, Faculty of Statistical Sciences, Department of Economics and Management “Marco Fanno”, Via del Santo, 33, 35123 Padova, Italy – email: massimiliano.caporin@unipd.it – ph. +39-049-827-4258, fax +39-049-827-4211

1. Introduction

Wind conditions represent a relevant source of risk for every wind farm. The risk exposure can be associated with two elements that characterize the wind: the overall wind speed, or wind speed intensity, and the wind direction. However, most turbines currently built have mechanisms for automatic rotation of blades in the appropriate wind direction. Therefore, in practice, the most relevant weather exposure of wind farms can be measured by analyzing only the wind speed intensity (an approximation is included since blade rotation is not immediate). In order to mitigate this wind risk, we often use insurance contracts or wind derivatives, making a direct connection between atmospheric elements, financial markets, and economic implications. For a survey of the literature on wind risk evaluation and modelling, see Brix et al. (2005).

From a meteorological point of view, wind speed intensity is just one of the variables generally modelled to provide robust short-term weather forecasts. However, for the purpose of wind risk evaluation, on the one hand, we need long-term forecasts and simulations, while, on the other hand, the interest in weather variables other than wind may be limited. As a result, meteorological models are not appropriate for long-range wind speed intensity forecast and simulation, while statistical approaches are viable. The scientific literature includes some contributions on that topic, for example, Brown et al. (1984), Castino et al. (1998), and Aillot et al. (2006), among others. However, few authors have yet considered the economic and financial point of view; that is, studying problems associated with the modelling of wind speed intensity for evaluating wind risks, and for the pricing of wind derivatives (see, for instance, Leroy (2004) and Yamada (2008)).

In this framework, the main interest is in the evaluation (even by simulated methods) of a wind index, which is a transformation of wind speed intensity into a quantity proportional to the energy power potentially produced by a turbine over a given period. Simple and naïve approaches focus on low frequency (from monthly to yearly, usually) wind data time series. These data are fitted with an unconditional density, which we then use to simulate wind speed paths. In turn, these paths determine simulated wind indices. However, these unconditional approaches fail to consider the periodic evolution of wind speed intensity, or the presence of serial correlation. We can capture both of these

features by alternative methods based on higher frequency data. Moreover, the normality assumption of wind indices is questionable.

Our present paper contributes to the weather derivatives and environmental finance literature in several ways. First, we consider a set of models that capture serial correlation; these have already appeared in scientific contributions, but generally in different contexts. We consider the ARFIMA-FIGARCH model of Beine and Laurent (2003), as used in many frameworks including financial time series analysis, as well as for the study of average temperature sequences, see Caporin and Prés (2009). We adapt the ARFIMA-FIGARCH model to the logarithms of wind speed intensities, adding periodic deterministic components to capture the yearly seasonal cycle. We then consider two competing approaches that introduce serial correlation in Gamma densities: the Gamma Auto Regressive (GAR) model of Tol (1997) (already used by the author to model wind speed intensity) and the Auto Regressive Gamma (ARG) of Gouriéroux and Jasiak (2006), which has been proposed for modelling intertrade durations.

Our second contribution is to provide a methodology for comparing wind speed models by focusing on both in-sample and out-of-sample model performances. In-sample, we compare the model fit with traditional statistical methods. Out-of-sample, we compare the alternative models and specifications using both one-step-ahead wind speed point and density forecasts, as well as in terms of the model's ability to simulate the evolution of wind indices. This last element will provide fundamental results, given that wind indices are the relevant element for the pricing of contracts covering wind risks. From a different point of view, the ability to simulate wind indices will also be interesting for the identification of optimal locations for wind farm construction. All of the considered models, approaches, and methods we will compare using a number of daily wind speed intensity time series obtained from meteorological stations based in Poland. By simulation based methods we will verify the model performance, and we will provide theoretical prices of wind derivatives, comparing these to the ones provided by simpler unconditional methods.

This paper differs from other contributions on wind speed modelling in a number of ways. Firstly, we do not base the models on data transformations, such as the Box-Cox in Brown et al. (1984), but on actual original data, for which we postulate a specific stochastic structure whose components have a direct interpretation. In detail, we assume that wind speed intensities are function of periodic deterministic components and of

stochastic components. As a result, we incorporate serial correlation in stochastic processes, defined over a positive support. Secondly, the models we propose are able to capture the yearly seasonal evolution of wind speed and, with generalization, can capture intra-daily periodicity created, for instance, by the differences between day and night wind speed intensities. However, the dataset we consider in the empirical analysis contains only average daily wind speed intensities, and thus we are not able to recover intra-daily patterns. The empirical results of our work show how alternative models appropriately fit historical data, while their forecasting performances may differ.

The paper proceeds as follows: Section 2 introduces the modelling approaches and the techniques we consider for model comparison. Section 3 presents the empirical estimates of the models and a first comparison across these. Section 4 compares the fitted model within a wind risk management framework. Section 5 concludes the paper.

2. Modelling wind speed intensity

Meteorological stations measure wind speed intensities and direction at regular frequencies, typically in minutes. These data present a superimposition of several elements: the stochastic nature of the series evolution, the intra-daily periodic evolution governed by the alternation of day and night and the long-term seasonal component driven by the sequence of seasons. The models we propose in this paper appropriately capture these components. All of the models we consider presume that the periodic components are purely deterministic and we could filter these out a priori, before modelling the underlying stochastic process. We note that we could follow alternative approaches; for instance, by assuming a stochastic nature of the periodic components, we could specify a GARMA processes as in Guégan and Diongue (2009). In the following section, we describe the structure of the periodic evolution of wind speed intensities and show possible ways to estimate and filter out this component. The subsequent sections will then introduce the alternative models (dealing with their structure, estimation and simulation), and the approaches for their comparison.

2.1. Periodic component

We consider two alternative specifications of the periodic deterministic component: the first is multiplicative and we apply it to wind speed levels (we can also recast it into an additive periodic function using logs); the second models the wind speed logarithm with additive and multiplicative periodic effects. Note that the purely additive approach provides a filtered series with support over \mathbb{R}^+ as an output, while the additive and multiplicative method gives a filtered series with support over \mathbb{R} .

If x_t denotes the positively valued random variable measuring the average wind speed over a relatively short time period (say 1 hour), the multiplicative periodic component $s(t)$ affects x_t as follows:

$$x_t = e^{s(t)} y_t . \quad (1)$$

where y_t is the “seasonally adjusted” series. We assume that the function $s(t)$ is deterministic, as we previously stated, and has the following structure:

$$s(t) = \beta_0 + \sum_{i=1}^w \beta_i t^i + \sum_{j=1}^q \left(\varphi_j \cos(2\pi j f(t)) + \psi_j \sin(2\pi j f(t)) \right), \quad (2)$$

where $f(t)$ is a yearly periodic sequence adapted to the data frequency. For instance,

when daily data are used, it has the following form $\left\{ \frac{1}{365}, \frac{2}{365}, \dots, \frac{364}{365}, 1, \frac{1}{365}, \frac{2}{365}, \dots \right\}$,

assuming values $\frac{1}{366}, \frac{2}{366}, \dots, \frac{365}{366}, 1$ for leap years. Similarly, for hourly data, we would

have $\left\{ \frac{j}{365 \times 24} \right\}_{j=1}^{365 \times 24}$ for regular years and $\left\{ \frac{j}{366 \times 24} \right\}_{j=1}^{366 \times 24}$ for leap years.

The periodic component is formed by a constant, a polynomial trend, and a combination of harmonics. Note that the multiplicative periodic component affects the wind speed levels through an exponential transformation, which ensures that it has a positive value, and its interpretation is as a multiplicative seasonal factor.

The standard least squares approaches on the log-transformation ideally estimate the parameters in equation (2) (we assume that the probability of having zero average wind speed over a not too short period is zero or negligible):

$$\ln(x_t) = s(t) + \ln(y_t) . \quad (3)$$

We then identify the orders w and q , as well as of the significant coefficients, based on information criteria and using Heteroskedastic and Autocorrelation Consistent (HAC) standard errors, due to the possible presence of serial correlation and heteroskedasticity over the $\ln(y_t)$ series. Note that the multiplicative representation in (1) allows modelling y_t with positive-valued stochastic processes.

As previously stated, we also consider an alternative approach for removing the periodic component. In this case we replace (1) with the following equation

$$x_t = e^{s(t) + \nu(t)y_t} , \quad (4)$$

which adds a second periodic function $\nu(t)$. Taking the logarithms of (4), we obtain:

$$\ln(x_t) = s(t) + \nu(t)y_t . \quad (5)$$

In equation (5), $s(t)$ follows (2) and we estimate this by OLS, as in the previous case. In contrast, the extraction of $\nu(t)$ requires a preliminary transformation. Given the fitted $\hat{s}(t)$ values, we can compute the residuals $\ln(x_t) - \hat{s}(t)$, square them and take the logarithms. Theoretically, the resulting quantity obeys:

$$\ln\left([\ln(x_t) - \hat{s}(t)]^2\right) = \ln(\nu(t)^2) + \ln(y_t^2) . \quad (6)$$

Therefore, we could recover $\nu(t)$ by assuming a representation similar to that of $s(t)$:

$$\ln\left(v(t)^2\right) = \alpha_0 + \sum_{i=1}^m \alpha_i t^i + \sum_{j=1}^p \left(\delta_j \cos(2\pi j f(t)) + \gamma_j \sin(2\pi j f(t)) \right). \quad (7)$$

We could estimate the parameters in (7) by OLS using HAC standard errors. Removal of both periodic components allows modelling of the y_t series using stochastic processes with support over the real line. Note that under an assumption of Gaussianity for y_t , the wind speed intensity x_t follows a Log-Normal distribution. For this reason, we could recover parameter estimates for the periodic components by maximum likelihood methods.

