

A seasonal copula mixture for hedging the clean spark spread with wind power futures

Christensen, Troels Sønderby; Pircalabu, Anca; Høg, Esben

Published in:
Energy Economics

DOI (link to publication from Publisher):
[10.1016/j.eneco.2018.11.002](https://doi.org/10.1016/j.eneco.2018.11.002)

Creative Commons License
CC BY-NC-ND 4.0

Publication date:
2019

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Christensen, T. S., Pircalabu, A., & Høg, E. (2019). A seasonal copula mixture for hedging the clean spark spread with wind power futures. *Energy Economics*, 78(February 2019), 64-80.
<https://doi.org/10.1016/j.eneco.2018.11.002>

General rights

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

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


Take down policy

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

Accepted Manuscript

A seasonal copula mixture for hedging the clean spark spread with wind power futures

Troels Sønderby Christensen, Anca Pircalabu, Esben Høg



PII: S0140-9883(18)30442-0

DOI: <https://doi.org/10.1016/j.eneco.2018.11.002>

Reference: ENEECO 4211

To appear in: *Energy Economics*

Received date: 30 August 2017

Revised date: 2 October 2018

Accepted date: 2 November 2018

Please cite this article as: Troels Sønderby Christensen, Anca Pircalabu, Esben Høg , A seasonal copula mixture for hedging the clean spark spread with wind power futures. Eneeco (2018), <https://doi.org/10.1016/j.eneco.2018.11.002>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

A seasonal copula mixture for hedging the clean spark spread with wind power futures

Troels S nderby Christensen^{a,b,*}, Anca Pircalabu^{a,b}, Esben H g^a

^a*Department of Mathematical Sciences, Aalborg University, Skjernvej 4A, 9220 Aalborg  st, Denmark*

^b*Quantitative Analytics, Neas Energy, Skelagervej 1, 9000 Aalborg, Denmark*

Abstract

The recently introduced German wind power futures have brought the opportunity to address the problem of volume risk in wind power generation directly. In this paper, we study the hedging benefits of these instruments in the context of peak gas-fired power plants, by employing a strategy that allows trading in the day-ahead clean spark spread and wind power futures. To facilitate hedging decisions, we propose a seasonal copula mixture for the joint behavior of the day-ahead clean spark spread and the daily wind index. The model describes the data surprisingly well, both in terms of the marginals and the dependence structure, while being straightforward and easy to implement. Based on Monte Carlo simulations from the proposed model, the results indicate that significant benefits can be achieved by using wind power futures. Moreover, a comparison study shows that accounting for asymmetry, tail dependence, and seasonality in the dependence structure is especially important in the context of risk management.

Keywords: Clean spark spread, Wind power futures, Copula models, Time-varying dependence, Hedging

JEL: C22, C58, G11, G17, Q40

1. Introduction

The sudden change in German energy policy that followed the Fukushima nuclear accident marked a new era for the German power market. Since the nuclear shutdown and the shift to renewables, Germany has experienced an impressive growth in both wind and solar power, and has reached a level that far exceeds the Ky to climate obligations. This change has undoubtedly brought benefits on several fronts, however, the non-programmable nature of wind and solar electricity production has resulted in a large share of weather-dependent supply of electricity. From a financial point of view, the cash-flows from such non-programmable power plants can be incredibly volatile, not only due to price uncertainty, but also due to the uncertainty associated with the volume produced. While renewable generators are clearly affected by the uncertain volume, they are not the only ones; by market design and economics principle, the presence of renewables in the bid stack will always force conventional generators to produce less. In Germany, where the share of renewable energy is especially high, the conventional producers' competitiveness on e.g. the spot and forward markets has deteriorated, which has in turn invoked the need for far more intricate operation patterns and strategies.

In light of the advancements concerning renewables in Germany and the challenges imposed by volume risks for many different market players, the European Energy Exchange (EEX) recently introduced a financial instrument to mitigate the volume uncertainty associated with wind power generation. This instrument is referred to as a *wind power futures*, and its underlying is the German wind index. Representative agents for the sell and buy sides of wind power futures are the wind electricity producers and the conventional electricity producers, respectively. On one hand, low wind scenarios are unfavorable for wind electricity

*Corresponding author. E-mail: troelsc@math.aau.dk

producers, since they have a lowering effect on cash-flows; on the other hand, conventional generators are exposed to high wind scenarios, since a large share of wind power in the electrical grid displaces the costlier sources.

In this paper, we study the risk-reducing benefits of wind power futures in the context of conventional generators that operate in the day-ahead market whenever profitable. As a representative agent for the conventional generator, we consider the case of a peak gas-fired power plant whose profit per unit of electricity produced is measured in terms of the day-ahead clean spark spread. Since the dependence between the day-ahead clean spark spread and the wind index is essential for assessing the benefits of wind power futures, the contribution of this paper is twofold.

First, we propose a seasonal copula mixture to model the joint behavior of the day-ahead clean spark spread and the daily wind index. The model is fitted to four years of German data, and captures the marginal behavior of the individual variables and also the seasonality in the dependence between the variables very well. Second, we employ the proposed seasonal copula mixture to facilitate hedging decisions and showcase the effectiveness of wind power futures. To highlight the benefits of the seasonal copula mixture, we perform a study where the proposed model is compared against alternative models.

Owing to the recent introduction of the German wind power futures, the related literature is very scarce. The first related study is that of Gersema and Wozabal (2017), where the authors focus mainly on the pricing of wind power futures and explaining risk premia, for which an equilibrium pricing model is proposed. Also concentrating on the pricing aspect is the work of Benth and Pircalabu (2018), who apply a no-arbitrage approach to the pricing of wind power futures, and obtain results concerning the sign of risk premia that support the conclusions drawn in Gersema and Wozabal (2017). In contrast to the two existing studies, which focus mainly on pricing and less on hedging and risk management, we take a simplistic approach to pricing but study in detail aspects related to the risk-reducing ability of wind power futures. Nevertheless, we acknowledge that some of the results in Gersema and Wozabal (2017) and Benth and Pircalabu (2018) are very relevant in the context of the present study, and they shall thus be included in our discussion.

Turning to applications of copulas in energy markets, we mention that these models have gained substantial interest over the past years and have become a popular tool to model the non-linear dependence between different commodities. Some examples concerning applications of bivariate copulas are Börger *et al.* (2009), Benth and Kettler (2011), Grothe and Schneiders (2011), Avdulaj and Barunikl (2015), and Elberg and Hagspiel (2015). For applications beyond bivariate copulas, we mention the study of Pircalabu and Jung (2017), and that of Aeppli *et al.* (2017). The present paper contributes to this stream of literature in terms of the application, which to the best of our knowledge has not yet been considered, and also in terms of modeling approach, by proposing an extension that deals with seasonality in the dependence structure.

The remaining of this paper is structured as follows: In Section 2, we introduce the data and elaborate on the construction of the variables. In Section 3, we describe the modeling framework and report estimation results. Section 4 introduces the seasonal copula mixture model and provides evidence for its quality of fit. In Section 5, we employ the proposed model to study the benefits of wind power futures, and perform various comparison studies. Section 6 concludes.

2. Background and data

To investigate the benefits of wind power futures for a gas-fired power plant (GFPP), two data components are of interest in the analysis performed in this paper: The day-ahead clean spark spread and the daily wind index. In this section, we address each of these in turn, commenting on their construction.

2.1. Clean spark spread

As an indicator for the profit per unit of electricity generated by a GFPP, we consider the day-ahead clean spark spread (CSS). This measure depends on electricity, gas, and emission prices, and also on the heat rate and the emission factor. The heat rate represents the required number of natural gas MWhs to produce one MWh electricity, i.e., the efficiency at which the GFPP transforms gas to electricity. Further, the emission factor represents the number of tons of CO₂ emitted by producing one MWh electricity.

With GFPPs being mainly peak-operated power plants—that is, power plants dispatching during the peak hours between 8 AM to 8 PM on weekdays and non-holidays—we consider the peak electricity price. Specifically, let S_t^E denote the day-ahead peak load electricity price, S_t^G the day-ahead gas price, and S_t^C the day-ahead emission price, with the subscript t indicating time measured in days. Further, let h be the heat rate and e the emission factor. We define the day-ahead CSS on day t as

$$CSS_t = S_t^E - hS_t^G - eS_t^C, \quad (1)$$

where S_t^E and S_t^G are measured in EUR/MWh, and S_t^C is measured in EUR/tCO₂.

2.1.1. Data preparation for the clean spark spread

To construct a time series for the day-ahead CSS, we consider the following time series:

- S_t^E : The German electricity price, which is computed as the average of all hourly electricity prices between 8 AM to 8 PM on weekdays and non-holidays. The source of this data is EEX.
- S_t^G : The day-ahead gas price for NetConnect Germany (NCG), which corresponds to the closing price. The source of this data is EEX.
- S_t^C : The EU Allowance unit of one tonne of CO₂ (EUA) phase 3 daily futures price. This data is collected from the Intercontinental Exchange, and represents the closing price. For more information regarding the EU emissions trading system, we refer the interested reader to European Commission (2017).

