Short-term Density Forecasting of Wave Energy Using ARMA-GARCH Models and Kernel Density Estimation

Jooyoung Jeon*

School of Management, University of Bath

James W. Taylor

Saïd Business School, University of Oxford

Forthcoming in the Special Issue of the *International Journal of Forecasting* on Probabilistic Energy Forecasting

* Address for Correspondence:
Jooyoung Jeon
School of Management
University of Bath
Bath, BA2 7AY, UK
Tel: +44 (0)1225 386 742
Fax: +44 (0)1225 386 473
Email: j.jeon@bath.ac.uk

Short-term Density Forecasting of Wave Energy Using ARMA-GARCH Models and Kernel Density Estimation

Abstract

Wave energy has great potential as a renewable source of electricity. Installed capacity is increasing, and with developments in technology, wave energy is likely to play an important role in the future mix of electricity generation. The short-term forecasting of wave energy is needed for the efficient operation of wave farms and power grids, as well as for energy trading. The intermittent nature of wave energy motivates the use of probabilistic forecasting. In this paper, we evaluate the accuracy of probabilistic forecasts of wave energy flux from a variety of methods, including unconditional and conditional kernel density estimation, univariate and bivariate autoregressive moving average generalised autoregressive conditional heteroskedasticity (ARMA-GARCH) models, and a regression-based method. The bivariate ARMA-GARCH models are implemented with different pairs of variables, such as (1) wave height and wave period, and (2) wave energy flux and wind speed. Our empirical analysis uses hourly data from the FINO1 research platform in the North Sea to evaluate density and point forecasts, up to 24 hours ahead, for the wave energy flux. The empirical study indicates that a bivariate ARMA-GARCH model for wave height and wave period led to the greatest accuracy overall for wave energy flux density forecasting, but as the lead time increases, its usefulness reduces for point forecasting. This model also performed well for wave power data that had been generated from wave height and wave period observations using a conversion matrix.

Key words: Wave energy; Probability density; ARMA-GARCH; Kernel density estimation.

1. Introduction

The growing demand for energy and sustained efforts to reduce carbon emissions, in light of the threat of global climate change, have led to the development of new technologies for harnessing energy from renewable sources. Among renewable energy sources, ocean energy has relatively high energy density in the form of currents, waves, heat and tides, with the benefit of reduced visual impact (Brekken et al., 2009). Estenban and Leary (2012) predict ocean energy could deliver around 7% of global energy production by 2050. Among the ocean energy sources, which include tidal and thermal, wave energy is known to have the greatest capacity (Clément et al., 2002; Falnes, 2007).

Wave energy forecasts from a few hours to several days ahead are of interest for the management of power grids (Pinson et al., 2012; Reikard et al., 2015). When the technology is fully commercialised, inaccurate forecasts for wave power are likely to induce penalties for wave generators that reflect over- or under-production, and increase the cost of the spinning reserve needed when the wave energy production is integrated into the power system. To ensure reliable grid operation, system operators are interested in probabilistic forecasts, such as probability density forecasts, as these deliver a fuller description of the future energy.

Ocean waves are generated mainly by the wind blowing over the sea surface, which propagates its energy horizontally on the surface. The rate of transfer of the energy is called the wave energy flux, and it is a nonlinear function of wave height and period. Due to its nonlinear dependence on meteorological and ocean variables, forecasting the wave energy flux is challenging, even for a lead time of just a few hours. As data on the power output from wave farms has not been available, researchers have largely focused on the wave energy flux, which is the energy that theoretically could be obtained. In this paper, we follow the literature by having the wave energy flux as our main focus. However, as the electricity obtained in practice from wave power devices can differ notably from the wave energy flux, we also

consider the prediction of wave power data, which we generate from wave height and wave period observations using a theoretical conversion matrix.

Studies of wave energy forecasting have generally involved the modelling of wave height and period using statistical and physics-based approaches. The statistical models used for this application include time-varying parameter regressions (see Reikard, 2009, 2013), unconditional kernel density estimation (see Ferreira and Guedes-Soares, 2002), neural networks (see, for example, Zamani et al., 2008), and autoregressive models (see, for example, Guedes-Soares and Ferreira, 1996; Guedes-Soares and Cunha, 2000; Fusco and Ringwood, 2010). Physics-based models are used by Hasselmann et al. (1976, 1980, 1985) and Janssen (1991, 2007). Combined models using both time series and physics approaches are used by Roulston et al. (2005), Woodcock and Engel (2005), Woodcock and Greenslade (2006), Durrant et al. (2008), Reikard et al. (2011) and Pinson et al. (2012).

In terms of probabilistic forecasting for wave energy, we are aware of only one published paper. Pinson et al. (2012) model wave energy flux under the assumption that the density is log-normal. Using time-varying recursive regression, they model the errors from atmospheric model forecasts for wave height and wave period up to 48 hours ahead.

In this paper, we provide an empirical comparison of statistical approaches to producing density forecasts of wave energy flux. We consider a regression-based method, conditional kernel density estimation, and ARMA-GARCH models. Our study is the first that we are aware of to use ARMA-GARCH models in the wave energy forecasting context. As statistical models have been found to be more useful than atmospheric models for shorter lead times (Reikard and Rogers, 2011; Reikard, 2013), we consider lead times from 1 to 24 hours.

Section 2 describes our data. Section 3 presents the forecasting methods that we consider in our empirical analysis of Section 4 for wave energy flux. In Section 5, we present empirical results for wave power. Section 6 summaries, and provides concluding comments.

2. THE DATA FOR WAVE ENERGY FORECASTING

Wave energy flux is the average horizontal transport momentum of a wave, per unit of wave-crest length. It is a function of wave height and wave period as in the following:

$$E_t \backsimeq (\rho g^2/64\pi) H_t^2 P_t \quad (1)$$

E_t is the wave energy flux, which is expressed in kilowatts per metre of wave-crest length; H_t is the significant wave height, which is defined as the average height of the highest third of the waves, where height is measured in metres from trough to crest; P_t is the mean wave period in seconds, which is the average of the time taken for two successive wave crests to pass a given point; g is the acceleration due to gravity; and ρ is ocean water density, which is approximately 1025 kg/m^3 at the sea surface.

We used data recorded at the FINO1 research platform, which is located in the North Sea approximately 45km to the north of Borkum, Germany. FINO1 was brought into service to facilitate the research in offshore wind energy by providing meteorological and oceanographic data. Raw wind speed data (m/s) are measured at FINO1 every second, and then averaged every 10 minutes. Raw wave height and wave period data are measured every half second from a buoy, which follows the movement of the sea surface and is located 200 metres away from FINO1. The significant wave height and peak wave period, which is defined as the wave period where the highest energy occurs among the individual wave periods, are derived every 30 minutes following the definitions below expression (1). As mean wave period data was not available, we instead used peak wave period observations. In our study, wind speed, significant wave height and peak wave period data were averaged to give hourly observations. We used hourly values of significant wave height and peak wave period to obtain hourly values of wave energy flux based on expression (1).

To avoid long periods of missing observations, two periods were chosen: (a) 1 June 2011 to 30 June 2012, inclusive, and (b) 1 September 2012 to 31 October 2013, inclusive. We

[1] http://www.fino1.de/

name these periods FINO1a and FINO1b, respectively. For each period, the last two months are used for post-sample evaluation, with the earlier observations used for model fitting. We evaluated forecasts produced using each period of the evaluation sample as forecast origin.

In FINO1a, one observation for wave height and wave period is missing from the estimation sample, and none from the evaluation sample. In FINO1b, seven separate observations for wave height and wave period are missing from the estimation sample, and one from the evaluation sample. For these periods with missing observations, we used linear interpolation to generate values. We did not evaluate forecast accuracy for these periods when they occurred in the evaluation sample. We used wind speed observations from a height of 33m, which was the lowest height available. Around 3.5% of the measurements at this height were missing. Any missing wind speed observations were replaced with the measurements at 50m, 60m or 70m height after level adjustment, which uses linear scaling to minimise the bias from the different dynamics of wind speed at different heights.

Fig. 1 shows the time series of wind speed, wave height, wave period, and wave energy flux for FINO1b. All four series are highly volatile, bounded below by zero, and skewed to the right. This is shown in the first column of histograms in Fig. 2. The other histograms in Fig. 2 relate to transformations of the series, and we discuss this further in Section 3.3.1. It is interesting to see from Fig. 1 that the variation in wave height appears to be related to the variation in wind speed. The correlation between these two series is 0.74. The wave energy flux seems to be more related to the wave height than the wave period. It is not clear that the wave period has any similar movement with the wind speed, and this is supported by the correlation between wind speed and wave period, which is -0.04.

3. DENSITY FORECASTING METHODS FOR WAVE ENERGY FLUX

In this section, we describe the methods that we implement in our empirical study. We present a regression approach, kernel density estimators, and ARMA-GARCH models.

3.1. Regression models

Reikard (2009) finds that the best approach to the point forecasting of wave energy flux, up to 4 hours ahead, is to use independently estimated least squares regression models for wind speed, wave period and wave height, as shown in expressions (2) to (4).

$$\log S_t = \theta_{0t} + \theta_{1t} \log S_{t-1} + \theta_{2t} \log S_{t-2} + \theta_{3t} \log S_{t-3} + \theta_{24t} \log S_{t-24} + \epsilon_t \quad (2)$$

$$\log P_t = \nu_{0t} + \nu_{1t} \log P_{t-1} + \nu_{2t} \log P_{t-2} + \nu_{3t} \log P_{t-3} + \nu_{4t} \log P_{t-4} + \epsilon_t \quad (3)$$

$$\log H_t^2 = \delta_{0t} + \delta_{1t} \log H_{t-1}^2 + \delta_{2t} \log H_{t-2}^2 + \delta_{3t} \log H_{t-3}^2 + \delta_{4t} \log H_{t-4}^2 + \delta_{5t} \log P_t + \delta_{6t} \log S_t + \epsilon_t \quad (4)$$

P_t and H_t were defined in Section 2, and S_t denotes wind speed at time t. θ_{it}, ν_{it} and δ_{it} are parameters estimated separately for each expression and for each forecast origin using ordinary least squares. ϵ_t is assumed to be Gaussian white noise, implying that the energy flux E_t of expression (1) follows a conditional log-normal distribution, which was also assumed by Pinson et al. (2012). The resulting forecasts for wave period and wave height are plugged into expression (1) to deliver a wave energy flux forecast.