Note that we do not consider the estimation of both additive and multiplicative periodic components over the wind speed levels, as this approach could create problems in the simulation of wind speed trajectories. In fact, use of an additive component over the wind speed mean results in the residuals having support over the real line, and the simulation of wind speed patterns cannot be easily constrained to assume only positive values.

Summarizing, we will consider two alternative specifications for the periodic components affecting wind speed time series. We represent the first by equations (1)-(2) and the second by (4), (2) and (7). Note that if the estimated periodic component in (7) collapses to a constant, the specification in (4) becomes equivalent to that in (1).

2.2. Modelling wind speed intensity

This section presents a set of competing models that are usable for capturing the dynamic evolution of wind speed. We will compare these models over real data in the subsequent sections. We describe here the model structure, the estimation approaches as well as the simulation methods that are fundamental for generating the simulated distribution of wind speed series and for pricing wind speed derivatives. We will propose all of the models for series filtered from the periodic deterministic component, which we denote by y_t , independently from the approach used for recovering the periodic component. As a result, the support for the y_t densities could be different and will be clear in the context of each proposed model.

2.2.1 ARFIMA-FIGARCH

If we remove period components following equation (4), the residual series y_t could be modelled with an ARFIMA-FIGARCH process (see Beine and Laurent, 2003), allowing for the possible presence of long-memory both in the mean and in the variance. We consider this double long-memory model since preliminary analyses we verified that “seasonally adjusted” wind speed series could present mild long-memory behaviours. In the empirical applications, we estimate the ARFIMA-FIGARCH process, and when the data do not support the presence of long-memory in mean and/or in variance, we could consider short-memory specifications. In fact, the ARFIMA-FIGARCH model nests both ARMA mean specifications as well as GARCH structures in the conditional variances.

The following equations characterize an ARFIMA-FIGARCH process for the series y_t :

$$\begin{aligned} \Phi(L)(1-L)^d y_t &= \Theta(L)\varepsilon_t \quad \varepsilon_t | I^{t-1} \sim D(0, \sigma_t^2), \\ \sigma_t^2 &= \omega + \beta(L)\sigma_t^2 + \left[1 - \beta(L) - \Psi(L)(1-L)^\lambda\right] \varepsilon_t^2, \end{aligned} \quad (8)$$

where I^{t-1} is the information set up to time $t-1$, L is the lag operator, D is an un-specified conditional density, $0 \leq d \leq 0.5$, $0 \leq \lambda \leq 1$ are the long-memory coefficients for mean and

variance, respectively, $\omega > 0$, $\Phi(L) = 1 + \sum_{i=1}^p \phi_i L^i$, $\Theta(L) = 1 + \sum_{i=1}^q \theta_i L^i$, $\beta(L) = \sum_{i=1}^m \beta_i L^i$,

$\Psi(L) = 1 + \sum_{i=1}^n \psi_i L^i$, and all polynomials satisfy the constraints ensuring stationarity,

invertibility, and positivity of conditional variances. The ARFIMA model in (8) does not include an intercept since it must be equal to zero by construction, as we filter the wind speed logarithms following (5).

Under an assumption of normality, we derive the one-step-ahead conditional mean and variance of wind speed as follows. At first, we determine the one-step-ahead conditional mean and variance of the y_t series:

$$E_{t-1}[y_t] = \sum_{i=1}^M \pi_i y_{t-i} + \sum_{i=1}^q \theta_i \varepsilon_{t-i}, \quad Var_{t-1}[y_t] = \sigma_t^2, \quad (9)$$

where $E_{t-1}[\cdot]$ and $Var_{t-1}[\cdot]$ are the expectation and the variance conditional to time $t-1$ information set, respectively, and $\Phi(L)(1-L)^d = 1 - \sum_{i=1}^{\infty} \pi_i L^i$. The conditional mean depends on the past values of y_t with a truncation lag M , which is normally set to one thousand (or to all the available information if the sample dimension is lower). We then obtain the mean and variance of wind speed logarithms by:

$$E_{t-1}[\ln(x_t)] = s(t) + v(t)E_{t-1}[y_t], \quad Var_{t-1}[\ln(x_t)] = v(t)^2 \sigma_t^2. \quad (10)$$

Finally, the normality assumption of y_t translates into a log-normality assumption of wind speed intensity x_t whose conditional mean and variance follows from standard formulae. We suggest the use of Maximum likelihood (ML) approaches or Quasi Maximum Likelihood (QML) methods to estimate the parameters of the ARFIMA-FIGARCH model in equations (4), (2), (7) and (8). The estimation of the entire model could be performed by a three-step procedure as follows: i) estimate the parameters of the additive periodic component by OLS over equation (5) when the periodic component follows (2); ii) estimate the parameters of the multiplicative periodic component by OLS in equation (6) when the periodic component is that in (7); iii) estimate the ARFIMA-FIGARCH parameters in (8) using a normal likelihood function. Alternatively, we may estimate the full model using a normal likelihood function over the wind speed logarithms and using the expected mean and variances reported in (10).

We could perform the simulation of wind speed under the model in (4), (2), (7) and (8) by extracting the standardized normal residuals $z_t = \sigma_t^{-1} \varepsilon_t$ from a normal density or by resampling them from the in-sample residuals. Use of (9)-(10) allows easy derivation of the recursion for the simulation of wind speed patterns.

2.2.2 Gamma Auto Regressive (GAR) models

The statistical literature includes a number of works combining autoregressive behaviours and Gamma distributions. The Gamma Auto Regressive model of Gaver and Lewis (1980) provides a relevant example (see also Lawrence and Lewis, 1981, and

Smith and Miller, 1986). However, the approach of Gaver and Lewis (1980) is not appropriate in our setup since it assumes homoscedasticity and a probability mass at zero which is not necessarily present in wind speed data.

Tol (1997) used a different approach to model wind speed data, proposing the following two-parameter conditional Gamma process for the modelling of series y_t :

$$f(y_t) = \frac{e^{-c_t^{-1}y_t} y_t^{\delta_t-1}}{c_t^{\delta_t} \Gamma(\delta_t)}, \quad (11)$$

where $\Gamma(\cdot)$ is the Gamma function, $c_t > 0$ is the scale parameters and $\delta_t > 0$ is the shape parameter. The two time-varying parameters have the following dynamic representation:

$$\delta_t = \frac{\mu_t^2}{\sigma_t^2}, \quad c_t = \frac{\sigma_t^2}{\mu_t}, \quad (12)$$

$$\mu_t = \gamma_0 + \gamma(L)y_t, \quad (13)$$

$$\sigma_t^2 = \omega + \beta(L)\sigma_t^2 + \alpha(L)(y_t - \mu_t)^2. \quad (14)$$

where $\gamma(L) = \sum_{i=1}^m \gamma_i L^i$, $\beta(L) = \sum_{i=1}^p \beta_i L^i$, and $\alpha(L) = \sum_{i=1}^q \alpha_i L^i$. The conditional first order moment, and the conditional variance of y_t are then equal to $E_{t-1}[y_t] = \delta_t c_t = \mu_t$ and $Var_{t-1}[y_t] = \delta_t c_t^2 = \sigma_t^2$, respectively. Note that (13) represents the conditional mean of the Gamma process, while (14) is a conditional variance in the spirit of the GARCH process of Engle (1982) and Bollerslev (1986). The GAR model structure previously outlined does not include long-range persistence, neither in the mean nor in the variances. In principle, we could generalize both (13) and (14) by introducing long-memory polynomials, but we will not pursue this objective in the current paper, since the empirical analyses we will present show that the simple GAR model could appropriately capture the dynamic structure of wind speed intensity.

We can estimate the complete model in (1)-(3), (11)-(14) by a two-step procedure: i) estimate the multiplicative periodic component in (3) by OLS; and ii) estimate by ML

methods the other parameters under a Gamma density. Alternatively, we could estimate jointly the parameters by using a Gamma density replacing the scale parameter c_t with the scale parameter $e^{s(t)}c_t$, and therefore using the following log-likelihood function:

$$L(\theta) = \sum_{t=1}^T \left(-\delta_t s(t) - \delta_t \ln(c_t) - \ln(\Gamma(\delta_t)) - e^{-s(t)} c_t^{-1} x_t + (\delta_t - 1) \ln(x_t) \right), \quad (15)$$

where θ contains the parameters in (2), (13) and (14). Note that in this last case, the Gamma parameters, its mean and variance follow:

$$\delta_t = \frac{\mu_t^2}{\sigma_t^2}, \quad c_t = \frac{\sigma_t^2}{\mu_t} e^{s(t)}, \quad E_{t-1}[x_t] = \mu_t e^{s(t)}, \quad \text{and} \quad \text{Var}_{t-1}[x_t] = \sigma_t^2 e^{2s(t)}, \quad (16)$$

highlighting that the same periodic component affects both the mean and the standard deviation.

The simulation of the Gamma process in (11)-(14) goes through the generation of random numbers from a two-parameter Gamma density, where parameters are time-varying.

2.2.3 Auto Regressive Gamma (ARG) models

Gourieroux and Jasiak (2006) (GJ henceforth) proposed an alternative conditional Gamma process, called Auto Regressive Gamma (ARG). Their proposal builds on the three parameters Non-central Gamma density, represented as:

$$f(x_t; c, \delta, \beta) = \sum_{i=0}^{\infty} \frac{e^{-\beta} \beta^i}{i!} \frac{e^{-c^{-1}x_t} x_t^{\delta+i-1}}{c^{\delta+i} \Gamma(\delta+i)}, \quad (17)$$

where c is the scale parameter, β is the non-centrality parameter and δ is the shape parameter (or degree of freedom). All parameters are non-negative, and the random variable assumes values on the positive real line. The Non-centred Gamma arises as combination of centred two-parameter Gamma densities with weights given by a Poisson

process driving the non-centrality. The first order moment and the variance of the density in (17) are given by $E[x_t] = c\delta + c\beta$ and $Var[x_t] = c^2\delta + 2c^2\beta$, respectively.