All time series above span 1030 observations in the period from 3 January 2013 to 30 December 2016, and cover weekdays that are non-holidays. To provide a sense of the data, we plot in Fig. 1 the time series corresponding to each of the three data sources described above. Moreover, we plot the CSS obtained by applying Eq. (1), and using the values for h and e reported in Table 1. These numbers are based on ICIS (2016), and shall be used in the remaining of this paper unless explicitly stated otherwise.

Heat rate h	Emission factor e
2.035	0.375

Table 1: Heat rate (MWhs natural gas per MWh electricity) and emission factor (tCO₂ per MWh electricity) based on ICIS (2016). The chosen heat rate corresponds to an efficiency of 49.13%.

2.2. Wind index

Since the German wind power futures (WPF) were introduced only recently, we find it relevant to provide a brief description of these products and to clarify their payoff structure. WPF contracts are written on the average wind index in Germany, and can be traded at the European Energy Exchange (EEX) and Nasdaq OMX. In this paper, we shall restrict our attention to the WPF traded at EEX.

The German wind index is obtained as the ratio between the total wind power generation and the total available installed wind power capacity. Hence, the index is bounded between zero and one, and provides a measure of the German wind utilization. Currently, delivery periods for WPF correspond to weeks, months, quarters and years, and only trading the base load profile is possible. Compared to the definition of the day-ahead CSS data in Eq. (1), there is clearly a mismatch between delivery periods, with wind power futures hedging all hours of every day, and gas turbines generating output during peak hours. However, this reflects the present market conditions, where the volume risk of a GFPP can only be imperfectly hedged.




Figure 1: Historical evolution of the daily day-ahead CSS (peak load), the German day-ahead electricity price (peak load), the NCG day-ahead gas closing price, and the EUA day-ahead closing price, from 3 January 2013 to 30 December 2016. The applied heat rate and emission factor to construct the day-ahead CSS are given in Table 1.

Assuming a delivery period $[T_1, T_2]$ consisting of H hours, the payoff corresponding to a long position in one WPF contract is given by

$$R^{WPF} = H \left(\underbrace{\frac{1}{T_2 - T_1 + 1} \sum_{t=T_1}^{T_2} W_t - W_{t_0}}_{=\bar{W}} \right) X, \quad (2)$$

where $W_t \in [0, 1]$ is the daily wind index, $\bar{W} \in [0, 1]$ is the realized average wind index over the delivery period, and $W_{t_0} \in [0, 1]$ can be thought of as the “futures price”, i.e., the index set at t_0 when entering the contract. Further, X is a pre-specified constant tick size which is used to convert the index differences into monetary measures. According to EEX, $X = 100$ EUR. We see from Eq. (2) that a short position in WPF will generate a profit in low-wind scenarios, making it a useful hedging instrument for the wind power producer. Conversely, a long position will generate a profit in the high-wind scenarios, to which the GFPPs are exposed.

2.2.1. Data preparation for the wind index

The index that a WPF contract is settled against is externally provided by EuroWind. Since trading in WPF started only recently, the amount of data available on the spot wind index provided by EuroWind is limited. To obtain a longer time series, we consider instead a proxy wind index constructed using the wind power production in Germany on a daily basis, and monthly recordings of the German installed wind power capacity, which are updated at the start of each month. The wind power production data is collected from the four different transmission system operators in Germany, and the source of the installed capacity data is PointConnect.

Specifically, the daily German wind index is constructed as

$$W_t = \text{Daily wind index} = \frac{\text{Daily wind power generation (MWh)}}{H_t \cdot \text{Installed capacity (MW)}},$$

where H_t denotes the number of hours in day t , and the installed capacity on a daily basis is obtained by linear interpolation. In order to unify the length of the day-ahead CSS and the wind index, we omit weekends and holidays for the wind index data. Hence, the constructed index spans the period from 3 January 2013 to 30 December 2016, a total of 1030 observations, and is plotted in Fig. 2(a).

To provide some evidence for how the constructed wind index matches the true settlement data, we plot in Fig. 2(b) our proxy together with the one year of actual data from EuroWind that we have available. The time series plot reveals an acceptable resemblance, and to provide a quantitative indication, we compute the mean absolute error to 0.020.

3. Model construction and fit

To model the joint behavior of the day-ahead CSS and the daily wind index (henceforth referred to as simply CSS and wind index, respectively), we consider copula models. Restricting our presentation to the two-dimensional case, a copula is the joint distribution of the random variables U_1 and U_2 , where each variable is marginally uniformly distributed as $\text{Unif}(0,1)$. Since our data exhibits seasonality and autocorrelation, we wish to filter out these effects before applying the copula. Therefore, we are here considering the conditional copula.

Let $F(\cdot|\mathcal{F}_{t-1})$ denote the conditional joint distribution function of the random vector $\mathbf{Y}_t = (Y_{1t}, Y_{2t})$, and let $F_1(\cdot|\mathcal{F}_{t-1})$ and $F_2(\cdot|\mathcal{F}_{t-1})$ denote the conditional continuous marginal distribution functions of Y_{1t} and Y_{2t} , respectively. Then, according to Sklar’s theorem [Sklar (1959)] for conditional distributions, there exists a unique copula C such that F can be decomposed as

$$F(y_{1t}, y_{2t}|\mathcal{F}_{t-1}) = C(F_1(y_{1t}|\mathcal{F}_{t-1}), F_2(y_{2t}|\mathcal{F}_{t-1})|\mathcal{F}_{t-1}). \quad (3)$$




Figure 2: Historical evolution of the German wind index (W_t) on a daily basis from 3 January 2013 to 30 December 2016, and a comparison of W_t with the actual wind index provided by EuroWind for the year 2016.

The converse also holds, meaning that given two univariate distributions F_1, F_2 and a copula C , F as defined in Eq. (3) is the joint distribution with margins F_1, F_2 . Thus, Sklar's theorem not only provides a way of *decomposing* a joint distribution function, but also a way of *composing* it given marginal distributions and a copula, both of which are very useful in practical applications. For the proof of Sklar's theorem for conditional distributions, we refer to Patton (2006(a)).

Recalling the probability integral transform, we note that $U_{it} := F_i(Y_{it}|\mathcal{F}_{t-1}) \sim \text{Unif}(0, 1)$, for $i = 1, 2$. Differentiating both sides of Eq. (3) with respect to (y_{1t}, y_{2t}) thus yields

$$f(y_{1t}, y_{2t}|\mathcal{F}_{t-1}) = c(u_{1t}, u_{2t}|\mathcal{F}_{t-1}) \cdot f_1(y_{1t}|\mathcal{F}_{t-1}) \cdot f_2(y_{2t}|\mathcal{F}_{t-1}), \quad (4)$$

where f denotes the joint density function, c is the copula density, and f_1, f_2 denote marginal density functions.

In our context, copula models are advantageous for various reasons: First, being able to capture dependence beyond the linear correlation can be of utmost importance when illustrating the hedging benefits of WPF, and this can be achieved with copulas. Second, we can separate the treatment of the dependence structure from that of the marginal behavior of the individual variables (cf. Eq. (4)), since the dependence structure is fully contained in the copula. Third, selecting one type of marginal distribution for the first variable does not restrict our choice of marginal distribution for the second variable.

Turning to the estimation of the model parameters, we let T denote the sample size, θ_c the copula parameters, and θ_1 and θ_2 the parameters of the marginal models. From Eq. (4) it follows that the log-likelihood function is

$$\begin{aligned} \log \mathcal{L} = & \sum_{t=1}^T \log c(u_{1t}, u_{2t}|\mathcal{F}_{t-1}, \theta_c) + \sum_{t=1}^T \log f_1(y_{1t}|\mathcal{F}_{t-1}, \theta_1) \\ & + \sum_{t=1}^T \log f_2(y_{2t}|\mathcal{F}_{t-1}, \theta_2). \end{aligned}$$

Here, we consider multi-stage maximum likelihood (MSML) estimation. This provides a far less complicated estimation procedure relative to one-stage MLE. Moreover, the studies of Joe (2005) and Patton (2006(b)) suggest that the efficiency loss is not substantial.

In the following two sections, we present in detail the marginal models and the constant copulas considered in this paper, and provide empirical evidence for the fit of these models to our data.

3.1. Marginal models

Since both the CSS and the wind index exhibit seasonality, we start the marginal treatment of the individual variables by applying suitable seasonal functions to remove the deterministic seasonal component.

For the CSS, we consider the seasonal function

$$f_t = a_1 + b_1 t + c_1 \sin(2\pi t/K) + c_2 \cos(2\pi t/K),$$

where a_1 is a constant, b_1 is the trend coefficient, and c_1 and c_2 are coefficients for the annual cycle. We have on average approximately $K = 258$ observations per year.

Not surprisingly, the seasonality function for the CSS resembles a seasonality function that would typically be considered for the day-ahead electricity price (see e.g. Haldrup and Nielsen (2006), Benth and Šaltytė Benth (2011), and Härdle and López Cabrera (2012)). This resemblance is caused by the magnitude of the electricity price compared to the gas and emission price, cf. Fig. 1, causing the former to have the dominant effect. Aside from electricity prices usually exhibiting a yearly seasonality, a strong within-week seasonality is also often observed. However, with the exclusion of weekends from our data, adding a term that addresses the weekly seasonality (e.g. day-of-week dummies) is unnecessary. Furthermore, adding more trigonometric terms (based on the periodogram) does not improve the fit of the seasonal function substantially.