3.2. Kernel density estimation

Kernel density estimation is a nonparametric approach, which has the appeal of avoiding the need for a distributional assumption. We used both unconditional and conditional kernel density estimation in our modelling. In both approaches, the kernel bandwidth parameters were chosen by minimising the mean continuous ranked probability score (CRPS) calculated for the in-sample period. The CRPS, which is described by Gneiting et al. (2007), assesses the calibration and sharpness properties of density forecasts. They explain that calibration measures the statistical consistency between the predicted density and the observed value, while sharpness refers to the concentration of the density forecast, which is a property of the density forecast alone. By contrast with the other methods that we considered, we applied the kernel density estimation methods to untransformed data.

3.2.1. Unconditional kernel density estimation (UKDE)

As a relatively simple benchmark method, we used kernel density estimation. The unconditional kernel density estimator of wave energy flux (E-UKDE) is defined as:

$$f(e) = \sum_{t=n-k+1}^{n} K_{h_e}(E_t - e)/k \ ,$$

where e is a value of wave energy flux for which a density is to be estimated; n is the forecast origin; k is the length of the sliding window used for the estimation; and K is a Gaussian kernel function, with bandwidth h_e, which dictates the smoothness of the estimated density. Having observed that a relatively small sliding window performs better for short-term forecasting of wave energy flux, we considered two versions of the E-UKDE approach using the following sliding window lengths, k: (a) 4 hours and (b) 24 hours.

3.2.2. Conditional kernel density estimation (CKDE)

Given that wind passing over the sea surface generates wave energy, we also considered the kernel density estimation of wave energy flux conditional on wind speed. We implemented the two-step conditional kernel density estimation approach of Jeon and Taylor (2012), which allows for a stochastic conditioning variable. As shown in Fig. 3(a), for our data, the wave energy flux was most highly correlated with the third lag of wind speed, and so we conditioned on this lag. Miller (1958) and Rieder (1997) find the time lag between the rise of wind speed and the rise of wave height varied between several hours and twenty hours depending on location and the persistence of wind speed and direction. Using an exponential decay parameter τ ($0< \tau \leq 1$), and an additional kernel for wind speed with bandwidth, h_s, the density estimate of wave energy flux, conditional on wind speed (E-CKDE), is given as:

$$f(e|s) = \frac{\sum_{t=1}^{n} \tau^{n-t} K_{h_s}(S_{t-3} - s) K_{h_e}(E_t - e)}{\sum_{t=1}^{n} \tau^{n-t} K_{h_s}(S_{t-3} - s)}.$$

This estimator can be viewed as a weighted average of the kernels $K_{h_e}(E_t - e)$, where the weights are larger for more recent data and for observations for which the wind speed at time

t was closer to the conditioning wind speed s. The approach requires density forecasts of wind speed. We generated these using a univariate ARMA-GARCH model, of the type discussed in Section 3.3.2, with no exogenous variable, fitted to the wind speed series.

3.3. Univariate and multivariate ARMA-GARCH models

3.3.1. Data transformation

Fig. 2 and Table 1 show that positive skewness is a feature of the unconditional distributions of wind speed, wave height, wave period, and especially wave energy flux. In Table 1, the kurtosis values for wave height, wave period and wave energy flux indicate fat tails relative to a Gaussian distribution. For time series models, transformations are often used prior to model fitting. We considered the log, square root and Box-Cox transformations.

Previous studies of wave energy forecasting have involved the use of the log transformation for wave energy flux, wave height, wave period and wind speed. We were able to apply the log transformation to each of our time series, because they contained no zero values. Table 1 and Fig. 2 show that wave height, wave period and wave energy flux are each closer to being Gaussian when the log transformation is applied.

Taylor et al. (2009) find the square root transformation useful in modelling hourly wind speed data. This transformation has the benefit that it can be used for data with zero values. Table 1 and Fig. 2 show that the square root transformation is more useful for wind speed than the log transformation, but clearly not for wave period and energy flux. For wave height, the square root transformation is slightly better for kurtosis, but not for skewness.

A third transformation that we considered is the single parameter form of the Box-Cox transformation (Box and Cox, 1964), which is given as:

$$BC(y,\lambda) = \left(y^{\lambda} - 1\right)/\lambda \quad (if\ \lambda \neq 0, y > 0) \tag{5}$$

$$= \log(y) \qquad (if\ \lambda = 0, y > 0). \tag{6}$$

This transformation has been used in modelling wave height and period (see, for example, Galiatsatou and Prinos, 2007; Ferreira and Guedes-Soares, 2002). The λ parameter can be optimised using maximum likelihood. Table 1 shows that the optimised value of λ was close to zero for wave height, wave period and wave energy flux, implying that the transformation is very similar to the log transformation.

In summary, the log transform seems to be suitable for wave height, wave period and wave energy flux, and this transformation is also reasonable for wind speed. In Section 4, we present post-sample forecasting results comparing the four different transformations.

3.3.2. Univariate ARMA-GARCH

ARMA-GARCH models are widely used for capturing the autocorrelation in the conditional mean and variance. In this paper, for wave energy flux, we use the ARMA(*r*,*m*)-GARCH(*p*,*q*) model with exogenous variables presented in expressions (7)-(9):

$$y_t = s(\boldsymbol{\mu}, t) + \sum_{i=1}^{r} \varphi_i Y_{t-i} + \sum_{j=1}^{m} \psi_j \varepsilon_{t-j}, \quad (7)$$

$$\sigma_t^2 = s(\boldsymbol{\omega}, t) + \sum_{i=1}^{p} \alpha_i \sigma_{t-i}^2 + \sum_{j=1}^{q} \beta_j \varepsilon_{t-j}^2, \quad (8)$$

$$\varepsilon_t = \sigma_t \eta_t, \quad (9)$$

where y_t is an observation of wave energy flux at time *t*; ε_t is an error term; η_t is white noise; σ_t is the conditional standard deviation (volatility); φ_i, ψ_i, α_i and β_i are the coefficients of the AR, MA, GARCH and ARCH components, and their orders are defined by non-negative integer valued constants *r*, *m*, *p* and *q*, respectively; $\boldsymbol{\mu}$ and $\boldsymbol{\omega}$ are vectors of parameters; and $s(\boldsymbol{\mu}, t)$ and $s(\boldsymbol{\omega}, t)$ are functions of exogenous variables that have an effect on the mean and the volatility, respectively. We imposed restrictions on α_i and β_i to ensure positivity of σ_t^2. For η_t, we considered Gaussian, Student *t*, and skewed *t* distributions, as they have often been considered in the GARCH modelling of daily financial returns data, and Table 1 shows a degree of skewness and high kurtosis in some of the variables.

Given that wave energy flux is a function of wave height and wave period, it is natural that wave energy flux will have annual seasonality (see, for example Jardine and Latham, 1981; Guedes-Soares and Cunha, 2000). However, as our time series are not sufficiently long to capture this, the only cyclicality that we model is the diurnal cycle. We allow for this in the level and volatility, using expressions (10) and (11), respectively:

$$s(\boldsymbol{\mu}, t) = \mu_0 + \sum_{i=1}^{N_\mu} \left[\mu_{i,1} \sin\left(2i\pi \frac{h(t)}{24}\right) + \mu_{i,2} \cos\left(2i\pi \frac{h(t)}{24}\right) \right], \quad (10)$$

$$s(\boldsymbol{\omega}, t) = \omega_0 + \sum_{i=1}^{N_\omega} \left[\omega_{i,1} \sin\left(2i\pi \frac{h(t)}{24}\right) + \omega_{i,2} \cos\left(2i\pi \frac{h(t)}{24}\right) \right], \quad (11)$$

where $h(t)$ is the hour of the day; and N_μ and N_ω are positive integers. To emphasise our use of wave energy flux as the target variable y_t, we denote the model as E-ARMA-GARCH. We also built a wave energy flux model with only AR components, namely E-AR, to test the usefulness of the MA and GARCH terms. We used AR lags from 1 to 4 in this simple model.

3.3.3. ARFIMA-FIGARCH

When a time series shows a slowly decaying persistence in the autocorrelation, this pattern is called 'long memory' dependence, and it can be modelled by a fractional integrated model. Reikard (2009) observes that wave energy flux is characterised by long memory, and this pattern appears more clearly in deep water locations than coastal sites. Fig. 3(b) shows that our wave energy flux data possesses long-memory, because there is significant autocorrelation at long lags. This prompted us to consider a fractionally integrated model.

Long memory in the level of a series can be modelled by the autoregressive fractionally integrated moving average (ARFIMA) model proposed by Granger and Joyeux (1980) and Hosking (1981). Long memory in the volatility can be captured with the fractionally integrated generalized autoregressive conditionally heteroscedastic (FIGARCH) model of Baillie et al. (1996). Before the ARFIMA-FIGARCH is presented, we rewrite the ARMA-GARCH model of expressions (7) and (8) as expressions (12) and (13), respectively:

$$\varphi(L) y_t = s(\boldsymbol{\mu}, t) + \psi(L) \varepsilon_t, \quad (12)$$

$$\alpha(L)\sigma_t^2 = s(\boldsymbol{\omega}, t) + \beta(L)\varepsilon_t^2, \quad (13)$$

where $\varphi(L), \psi(L), \alpha(L)$ and $\beta(L)$ are polynomial functions of the lag operator L. Using the fractional differencing parameter d, the ARFIMA and FIGARCH processes are defined in expressions (14) and (15), respectively:

$$\varphi(L)(1-L)^{d_1} y_t = s(\boldsymbol{\mu}, t) + \psi(L)\varepsilon_t, \quad (14)$$

$$\alpha(L)\sigma_t^2 = s(\boldsymbol{\omega}, t) + [1 - \alpha(L) - \zeta(L)(1-L)^{d_2}]\varepsilon_t^2, \quad (15)$$

$$(1-L)^d = \sum_{i=0}^{\infty} \frac{\Gamma(\mathrm{d}+1)}{\Gamma(\mathrm{i}+1)\Gamma(\mathrm{d}-\mathrm{i}+1)}(-1)^i L^i, \quad (16)$$

where $\zeta(L)$ is a polynomial function of L, $\Gamma(\cdot)$ is the gamma function, and d_1 and d_2 are the fractional differencing parameters, which determine the degree of long memory dependence in the level and volatility processes, respectively. If $d_i = 0$, the process follows short memory dependence. The time series has long memory dependence if $0 < d_i < 0.5$ and moderate long memory dependence if $-0.5 < d_i < 0$. When $d_i = 1$, the ARFIMA process is non-stationary.