GJ suggest expressing the non-centrality parameter as a function of past observations of the modelled variable, thus replacing β with $\sum_{j=1}^p \beta_j x_{t-j}$. This makes the density conditional to the information set at time t-1, and the one-step-ahead first order moment and variance becomes:

$$E_{t-1}[x_t] = c\delta + c\left(\sum_{j=1}^p \beta_j x_{t-j}\right), \quad Var_{t-1}[x_t] = c^2\delta + 2c^2\left(\sum_{j=1}^p \beta_j x_{t-j}\right). \quad (18)$$

GJ also discuss the stationarity of the model, which is associated with parameter combinations satisfying the nonlinear constraint $c\left(\sum_{i=1}^p \beta_i\right) < 1$. Furthermore, they present alternative approaches for the parameter estimation and the model simulation. In their work, GJ allow for a time varying scale parameter, but do not discuss in detail its possible forms and interpretation. If the wind speed x_t follows the Non-centred Gamma density in (17), denoted as $x_t \sim \gamma(\delta, \beta, c)$, then, the ratio $x_t c^{-1}$ follows the two-parameter Non-centred Gamma density $x_t c^{-1} \sim \gamma(\delta, \beta)$ (the density function can be easily obtained from 17). The scale parameter thus affects the density in a multiplicative way, similarly to the periodic component in (1). We note that we may interpret the scale parameter as the periodic component of wind speed; in fact, we may assume a time varying scale parameter equal to:

$$c_t = e^{s(t)}, \quad (19)$$

where the periodic component is that of equation (2). In principle, we could have introduced periodic elements in the time-varying non-centrality parameters, for instance, making it a function of harmonics. However, we would have required appropriate constraints to ensure positivity. Alternatively, we could use a representation similar to (19) to introduce time dependence in the non-centrality parameter in (18). However, we

do not consider this extension here, as some preliminary estimates (not reported here to save space) show that this approach does not provide any improvement over the previous models.

Given that we can interpret the scale parameter as the periodic component of wind speed series, we suggest filtering out the periodic component as in (1)-(3) and model the residual series with a conditional two-parameter non central Gamma process. The density we consider for y_t , the seasonally adjusted wind speed, is a special case of that in GJ, and has the following structure:

$$f(y_t; \delta, \beta) = \sum_{i=0}^{\infty} \frac{\exp\left(-\sum_{i=1}^p \beta_i y_{t-i}\right) \left(\sum_{i=1}^p \beta_i y_{t-i}\right)^i}{i!} \frac{e^{-y_t} y_t^{\delta+i-1}}{\Gamma(\delta+i)}. \quad (20)$$

As a result, the time-varying nature of both the mean and the variance of x_t , the wind speed, will depend on two elements: the Auto Regressive behaviour of the non-centrality parameter, and the dynamic evolution of the scale parameter. We can show the two moments by:

$$E_{t-1}[x_t] = c_t \delta + c_t \left(\sum_{j=1}^p \beta_j x_{t-j} \right), \quad \text{Var}_{t-1}[x_t] = c_t^2 \delta + 2c_t^2 \left(\sum_{j=1}^p \beta_j x_{t-j} \right). \quad (21)$$

Note that we know the scale parameter at time t since it is a function of deterministic components only. Finally, we highlight that the ARG model is consistent with a long-memory behaviour of the modelled series when the parameters satisfy $c \left(\sum_{i=1}^p \beta_i \right) = 1$, see GJ for details.

We can also use the ARG model in (1)-(3) and (20) for modelling intertrade durations, as in GJ, since intraday durations may show periodic behaviours. We also highlight the close relationship between the ARG model and the Autoregressive Conditional Duration model of Engle and Russell (1998) (see Pacurar, 2008, for a survey of duration models). In fact, we may specify an ACD model with a periodic conditional duration as in

equation (19), where only the time-varying scale parameter governs the dynamic behaviour and the innovation density is a two-parameter non-central Gamma.

The estimation of ARG in (1), (2) and (20) could follow the QML approach suggested by GJ. They propose using the normal likelihood function for wind speed intensity, with the expected mean and variances reported in (21). Alternatively, again as suggested by GJ, we could estimate the model by maximizing the truncated log-likelihood:

$$L(\theta) = \sum_{t=1}^T \left(\ln \sum_{i=0}^k \frac{e^{-\beta_t} \beta_t^i}{i!} \frac{e^{-c_t^{-1} x_t} x_t^{\delta+i-1}}{c_t^{\delta+i} \Gamma(\delta+i)} \right) \quad (22)$$

where k should not be too small (we verified through simulations that reasonable values of k should be chosen between 10 and 20; higher value do not improve the evaluation of the likelihood function and are less efficient from a computational point of view).

Note that we can also estimate the model using a two-step procedure similar to that of GAR: i) filter out the periodic component by OLS applied to (3); and ii) estimate the remaining parameters with QML using the normal likelihood and setting to 1 the scale parameter. To simulate the ARG process, we adapt the approach of GJ, by sampling from Gamma and Poisson random variables. In particular, we generate y_t values from:

$$y_t = \sum_{j=1}^{Z_t} W_{j,t} + \varepsilon_t \quad (23)$$

where $W_{j,t}$ are standard Gamma random variables extracted from the two-parameter centred Gamma density $\gamma(1,1)$, ε_t follows a two-parameter centred Gamma density

$\gamma(\delta,1)$, and Z_t is extracted from a Poisson density $P\left(\sum_{j=1}^p \beta_j y_{t-j}\right)$. Given a simulated

trajectory of y_t we obtain the corresponding simulated path of x_t by multiplication with the deterministic path of the scale parameter c_t that we generate using (2).

2.3. Model comparison

In the previous subsections we introduced three alternative models that we could use for fitting wind speed intensity series. The models differ in the distributional hypothesis, and we should adapt the traditional diagnostic checking tools to this framework. For instance, using the expected conditional mean and variance of the models, we can compute standardized wind speed series. In turn, we could analyze these for serial correlation and heteroskedasticity for a first comparison of the models. However, given that the main purpose of our study is to forecast wind speed, we chose to compare the models on the basis of their one-step-ahead point forecasts and density forecasts, as well as in their ability to simulate wind speed sequences.

We assume, in the following, that the researcher's purpose is to compare m alternative wind speed models over h one-step-ahead forecasts. We will estimate the models with an expanding window approach from the starting sample, from I to T . We denote model m time $T+i$ one-step-ahead forecast by $x_{T+i}^{f,m} = g(\hat{\theta}; I^{T+i-1})$, where $\hat{\theta}$ is the vector of estimated parameters entering the conditional density of x_{T+i} (estimated using the information set up to time $T+i-1$). In contrast, we indicate one step-ahead density forecasts as $f_m(x_{T+i} | \hat{\theta}, I^{T+i-1})$, where the density could be a Gamma, Non-centred Gamma, or Log-Normal.

We can draw a preliminary comparison of the model forecast from the analysis of standard quantities, such as the Root Mean Squared Forecast Error (RMSFE):

$$RMSFE(h, m) = \sqrt{\sum_{i=1}^h (x_{T+i} - x_{T+i}^{f,m})^2}, \text{ and the Mean Absolute Deviation (MAD)}$$

$$MAD(h, m) = \sum_{i=1}^h |x_{T+i} - x_{T+i}^{f,m}|. \text{ We also implement more formal approaches proposed in}$$

the statistical literature. We compute the Weighted Likelihood Ratio (WLR) test of Amisano and Giacomini (2007), since we can focus this on the tail of the forecast density. In fact, the following will clarify that too low or too high wind speed values are associated with an absence of power production in wind farms, with a relevant impact on the wind risk management process. The WLR test compares two models by contrasting

over a given horizon their density forecast. Given models m_l and m_j , and their density forecasts $f_{m_l}(x_{T+i} | \hat{\theta}, I^{T+i-1})$ and $f_{m_j}(x_{T+i} | \hat{\theta}, I^{T+i-1})$, the test is based on the quantity

$$WLR_{T+i} = w(x_{T+i}) \left[\log \left(f_{m_l}(x_{T+i} | \hat{\theta}, I^{T+i-1}) \right) - \log \left(f_{m_j}(x_{T+i} | \hat{\theta}, I^{T+i-1}) \right) \right], \quad (24)$$

where $w(x_{T+i})$ is a weighting function which can be calibrated to select specific regions of the distribution of x_{T+i} . For the purposes of this paper, we consider the following cases: i) $w(x_{T+i})=1$, which correspond to a test considering the entire distribution; and

ii) $w(x_{T+i})=1-\phi\left(\frac{x_{T+i}-E[x_t]}{\sqrt{\text{Var}[x_t]}}\right)\phi(0)^{-1}$, where $\phi(\cdot)$ is the Normal PDF. The second case

focuses the test on the tails of the distribution of wind speed intensity. The null hypothesis of model equivalence is associated with the condition $E[WLR_{T+i}]=0$, and the test statistic is $\overline{WLR}_{T+i}\sqrt{h}\sigma_h^{-1}$, where \overline{WLR}_{T+i} is the sample mean of WLR_{T+i} over the forecast horizon h , and σ_h^{-1} is the heteroskedasticity and autocorrelation consistent estimation of $\text{Var}[\sqrt{h}WLR_{T+i}]$; see Amisano and Giacomini (2007) for further details.