Turning to the wind index, recall that this series is bounded between 0 and 1, cf. Fig. 2. Following Pircalabu and Jung (2017), we apply the logit-transform to the wind index¹, and consider the following seasonal function for the logit wind index (LWI):

$$f_t = a_1 + c_1 \sin(2\pi t/K) + c_2 \cos(2\pi t/K),$$

which is motivated by the prominent annual cycles we observe in the sample autocorrelation of the LWI. Also here, different meaningful extensions of the seasonal function were experimented with, without yielding a significant improvement.

The seasonality functions are fitted to the data by ordinary least squares, and Table 2 summarizes the results obtained for the CSS and the LWI.

	\hat{a}_1	\hat{b}_1	\hat{c}_1	\hat{c}_2
CSS	-10.131 (0.579)	0.014 (0.001)	-3.611 (0.411)	2.990 (0.404)
LWI	-1.828 (0.030)	- -	0.066 (0.042)	0.472 (0.043)

Table 2: OLS estimates for parameters of the seasonal functions for the CSS and the LWI. Standard errors are reported in parenthesis, and are based on a naive OLS calculation.

Next, we apply ARMA-GARCH filters to the deseasonalized data. Given a time series of data y_t , an ARMA(p, q)-GARCH(h, k) model is defined by

$$\begin{aligned} y_t &= \sum_{i=1}^p \phi_i y_{t-i} + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + \varepsilon_t, \\ \varepsilon_t &= \sigma_t \eta_t, \\ \sigma_t^2 &= \omega + \sum_{i=1}^h \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^k \beta_j \sigma_{t-j}^2, \end{aligned}$$

where $\eta_t \sim iid N(0, 1)$. It was Engle (1982) who introduced the ARCH model, and later Bollerslev (1986) who extended the variance equation to include lagged values of σ_t^2 . For a review of ARMA and GARCH

¹The logit function is given by $\text{logit}(x) = \log(x) - \log(1 - x)$.

models, see for example Shumway and Stoffer (2007). In the following, we denote by η_t^{CSS} and η_t^{LWI} the standardized residuals resulting from applying the ARMA–GARCH models to the CSS and the LWI, respectively.

Model selection is based on the Bayesian Information Criterion (BIC), and we consider ARMA(p, q)–GARCH(h, k) models for all possible combinations of p, q, h and k , for $p = 0, \dots, 7$, $q = 0, \dots, 7$, $h = 0, 1, 2$, and $k = 0, 1, 2$. The optimal order of the models and the corresponding estimated parameters are reported in Table 3.

	CSS	LWI
Model	ARMA(2,1) – GARCH(1,1)	ARMA(1,1)
Conditional mean		
AR1 $\hat{\phi}_1$	1.313 (0.077)	0.360 (0.055)
AR2 $\hat{\phi}_2$	-0.360 (0.058)	-
MA1 $\hat{\theta}_1$	-0.818 (0.064)	0.208 (0.055)
Variance $\hat{\sigma}^2$	-	0.677 (0.030)
Conditional variance		
Constant $\hat{\omega}$	3.611 (0.636)	-
ARCH $\hat{\alpha}_1$	0.110 (0.022)	-
GARCH $\hat{\beta}_1$	0.827 (0.027)	-

Table 3: Type and order of marginal models, parameter estimates and corresponding standard errors in parenthesis.

Considering the goodness-of-fit of the normal distribution, we find a satisfactory fit in the case of $\hat{\eta}_t^{LWI}$, cf. Figs. 3(e) and 3(f). This is however not the case for $\hat{\eta}_t^{CSS}$. Consequently, we relax the normality assumption for the CSS, and consider instead the normal-inverse Gaussian (NIG) distribution. The probability density function of the NIG distribution is given by

$$g(x|\alpha, \beta, \mu, \delta) = \frac{\alpha \delta G_1 \left(\alpha \sqrt{\delta^2 + (x - \mu)^2} \right)}{\pi \sqrt{\delta^2 + (x - \mu)^2}} e^{\delta \sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)},$$

where

$$G_1(x) = \frac{1}{2} \int_0^\infty e^{-\frac{1}{2}x(t+t^{-1})} dt$$

is the modified Bessel function of third kind and index 1. The NIG distribution is a popular choice in the financial literature (for some examples, see Barndorff-Nielsen (1997a), Rydberg (1997), Barndorff-Nielsen (1997b), and Jensen and Lunde (2001)), and is also often able to provide a good description of commodity data, see e.g. Benth and Šaltytė Benth (2004) and Benth and Kettler (2011). The NIG distribution is fitted to the residuals from the ARMA(2,1)–GARCH(1,1) model cf. Table 3 by maximum likelihood, and the parameter estimates are reported in Table 4. As it appears from the histogram and quantile plots displayed in Figs. 3(a) and 3(b), the NIG distribution provides a satisfactory fit to the CSS data.

$\hat{\alpha}$	$\hat{\beta}$	$\hat{\mu}$	$\hat{\delta}$
1.584 (0.307)	-0.189 (0.144)	0.189 (0.126)	1.534 (0.281)

Table 4: Maximum likelihood estimates obtained by fitting the NIG distribution to $\hat{\eta}_t^{CSS}$. Corresponding standard errors are given in parenthesis.

To provide further evidence for the appropriateness of the chosen marginal distributions, we perform the Kolmogorov-Smirnov (K-S) and the Cramer-von Mises (CvM) goodness-of-fit tests. To obtain critical values for the tests, we employ the simulation-based method described in detail in Patton (2013). In the

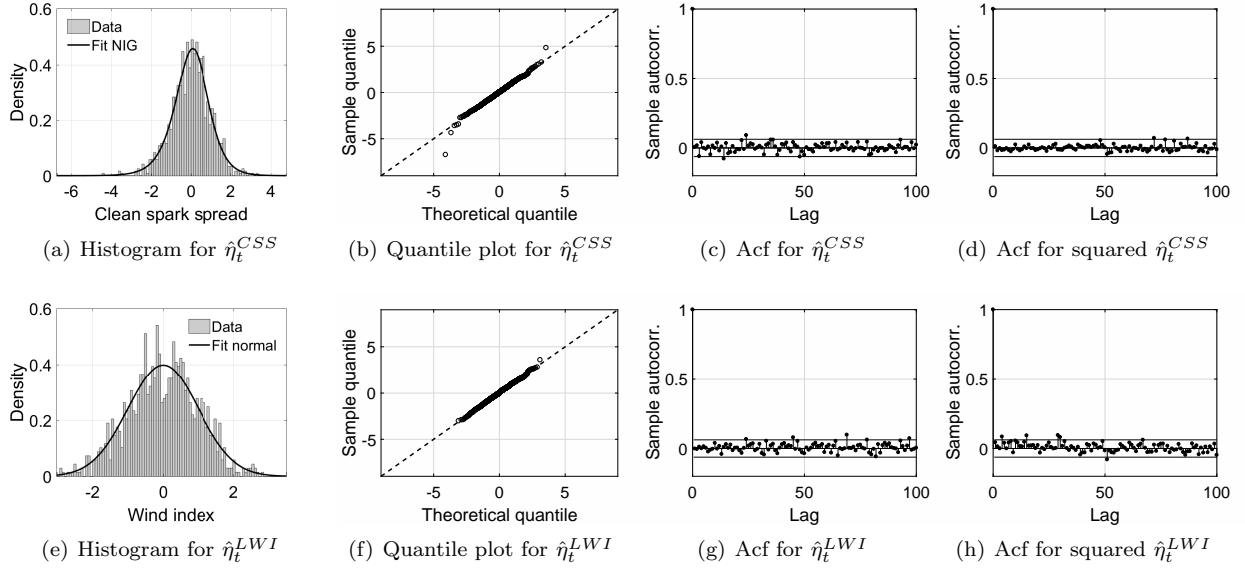


Figure 3: Diagnostics for the standardized residuals $\hat{\eta}_t^{CSS}$ (first row) and $\hat{\eta}_t^{LWI}$ (second row).

CSS case we obtain p -values of 0.627 and 0.785 for the K-S and CvM test, respectively, and can thus not reject the null that the NIG distribution is well-specified. This is also the conclusion in the LWI case, where we test the goodness-of-fit of the normal distribution. Here, the resulting p -values are 0.746 and 0.915 for the K-S and CvM test, respectively.

Aside from providing evidence for the goodness-of-fit of the marginal distributions, the sample autocorrelations provided in Fig. 3 suggest that no considerable serial dependence is left in the conditional mean and variance, for either variable. Having verified that the models proposed here are suitable for describing the marginal behavior of the CSS and the LWI, we proceed in the next section to the modeling of the dependence structure.

3.2. Constant copula models

Let F^{NIG} and F^N denote the cumulative distribution functions for the NIG and standard normal distribution, respectively. To obtain the approximately uniforms that are the input variables to the copula function, we apply the probability integral transform, i.e.,

$$\begin{aligned}\hat{u}_t^{CSS} &= F^{NIG}(\hat{\eta}_t^{CSS} | \mathcal{F}_{t-1}, \hat{\alpha}, \hat{\beta}, \hat{\mu}, \hat{\delta}), \\ \hat{u}_t^{LWI} &= F^N(\hat{\eta}_t^{LWI} | \mathcal{F}_{t-1}),\end{aligned}$$

for $t = 1, \dots, T$. In Fig. 4 we plot the resulting probability integral transforms against each other, revealing that the variables are negatively related. This finding is not surprising considering the negative dependence between electricity prices and the wind index. A high wind penetration in the electricity grid puts a downward pressure on day-ahead electricity prices owing to the process of day-ahead price formation, which prioritizes cheap electricity producers. With everything else being equal, this lowering effect on the electricity price is then reflected in the CSS, which is also lowered, cf. Eq. (1). Similar arguments apply to the case of a low wind scenario, where electricity prices are typically pushed upwards.