3.3.4. VARMA-MGARCH

Soares and Cunha (2000) observe correlation between wave height and wave period, and fit a bivariate vector autoregressive model in order to preserve the covariance structure. In this paper, we jointly model wave height and wave period, and then convert the resulting forecasts into predictions of wave energy flux using expression (1). (In Section 5, we consider the conversion of wave height and wave period forecasts to wave power.)

We were also curious to investigate a joint model involving the variable wave energy flux. In such a model, it makes no sense to include either wave height or wave period, because we have, in this paper, used these two variables to construct the wave energy flux series. Instead, we implemented a joint model for wave energy flux and wind speed.

We, therefore, implemented a bivariate model for two pairs of variables. One pair was wave energy flux and wind speed, and the other was wave height and wave period. The benefit of modelling wave height and wave period is that their forecasts can be conveniently

plugged into the formula in expression (1) for conversion to wave energy flux, or converted to wave power as described in Section 5.

To model the dynamics of the conditional variance and covariance of the pairs of variables, we implemented vector ARMA models with multivariate GARCH terms (VARMA-MGARCH). Although similar models have been used by Cripps and Dunsmuir (2003) and Jeon and Taylor (2012) to model wind velocity, we are not aware of the previous use of ARMA-GARCH models in the wave energy context.

We used the VEC type VARMA-MGARCH model of Bollerslev et al. (1988) with diurnal cyclical terms, as given in expressions (17)-(19):

$$\boldsymbol{y}_t = s(\boldsymbol{\mu}, t) + \sum_{i=1}^{r} \boldsymbol{R}_i \boldsymbol{y}_{t-i} + \sum_{j=1}^{m} \boldsymbol{M}_j \boldsymbol{\varepsilon}_{t-j}, \tag{17}$$

$$vech(\boldsymbol{V}_t) = s(\boldsymbol{\omega}, t) + \sum_{i=1}^{p} \boldsymbol{P}_i vech(\boldsymbol{V}_{t-i}) + \sum_{j=1}^{q} \boldsymbol{Q}_j vech(\boldsymbol{\varepsilon}_{t-j}\boldsymbol{\varepsilon}'_{t-j}), \tag{18}$$

$$\boldsymbol{\varepsilon}_t = \boldsymbol{\eta}_t, \tag{19}$$

where $\boldsymbol{y}_t$ is a vector of (1) wave height and wave period or (2) wave energy flux and wind speed; $\boldsymbol{\varepsilon}_t$ is a vector of error terms; $\boldsymbol{V}_t$ is the conditional covariance matrix of $\boldsymbol{\varepsilon}_t$; $\boldsymbol{\eta}_t$ is a vector of white noise, for which multivariate Gaussian, Student *t* or skewed *t* distributions are considered in our empirical study; $vech(\cdot)$ denotes the column stacking operator of the lower triangular part of its argument symmetric matrix; $\boldsymbol{R}_i$ and $\boldsymbol{M}_i$ are (2×2) matrices of parameters; $\boldsymbol{P}_i$ and $\boldsymbol{Q}_i$ are (3×3) matrices of parameters; *r*, *m*, *p* and *q* are the order of $\boldsymbol{R}_i$, $\boldsymbol{M}_i$, $\boldsymbol{P}_i$ and $\boldsymbol{Q}_i$ respectively, selected by the SBC. Among various forms of multivariate skewed *t* distributions, we used the definition by Azzalini and Genton (2008). In our empirical study, we imposed restrictions on $\boldsymbol{P}_i$ and $\boldsymbol{Q}_i$ using the sufficient condition for the positivity of $\boldsymbol{V}_t$ proposed by Gourieroux (1997). We also implemented the Baba-Engle-Kraft-Kroner VARMA-MGARCH model (see Engle and Kroner, 1995), but this did not lead to improved post-sample forecasting results, and so we do not discuss it further in this paper.

In addition to the standard VEC approach, which we refer to as MGARCH, we also implemented the approach with $\boldsymbol{P}_i$ and $\boldsymbol{Q}_i$ restricted to be diagonal matrices. We refer to this

as MGARCH-DG and implemented this for the model for wave height and wave period (H-P-MGARCH-DG), and for the model for wave energy flux and wind speed (E-S-MGARCH-DG). The diagonal matrices ensure that $\boldsymbol{V}_t$ is positive definite for all *t* (Bollerslev et al., 1988), although this is perhaps overly restrictive, as it allows no interaction between the conditional variances and covariances. For the joint model of wave energy flux and wind speed, to avoid wind speed being modelled in terms of wave energy flux or its lags, we restricted $\boldsymbol{P}_i$ and $\boldsymbol{Q}_i$ to be upper triangular. We refer to this as E-S-MGARCH-UP.

As a relatively simple VAR benchmark model, we constructed a model for wave energy flux and wind speed with lags from 1 to 4, assuming a constant variance. We call this E-S-VAR. Similarly, H-P-VAR is the same model fitted to wave height and wave period.

3.3.5. Orders of the various (V)ARMA-(M)GARCH models

For the (V)ARMA-(M)GARCH models, we used the SBC to select the orders, and also to select terms (values of *i*) to use in the summations of expressions (10) and (11), which capture the diurnality. Table 2 summarises the resulting orders and values for models with Gaussian noise terms fitted to the in-sample FINO1b data. We consider only the Gaussian models here, and in the rest of the paper, because the post-sample results for models fitted with the Student *t* and skewed *t* distributions were no better.

Table 3 presents the d_1 and d_2 parameters estimated for the ARFIMA-GARCH and ARFIMA-FIGARCH models applied to the log transformation of wave energy flux for the two in-sample periods. As explained in Section 3.3.3, a parameter between -0.5 and 0.5 indicates the existence of long memory. The level parameter d_1 indicates the level process does not have long memory, while the volatility parameter, d_2, indicates there are long memory effects when the models are fitted to FINO1a, but not when fitted to FINO1b.

4. EMPIRICAL POST-SAMPLE RESULTS FOR WAVE ENERGY FLUX

As we explained in Section 2, for the final two months of the FINO1a and FINO1b periods, we produced 1 to 24 hour-ahead post-sample density forecasts for wave energy flux, using each period of the evaluation sample as forecast origin. For the VARMA-MGARCH models, we felt that it is not practical to re-optimise repeatedly the parameters for a sliding window of observations, and so, for each of the two periods (FINO1a and FINO1b), we estimated the parameters just once. For consistency, we followed the same approach with the other methods, although we acknowledge that the ranking of methods may change parameters were re-optimised. In Sections 4.1 and 4.2, we use the mean of the CRPS to evaluate density forecasting accuracy, which is the main focus in this paper. In Section 4.3, we consider point forecasting. As statistical methods have been shown to be more competitive with atmospheric models for short forecast horizons, our analysis provides more detail for the earlier horizons.

4.1. Evaluation of the transformations for use with ARMA-GARCH

Table 4 presents post-sample CRPS density forecasting results, averaged over the FINO1a and FINO1b periods, for the univariate ARMA-GARCH models in Section 3.3.2 fitted to wave energy flux using the transformations described in Section 3.3.1. The table indicates that using any transformation was preferable to using none. The square root was not as useful as the log and Box-Cox transformations, which is consistent with the results for wave energy flux in Table 1. The log and Box-Cox transformations delivered similar results, and as the log transformation is simpler, in the rest of this paper, we report results for all the (V)ARMA-(M)GARCH models applied to variables that were logged prior to model fitting.

4.2. Density forecasting results for wave energy flux

Table 5 compares the accuracy of density forecasts from the ARMA-GARCH, ARFIMA-GARCH and ARFIMA-FIGARCH models applied to log transformed wave energy

flux. The table shows that the models with fractional integration were slightly outperformed by the ARMA-GARCH model. It is likely that the forecast lead times that we consider are not sufficiently long for models with fractional integration to be of benefit. In view of this, in the rest of the paper, we do not report results for the fractionally integrated models.

Table 6 and Fig. 4 compare the density forecast accuracy from the regression-based approach, the KDE methods, and the ARMA-GARCH models. Table 6 shows that the regression method produced less accurate density forecasts than each of the ARMA-GARCH models beyond four hours ahead. Both the UKDE and CKDE methods did not perform well, particularly for the shorter lead times. The CKDE approach allows exponential weighting, but there is little weight decay as the optimal values of the exponential decay factor τ were 0.998 and 1.000 for FINO1a and FINO1b, respectively. We experimented with weight decay in the UKDE, but the optimised decay parameter was close to zero, implying very large weight on the most recent period, which had little appeal, and so we did not consider the method further.

The (V)ARMA-(M)GARCH models used three different combinations of data, namely wave energy flux (E) alone, wave energy flux and wind speed (E-S), and wave height and wave period (H-P). Table 6 shows no great difference in the results of the methods, with the H-P models performing slightly better than the others. For both the E-S and H-P combinations of data, the MGARCH-DG model, which is a diagonal form of multivariate GARCH, delivered slight improvement over standard MGARCH. This is a useful result because this simplified model has fewer parameters, and so is easier to estimate.

To further evaluate the density forecasts, histograms of the probability integral transform (PIT) (see Gneiting et al., 2007) are provided in Fig. 5 for FINO1b. The figures show results, for lead times of 1, 4, 12 and 24 hours ahead, for the following four methods: the regression-based method, E-UKDE (4hour), E-CKDE, and H-P-VARMA-MGARCH-DG. The ideal shape of a PIT histogram is a uniform distribution. For the regression-based method, E-UKDE (4 hour) and E-CKDE, the PIT histograms are far from uniform. As the lead time

increases, the peaks in each tail become larger, indicating that the density forecasts were overly wide. The PIT histograms are closer to uniform for H-P-VARMA-MGARCH-DG.