Note that density forecasts could be defined using a common function across all models, such as the normal density where the first and second order conditional moments are determined in accordance with the alternative specifications. In this case, we compare only mean and variance forecasts, with a minor impact of the shape of the distribution and of higher order moments.

Furthermore, following the contribution of Diebold et al. (1998), we evaluate density forecasts using the inverse probability transform approach of Rosenblatt (1952). Campbell and Diebold (2005) used similar methods in the context of weather forecasts for weather derivative pricing, as did Allen et al. (2009) for the evaluation of density forecast of financial durations, which, similarly to wind speed intensities, they characterized by distributions with a positive support.

Given the one-step-ahead density forecasts $f_m(x_{T+i} | \hat{\theta}, I^{T+i-1})$, we compute the cumulated density forecasts as:

$$F_{T+i,m}(x_{T+i}) = \int_{-\infty}^{x_{T+i}} f_m(v | \hat{\theta}, I^{T+i-1}) dv \quad (25)$$

If a model is correctly specified, the sequence $\{F_{T+i,m}(x_{T+i})\}_{i=1}^h$ should distribute as a Uniform in the interval $[0,1]$. Following Campbell and Diebold (2005), and Allen et al. (2009), we suggest analyzing the histograms of (25), as well as its ACF, and to test for the null hypothesis of a Uniform distribution.

3. Empirical analysis

In this section, we fit and compare the previous models on a set of three wind speed intensity series. The data, available in the range 1st January 1986 to 31st December 2008, refer to three meteorological stations located in Poland at Wielun (WMO 12455), Lodz (WMO 12465), and Sulejow (WMO 12469). Within parentheses, we report the World Meteorological Organization (WMO) code for each station. We use the historical time series of daily average wind speed provided by National Climatic Data Centre, measured at 10 meters above the ground level (10 magl). The series have been checked for homogeneity and coherency, see Boissonade et al. (2002). The series do not allow capturing the intra-daily periodicities, and will require some assumptions for the evaluation of power production at turbine height. We will discuss these last elements in a subsequent section. We estimate the models using data up to the 31st of December 2007, while 2008 data will serve for out-of-sample model comparison.

Strong periodic components characterize these time series. These are evident in the graph of the original series, and clearly appear both in the correlogram and in the periodogram (which peaks at zero and annual frequencies). The overall distribution of the data could suggest Gamma, Weibull, and Log-Normal densities. Figure 1 reports an example for Wielun station. Similar graphs are obtainable for the other time series.

[FIGURE 1]

3.1. Periodic components estimation

We estimate the periodic components following the approaches described in Section 2.1. If we consider a multiplicative model as in equations (1)-(2), all series have a strong yearly pattern, and some other minor periodic components. If we also add a second periodic function, and thus following model (2)-(4)-(7), we note this has a minor impact. Table 1 summarizes the coefficients associated with the estimated periodic components, while Figure 2 plots the deterministic periodic components of Wielun time series.

Once periodic components are filtered out with a purely multiplicative approach (equations (1) and (2)), the seasonally adjusted series evidences serial correlation, which might be associated with mild long-memory behaviours. In fact, the GPH and Whittle tests for fractional integration suggest, but only at the 5% level, the presence of long-memory, with estimated integration coefficients close to 0.1 (we have not reported these tests for brevity). Figure 3 reports the graphical analysis of Wielun seasonally adjusted series after the estimation of a multiplicative periodic component. Results for the other stations are similar, as well as for the case of additive and multiplicative periodic components, apart the obvious difference in the kernel density.

[FIGURES 2 AND 3]

[TABLE 1]

3.2. ARG, GAR and ARFIMA model estimation

On the seasonally adjusted series, we fit the dynamic models presented in section 2.2. Tables 2, 3 and 4 report the estimates for ARG, GAR, ARFIMA-GARCH and ARFIMA-FIGARCH (GARCH variances have been introduced to verify the advantages of long-memory in variances), respectively.

With respect to the GAR estimates, we note that these require the inclusion of higher order lags. In fact, we have included lags up to order 10 (selected using a general to specific approach), suggesting that the presence of mild long-memory effects (however, standardized GAR residuals do not show serial correlation, supporting the estimate of short-memory GAR specifications). Parameters evidencing a relevant persistence characterize the conditional variances provided by the GAR model, similar to what is

observed for financial assets. The ARG model does not require lags higher than 5. Notably, the parameters of ARG are far from the long-memory limit (the β should sum to 1).

In contrast, ARFIMA mean models return estimated memory coefficients ranging from 0.17 (Wielun) to 0.2 (Sulejow), while only an MA(1) term is needed to capture short-memory behaviours. In the variances, GARCH and FIGARCH specifications are close one to the other, with a clear preference for short-memory structures for Wielun and long-memory for Lodz, while Sulejow suggests the need for long-memory, but only at the 5% level (on the basis of Likelihood Ratio tests). For ARFIMA models, we also test the normality hypothesis over residual. Several tests are all concordant and rejecting the null (not reported).

In Figures 4 and 5, we graphically compare the fitted mean and variances of wind speed intensity (thus also including periodic components) provided by the four alternative specifications for Wielun data. It clearly appears from the scatter plots that the models provide extremely close fit for the mean, while some differences appear for the fitted conditional second order moment. In particular, ARFIMA specifications have the closer fitted variances, and they sensibly differ from those of ARG and GARCH models. These last two models also have different variance patterns. We obtain similar results for Sulejow and Lodz. This fact could have relevant impact for the pricing of contracts based on the wind power production.

[TABLES 2, 3 AND 4]

[FIGURES 4 AND 5]

3.3. One-step-ahead model forecast

We now evaluate the one-step-ahead model forecasts. The evaluation sample consists of 366 days, i.e., the entire year 2008. In Table 5, we report the MAE and RMSFE measures for the wind speed intensity point forecasts across all models and stations. The two indices are concordant in suggesting a preference of ARG model for Wielun and Sulejow, while ARFIMA-FIGARCH is preferred for Lodz. We note that the models are all very close to one another. Within ARFIMA-GARCH and ARFIMA-FIGARCH, the two closest specifications, there is always a marginal preference for the second.

We search for confirmation of this result by evaluating the density forecasts produced by the different models. In this case, we consider both the Amisano-Giacomini test, and the inverse probability transform. We stress that the evaluation of density forecast assumes a central role, given the subsequent use of models for wind risk management, as we will discuss in the following section.

Table 6 reports the Amisano-Giacomini robust test statistic for all pairs of models with weights focusing on the tails, and using, as forecast density, the normal calibrated with conditional mean and variance forecasts produced by the four alternative models. Amisano-Giacomini results confirm the previous findings with respect to density forecasts: in most cases, the models provide density forecasts that are not statistically different. We observe few pairs of models with marginally different forecasting ability, in particular for Sulejow. We obtained substantially similar results with weights set all to 1 (not reported to save space).

If we consider the inverse probability transformation, the outcomes are different. In fact, as shown in Table 7, ARG provide the worst results, accepting the null of Uniform distribution only in a few cases. The other models are almost equivalent for Sulejow, while Wielun and Sulejow show a preference for GAR (by the fewer rejections of the null). The result for ARG depends on the overconcentration of the forecasts around the mean. Summarizing, the results based on the one-step-ahead forecasts suggest that the models are quite close to one another, with the exception of ARG, whose density forecasts are unsatisfactory.

[TABLES 5, 6 AND 7]

4. Wind risk management

Wind farms represent long-term investment projects for their owners. Guaranteeing a profit in the medium/long-term requires careful choice of locations. In fact, one of the most relevant features of wind is its large volatility over both intensity and direction, and the interdependence of these quantities with the localization and structure of the land (roughness). As a result, before wind farm installation, locations are analysed and monitored for at least one year, using weather stations that measure wind speed intensity

and direction at different heights. The collected data allow the development of an appropriate budget, including the expected wind power production. Alternatively, when historical data for specific locations are not available, close WMO stations can be potential proxies. However, wind speed is not a monetary quantity and we need to transform it into electricity, for which the market provides prices.

The process of transforming wind kinetic energy into electric energy is nonlinear and stochastic, due to the combined effect of wind speed intensity, wind direction, and the time taken for turbines to rotate to follow the wind direction. As an example, Figure 6 shows the observed wind speed and energy production for an already existing turbine. The empirical power production oscillates around a non-linear path that represents the theoretical power curve, i.e., the non-linear function that maps wind speed intensity into produced power (power curves are turbine-specific and provided by turbine manufacturers). Power production starts above a minimum wind speed cut-off level (generally 2 meters-per-second, m/s), gradually increases up to a maximum level (turbine full production is reached), and then stops at an upper wind speed cut-off value (depending on turbines, but generally between 25 and 35 m/s) in order to avoid damage to the turbine from excessive wind speeds (in extreme weather conditions).

[FIGURE 6]

We must consider an additional aspect in the monetary evaluation of wind power: historical data are not generally available at the turbine height (it could vary between 50 and 110 magl), but at lower heights (in most cases 10 magl for WMO stations). Therefore, we must convert the wind speed historical data into wind speed time series at the desired turbine height. One approach available in the meteorological literature is the so-called Hellman's formula:

$$x_h = x_m \left(\frac{h}{m} \right)^a \quad (26)$$

where x_h is the wind speed at the desired height h , x_m is the observed wind speed at height m , and a is a coefficient associated with the land roughness. This parameter can be

either calibrated or estimated using closer or similar locations for which data are available at two different heights. Note that the literature contains alternative approaches; see Leroy (2004) for an example.