The dependence structure seems to be slightly asymmetric, with the north west corner of Fig. 4 exhibiting more concentration and being sharper in shape compared to the south east corner. That is, there seems to be more probability of observing the combination of extremely high wind / extremely low CSS than the reverse. Since non-zero dependence in extreme events could have notable implications for the benefits of WPF, being able to capture such behavior in a model for the dependence structure must be considered.




Figure 4: Empirical copula density.

Luckily in the context of copulas, such extreme events can be easily captured by considering certain copula families that allow for non-zero *tail dependence*. More specifically, the lower and upper tail dependence can be defined as

$$\begin{aligned}\Lambda^l &= \lim_{q \rightarrow 0^+} P(u_t^{CSS} \leq q | u_t^{LWI} \leq q) \\ \Lambda^u &= \lim_{q \rightarrow 1^-} P(u_t^{CSS} > q | u_t^{LWI} > q),\end{aligned}$$

where q denotes the quantile. Clearly, since our data is characterized by negative dependence, computing Λ^l and Λ^u as defined above is not meaningful. This can however be resolved by performing suitable rotations of the data, which shall be discussed in more detail shortly.

To investigate which copula best describes the dependence structure illustrated in Fig. 4, we consider first the following standard copulas, which are often employed in the related literature: Gaussian, Gumbel, rotated Gumbel (RGumbel), Clayton, rotated Clayton (RClayton), Frank, symmetrized Joe-Clayton, and Student t .² These copula models cover a wide range of dependency structures, with some models being able to capture asymmetric dependence, and also upper and lower tail dependence, i.e., a non-zero probability of extreme events happening simultaneously. In the interest of brevity, we shall not go into detail with the properties of each copula model here, and refer instead to McNeil *et al.* (2005), Nelsen (1999), and Patton (2006(a)) for a comprehensive description.

To allow for further flexibility compared to the standard copulas enumerated above, we also consider copula mixtures. As in e.g. Rodriguez (2007) and Dias and Embrechts (2009), for a given t we mix copula a , having copula density $c^a(\cdot | \mathcal{F}_{t-1}, \theta^a)$, with copula b , having copula density $c^b(\cdot | \mathcal{F}_{t-1}, \theta^b)$, by using a mixing parameter $0 < \lambda < 1$ and the following form:

$$c^m(\cdot | \mathcal{F}_{t-1}, \theta^a, \theta^b, \lambda) = \lambda c^a(\cdot | \mathcal{F}_{t-1}, \theta^a) + (1 - \lambda) c^b(\cdot | \mathcal{F}_{t-1}, \theta^b). \quad (5)$$

As expected, a mixture copula inherits characteristics from its mixing components. In the following proposition, we present an especially useful result relating to the tail dependence of a mixture copula, which we shall use shortly. Notice that we omit the conditioning to ease the notation.

²By rotated, we mean a 180 degree rotation of the data.

Proposition 1. Let $U_i \sim \text{Unif}(0,1)$ for $i = 1, 2$, and let C^m denote the bivariate copula of (U_1, U_2) . Further assume C^m is given as the mixture

$$C^m(u_1, u_2) = \lambda C^a(u_1, u_2) + (1 - \lambda) C^b(u_1, u_2),$$

where C^a and C^b are two bivariate copulas, and $0 < \lambda < 1$. Then, the lower tail dependence Λ^l and the upper tail dependence Λ^u for the mixture C^m are given as

$$\Lambda^l = \lambda \Lambda^{l,a} + (1 - \lambda) \Lambda^{l,b},$$

and

$$\Lambda^u = \lambda \Lambda^{u,a} + (1 - \lambda) \Lambda^{u,b},$$

where $\Lambda^{l,a}$, $\Lambda^{l,b}$, $\Lambda^{u,a}$, and $\Lambda^{u,b}$ are the respective tail dependence measures for C^a and C^b .

Proof. See Appendix A.1.

Moving on to the estimation aspect, we let $c(\cdot | \mathcal{F}_{t-1}, \boldsymbol{\theta})$ denote the conditional copula density with parameter vector $\boldsymbol{\theta} \in \mathbb{R}^l$, where $l \in \mathbb{N}$ is the number of parameters in the copula. For each copula model, we obtain an estimate for $\boldsymbol{\theta}$ by maximizing the copula log-likelihood, i.e.,

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmax}} \sum_{t=1}^T \log c(\hat{u}_t^{CSS}, \hat{u}_t^{LWI} | \mathcal{F}_{t-1}, \boldsymbol{\theta}). \quad (6)$$

We note that it is only the Gaussian and Student t copulas that allow for negative dependence. To fit the remaining copulas to our data, we perform suitable rotations of the data. Specifically, we rotate around the LWI variable for the case of the Gumbel, Clayton, Frank, and symmetrized Joe-Clayton copulas, and consider thus the pair $(\hat{u}_t^{CSS}, 1 - \hat{u}_t^{LWI})$ as input to Eq. (6) for these models. Regarding tail dependence, the rotation of data implies that lower tail dependence for the estimated copulas corresponds to high wind index / low CSS scenarios (north west corner of Fig. 4), whereas upper tail dependence for the estimated copulas corresponds to low wind index / high CSS scenarios (south east corner of Fig. 4). To fit the RGumbel and the RClayton, we note that a further 180 degree rotation of the pair $(\hat{u}_t^{CSS}, 1 - \hat{u}_t^{LWI})$ is performed.

In Table 5, we report the estimation results for all standard copula models and three selected mixtures. Other copula mixtures aside from those reported in Table 5 were considered, but we found no increase in performance. As a model selection criterion, we employ the Akaike Information Criterion (AIC). According to the AIC, the preferred model is the mix of Frank and RGumbel (hereafter denoted FRG copula), confirming the presence of slight asymmetry in the dependence structure illustrated in Fig. 4.

Considering the FRG copula in more detail, its first mixing component, the Frank copula, imposes symmetric dependence and a zero tail dependence. Its second mixing component, the RGumbel, imposes an asymmetric dependence structure, with zero upper tail dependence and lower tail dependence given by

$$\Lambda^{l, RG} = 2 - 2^{1/\theta^{RG}},$$

where θ^{RG} is the parameter for the RGumbel copula. Recalling Prop. 1, we thus have that the upper and lower tail dependence for the FRG copula are

$$\begin{aligned} \Lambda^{u, FRG} &= 0, \\ \Lambda^{l, FRG} &= (1 - \lambda) \Lambda^{l, RG}. \end{aligned} \quad (7)$$

The fit produced by the FRG translates into a tail dependence coefficient of approximately 0.359 when considering the north west corner of Fig. 4, and hence a rather high probability of extremely high wind index / low CSS happening simultaneously. To illustrate the shape of the FRG copula and how it deviates from the shapes of the individual copulas in the mixture, we plot in Fig. 5 simulations from the fitted Frank, RGumbel and FRG copulas. The simulations reveal that while the fitted Frank copula is too symmetric and the fitted RGumbel is too asymmetric compared to the observed dependence in Fig. 4, the fitted FRG mixture is able to dampen the individual effects, hence providing a better resemblance to the observed dependence structure.

Copula model		Param. (s.e.)	$\log \mathcal{L}_c$	AIC
Gaussian	$\hat{\rho}$	-0.636 (0.019)	266.503	-531.006
Gumbel*	$\hat{\theta}$	1.675 (0.049)	226.083	-450.165
RGumbel*	$\hat{\theta}$	1.773 (0.053)	277.829	-553.657
Clayton*	$\hat{\theta}$	1.233 (0.083)	245.947	-489.894
RClayton*	$\hat{\theta}$	0.911 (0.071)	167.862	-333.724
Frank*	$\hat{\theta}$	5.029 (0.247)	267.805	-533.611
Sym. Joe-Clayton*	$\hat{\Lambda}^u$	0.274 (0.048)	265.777	-527.553
	$\hat{\Lambda}^l$	0.539 (0.026)		
Student t	$\hat{\rho}$	-0.646 (0.020)	274.596	-545.192
	$\hat{\nu}$	9.873 (15.008)		
Mix of Gumbel and RGumbel*	$\hat{\theta}_1$	1.964 (0.500)	285.060	-564.120
	$\hat{\theta}_2$	1.797 (0.101)		
	$\hat{\lambda}$	0.219 (0.081)		
Mix of Frank and RGumbel*	$\hat{\theta}_1$	4.552 (1.309)	286.419	-566.837
	$\hat{\theta}_2$	1.920 (0.169)		
	$\hat{\lambda}$	0.365 (0.104)		
Mix of Gaussian and RGumbel*	$\hat{\rho}$	-0.494 (0.092)	285.152	-564.304
	$\hat{\theta}$	2.095 (0.234)		
	$\hat{\lambda}$	0.357 (0.130)		

Table 5: Estimation results for 11 selected copula models. The maximized value of the copula log-likelihood is denoted $\log \mathcal{L}_c$. For the functional forms of the considered copulas and other characteristics, we refer to McNeil *et al.* (2005), Nelsen (1999) and Patton (2006(a)). A copula marked by an asterisk has been estimated using a suitable rotation of the data. Standard errors are based on 999 simulations.