4.3. Point forecasting results for wave energy flux

Although density forecasting is our primary concern, the evaluation of point forecasting is also of interest. Table 7 and Fig. 6 present the root mean squared error (RMSE) results, averaged over the FINO1a and FINO1b periods, for point forecasts produced by the different methods. The table shows that H-P-VARMA-MGARCH-DG produced the best results overall. Indeed, this method was not outperformed by any other method at any lead time. The regression method also performed very well. These findings show that modelling wave height and wave period, albeit separately, led to better results than directly modelling wave energy flux.

For the longer lead times, the regression method and CKDE were much more competitive in terms of point forecasting than they were in Table 6 for density forecasting. The UKDE methods did not perform well in terms of point forecasting.

With regard to the relative performances of the (V)ARMA-(M)GARCH models in Table 7, we can make a number of points. Firstly, the bivariate (E-S) models for wave energy flux and wind speed seem to offer very little over the univariate (E) models for wave energy flux. Secondly, all of the bivariate (H-P) models for wave height and wave period are, overall, more accurate than the univariate (E) models for wave energy flux and the bivariate (E-S) models for wave energy flux and wind speed. Thirdly, with regard to the (H-P) models for wave height and wave period, up to about 8 hours-ahead, there does seem to be benefit in the increased complexity of the VARMA-MGARCH-DG model over the much simpler H-P-VAR model, and the diagonal (DG) version of the VARMA-MGARCH model does seem preferable to the more highly parameterised VARMA-MGARCH model.

5. EMPIRICAL POST-SAMPLE RESULTS FOR WAVE POWER

We generated wave power density and point forecasts by converting wave height and wave period to wave power using a conversion matrix for the Pelamis P2 device (see Henderson, 2006; Retzler, 2006; Yemm et al., 2012), which is an established wave power technology. The matrix is presented and used by Reikard (Fig. 1, 2013). The Pelamis P2 wave energy converter consists of semi-submerged multiple cylindrical sections. As waves pass along the length of the device, the differences in buoyancy make the joints of the cylinders bend, and this induces hydraulic cylinders to pump high pressure oil through hydraulic motors, which drives electrical generators to produce electricity.

Since our wave period data ranges up to 20 seconds, which is higher than the upper limit of the conversion matrix, we extrapolated the conversion matrix using the inpaint_nan function by D'Errico (2012), which is based on sparse linear algebra and PDE discretizations, to give the conversion function of Fig. 7. It is notable from this figure that wave power from the Pelamis P2 device has an upper bound, and that, regardless of the value of wave height, wave power is at its highest when wave period is approximately 7.5 seconds, which is consistent with the finding of Retzler (Fig. 4, 2006) that the power capture of the device is highest when the frequency is around 0.13 Hz. Due to the shape of the nonlinear conversion function, the resulting wave power times series for the FINO1b data series in Fig. 8 exhibits less extreme spikes than the wave energy flux series of Fig 1.

Fig. 9 shows that none of the log, square root and Box-Cox transformations are able to change the strong skewness in wave power. Consequently, there was no appeal in performing direct modelling of wave power using a univariate ARMA-GARCH model, or using this variable along with another in a bivariate VARMA-MGARCH model.

In terms of modeling wave power directly, we applied the kernel density estimation methods of Section 3.2 to wave power. In addition, we generated wave power density forecasts by using the function of Fig. 7 to convert wave height and wave period density

forecasts produced by (a) the regression method of Section 3.1, and (b) the bivariate VARMA-MGARCH models of wave height and wave period, discussed in Section 3.3.4.

For a selection of the methods, Figs. 10 and 11 present the CRPS and RMSE for the wave power density and point forecasts, respectively. Both figures show the KDE approaches performing relatively poorly. By contrast the results of the CKDE approach are competitive, and comparable with the regression approach. Overall, the best CRPS results correspond to the H-P-VARMA-MGARCH-DG method, although the regression method is as accurate for lead times less than about 8 hours. Fig. 11 shows that the point forecasting results were similar for these two methods and the CKDE method.

6. SUMMARY

In this paper, we evaluated density forecasts of wave energy flux and wave power produced by a regression method, UKDE methods, a CKDE approach, and univariate and multivariate ARMA-GARCH models. Our results showed the following:

(i) Although the regression method performed well in terms of point forecasting for the longer lead times, overall, the best point and density forecast accuracy were produced by the ARMA-GARCH models. We found that the GARCH component was useful only for lead times up to about 8 hours ahead. Our results do not support the use of a Student *t* or skewed *t* distribution instead of a Gaussian distribution.

(ii) Bivariate ARMA-GARCH modelling of the log transformed wave height and wave period produced the best result for both wave energy flux in Section 4 and wave power in Section 5. For energy flux, it was interesting that this was preferable to forecasting the energy flux directly.

(iii) Despite evidence of long-memory in wave data, we could not find any clear evidence to support the use of fractionally integrated models.

(iv) Kernel density estimation was not particularly competitive.

ACKNOWLEDGEMENT

We thank the BMU (Bundesministerium fuer Umwelt, Federal Ministry for the Environment, Nature Conservation and Nuclear Safety) and the PTJ (Projekttraeger Juelich, Project Executing Organization) for making the FINO1 data available. We are very grateful to Gordon Reikard for his advice regarding both the modelling and literature on modelling wave energy. We are also grateful for the useful comments of the reviewers.

REFERENCES

Azzalini, A. & Genton, M. G. (2008). Robust Likelihood Methods Based on the Skew-t and Related Distributions. International Statistical Review, 76, 106-129.

Bollerslev, T., Engle, R.F. & Wooldridge, J.M. (1988). A capital asset pricing model with time varying covariances. Journal of Political Economy 96, 116-131.

Box, G.E.P. & Cox, D.R. (1964). An analysis of transformations. Journal of the Royal Statistical Society, 26, 211-252.

Brekken, T. K. A., von Jouanne, A. & Han, H.Y. (2009) Ocean wave energy overview and research at oregon state university. In Proceedings of Power Electronics and Machines in Wind Applications, June 24-29, 2009, 1-7.

Clément, A., McCullen, P., Falcao, A. et al. (2002). Wave energy in Europe: current status and perspectives. Renewable and Sustainable Energy Review, 6, 405-431.

Cripps, E. & Dunsmuir, W.T.M. (2003). Modelling the variability of Sydney harbour wind measurements. Journal of Applied Meteorology, 42, 1131-1138.

D'Errico, J. (2012). inpaint_nans: Interpolates (& extrapolates) NaN elements in a 2d array [Computer program, http://de.mathworks.com/matlabcentral/fileexchange/4551-inpaint-nans. Accessed on 9 June 2015.

Durrant, T.H., Woodcock, F. & Greenslade, D.J.M. (2008). Consensus forecasts of modeled wave parameters. Weather and Forecasting, 24, 492-503.

Engle, R.F. & Kroner, K.F. (1995). Multivariate simultaneous generalized ARCH. Econometric Theory, 11, 122-150.

Falnes, J. (2007). A review of wave-energy extraction. Marine Structures, 20, 185-201.

Ferreira, J.A. & Guedes-Soares, C. (2002). Modelling bivariate distributions of significant wave height and mean period. Applied Ocean Research, 24, 31-45.

Fusco, F. & Ringwood, J. (2010). Short-term wave forecasting for real-time control of wave energy converters. IEEE Transactions on Sustainable Energy, 1, 99-106.

Galiatsatou, P. & Prinos, P. (2007). Bivariate models for extremes of significant wave height and period. An application to the Dutch Coast. In Proceedings of 2nd IMA International Conference on Flood Risk Assessment, University of Plymouth, UK, September 4-5.

Gneiting, T., Balabdaoui, F. & Raftery, A.E. (2007). Probabilistic forecasts, calibration and sharpness. Journal of the Royal Statistical Society. Series B: Statistical Methodology, 69, 243-268.

Gourieroux, C. (1997). ARCH models and financial applications, New York: Springer.

Granger, C.W.J. & Joyeux, R. (1980). An introduction to long-memory time series models and fractional differencing. Journal of Time Series Analysis, 1, 15-29.

Hasselmann, K., Sell, W., Ross, D.B. & Müller P. (1976). A parametric wave prediction model. Journal of Physical Oceanography, 6, 200-228.

Hasselmann, D.E., Dunckel, M. & Ewing, J.A. (1980). Directional wave spectra observed during JONSWAP 1973. Journal of Physical Oceanography, 10, 1264-1280.

Hasselmann, S., Hasselmann, K., Allender J.H. & Barnett, T.P. (1985). Computations and parameterizations of the non-linear energy transfer in a gravity wave spectrum. Part II: Parameterizations of the non-linear energy transfer for application in wave models. Journal of Physical Oceanography, 15, 1378-1391.

Henderson, R. (2006). Design, simulation, and testing of a novel hydraulic power take-off system for the Pelamis wave energy converter. Renewable Energy, 31, 271-283.

Hosking, J.R.M. (1981). Fractional differencing. Biometrika, 68, 165-176.

Janssen, P.A.E.M. (1991). Quasi-linear theory of wind-wave generation applied to wave forecasting. Journal of Physical Oceanography, 21, 1631-1642.

Janssen, P.A.E.M. (2007). Progress in ocean wave forecasting. Journal of Computational Physics, 227, 3572-3594.

Jardine, T.P. & Latham, F.R. (1981). An analysis of wave height records for the N.E. Atlantic. Quarterly Journal of the Royal Meteorological Society, 107, 415-426.

Jeon, J. & Taylor, J.W. (2012). Using conditional kernel density estimation for wind power forecasting, Journal of the American Statistical Association, 107, 66-79.

Miller, A.R. (1958). The effects of winds on water level on the New England coast. Limnology and Oceanography, 3, 1-14.

Pinson, P., Reikard, G. & Bidlot, J.R. (2012). Probabilistic forecasting of the wave energy flux. Applied Energy, 93, 364-370.