Given the wind speed at a desired height, we determine the (theoretical or historical) power production through the power curve. However, due to the several elements affecting power production (changes in the direction and intensity over the day, hitting of the upper or lower cut-off values...) cumulative wind speed indices (CWSI) are preferred instead of power productions to price wind derivatives. We determine the CWSI as simple sums of average daily wind speed intensities over the target period, filtered from values below the lower cut-off and above the upper cut-off. For instance, with observed average daily wind speed intensities x_t from 1 to T, we can obtain the CWSI for the last year by:

$$CWSI_{T|T-364} = \sum_{t=T-364}^T x_t I(l \leq x_t \leq u) \quad (27)$$

where $I(\cdot)$ is the indicator function, while l and u are the lower and upper cut-off values, respectively. Despite the non-linear correlation between average daily wind speed and effective power production, as evidenced by Brix et al. (2005), we generally base wind contracts on wind indices determined in accordance with (27).

Wind indices represent the most relevant elements for the assessment of the economic value of a new wind farm, or to price wind related derivative contracts and insurances. Owners of wind farms may consider buying these instruments in order to offset part of the risks associated with the large variability of wind speed. From a different point of view, determination of the future theoretical wind production for specific locations could use wind indices.

Recently, the introduction of green certificates by the EU introduced a second significant source of revenues for wind farms. In fact, every megawatt of energy produced by a wind farm provides one green certificate that has a sales value in the market. Given the purposes of the actual work, we will not consider the effects of green certificates on the risk evaluation of wind farm, but will focus only on the wind energy production.

We can apply the models proposed in this paper for the simulation of the future evolution of wind speed intensities, and thus simulate the future density of wind speed indices. Before presenting some empirical evidence, we briefly introduce the approaches used to price a contract based on a wind index.

4.1. Pricing wind-related contracts

The market for weather derivatives that depend on wind indices is highly illiquid, due to the location-specific features of the contracts and the limited number of subjects trading them (the main operators are reinsurance companies). As a result, the basis of the pricing process should arise from a suitable model, following the approach of VanderMarck (2003). The standard Black and Scholes (1973) arbitrage-free pricing is unfeasible given, in that the underlying stochastic variables do not follow the geometric Brownian motion, they strongly deviate from normality, and they are not directly tradable in the market, see Dischel (2002). Pricing generally follows an actuarial approach based on the simulated distribution of contract outcomes. Simulations are conditional on historical data and, if available, on short and medium term weather forecasts, see Cao and Wei (2000, 2003), Zeng (2000), Davis (2001), and Brix et al. (2005), among others. The actuarial price is thus the discounted expected payoff of a ‘neutral’ or ‘fair’ value plus some margin that includes fixed costs and risk loading factor for the contract writer, see Henderson (2002). The simulations may follow two alternative approaches: the direct generation of wind indices, as done by Historical Burn Analysis (HBA) and Index Modelling, the simulation of the underlying wind speed intensity, as advocated by Daily Modelling. The latter method avoids some of the drawbacks of index-based approaches, as pointed out by Nelken (2000). In fact, Daily Modelling is consistent with the features of the underlying variables, such as serial correlation, long-memory, and non-normal distribution. Moreover, it is much more flexible, and can easily incorporate short-term forecasts. Finally, its basis comes from higher frequency data compared to HBA and Index Modelling, and thus it could reduce the impact of estimation errors. In this paper, we clearly follow the Daily Modelling approach and use the proposed models to simulate the future evolution of wind speed intensity. We then use the simulations to construct the simulated density of a wind index, and to price an insurance contract.

4.2. Simulating wind speed intensities and CWSI

The ARG, GAR, ARFIMA-GARCH, and ARFIMA-FIGARCH models represent a selection of alternative methods for the simulation of future paths of wind speed intensity. Before using them out-of-sample, to generate the CWSI for the year 2008, we assess their performances in-sample, following the approach of Campbell and Diebold (2005), which employs the probability inverse transform approach similar to that presented in Section 2.3. In this case, the focus is on the simulated distribution of CSWI over a target period of length h conditionally to the information set at time T . We denote this simulated distribution as $f_T^S(CWSI_{T+h})$. The $CWSI_{T+h}$ is also directly observable (we are in-sample), and thus we define the Roseblatt transform as:

$$F_{T+h}^S(CWSI_{T+h}) = \int_{-\infty}^{CWSI_{T+h}} f_T^S(v) dv \quad (28)$$

Given that our out-of-sample target is the year, we set the in-sample h equal to 1 year (either 365 or 366 days). We thus split the time series into several 1-year periods, and evaluate (28) for each year excluding the first one, using every time $n=10,000$ simulations to recover $f_T^S(CWSI_{T+h})$. Under the assumption of correct model specification, $\left\{ F_{year}^S(CWSI_{year}) \right\}_{year=1987}^{2007}$ is distributed as a Uniform in $[0,1]$. This test allows comparison of models for daily wind speed with respect to their ability to simulate the distribution of wind speed indices.

For the simulation of CWSI, we assume that the turbine rotor is at 82 magl. We obtain the 82 magl wind speed by the Hellman's formula, with roughness coefficient $a=0.305$, which we estimated using the empirical daily data (wind speed at 82 magl) from one Polish wind farm, which is in the middle of the three WMO stations we are considering. Table 8 reports the p-values for the Uniformity tests over the 21 points in the sample. The tests show evidence of a preference of ARFIMA specifications over GAR and ARG models. The Gamma-based models are substantially equivalent, while the introduction of variance long-memory does not improve over the simpler ARFIMA-GARCH model for the logarithms of wind speed intensity. We then analysed the ACF of the series defined using (28), but found no significant serial correlation (results not reported), a further result supporting model adequacy. The different performances of the models strongly depend on the very different shapes of the simulated CWSI densities. As an example,

Figure 7 shows the CWSI densities for the year 2008 (similar results are obtained in-sample). The densities differ for the mean and the shape: ARFIMA models provide CWSI with lower mean and fatter tails. This is a by-product of the overconcentration of ARG models around the mean (thus, less simulated wind speed intensities below and above the cut-off values for energy production, and thus larger CWSI); while for GAR, the effect is due to the inappropriateness of the model in capturing tail events. Finally, there is a strong difference in the conditional variances of the models, as evidenced in section 3; a fact that plays a relevant role in the construction of CWSI. All models provide densities far from being normal, as confirmed by several tests (not reported for brevity). Table 9 reports the expected mean and standard deviations of 2008 CWSI for the three weather stations. Notably, the averages and standard deviations of ARFIMA models are closer in mean and variance to the historical moments. Summarizing, all the previous differences have a relevant impact in the wind risk management framework, and we thus conclude that ARFIMA specifications are preferred over ARG and GAR.

4.3. Wind risk management

In order to examine the pricing abilities of the presented wind models, we apply them for pricing a protective capped put option, a standard contract in the weather derivatives market. This type of contract would be of interest to a wind farm owner who wants to hedge the risk of a low power production over a given period (in our case, the year 2008). The underlying instrument of the put is the CWSI. We represent the payout of the contract as follows:

$$P = \min\left(\max\left((K - CWSI)\tau, 0\right), cap\right) \quad (29)$$

where CWSI is the observed wind speed index over the target period, K is the strike value of the CWSI, τ is the tick value (monetary value of one point of the CWSI), and cap is the cap value of the contract. In the example we set $\tau = 10,000\text{€}$, $cap = 1,000,000\text{€}$, and K at the average yearly CWSI of the last 10 years. The contract premium is determined as follows: First, we determine the expected payoff of (29), evaluated using the simulated CWSI densities for 2008. This quantity is then increased by a risk loading factor set at 6% of the Return-on-Value-at-Risk (RoVaR) of the CWSI

density at the 95% confidence level, by a 3% management and transaction costs and by fixed margin of 10,000 €. Finally, we discount the total future premium with a 3% risk-free rate. We then compare the theoretical prices obtained from the simulated CWSI to those obtained by HBA, which we apply following Jewson et al. (2005). Note that HBA, by using a limited number of data (in our case 21 observations), contains a huge estimation error, an element we try to minimize by daily modelling. We prefer HBA to index modelling as it is much more diffused among practitioners.

Table 9 contains all contract prices for the three stations. For all localizations, the put options, with strikes fixed as 10-year averages, appeared to be deeply out-of-money. ARG or GAR models always gave the lowest prices, because their simulated densities of CWSI indices overestimated the mean and underestimated the variance. As a result, by using these models to price wind options, we have received a final price that is slightly higher than the discounted value of the brokerage margin. ARFIMA models with short and long memory in conditional variances provide prices higher than ARG and GAR models, and lower than the HBA approach. The only exception was Sulejow, where considerably lower values in the mean of the simulated CWSI index, coupled with higher variances, made this option more expensive. Ex-post, comparing the realized 2008 CWSI with the simulated densities, we note how all contracts were not exercised.

5. Conclusion

In this paper, we compared three alternative approaches for the estimation and simulation of wind speed intensities for the pricing of wind options. We show how we could fit the models over real data and we compared them from a statistical point of view by contrasting their density forecasts. We also tackled the point of view of a weather risk manager and considered the models as tools for the pricing of wind options. Our study provides a framework for the comparison of these alternative models and, for the series we used, ARFIMA-based specifications provided better results.