Figure 5: T simulations from the fitted Frank, RGumbel and FRG copulas, cf. Table 5.

4. Time-Varying Dependence

Up until this point, we have assumed a static model for the dependence structure, which is seldom a realistic representation. Natural follow-up questions are therefore related to the presence and type of time variation in the dependence. In this section, we consider these questions in more detail.

To investigate the time-varying aspect we consider Spearman's ρ , which is a measure of concordance. In terms of a bivariate copula C , Spearman's ρ can be expressed as (see e.g. McNeil *et al.* (2005))

$$\rho = 12 \int_0^1 \int_0^1 C(u_1, u_2) du_1 du_2 - 3. \quad (8)$$

We compute Spearman's ρ between \hat{u}_t^{CSS} and \hat{u}_t^{LWI} based on a rolling window of 60 days. Fig. 6 displays the results, and reveals a strong seasonal pattern in the dependence structure. According to Fig. 6, the dependence is strongest around winter and weakest around summer. A possible explanation for this behaviour relates to the power generation mix in Germany and the import/export conditions. During winter, the increased wind power production has a direct lowering effect on the daily electricity price due to the mechanism of day-ahead electricity price formation. As argued in Section 3.2, this lowers the CSS. During summer, the lower wind power production does not have the same direct effect on the daily electricity price. If that were the case, prices should increase. The high photovoltaic production during peak periods combined with the high likelihood of being able to import cheap nuclear power from France prevents however prices from increasing. Consequently, this weakens the dependence between the wind power production and the CSS during the summer months.




Figure 6: Spearman's ρ between \hat{u}_t^{CSS} and \hat{u}_t^{LWI} based on a 60-days rolling window. The confidence interval is based on 999 bootstraps. Note that the date corresponding to each estimate refers to the last day in the 60-days period.

In light of these findings we consider next extending the static copula mixture, such that the yearly seasonality in the dependence measured by Spearman's ρ in Fig. 6 can be accounted for.

4.1. A seasonal copula mixture

Since the FRG copula is the best performing static copula cf. Table 5, we shall restrict our attention to this particular model in order to fix ideas. First, let us state a general result concerning Spearman's ρ for copula mixtures, which is particularly useful in our modeling context. Again here, we omit the conditioning for notational convenience.

Proposition 2. Let $U_i \sim \text{Unif}(0,1)$ for $i = 1, 2$, and let C^m denote the bivariate copula of (U_1, U_2) . Further, suppose C^m is given as

$$C^m(u_1, u_2) = \lambda C^a(u_1, u_2) + (1 - \lambda)C^b(u_1, u_2)$$

for two copulas C^a and C^b , and mixing parameter $0 < \lambda < 1$. Then Spearman's ρ implied by C^m can be expressed as

$$\rho^m = \lambda \rho^a + (1 - \lambda) \rho^b, \quad (9)$$

where ρ^a is Spearman's ρ corresponding to copula C^a , and ρ^b is Spearman's ρ corresponding to copula C^b .

Proof. See Appendix A.2.

It follows from Prop. 2 that Spearman's ρ for the copula mixture is simply a linear combination of the individual Spearman's ρ 's corresponding to the copulas comprised in the mixture. Thus, introducing time variation in ρ^a and ρ^b translates into time variation in ρ^m . Further, it is relatively easy to compute Spearman's ρ , even for copulas where no explicit relation between Spearman's ρ and the copula parameter is available. Considering the FRG copula, the relationship between the copula parameter and Spearman's ρ for both the Frank and RGumbel copula is shown in Fig. 7. By letting the superscripts F and RG indicate




Figure 7: Spearman's ρ as a function of parameter value for both the Frank and RGumbel copula.

their link to the particular copula, we note that Spearman's ρ is monotonically increasing as a function of the corresponding copula parameter, θ^F and θ^{RG} . Therefore, specifying time variation for ρ^F and/or ρ^{RG} will also uniquely determine values of θ^F and θ^{RG} . If we instead were to introduce time variation directly in Spearman's ρ for the FRG copula, we would not be able to identify θ^F and θ^{RG} .

Based on the discussion above and motivated by the pronounced yearly cycle in Fig. 6, we propose the following extension to the static FRG model. Specifically, we introduce a yearly cycle in Spearman's ρ corresponding to the RGumbel copula, i.e.,

$$\rho_t^{RG} = a^{RG} + b^{RG} \sin(2\pi t/K) + c^{RG} \cos(2\pi t/K), \quad (10)$$

where a^{RG} , b^{RG} , and c^{RG} are constant coefficients, and $K = 258$ as was the case with the seasonal functions in Section 3.1. Regarding the Frank contribution in the FRG copula, we keep the corresponding Spearman's ρ static. Consequently, the evolution equation for the overall Spearman's ρ implied by the *seasonal* FRG copula is

$$\rho_t^{SFRG} = \lambda^{SFRG} \rho^F + (1 - \lambda^{SFRG}) \rho_t^{RG}. \quad (11)$$

Given the seasonal specification in Eqs. (10)–(11), the model is estimated by maximizing the FRG copula loglikelihood. The estimation results are given in Table 6, revealing a clear improvement in AIC compared to the static FRG copula.

$\hat{\rho}^F$	\hat{a}^{RG}	\hat{b}^{RG}	\hat{c}^{RG}	λ^{SFRG}	$\log \mathcal{L}_c$	AIC
0.742	0.566	-0.132	0.208	0.316	309.894	-609.788
(0.072)	(0.041)	(0.040)	(0.043)	(0.092)		

Table 6: Maximum likelihood estimation results for the seasonal FRG copula described in Eqs. (10)–(11). The maximized value of the copula log-likelihood is denoted $\log \mathcal{L}_c$. Standard errors are reported in parenthesis and are computed following the simulation-based procedure described in detail in Patton (2013), where we note that the estimation error from the marginal models is taken into account. The seasonal FRG copula was fitted to a suitable rotation of the data, cf. Section 3.2.

Since the FRG copula has three parameters, there are of course different ways of incorporating yearly seasonality in the model. Some alternatives in terms of the FRG copula are discussed in Appendix B, where we also provide detailed estimation results to support the model specification stated above. On a different note, we stress that selecting the static FRG copula as the optimal model amongst static alternatives does not guarantee that the seasonal FRG copula will be preferred to time-varying extensions of other copula models. As a result, similar extensions as those proposed in Eqs. (10)–(11) were implemented for most of the copulas in Table 5 to ensure that the seasonal FRG is superior in terms of AIC.

To illustrate the fit of the proposed seasonal FRG model, we plot in Fig. 8(a) the empirical Spearman’s ρ together with $\hat{\rho}_t^{SFRG}$ implied by the seasonal FRG, using a 60-days moving window, as in Fig. 6. As a standard of comparison, we include the Spearman’s ρ implied by the static FRG. The results indicate that the dependence implied by the seasonal FRG follows the yearly cycle observed in the actual Spearman’s ρ rather well. Moreover, it appears from Fig. 8(a) that we would underestimate the strength of the dependence between the CSS and the LWI during autumn and winter with the static FRG. The reverse is observed during spring and summer, with the strength of the dependence being overestimated by the static FRG. To provide further support for the proposed seasonal FRG, we display in Fig. 8(b) a simulated path over a four-year horizon, which resembles the actual data nicely.




Figure 8: (left) Actual Spearman’s ρ and Spearman’s ρ implied by the static and seasonal FRG copula, based on a 60-days rolling window. The date corresponding to each estimate refers to the last day in the 60-days period. (right) A simulated path of Spearman’s ρ from the seasonal FRG copula, aggregated using a 60-days rolling window.

To complement Fig. 8(a), a clearer picture of the yearly shape of Spearman’s ρ implied by the seasonal FRG is given in Fig. 9(a), where we illustrate the fit at each time point during a year (i.e., no averaging of Spearman’s ρ is performed). Equally interesting to consider is the lower tail dependence implied by

the fitted seasonal FRG copula, which follows directly from applying Eq. (7). The results are plotted in Fig. 9(b), revealing that the lower tail dependence coefficient reaches its lowest value of approximately 0.2 during summer and its maximum value of approximately 0.50 during winter. This entails that there is a rather large difference between the probability of observing the event of extreme high wind index / low CSS during winter compared to summer. In Fig. 9(b), we again provide as benchmark the corresponding static estimate.




Figure 9: Spearman's ρ and lower tail dependence implied by the static and seasonal FRG copulas throughout the year. A 90% confidence interval is provided for the seasonal parts.