Rieder, K. F. (1997). Analysis of sea-surface drag parameterizations in open ocean conditions, Boundary Layer Meteorology, 82, 355-377.

Reikard, G. (2009). Forecasting ocean wave energy: Tests of time-series models. Ocean Engineering, 73, 168-178.

Reikard, G. (2013). Integrating wave energy into the power grid: Simulation and forecasting. Ocean Engineering, 36, 348-356.

Reikard, G., Pinson, P. & Bidlot, J.R. (2011). Forecasting ocean wave energy: The ECMWF wave model and time series methods. Ocean Engineering, 38, 1089-1099.

Reikard, G., Robertson, B. & Bidlot, J.R. (2015). Combining wave energy with wind and solar: Short-term forecasting. Renewable Energy, 81, 442-456.

Reikard, G. & Rogers, W.E. (2011). Forecasting ocean waves: Comparing a physics-based model with statistical models. Coastal Engineering, 58, 409-416.

Retzler, C. (2006). Measurements of the slow drift dynamics of a model Pelamis wave energy converter. Renewable Energy, 31, 257-269.

Roulston, M.S., Ellepola, J., von Hardenberg, J. & Smith, L.A. (2005). Forecasting wave height probabilities with numerical weather prediction models. Ocean Engineering, 32, 1841-1863.

Soares, C.G. & Cunha, C. (2000). Bivariate autoregressive models for the time series of significant wave height and mean period. Coastal Engineering, 11, 139-148.

Soares, C.G. & Ferreira, A.M. (1996). Representation of non-stationary time series of significant wave height with autoregressive models. Probabilistic Engineering Mechanics, 11, 139-148.

Taylor, J.W., McSharry, P.E. & Buizza, R. (2009). Wind power density forecasting using ensemble predictions and time series models. IEEE Transactions on Energy Conversion, 24, 775-782.

Woodcock, F. & Engel, C. (2005). Operational consensus forecasts. Weather and Forecasting, 20, 101-111.

Woodcock, F. & Greenslade, D.J.M. (2006). Consensus of numerical forecasts of significant wave heights. Weather and Forecasting, 22, 792-803.

Yemm, R., Pizer, D., Retzler, C. & Henderson, R. (2012). Pelamis: experience from concept to connection. Philosophical Transactions of the Royal Society, 370, 365-380.

Zamani, A., Solomatine, D., Azimian, A. & Heemink, A. (2008). Learning from data for wind–wave forecasting. Ocean Engineering, 35, 953-962.

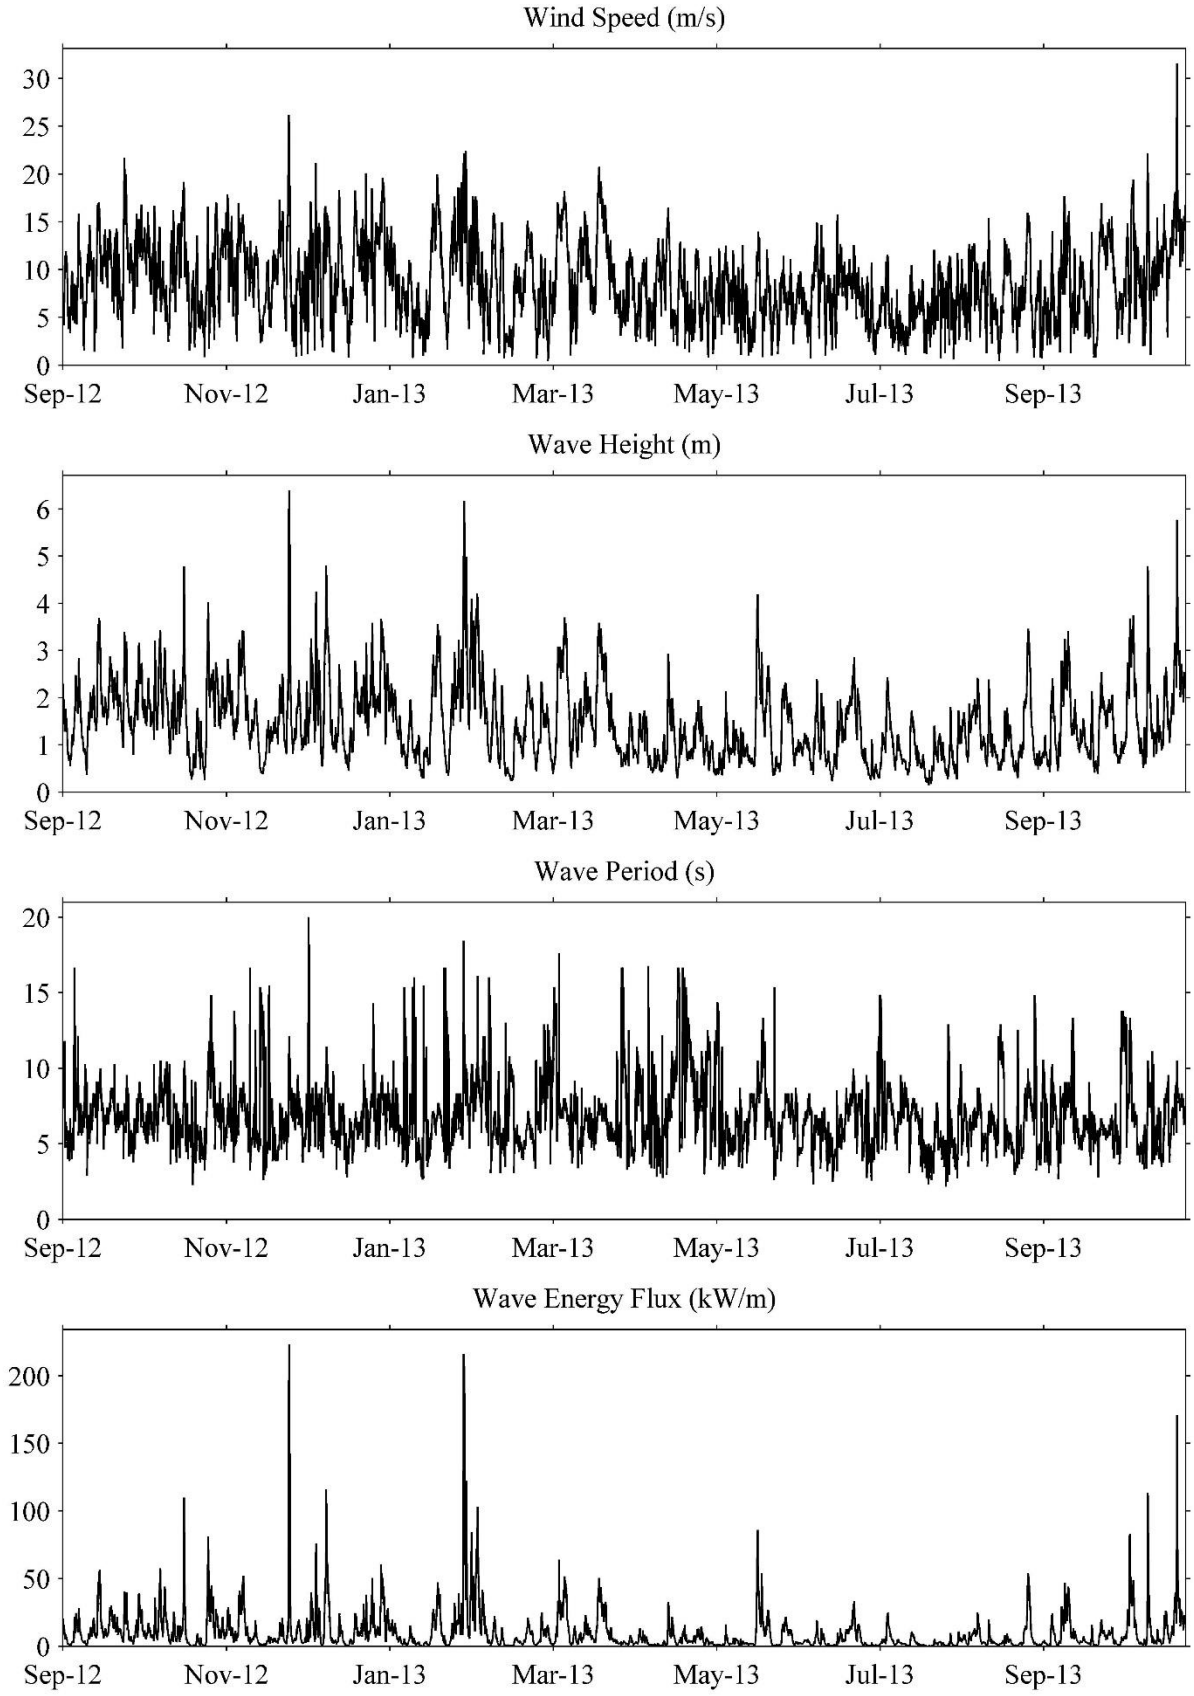

Fig. 1. Time series of wind speed, wave height, wave period and wave energy flux from the FINO1b dataset.