References

- Aillot, P., Monbet, V., and Prevosto, M., 2006, An autoregressive model with time-varying coefficients for wind fields, *Environmetrics*, 17, 107-117
- Allen, D., Lazarov, Z., McAleer, M., and Peiris, S., 2009, Comparison of alternative ACD models via density and interval forecasts: evidence from the Australian stock market, *Mathematics and Computers in Simulation*, 79, 2535-2555
- Amisano, G., and Giacomini, R., 2007, Comparing density forecasts via weighted likelihood ratio tests, *Journal of Business and Economic Statistics*, 25-2, 177-190
- Beine, M. and S. Laurent, 2003, Central bank interventions and jumps in double long memory models of exchange rates, *Journal of Empirical Finance*, 10, 641-660
- Black, F., and Scholes, M., 1973, The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 637-654.
- Boissonnade, A.C., Heitkemper, L.J. and Whitehead, D., 2002, Weather data: Cleaning and Enhancement, in *Climate Risk and the Weather Market*, Risk Books.
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 31, 307-327.
- Brix, A., Jewson, S. and Ziehmann, C., 2005, *Weather Derivative Valuation*, Cambridge: Cambridge University Press
- Brown, B.G., Katz, R.W., and Murphy, A.H., 1984, Time series models to simulate and forecast wind speed and wind power, *Journal of Climate and Applied Meteorology*, 23, 1184-1195
- 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., 2000, Pricing the weather. *Risk*, 13-5, 67-70.
- Cao, M. and Wei, J., 2003, Weather derivatives valuation and market price of weather risk, *Journal of Future Markets*, 24, 1065-1090
- Caporin, M., and Prés, J., 2009, Forecasting temperature indices with time-varying long-memory models, available in SSRN
- Castino, F., Festa, R., and Ratto, C.F., 1998, Stochastic modelling of wind velocities time series, *Journal of Wind Engineering and Industrial Aerodynamics*, 74-76, 141-151
- Davis, M., 2001, Pricing weather derivatives by marginal value, *Quantitative Finance*, 1, 1-14
- Diebold, F.X., Gunther, A., and Tay, S., 1998, Evaluating density forecasts with application to financial risk management, *International Economic Review*, 39, 863-883
- Dischel, R.S. (Ed.), 2002, *Climate risk and the weather market: financial risk management with weather hedges*, Risk Publications, London
- Engle, R.F., 2002, Dynamic conditional correlation, a simple class of multivariate GARCH models, *Journal of Business and Economic Statistics*, 20, 339-350.
- Engle R., and Russell, J., 1998, The autoregressive conditional duration model: a new model for irregularly spaced data, *Econometrica*, 66, 1127-1162.

- Gaver, D., and Lewis, P., 1980, First order autoregressive gamma sequences and point processes, *Advances in Applied Probability*, 12, 727-745
- Guégan, D., and Diongue, A.K., 2009, Temperature modeling and weather derivatives, *Proceedings of the Third International Conference on Computational and Financial Econometrics*, Cyprus
- Gourieroux, C., and Jasiak, J., 2006, Autoregressive Gamma Processes, *Journal of Forecasting*, 25, 129-152
- Henderson, R., 2002, Pricing Weather Risk, in *Weather Risk Management, Markets, Products and Applications*, ed. Banks E., 167-196.
- Lawrance, A, Lewis, P., 1981, A new autoregressive time series model in exponential variables, *Advances in Applied Probability*, 13, 826–845.
- Leroy, A., 2004, Design and valuation of wind derivatives, *White Papers*, Solvay Business School
- Nelken, I., 2000, *Weather Derivatives – Pricing and Hedging*, Super Computer Consulting, Inc. IL. USA , January.
- Pacurar, M., 2008, Autoregressive conditional duration models in finance : a survey of the theoretical and empirical literature, *Journal of Economic Surveys*, 22-4, 711-751
- Rosenblatt, M., 1952, Remarks on multivariate transformations, *Annals of Mathematical Statistics*, 23, 470-472
- Smith, S., Miller, J., 1986, A non-Gaussian state space model and application to prediction of records, *Journal of the Royal Statistical Society, Series B*, 48, 79–88.
- Tol, R.S.J., 1997, Autoregressive conditional heteroscedasticity in daily wind speed measurements, *Theoretical and Applied Climatology*, 56, 113-122.
- VanderMarck, P., 2003, Marking to model or to market?, *Environmental Finance*, 1, pp.36-38.
- Yamada, Y., 2008, Simultaneous optimization for wind derivatives based on prediction errors, *Proceedings of the American Control Conference 2008*, 11-13, June 2008, 350-355, DOI: 10.1109/ACC.2008.4586515
- Zeng, L., 2000, Weather derivatives and weather insurance: concept application and analysis, *Bulletin of the American Meteorological Society*, 81, 2075-2982

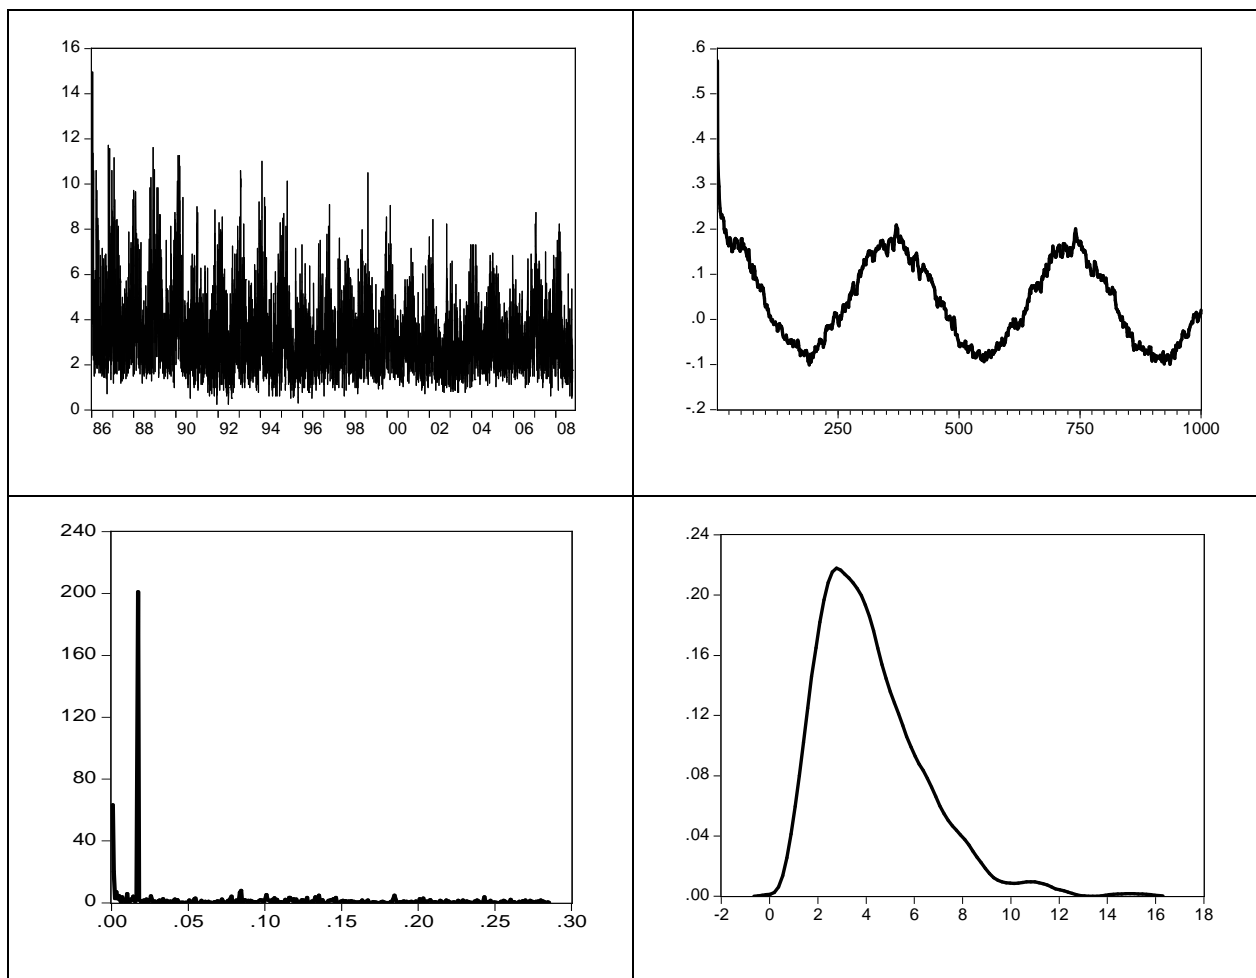


Figure 1. Wielun graphical analysis: Wielun data, ACF, Periodogram (selected frequencies up to $\pi/10$), and kernel density estimate.

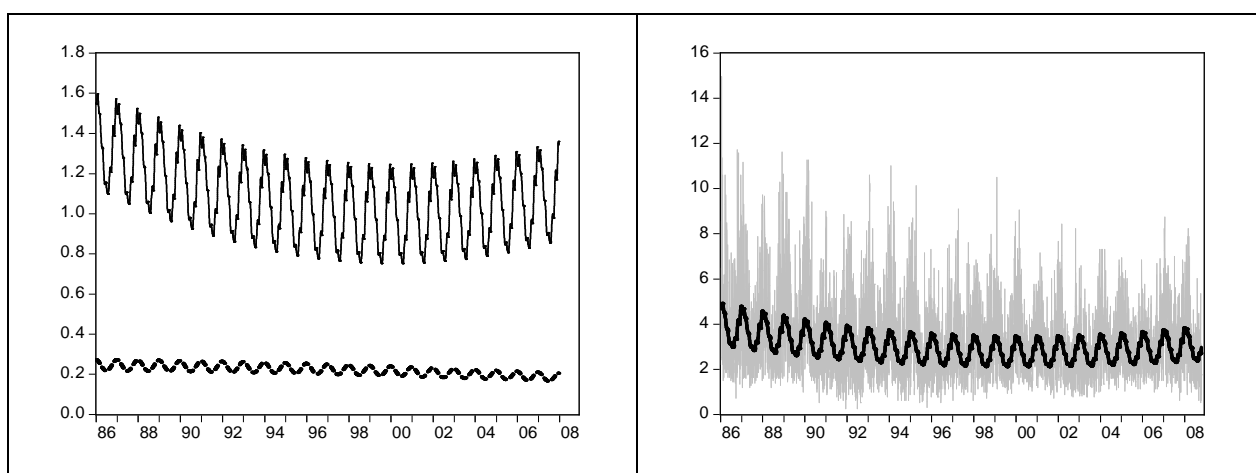


Figure 2. Periodic components for Wielun: left panel, additive (upper line) and multiplicative (lower line) components of Wielun log-transformed data; right panel, Wielun data (grey) superimposed to the exponential of the additive periodic component of the left panel.