4.1.1. Dynamic copula models

Before proceeding to quantifying the hedging benefits of wind power futures, we comment briefly on another class of models that have become popular in the related literature because of their broad applicability: The Generalized Autoregressive Score (GAS) models proposed by Creal *et al.* (2013). Extending the seasonal copula mixture proposed in Eqs. (10)–(11) as to allow for one or more parameters to evolve according to the GAS equation is possible. In the context of our study however, where we shall base hedging decisions on simulations over a long time horizon, we argue against such extensions. With the GAS model, we would introduce much complexity compared to the present straightforwardness of the seasonal copula, but not add that much value. Moreover, we stress that compared to the static dependence model, the addition of the simple seasonal extension not only provides a significant improvement, but is very easily interpretable, making it very appealing from a practical perspective.

For applications such as forecasting or short-term simulation, we acknowledge the added value of including a GAS dynamic to the seasonal copula mixture. Therefore, we include in Appendix B a description of the GAS model together with estimation results obtained by fitting two dynamic copula models to our data.

5. Application results

Having established a model for the joint behavior of the CSS and the wind index, we consider next the quantification of the benefits that WPF can offer GFPPs. A GFPP acting in the day-ahead market can decide from day to day whether to run or not, and thereby take advantage of the daily variation in the CSS. By the construction of the day-ahead wholesale electricity market, a GFPP will not run in times of a

negative CSS; its profit R^{CSS} for a period $t \in [T_1, T_2]$ can thus be represented as

$$R^{CSS} = \sum_{t=T_1}^{T_2} 12 \max(CSS_t, 0)s, \quad (12)$$

where s is the size of the GFPP measured in MW, and 12 is the number of peak load hours during a day. Recalling the payoff in Eq. (2), taking a position $\gamma \in \mathbb{Z}$ in WPF contracts yields the hedged profit of the GFPP, which we denote by R :

$$R = R^{CSS} + \gamma R^{WPF}. \quad (13)$$

We note that by excluding weekends and holidays from our analysis, these are not captured in R^{WPF} . We argue however that this does not alter the overall conclusions drawn below. To facilitate hedging decisions, we perform Monte Carlo simulations from the proposed model. Specifically, the marginal models fitted in Section 3.1 and the seasonal FRG copula fitted in Section 4.1 are employed to produce simulations of the joint behavior of the CSS and the wind index, i.e., the pair (CSS_t, W_t) . The “price” W_{t_0} affecting R^{WPF} in Eq. (13) is computed by averaging across all Monte Carlo simulations of W_t for the delivery $[T_1, T_2]$.³ While we recognize that this pricing approach is simplistic in that it assumes a zero market price of risk, it simplifies our hedging exercise somewhat, since the mean of the hedged profit R will not be affected by varying the quantity γ . Consequently, instead of the classical mean-variance objective, we can restrict ourselves to the variance minimization criterion in order to determine optimal positions in WPF contracts. Hence, we consider the following objective:

$$\min_{\gamma \in \mathbb{Z}} \text{Var}[R]. \quad (14)$$

5.1. Effectiveness of wind power futures

To illustrate the results obtained by applying the hedging approach described above, we fix $s = 200$ MW, $t_0 = 30$ December 2016 (the last date in our sample), and perform 20,000 Monte Carlo simulations of the pair (CSS_t, W_t) one year ahead. The resulting simulated paths are split into monthly periods, and WPF prices corresponding to monthly deliveries are computed as explained earlier. Then, monthly quantities for R^{CSS} and R^{WPF} are constructed for each simulated path, and the minimization in Eq. (14) is applied to each month in turn. The subdivision to monthly profits is motivated by the seasonal pattern observed in the dependence structure cf. Fig. 8(a), and allows us to investigate the effect of the yearly seasonality on hedging-related aspects.

In Fig. 10, we illustrate the simulated unhedged profit distribution R^{CSS} and the hedged profit distribution R obtained by solving Eq. (14) for the months July and October. We observe a compression of the profit distribution in both cases when applying the hedge, which entails that WPF have variance reducing effects. In fact, this finding applies to all 12 months, as will be illustrated shortly.

Perhaps unsettling is the fact that losses can occur when considering the hedged profit distribution in Fig. 10(a), whereas the unhedged profits cannot attain negative values by construction (see Eq. (12)). Nevertheless, we find that the probability of a loss when hedging with WPF is approximately 0.4% on average. The downside of performing the hedge is therefore quite small. In the pursuit to impair this

³When constructing the wind index data used in this paper, we considered the historical evolution of its two underlying data components, namely the wind power production and the installed capacity, as discussed in Section 2. This implicitly means that we have captured 1) the variations due to changes in wind speeds and 2) the variations caused by the increase in installed capacity and changes in the geographical distribution of wind turbines. While the latter aspect is important to capture in the modeling part of this paper, we argue that a different wind index series should be used in a pricing context. This is because today’s WPF price is not affected by the historical evolution of the installed capacity and the changes in the geographical distribution of turbines, but by the *present* conditions. We argue that this issue does not affect the conclusions drawn in this paper, but can have serious implications in other contexts. For more details, we refer the interested reader to Benth and Pircalabu (2018).




Figure 10: Examples of profit distributions before and after hedging with WPF, based on Monte Carlo simulations from the proposed seasonal FRG copula model.

concern even further, recall that the price of WPF is computed under the assumption of a zero market price of risk. In reality, the studies of Gersema and Wozabal (2017) and Benth and Pircalabu (2018) find evidence of a negative market price of risk in the German market for WPF, implying that a GFPP buys WPF at a discounted price compared to the one computed here. Accounting for this would shift the hedged profit distributions to the right, potentially excluding losses altogether.

Next, we consider in more detail the reduction in the variance of profit distributions attained by performing the hedge. The results are stated in Fig. 11 for all months of the year, and reveal considerable reductions; even for May and June, where we observe the lowest values, the variance reductions are above 10 %. Further, notice the connection between the yearly pattern of the reductions in Fig. 11 and the implied Spearman's ρ in Fig. 9(a): Not surprisingly, the stronger the dependence, the higher the variance reduction.




Figure 11: Variance reduction achieved by hedging with WPF, for each month of 2017. The results are based on Monte Carlo simulations from the proposed seasonal FRG copula model.

Also relevant to consider in this context is the impact on hedging effectiveness from changing the GFPP efficiency. To assess this, we allow the efficiency to vary from 43.13% to 55.13% with a step size of 2%, and let the emission factor vary according to

$$e = 0.184h,$$

which is based on ICIS (2016). Fig. 12 illustrates the variance reductions obtained with the different efficiencies, across all months of the year. It appears that increasing the efficiency (i.e., lowering the heat

rate) leads to an increase in the variance reductions for all months. The effect seems to be more pronounced during autumn and winter compared to spring and summer. From Eq. (1), it was already apparent that increasing the efficiency of a GFPP produces a higher CSS and hence increases profitability. The findings presented in Fig. 12 incentivize such action even further: Aside from the higher CSS, an increased hedging effectiveness of WPF can be achieved.




Figure 12: Variance reduction implied by hedging with WPF, for GFPP efficiencies spanning from 43.13% to 55.13% with a step size of 2%.

5.2. Comparison with alternative models

So far in the hedging application, we have focused on the results obtained with the preferred copula model, that is, the seasonal FRG. In this section, we wish to highlight the benefits of this copula compared to other less optimal alternatives. We consider the following natural progression in comparisons:

1. Frank copula versus FRG copula: In this comparison, we focus on the effect of asymmetry and tail dependence on the benefits of WPF. These features are captured by the FRG copula, as discussed in Section 3.2, but not by the Frank copula, which imposes symmetry and no tail dependence.
2. FRG copula versus seasonal FRG copula: Here, we concentrate on the effect of seasonal dependence on the hedging benefits.

To perform comparisons, we keep the marginal models proposed in Section 3.1 fixed, and repeat the simulations performed in Section 5.1 with the Frank, FRG, and seasonal FRG copulas, instead of only the seasonal FRG. We note that the same random seed was used to produce Monte Carlo simulations from the three models. Then, we compute optimal hedge quantities γ and associated variance reductions with each model, on a monthly basis.

5.2.1. The effects of asymmetry and tail dependence in the copula

Recalling Figs. 5(a) and 5(c), it is apparent that by employing the FRG copula instead of the Frank copula, we introduce a slight asymmetry and assign more probability to the extreme events where high wind and low CSS happen simultaneously. The resulting effects on hedging are depicted in Fig. 13, where we present the optimal hedge quantities and variance reductions produced by the two copulas. Regarding the former, we notice that the Frank copula generally suggests less WPF in the hedging portfolio. Further, the optimal hedge quantities vary across the year, which is a consequence of the seasonality captured in the marginal models.

Turning to the variance reductions, which are depicted in Fig. 13(b), we observe that the values implied by the Frank copula are generally lower compared to those implied by the FRG copula. This finding is expected, since GFPPs seek to cover their exposure to high wind / low CSS scenarios. By assigning more probability mass to precisely these events happening simultaneously, which is done by shifting from the Frank to the FRG copula, we increase the benefits of WPF. At the same time, due to the asymmetric

behavior of the FRG copula, we are not increasing the probability of observing the reverse combination of low wind / high CSS, and thus not counteracting the increased benefits of WPF. Briefly put, by believing in a dependence structure described by the Frank copula compared to the FRG copula, we would underestimate the risk-reducing power of WPF.




Figure 13: Comparison of hedging results implied by the Frank and FRG copulas.