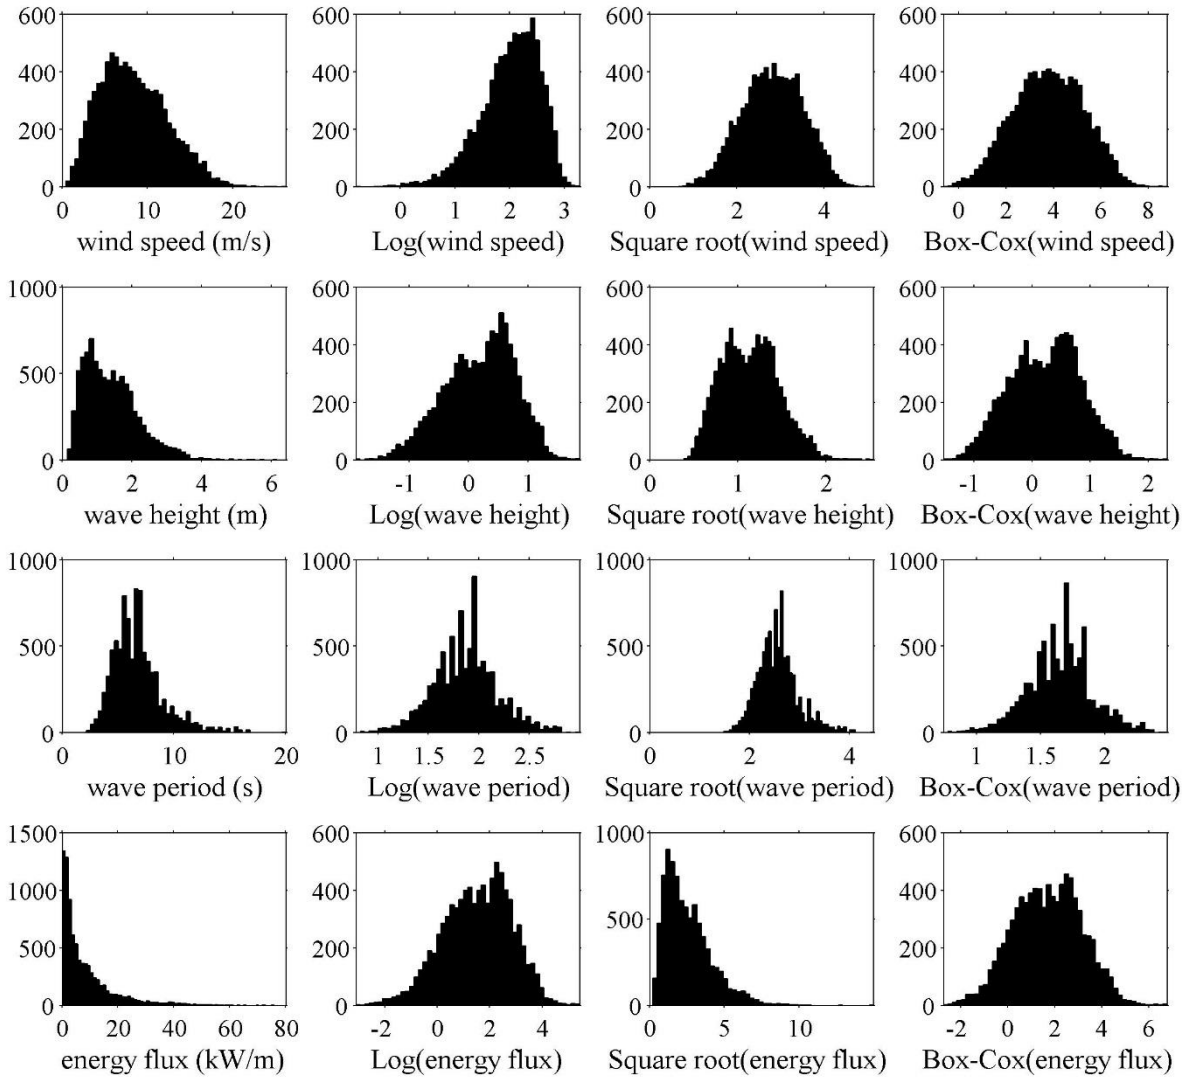

Fig. 2. Histograms of wind speed, wave height, wave period and wave energy flux, and their log, square root and Box-Cox transformed distributions from the FINO1b dataset.

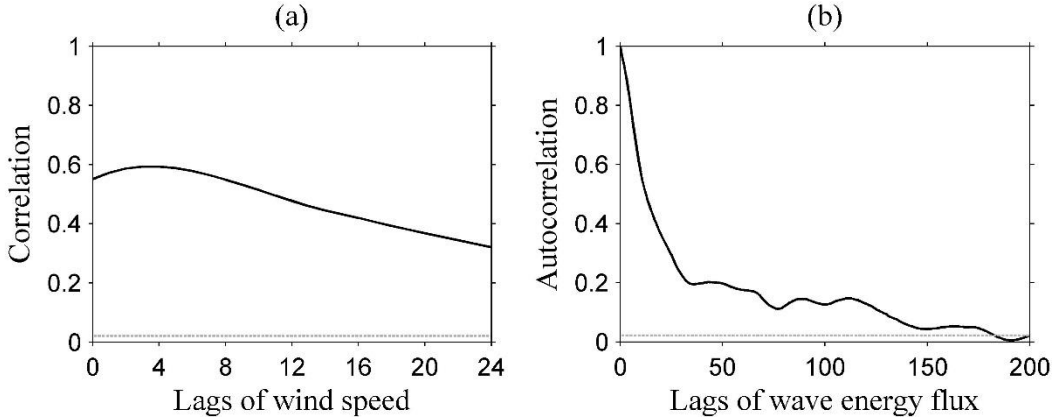

Fig. 3. (a) Correlations between wave energy flux and lags of wind speed up to 24 hours from FINO1b. (b) Autocorrelations in wave energy flux of the FINO1b dataset. The 95% significance level is indicated as a dotted line.

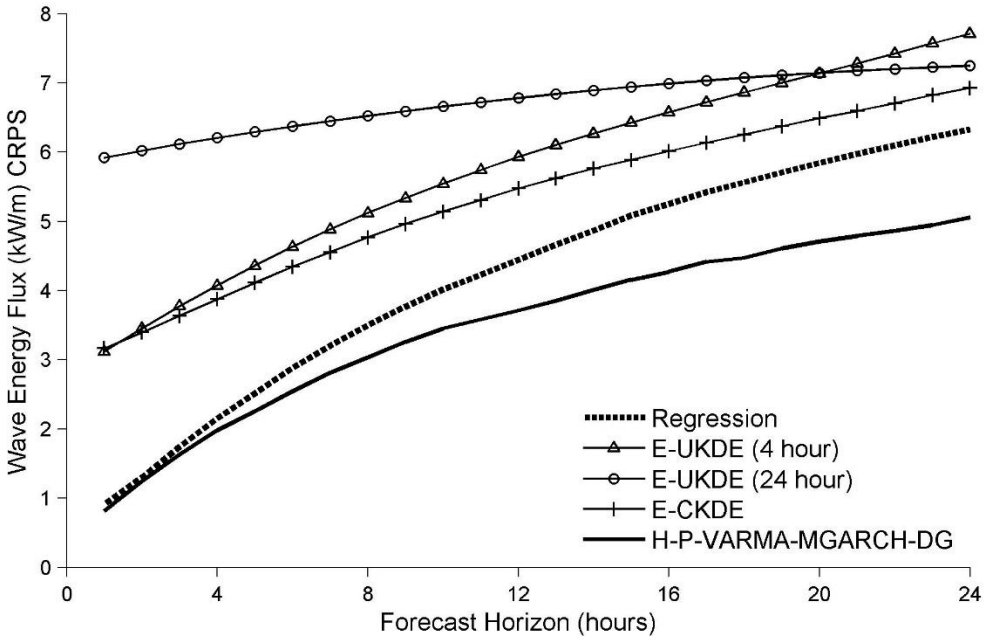

Fig. 4. CRPS evaluated for wave energy flux forecasts and averaged over FINO1a and FINO1b. Lower values are better.

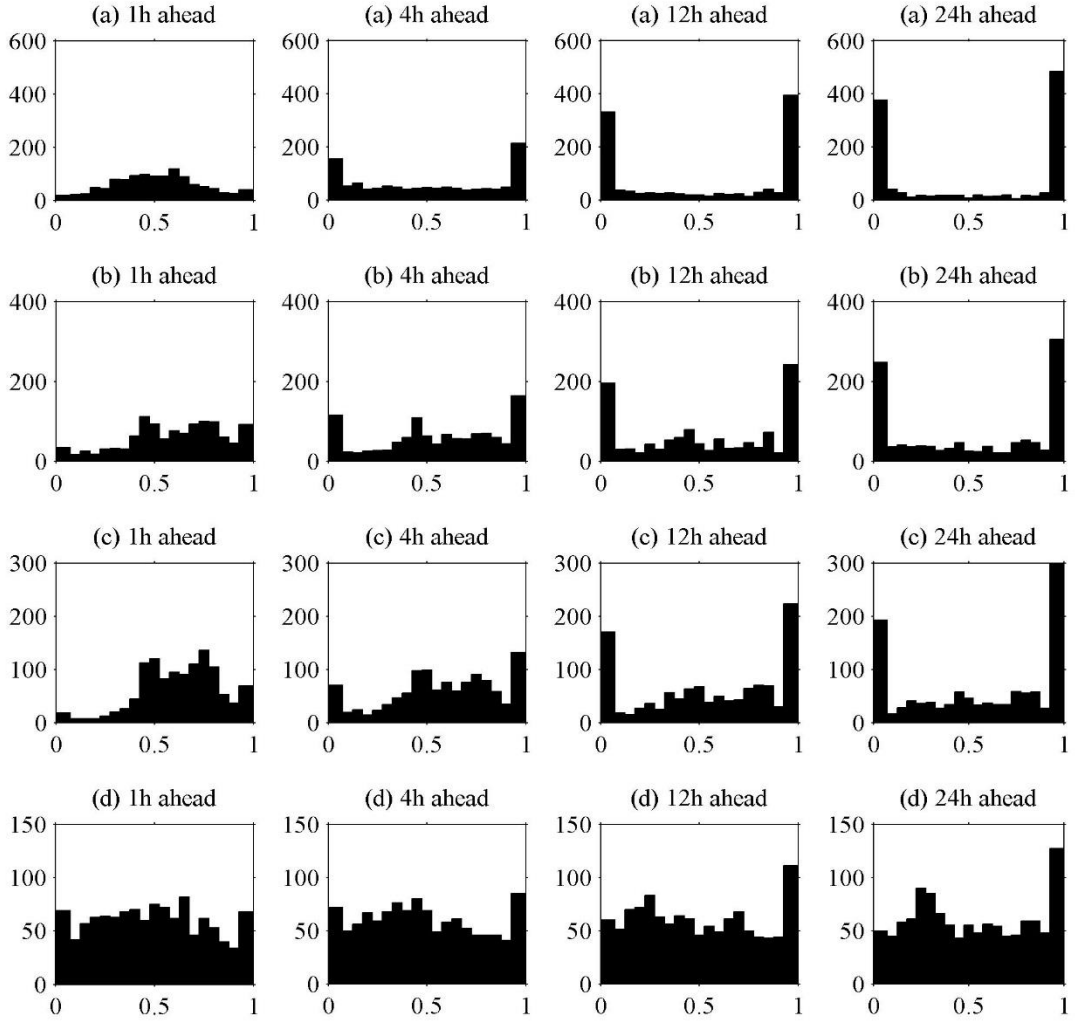

Fig 5. PIT histograms of 1, 4, 12 and 24 hours ahead wave energy flux forecasts for FINO1b using (a) Regression, (b) E-UKDE (4 hour), (c) E-CKDE and (d) H-P-VARMA-MGARCH-DG (Gaussian) methods.