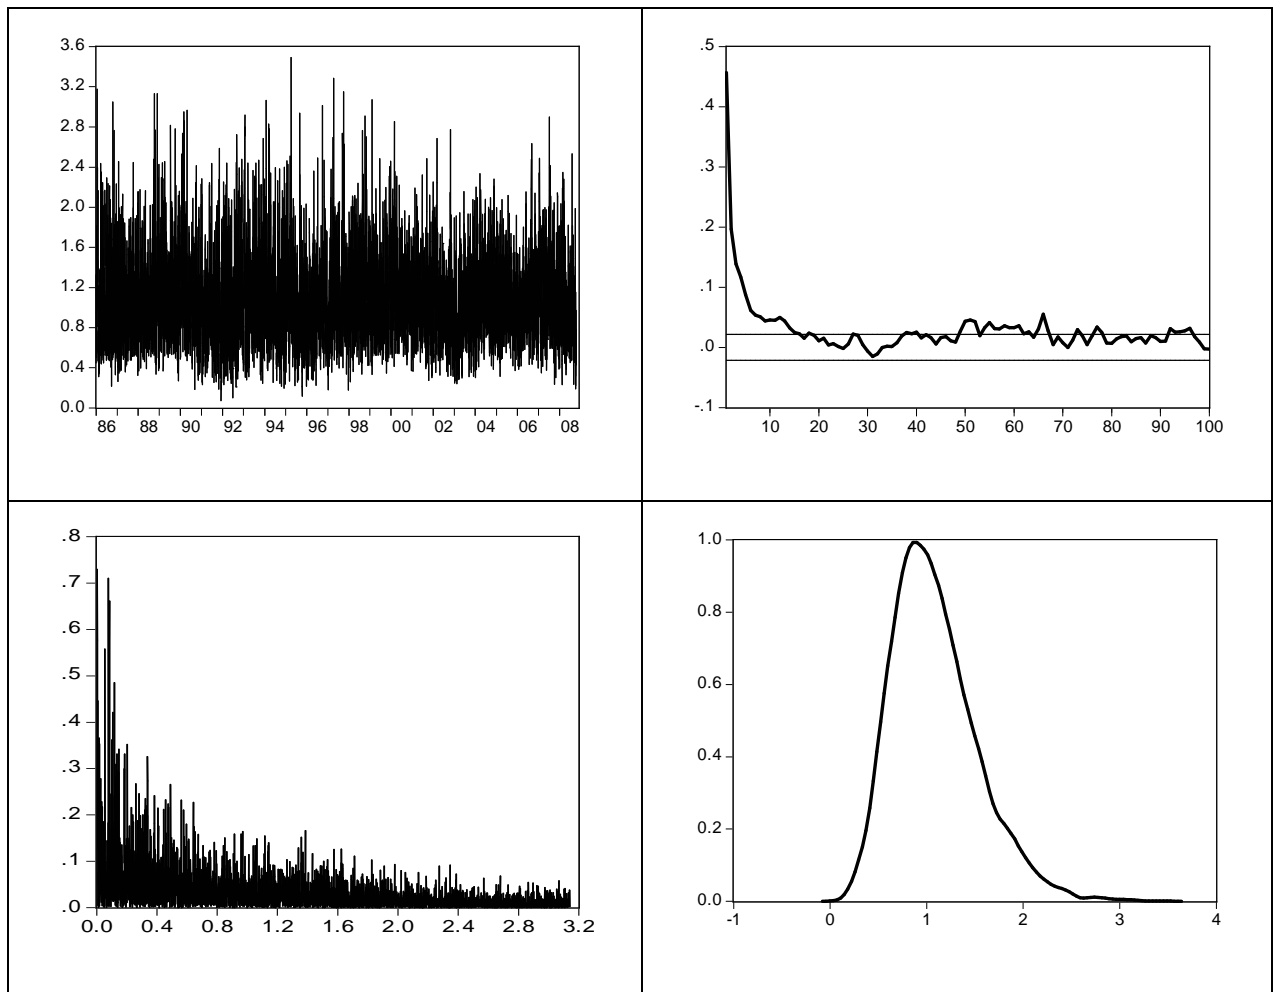


Figure 3. Wielun graphical analysis of "seasonally adjusted" series after removing a multiplicative periodic component: Wielun "seasonally adjusted" series, ACF (straight lines denotes the 5% confidence interval), Periodogram, and kernel density estimate.




Figure 4: scatter plot of the in sample fitted conditional mean for Wielun wind speed intensity provided by the ARFIMA-FIGARCH, ARFIMA-GARCH, ARG and GAR models




Figure 5: scatter plot of the in sample fitted conditional variances for Wielun wind speed intensity provided by the ARFIMA-FIGARCH, ARFIMA-GARCH, ARG and GAR models




Figure 6: Empirical power curve extracted from 10 minutes observations for a turbine located in Poland.

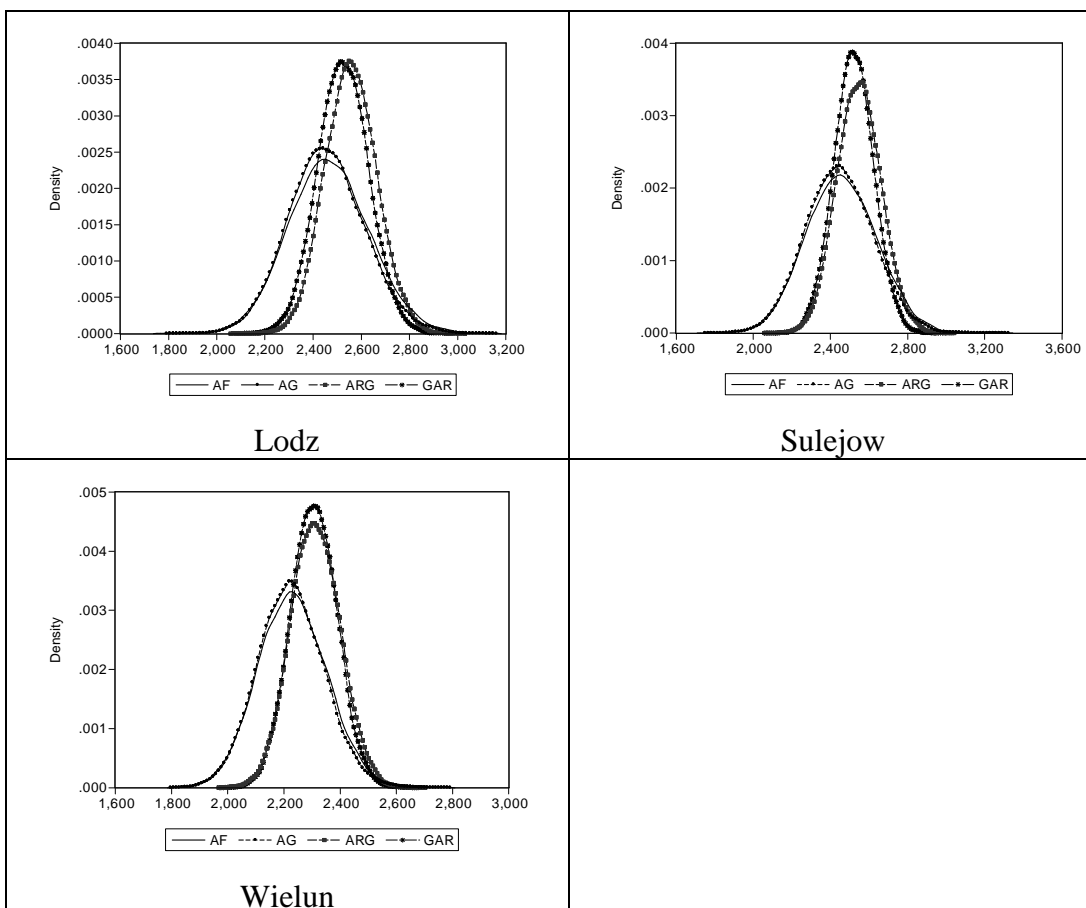


Figure 7: Simulated CWSI kernel density estimates for 2008 from ARFIMA-GARCH (AG), ARFIMA-FIGARCH (AF), ARG, and GARCH (GAR) models

Table 1: Parameter estimates of periodic components

	Wielun		Sulejow		Lodz	
	Coeffs.	T-stat	Coeffs	T-stat	Coeffs	T-stat
Coefficients of $s(t)$						
β_0	1.525	39.563	1.387	28.281	1.544	33.255
β_1	-0.174	-9.501	-0.104	-4.611	-0.200	-9.145
β_2	0.014	7.745	0.009	4.183	0.018	8.274
φ_1	0.213	16.731	0.153	9.710	0.157	9.499
φ_6	0.024	2.051				
ψ_1	0.071	5.971	0.074	4.952	0.086	5.985
ψ_7	-0.024	-2.165	-0.035	-2.604	-0.027	-2.081
ψ_8	-0.021	-1.968	-0.031	-2.392	-0.026	-2.091
Coefficients of $v(t)$						
α_0	-2.815	-58.441	-2.643	-54.648	-2.747	-47.482
α_2	-0.007	-5.888	-0.002	-1.902	-0.003	-2.320
δ_1	0.226	5.453	0.296	7.044	0.169	3.761

The table reports the significant coefficients and the corresponding robust T-statistics.