5.2.2. The effects of seasonal time variation in the copula

Proceeding to the comparison of the FRG copula with its seasonal version, we present in Fig. 14 results that are similar to those in Fig. 13. Regarding the optimal hedge quantities in Fig. 14(a), the FRG copula yields higher values than the seasonal FRG copula during spring and summer, while the situation reverses during autumn and winter. This alternating behavior is connected to that of the differences in Spearman's ρ implied by the FRG and the seasonal FRG, cf. Fig. 9(a). That is, the hedge quantities decrease (increase) with a decrease (increase) in absolute values of Spearman's ρ .

Considering Fig. 14(b), the results reveal a fairly constant level in the variance reduction produced by the FRG copula, compared to the variance reduction levels implied by the seasonal model. Hence, believing in static dependence can lead to very misleading conclusions when managing risks. Again in this context, we mention the link between the difference in percentage reductions and the difference between the Spearman's ρ implied by the two models (cf. Fig. 9(a)): The difference in reductions is largest in May/June and November/December, reflecting the fact that the dependence implied by the seasonal FRG model is weakest during May/June and strongest during November/December.

Having found clear evidence of seasonal dependence between the CSS and the wind index, we conclude this section by briefly addressing the error we would get by applying a hedge based on the static FRG model in a seasonal time-varying reality. To perform this analysis, we assess the optimal hedge quantities implied by the static FRG copula in a seasonal setting by using the simulated CSS and wind index from the seasonal model. The variance reductions obtained with this approach are then compared with the reductions implied by the seasonal model shown in Fig. 14(b). The results, presented in Fig. 15, reveal very small errors. The smallest and largest errors occur in February and June, respectively, which is connected to the findings presented in Fig. 14(a); the absolute difference in the optimal number of WPF in the static and seasonal case generates the pattern seen in Fig. 15.

With the small errors in mind, the real error one commits by believing in static dependence, is the belief in a wrong resulting variance reduction. Thus, while the static model creates a misleading picture in a risk management context, our results suggest that it could be employed to determine optimal hedging quantities.

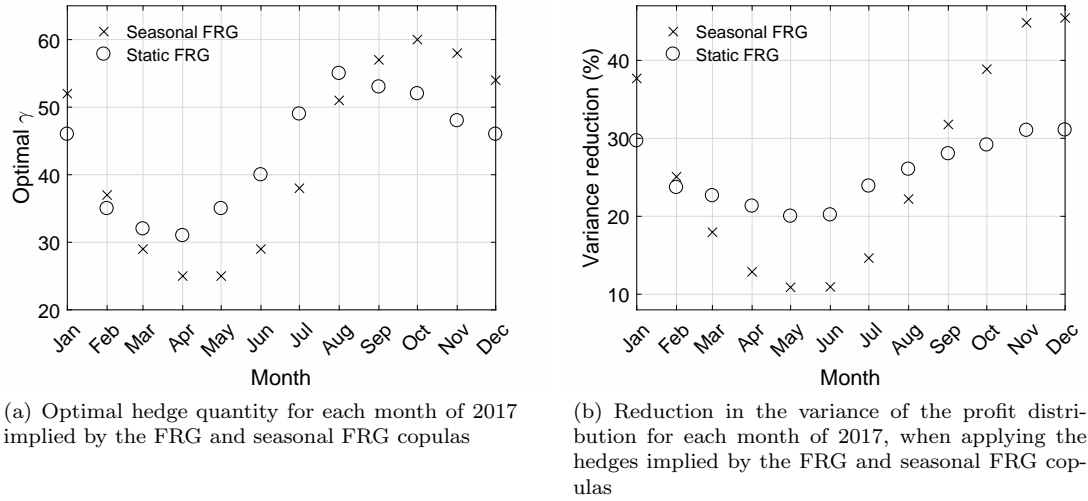


Figure 14: Comparison of hedging results implied by the FRG and seasonal FRG copulas.




Figure 15: The effects of using the “wrong” hedge quantity: Difference in variance reduction of the profit distribution when using hedge quantities obtained from 1) the FRG, and 2) the seasonal FRG, both evaluated using simulations from the seasonal FRG copula model.

5.3. Discussion of the proposed hedging strategy

To conclude this section, we turn briefly to the standard hedging principle often employed to hedge the day-ahead CSS. Usually, conventional generators remove their exposure to day-ahead price risk either completely or partially by entering a short position in standard power forwards, and a long position in fuel forwards and carbon credit forwards. In this paper, we have considered a different hedging approach with the purpose of determining the potential of the newly introduced wind power futures, but we stress that our strategy is not incompatible with the industry standard. In fact, our hedging portfolio consisting of wind power futures could be extended to include the additional forwards mentioned above. However, this would require us to switch from our bivariate modeling problem to a multivariate one, since the joint behavior of the wind index, the day-ahead CSS, and the different forward clean spark spreads should be considered. While this is outside the scope of the present paper, it is nevertheless an interesting perspective that has not been studied yet, and could possibly be approached with vine copulas. In the context of the effectiveness of the standard hedging principle that conventional generators usually employ, we mention the study of Charalampous and Madlener (2016).

6. Conclusion

In this paper, we propose a joint model for the day-ahead clean spark spread and the daily wind index that can facilitate hedging decisions for a gas-fired power plant. The modeling procedure is based on two steps: First, the marginal behavior of the variables is considered, where we apply seasonal functions and ARMA–GARCH filters to remove the seasonality and the serial dependence in the conditional mean and variance. While the usual Gaussian assumption for the innovation process works in the case of the daily wind index, the normal-inverse Gaussian distribution provides a better fit for the day-ahead clean spark spread. Second, the standardized residuals from the ARMA–GARCH models are connected through copulas. The data reveals a dependence structure that is slightly asymmetric, and also varying according to an annual cycle. To capture these empirical findings, we propose a seasonal copula mixture, where the mixing components are the rotated Gumbel and the Frank copulas.

Based on Monte Carlo simulations from the proposed model, we show that wind power futures have considerable risk-reducing benefits in the context of a gas-fired power plant operating in the day-ahead market. Further, their hedging effectiveness increases as a function of the efficiency of the gas-fired power plant. To highlight the importance of capturing asymmetry, tail dependence, and seasonality in the dependence structure, we perform comparison studies where the optimal model is compared to less optimal alternatives. Accounting for asymmetry and tail dependence (as opposed to imposing symmetry and zero tail dependence) leads to an increase in the effectiveness of wind power futures. Moreover, we find that the conclusions drawn with a static dependence model deviate to a large extent from those obtained with a seasonal dependence model. With static dependence, the variance reductions of the profit distributions attained by the hedge vary between 20% and 31%; in the seasonal case the corresponding reductions vary between 10% and 45%.

Although we have concentrated on the German market and the case of gas-fired power plants, the results are relevant for other markets, and are also transferable to other conventional electricity producers. Since the amount of electricity generated by wind turbines is expected to grow globally, the dependence between the day-ahead clean spark spread and the daily wind index in other market places will most likely be strengthened in the future. Hence, it is reasonable to assume that more weather-based instruments similar to the German wind power futures will be introduced, enabling similar analyses to be performed on other than the German market.

Acknowledgements

The authors would like to thank the Quantitative Analytics team at Neas Energy and the participants attending the Second Conference on the Mathematics of Energy Markets 2017 at the Wolfgang Pauli Institute for providing valuable comments and suggestions. Three anonymous referees are also thanked for their constructive criticism and suggestions, which improved the presentation of this paper.

Funding

Troels Sønderby Christensen is supported by the Innovation Fund Denmark under Grant 5189-00117B. Anca Pircalabu is supported by the Innovation Fund Denmark under Grant 4135-00082B.

Appendix A. Theoretical results for copula mixtures

Appendix A.1. Proof of Proposition 1

For positive dependent variables, the upper tail dependence for copula C can be written as (see e.g. (McNeil *et al.*, 2005, p. 209))

$$\Lambda^u = \lim_{u \uparrow 1} \frac{\hat{C}(1-u, 1-u)}{1-u} = \lim_{u \uparrow 1} \frac{1-2u+C(u, u)}{1-u}, \quad (\text{A.1})$$

where $\hat{C}(u, v) := P(U > u, V > v) = C(1-u, 1-v) + u + v - 1$ is the survival copula. Applying Eq. (A.1) to the copula mixture, we get that the upper tail dependence is

$$\begin{aligned} \Lambda^u &= \lim_{u \uparrow 1} \left[\frac{1-2u + \lambda C^a(u, u) + (1-\lambda)C^b(u, u)}{1-u} \right] \\ &= \lim_{u \uparrow 1} \left[\frac{1-2u}{1-u} + \lambda \frac{1-2u+C^a(u, u)}{1-u} + (1-\lambda) \frac{1-2u+C^b(u, u)}{1-u} - \lambda \frac{1-2u}{1-u} - (1-\lambda) \frac{1-2u}{1-u} \right] \\ &= \lambda \Lambda^{u,a} + (1-\lambda) \Lambda^{u,b}. \end{aligned}$$

Similarly, the lower tail dependence can be written in terms of copula C as (again for positive dependent variables)

$$\Lambda^l = \lim_{u \downarrow 0} \frac{C(u, u)}{u},$$

resulting in the following lower tail dependence for the copula mixture:

$$\begin{aligned} \Lambda^l &= \lim_{u \downarrow 0} \left[\frac{\lambda C^a(u, u) + (1-\lambda)C^b(u, u)}{u} \right] \\ &= \lambda \Lambda^{l,a} + (1-\lambda) \Lambda^{l,b}. \end{aligned}$$

Appendix A.2. Proof of Proposition 2

From Eq. (8) we have that Spearman's ρ implied by copula C^m is

$$\begin{aligned} \rho^m &= 12 \int_0^1 \int_0^1 C^m(u_1, u_2) du_1 du_2 - 3 \\ &= 12 \int_0^1 \int_0^1 (\lambda C^a(u_1, u_2) + (1-\lambda)C^b(u_1, u_2)) du_1 du_2 - 3(1-\lambda + \lambda) \\ &= \lambda \left(12 \int_0^1 \int_0^1 C^a(u_1, u_2) du_1 du_2 - 3 \right) + (1-\lambda) \left(12 \int_0^1 \int_0^1 C^b(u_1, u_2) du_1 du_2 - 3 \right) \\ &= \lambda \rho^a + (1-\lambda) \rho^b. \end{aligned}$$

Appendix B. Alternative models

As mentioned in Section 4.1, there are many ways of extending the FRG copula. Here, we present some alternatives to the model proposed in Eqs. (10)–(11), and the corresponding estimation results.