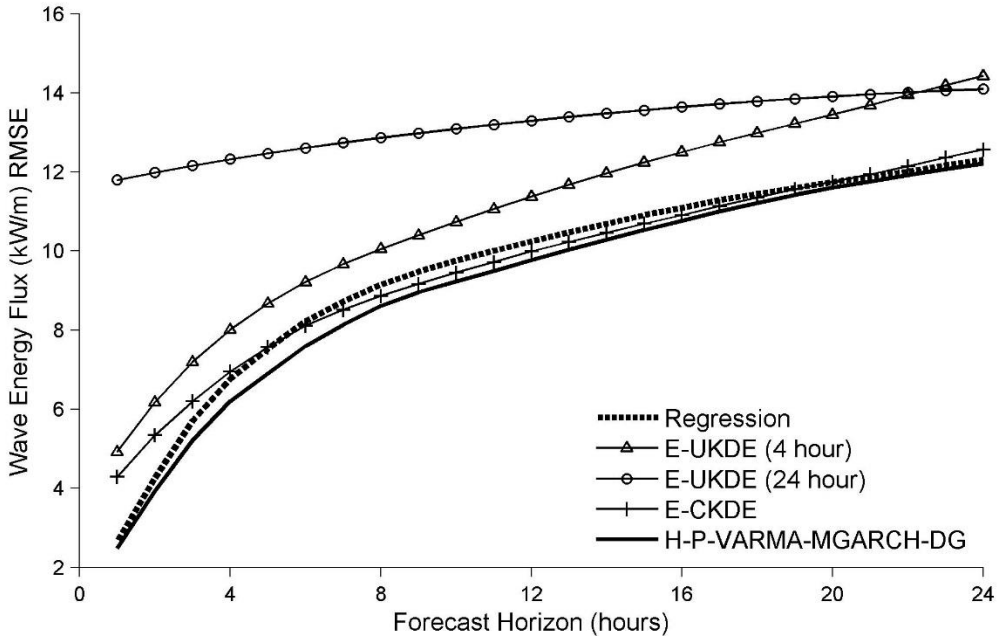

Fig. 6. For wave energy point forecasts, RMSE averaged over FINO1a and FINO1b.

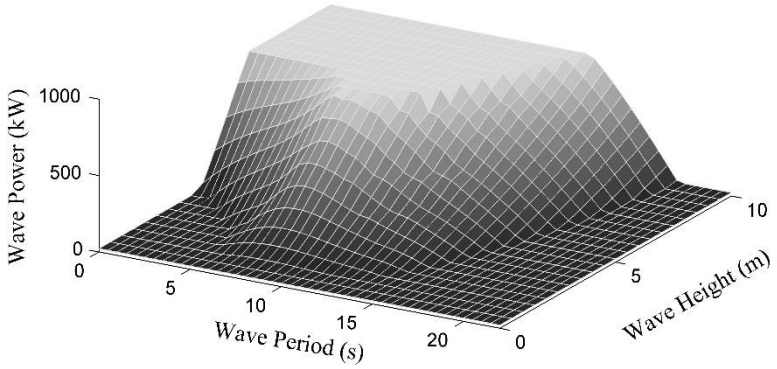

Fig. 7. Extended wave power conversation from wave period and wave height based on the conversion matrix for the Pelamis P2 device shown by Reikard (2013) and sparse linear algebra and PDE discretizations by D'Errico, J. (2012).

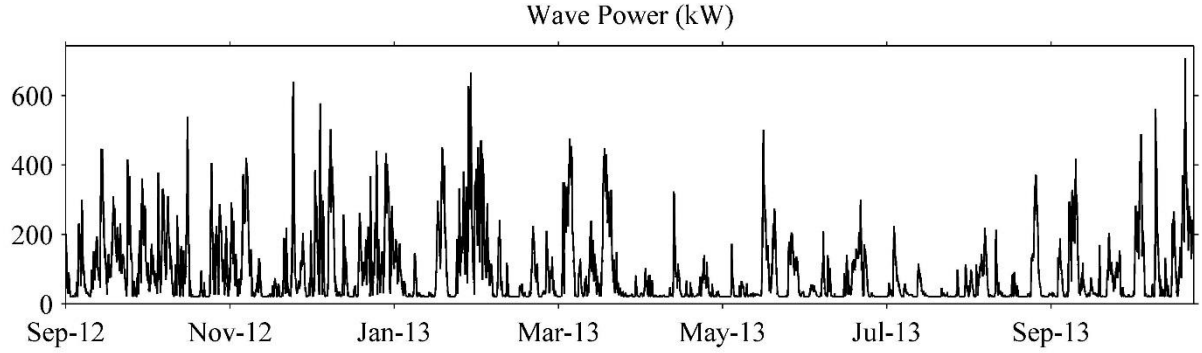

Fig. 8. Time series of wave power from the FINO1b dataset.

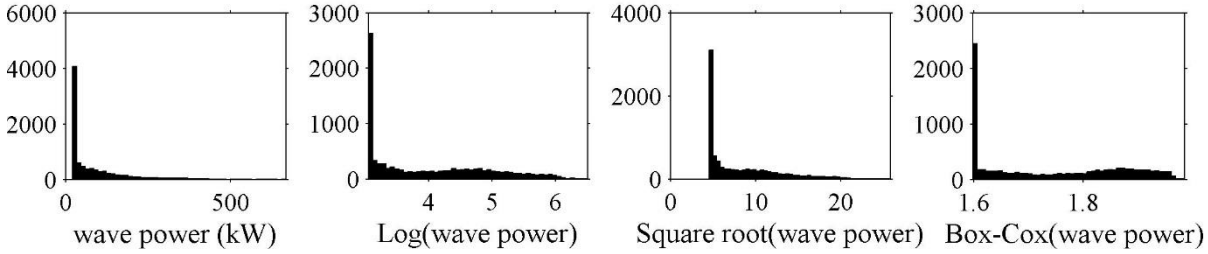

Fig. 9. Histograms of wave power, and its log, square root and Box-Cox transformed distributions from the FINO1b dataset.

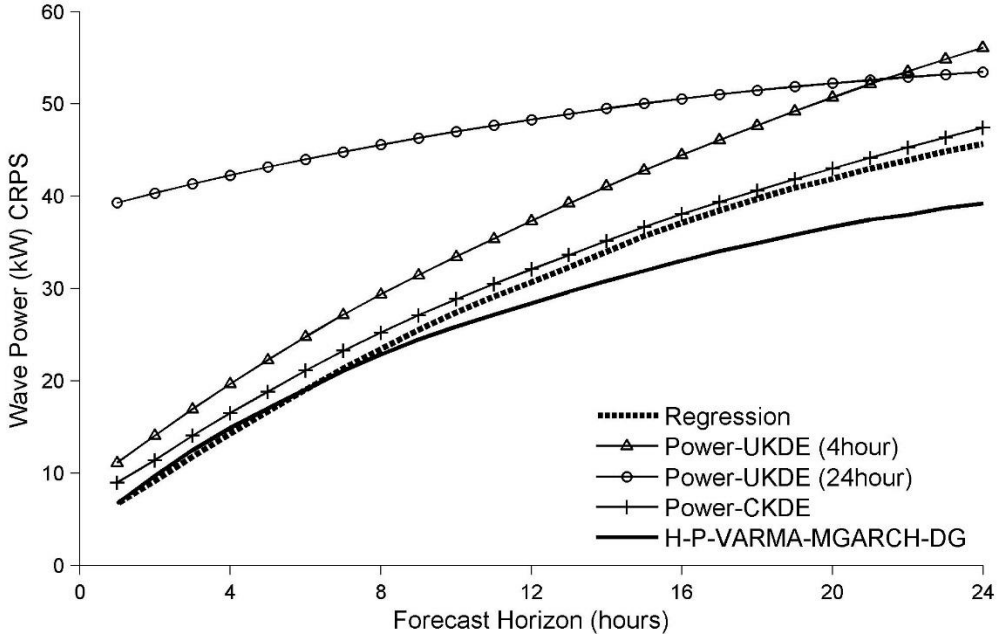

Fig. 10. CRPS evaluated for wave power forecasts and averaged over FINO1a and FINO1b. Lower values are better.

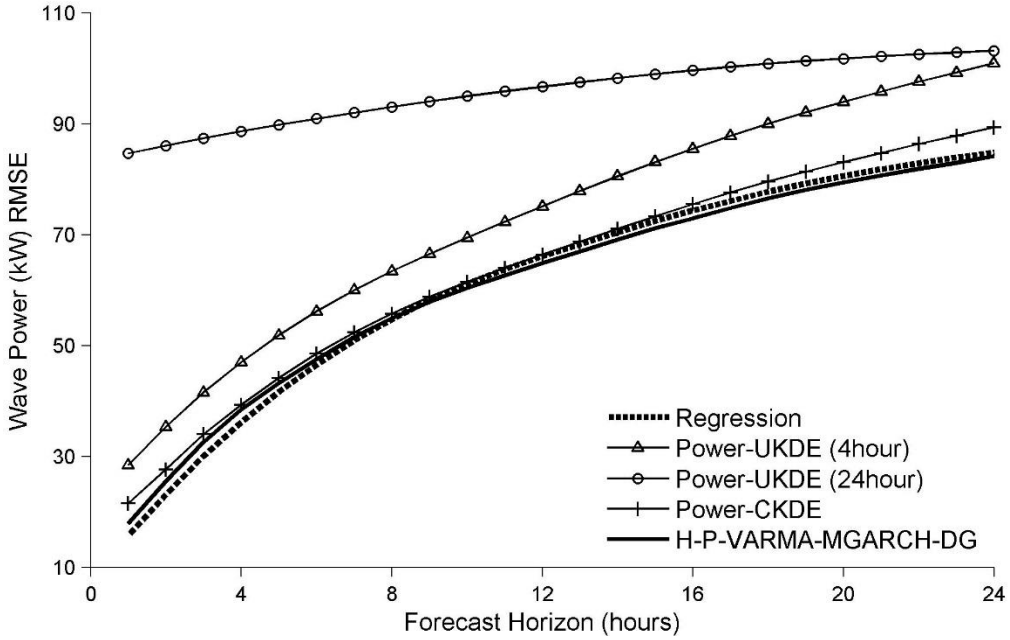

Fig. 11. RMSE evaluated for wave power forecasts and averaged over FINO1a and FINO1b.