Table 2: ARG estimation output

	Wielun		Sulejow		Lodz	
	Coeffs.	T-stat.	Coeffs	T-stat	Coeffs	T-stat
δ	5.633	37.033	4.116	32.758	4.689	31.342
β_1	0.447	37.893	0.490	41.157	0.460	39.620
β_3	0.034	2.725	0.049	3.956	0.034	2.722
β_4	0.038	3.061			0.027	1.997
β_5			0.031	2.690	0.035	2.943
c	0.093	46.543	0.105	48.638	0.117	46.076

The table reports the estimated coefficients and the maximum likelihood t-statistics.

Table 3: GAR estimation output

	Wielun		Sulejow		Lodz	
	Coeffs.	T-stat.	Coeffs	T-stat	Coeffs	T-stat
γ_0	0.550	28.726	0.528	30.173	0.517	24.762
γ_1	0.389	36.759	0.406	41.821	0.402	37.390
γ_3	0.031	2.847	0.042	4.451	0.038	3.991
γ_4	0.034	3.301	0.029	3.418	0.033	3.340
γ_7	0.017	2.067				
γ_8			0.019	1.911	0.025	2.398
γ_9	0.024	2.748				
γ_{10}			0.028	2.978	0.033	3.288
ω	0.003	0.743	0.023	3.187	0.028	1.818
α_1	0.024	1.377	0.061	5.347	0.064	3.442
β_1	0.956	21.721	0.824	18.047	0.793	8.235

The table reports the estimated coefficients and the maximum likelihood t-statistics.

Table 4: ARFIMA-FIGARCH and ARFIMA-GARCH estimation output

	Wielun		Sulejow		Lodz	
	Coeffs.	T-stat.	Coeffs	T-stat	Coeffs	T-stat
ARFIMA						
d	0.170	13.391	0.201	15.601	0.179	14.939
θ_1	-0.266	-18.090	-0.255	-17.298	-0.261	-18.015
GARCH						
ω	0.010	1.408	0.051	1.314	0.049	1.186
α_1	0.008	3.770	0.014	2.477	0.020	2.102
β_1	0.989	246.906	0.967	50.803	0.964	43.119
FIGARCH						
ω	0.343	3.377	0.368	3.611	0.333	2.824
ψ_1	0.719	10.502	0.694	11.533	0.677	8.063
λ	0.153	5.718	0.147	4.519	0.183	5.090
β_1	0.616	8.816	0.585	8.913	0.560	6.380
LR test	0.381		0.025		$<10^{-4}$	

The table reports the estimated coefficients and the maximum likelihood t-statistics. The estimated GARCH(1,1) conditional variances follow $\sigma_t^2 = \omega + \beta\sigma_{t-1}^2 + \alpha\varepsilon_{t-1}^2$. The last row report the p-value of the likelihood ratio test for the null hypothesis of no long-memory in conditional variances. Under the null, the estimated FIGARCH models collapse on GARCH specifications with $\alpha = \psi - \beta$. The test statistic has a Chi-square distribution with one degree of freedom.

Table 5: MAE and RMSFE

Weather Station	Model/Indicator	ARFIMA GARCH	ARFIMA FIGARCH	ARG	GAR
Wielun	MAE	0.911	0.910	0.909	0.912
	RMSFE	1.174	1.171	1.170	1.175
Sulejow	MAE	1.115	1.113	1.097	1.181
	RMSFE	1.473	1.471	1.431	1.550
Lodz	MAE	1.097	1.096	1.116	1.156
	RMSFE	1.440	1.437	1.465	1.492

Bold values identify the preferred model.

Table 6: Amisano-Giacomini test statistics

Weather Station	Model	ARG	GAR	ARFIMA GARCH
Lodz	ARFIMA-FIGARCH	-0.721	0.571	1.474
	ARFIMA-GARCH	-1.231	0.029	
	GAR	-1.393		
Sulejow	ARFIMA-FIGARCH	0.388	1.923	2.099
	ARFIMA-GARCH	0.469	1.630	
	GAR	-1.641		
Wielun	ARFIMA-FIGARCH	-0.654	-0.117	2.898
	ARFIMA-GARCH	-1.752	-1.077	
	GAR	-0.582		

Amisano-Giacomini robust test statistic for equal forecasting ability. The test statistic is distributed as a standardized normal. The test is based on the log-score of the row model minus the log-score of the column model. Positive (negative) statistically significant test statistics should be interpreted as a preference for the row (column) model. In bold we report non significant test statistics at the 1% confidence level (one sided).

Table 7: testing Uniform distribution of inverse probability transform

	Lodz	Sulejow	Wielun	Lodz	Sulejow	Wielun
	ARFIMA-FIGARCH			ARFIMA-GARCH		
Kolmogorov-Smirnov (D+)	0.542	0.474	0.126	0.573	0.461	0.160
Kolmogorov-Smirnov (D-)	0.054	0.208	0.053	0.054	0.206	0.053
Kolmogorov-Smirnov (D)	0.107	0.411	0.105	0.107	0.409	0.106
Kuiper (V)	0.045	0.177	0.002	0.052	0.166	0.003
Cramer-von Mises (W2)	0.160	0.554	0.117	0.156	0.577	0.156
Watson (U2)	0.020	0.244	0.004	0.019	0.269	0.009
Anderson-Darling (A2)	0.141	0.426	0.015	0.150	0.473	0.038
	GAR			ARG		
Kolmogorov-Smirnov (D+)	0.412	0.293	0.299	0.142	0.331	0.283
Kolmogorov-Smirnov (D-)	0.137	0.741	0.805	0.011	0.028	0.016
Kolmogorov-Smirnov (D)	0.274	0.570	0.583	0.021	0.056	0.032
Kuiper (V)	0.078	0.575	0.673	0.000	0.005	0.002
Cramer-von Mises (W2)	0.322	0.570	0.367	0.020	0.059	0.038
Watson (U2)	0.116	0.373	0.544	0.001	0.002	0.000
Anderson-Darling (A2)	0.421	0.569	0.157	0.001	0.003	0.006

The table reports the p-values of the various tests included in the first column. We report in bold cases where the null hypothesis of uniform distribution is rejected at the 5% level.

Table 8: testing Uniform distribution of inverse probability transform for CWSI

	Lodz	Sulejow	Wielun	Lodz	Sulejow	Wielun
	ARFIMA-FIGARCH			ARFIMA-GARCH		
Kolmogorov-Smirnov (D+)	0.854	0.295	0.772	0.579	0.132	0.340
Kolmogorov-Smirnov (D-)	0.477	0.610	0.301	0.843	0.844	0.636
Kolmogorov-Smirnov (D)	0.854	0.574	0.586	0.948	0.264	0.654
Kuiper (V)	0.923	0.399	0.612	0.968	0.375	0.497
Cramer-von Mises (W2)	0.949	0.753	0.865	0.818	0.482	0.654
Watson (U2)	0.929	0.502	0.703	0.884	0.412	0.669
Anderson-Darling (A2)	0.961	0.666	0.837	0.907	0.437	0.775
	GAR			ARG		
Kolmogorov-Smirnov (D+)	0.278	0.241	0.064	0.391	0.562	0.104
Kolmogorov-Smirnov (D-)	0.138	0.173	0.082	0.261	0.103	0.073
Kolmogorov-Smirnov (D)	0.275	0.344	0.128	0.514	0.206	0.146
Kuiper (V)	0.033	0.036	0.001	0.160	0.098	0.002
Cramer-von Mises (W2)	0.263	0.256	0.088	0.392	0.284	0.085
Watson (U2)	0.035	0.035	0.001	0.102	0.098	0.001
Anderson-Darling (A2)	0.000	0.110	0.027	0.000	0.155	0.000

The table reports the p-values of the various tests included in the first column. We report in bold cases where the null hypothesis of uniform distribution is rejected at the 10% level (we chose a more conservative confidence level with respect to that of Table 7 given the limited number of observations in our sample (the sample includes 21 observations)).

Table 9: Estimated contract prices

		Wielun	Sulejow	Lodz
Historical CSWI average		2238.1	2404.6	2271.5
Historical CWSI standard deviation		262.9	189.7	294.5
Strike price (CWSI points)		2097	2353	2128
Expected mean	Historical Burn Analysis	2255.7	2534.1	2510.2
	ARG model	2313.1	2548.2	2557.4
	GAR model	2307.0	2520.7	2525.2
	AG model	2226.5	2448.6	2456.1
	AF model	2231.2	2457.1	2467.1
Standard deviation	Historical Burn Analysis	120.1	150.5	199.6
	ARG model	87.9	114.4	107.1
	GAR model	81.9	103.1	105.2
	AG model	116.6	173.4	157.4
	AF model	122.1	182.4	167.4
Contract price	Historical Burn Analysis	142,088 €	47,856 €	21,233€
	ARG model	11,377 €	26,271 €	9,746 €
	GAR model	10,757 €	31,222 €	9,723 €
	AG model	108,111 €	272,727 €	16,575 €
	AF model	114,166 €	271,004 €	18,425 €
Realized CWSI		2205.1	2556.6	2625.3

The table reports the simulated CWSI mean and standard deviations and the price for a capped put option with the strike value reported in the first row.