Table 1. Skewness and kurtosis for wind speed, wave height, wave period and wave energy flux, and log, square root and Box-Cox transformations of each variable. Statistics calculated for the in-sample period of FINO1b. In each row, value closest to Gaussian (skewness=0, kurtosis=3) is in bold.

	Original	Log	Square Root	Box-Cox
Wind speed				
skewness	0.48	-0.85	-0.11	**-0.05** (*λ*= 0.54*)*
kurtosis	**2.94**	3.89	2.58	2.56 (*λ*= 0.54*)*
Wave height				
skewness	1.07	-0.32	0.35	**-0.02** (*λ*= 0.22*)*
kurtosis	4.76	2.65	**2.79**	2.52 (*λ*= 0.22*)*
Wave period				
skewness	1.27	0.15	0.70	**0.00** (*λ*=-0.14*)*
kurtosis	5.51	3.32	3.95	**3.30** (*λ*=-0.14*)*
Wave energy flux				
skewness	5.11	-0.25	1.52	**-0.01** (*λ*= 0.07*)*
kurtosis	49.46	**2.79**	7.16	2.68 (*λ*= 0.07*)*

Table 2. Lags, and terms in the diurnal expressions (10) and (11), selected by the SBC criterion for (V)ARMA-(M)GARCH models fitted to the in-sample period of FINO1b.

Lags	AR	MA	Diurnal in mean	ARCH	GARCH	Diurnal in volatility
Univariate models for log wave energy flux						
E-AR	[1,2,3,4]	no	no	no	no	no
E-ARMA-GARCH	[1,2,3,4]	[1,2,24]	no	[1,2,24]	[1,24]	[2]
E-ARFIMA-GARCH	[1,2,3,24]	no	no	[1,2,3]	[1,24]	[2,4,8]
E-ARFIMA-FIGARCH	[1,2]	[1,2,24]	no	no	[1]	[2,4,8]
Bivariate models for log wave energy flux and log wind speed						
E-S-VAR	[1,2,3,4]	no	no	no	no	no
E-S-VARMA-MGARCH	[1,2,3,4]	no	[2,4,8]	[1,2,3,24]	[1]	[2,4,6]
E-S-VARMA-MGARCH-UP	[1,2,3,24]	[1,2]	no	[1]	[1]	[2]
E-S-VARMA-MGARCH-DG	[1,2,3,4]	[1,2]	[2,4]	[1,24]	[1]	no
Bivariate models for log wave height and log wave period						
H-P-VAR	[1,2,3,4]	no	no	no	no	no
H-P-VARMA-MGARCH	[1]	[1,2,3,4,5,6]	[2,4,6]	[1]	[1]	no
H-P-VARMA-MGARCH-DG	[1,2,3,4]	[1,2,24]	[2,4,6]	[1,2]	[1,2,3]	[2]

Table 3. Coefficients of the fractional integration models fitted to the log-transformed wave energy flux for the in-sample period of FINO1a and FINO1b.

	ARFIMA d_1	FIGARCH d_2
FINO1a		
E-ARFIMA-GARCH	0.62	
E-ARFIMA-FIGARCH	0.62	0.17
FINO1b		
E-ARFIMA-GARCH	0.92	
E-ARFIMA-FIGARCH	0.92	0.58

Table 4. Evaluation of transformation methods in terms of post-sample wave energy flux density forecast accuracy using CRPS (in kW/m) averaged for FINO1a and FINO1b.

Lead Time (hours):	1	2	3-4	5-6	7-8	9-12	13-18	19-24	1-24
ARMA-GARCH type with Gaussian for wave energy flux									
No transformation	**0.9**	1.4	2.1	2.8	3.4	4.1	5.1	5.8	4.2
Log transformation	**0.9**	**1.3**	**1.9**	**2.5**	**3.1**	**3.6**	**4.4**	**5.1**	**3.7**
Square root transformation	**0.9**	**1.3**	**1.9**	2.6	3.2	3.8	4.8	5.7	4.0
Box-Cox transformation	**0.9**	**1.3**	**1.9**	**2.5**	**3.1**	**3.6**	**4.4**	**5.1**	**3.7**

NOTE: Smaller values are better. The best performing method at each horizon is in bold.

Table 5. Evaluation of the fractional integration models in terms of post-sample wave energy flux density forecast accuracy using CRPS (in kW/m) averaged for FINO1a and FINO1b.

Lead Time (hours):	1	2	3-4	5-6	7-8	9-12	13-18	19-24	1-24
Univariate models for log wave energy flux									
E-ARMA-GARCH	**0.9**	**1.3**	**1.9**	**2.5**	**3.1**	**3.6**	**4.4**	**5.1**	**3.7**
E-ARFIMA-GARCH	**0.9**	**1.3**	**1.9**	2.6	**3.1**	3.7	4.5	5.3	3.8
E-ARFIMA-FIGARCH	**0.9**	**1.3**	**1.9**	2.6	3.2	3.8	4.6	5.4	3.9

NOTE: Smaller values are better. The best performing method at each horizon is in bold.

Table 6. Evaluation of post-sample wave energy flux density forecast accuracy using CRPS (in kW/m) averaged for FINO1a and FINO1b.

Lead Time (hours):	1	2	3-4	5-6	7-8	9-12	13-18	19-24	1-24
Regression	0.9	1.3	1.9	2.7	3.3	4.1	5.1	6.0	4.2
Kernel density estimation for log wave energy flux									
E-UKDE (4 hour)	3.1	3.5	3.9	4.5	5.0	5.6	6.5	7.4	5.8
E-UKDE (24 hour)	5.9	6.0	6.2	6.3	6.5	6.7	7.0	7.2	6.7
E-CKDE	3.2	3.4	3.8	4.2	4.7	5.2	5.9	6.7	5.3
Univariate models for log wave energy flux									
E-AR	0.9	1.3	1.9	2.6	3.1	3.6	4.3	5.0	3.7
E-ARMA-GARCH	0.9	1.3	1.9	2.5	3.1	3.6	4.4	5.1	3.7
Bivariate models for log wave energy flux and log wind speed									
E-S-VAR	0.9	1.3	1.9	2.5	3.0	3.6	4.3	4.9	3.6
E-S-VARMA-MGARCH	**0.8**	1.3	**1.8**	2.5	3.0	3.6	4.4	5.1	3.7
E-S-VARMA-MGARCH-DG	**0.8**	**1.2**	**1.8**	**2.4**	3.0	3.6	4.3	5.0	3.6
E-S-VARMA-MGARCH-UP	0.9	1.3	1.9	2.5	3.0	3.6	4.3	5.0	3.7
Bivariate models for log wave height and log wave period									
H-P-VAR	0.9	1.3	1.9	**2.4**	**2.9**	**3.5**	**4.2**	4.9	3.6
H-P-VARMA-MGARCH	**0.8**	1.3	**1.8**	**2.4**	**2.9**	**3.5**	4.3	4.9	3.6
H-P-VARMA-MGARCH-DG	**0.8**	**1.2**	**1.8**	**2.4**	**2.9**	**3.5**	**4.2**	**4.8**	**3.5**

NOTE: Smaller values are better. The best performing method at each horizon is in bold.

Table 7. Evaluation of post-sample wave energy flux point forecast accuracy using RMSE (in kW/m) averaged for FINO1a and FINO1b.

Lead Time (hours):	1	2	3-4	5-6	7-8	9-12	13-18	19-24	1-24
Regression	2.6	4.3	6.2	7.9	8.9	9.9	11.0	12.0	9.6
Kernel density estimation for log wave energy flux									
E-UKDE (4 hour)	4.9	6.2	7.6	8.9	9.9	10.9	12.4	13.8	11.0
E-UKDE (24 hour)	11.8	12.0	12.2	12.5	12.8	13.1	13.6	14.0	13.2
E-CKDE	4.3	5.3	6.6	7.8	8.7	9.6	10.8	12.1	9.6
Univariate models for log wave energy flux									
E-AR	2.8	4.3	6.2	7.8	8.9	9.9	11.3	12.6	9.8
E-ARMA-GARCH	2.7	4.3	6.1	7.8	8.9	9.9	11.2	12.5	9.8
Bivariate models for log wave energy flux and log wind speed									
E-S-VAR	2.7	4.2	6.1	7.7	8.9	9.9	11.1	12.4	9.7
E-S-VARMA-MGARCH	2.7	4.2	6.0	7.7	9.0	10.2	11.9	14.0	10.3
E-S-VARMA-MGARCH-DG	2.6	4.0	5.9	7.6	8.8	9.9	11.4	13.1	9.9
E-S-VARMA-MGARCH-UP	2.7	4.2	6.1	7.8	8.9	9.9	11.3	12.7	9.8
Bivariate models for log wave height and log wave period									
H-P-VAR	2.7	4.1	5.8	7.4	8.5	**9.4**	**10.6**	**11.8**	9.3
H-P-VARMA-MGARCH	2.6	4.1	5.9	7.5	8.6	9.7	11.1	12.3	9.6
H-P-VARMA-MGARCH-DG	**2.5**	**3.9**	**5.7**	**7.2**	**8.4**	**9.4**	**10.6**	**11.8**	**9.2**

NOTE: Smaller values are better. The best performing method at each horizon is in bold.