Appendix B.1. Seasonal copula alternatives

Instead of introducing a yearly cycle in the evolution equation for Spearman's ρ corresponding to the RGumbel copula cf. Eq. (10), one could consider this for Frank, i.e.,

$$\rho_t^F = a^F + b^F \sin(2\pi t/K) + c^F \cos(2\pi t/K). \quad (\text{B.1})$$

Alternatively, yearly seasonality could be introduced in both Spearman's ρ implied by the Frank and the RGumbel copula. These alternative models are stated in Table B.7, and corresponding estimation results are reported in Table B.8.

Seasonal FRG	Alternative 1
$\rho_t^{SFRG} = \lambda^{SFRG} \rho_t^F + (1 - \lambda^{SFRG}) \rho_t^{RG}$	$\rho_t^{SFRG} = \lambda^{SFRG} \rho_t^F + (1 - \lambda^{SFRG}) \rho_t^{RG}$
$\rho_t^{RG} = a^{RG} + b^{RG} \sin\left(\frac{2\pi t}{K}\right) + c^{RG} \cos\left(\frac{2\pi t}{K}\right)$	$\rho_t^F = a^F + b^F \sin\left(\frac{2\pi t}{K}\right) + c^F \cos\left(\frac{2\pi t}{K}\right)$
Alternative 2	
$\rho_t^{SFRG} = \lambda^{SFRG} \rho_t^F + (1 - \lambda^{SFRG}) \rho_t^{RG}$	
$\rho_t^{RG} = a^{RG} + b^{RG} \sin\left(\frac{2\pi t}{K}\right) + c^{RG} \cos\left(\frac{2\pi t}{K}\right)$	
$\rho_t^F = a^F + b^F \sin\left(\frac{2\pi t}{K}\right) + c^F \cos\left(\frac{2\pi t}{K}\right)$	

Table B.7: Model specifications for the seasonal FRG proposed in Section 4.1 and two seasonal alternatives.

	\hat{a}^F	\hat{b}^F	\hat{c}^F	\hat{a}^{RG}	\hat{b}^{RG}	\hat{c}^{RG}	λ^{SFRG}	$\log \mathcal{L}_c$	AIC
Seasonal FRG	0.742 (0.072)	- (-)	- (-)	0.566 (0.041)	-0.132 (0.040)	0.208 (0.043)	0.316 (0.092)	309.894	-609.788
Alternative 1	0.573 (0.048)	-0.122 (0.048)	0.280 (0.056)	0.689 (0.054)	- (-)	- (-)	0.532 (0.097)	309.657	-609.314
Alternative 2	0.693 (0.080)	0.050 (0.120)	0.156 (0.099)	0.580 (0.055)	-0.162 (0.081)	0.150 (0.071)	0.374 (0.109)	312.454	-610.908

Table B.8: Maximum likelihood estimation results for the models stated in Table B.7. The maximized value of the copula log-likelihood is denoted $\log \mathcal{L}_c$. Standard errors are reported in parenthesis and are computed following the simulation-based procedure in Patton (2013).

The results in Table B.8 reveal that the seasonal FRG model proposed in Section 4.1 is slightly better than Alternative 1 in terms of AIC. Considering Alternative 2, although its AIC is slightly lower than that of the other seasonal models, notice the large standard errors, which imply that the seasonal parameters corresponding to the Frank part are not statistically significant at a 5% level.

Appendix B.2. Dynamic copula alternatives

In terms of dynamic copula alternatives, we consider the Generalized Autoregressive Score (GAS) model proposed by Creal *et al.* (2013). Assuming a copula having one governing parameter which we denote by θ , the evolution equation for the GAS(1,1) is given by

$$z_{t+1} = \omega + \Psi z_t + \varsigma I_t^{-1/2} s_t, \quad (\text{B.2})$$

where

$$\begin{aligned} z_t &= f(\theta_t) \\ s_t &= \frac{\partial}{\partial \theta_t} \log c(u_{1,t}, u_{2,t} | \theta_t) \\ I_t &= \mathbb{E}[s_t^2], \end{aligned}$$

and ω , Ψ , and ς are constant coefficients.

Recalling Eq. (5) and restricting this analysis to the FRG copula, we let the (transformed) parameter of the RGumbel copula evolve according to the GAS(1,1) model. Specifically,

$$c_t^{FRG}(\cdot|\mathcal{F}_{t-1}, \theta^F, \theta_t^{RG}, \lambda^{FRG}) = \lambda^{FRG} c^F(\cdot|\mathcal{F}_{t-1}, \theta^F) + (1 - \lambda^{FRG}) c_t^{RG}(\cdot|\mathcal{F}_{t-1}, \theta_t^{RG}),$$

where $z_t = \log(\theta_t^{RG} - 1)$ follows a GAS(1,1). Aside from this, we consider two other alternatives which are stated in Table B.9; one where we only allow the transformed RGumbel copula parameter to vary according to a yearly cycle, and one where we combine the GAS dynamics with the yearly seasonality.

Alternative 3 (GAS)	Alternative 4 (Season)
$z_{t+1} = \omega^{RG} + \Psi^{RG} z_t + \varsigma^{RG} I_t^{-1/2} s_t$	$z_{t+1} = \omega^{RG} + b^{RG} \sin\left(\frac{2\pi t}{K}\right) + c^{RG} \cos\left(\frac{2\pi t}{K}\right)$
Alternative 5 (Season GAS)	
$z_{t+1} = \omega^{RG} + \Psi^{RG} z_t + \varsigma^{RG} I_t^{-1/2} s_t + b^{RG} \sin\left(\frac{2\pi t}{K}\right) + c^{RG} \cos\left(\frac{2\pi t}{K}\right)$	

Table B.9: Model specifications for three alternative models.

Parameter estimates for the models stated in Table B.9 are obtained by maximum likelihood, and are reported in Table B.10. Comparing Alternative 3 to Alternative 4, AIC decreases quite a bit when considering a deterministic yearly cycle instead of the GAS(1,1) specification. Further, combining the two effects (Alternative 5) yields the lowest AIC value across all models; notice however that two out of seven parameters are not significant at a 5% level.

Regarding Alternative 4, we stress that this model is very similar to the seasonal FRG: in Alternative 4, we let the transformed copula parameter vary according to a yearly cycle, whereas in the seasonal FRG, we introduce yearly seasonality in Spearman's ρ . Although the AIC corresponding to the former specification is slightly lower, we find the seasonal FRG to be more appealing, given that the yearly cycle is detected in the rolling Spearman's ρ cf. Fig. 6.

	$\hat{\theta}^F$	$\hat{\omega}^{RG}$	$\hat{\Psi}^{RG}$	$\hat{\varsigma}^{RG}$	\hat{b}^{RG}	\hat{c}^{RG}	$\hat{\lambda}^{FRG}$	$\log \mathcal{L}_c$	AIC
Alternative 3 (GAS)	6.132 (2.393)	0.001 (0.648)	0.993 (0.378)	0.092 (0.314)	- (-)	- (-)	0.315 (0.262)	300.349	-590.697
Alternative 4 (Season)	6.605 (1.776)	-0.391 (0.148)	- (-)	- (-)	-0.451 (0.144)	0.758 (0.152)	0.312 (0.090)	310.376	-610.752
Alternative 5 (Season GAS)	6.725 (2.831)	-0.421 (0.379)	-0.231 (0.567)	0.235 (0.107)	-0.503 (0.208)	0.842 (0.303)	0.270 (0.121)	315.832	-617.663

Table B.10: Maximum likelihood estimation results for the models stated in Table B.9. The maximized value of the copula log-likelihood is denoted $\log \mathcal{L}_c$. Standard errors are reported in parenthesis and are computed following the simulation-based procedure in Patton (2013).

References

- Aepli, M. D., Füss, R., Henriksen, T. E. S. and Paraschiv, F. (2017). Modeling the multivariate dynamic dependence structure of commodity futures portfolios. *Journal of Commodity Markets*, **6**, 66–87.
- Aydulaj, K. and Barunik, J. (2015). Are benefits from oil - stocks diversification gone? New evidence from a dynamic copula and high frequency data. *Energy Economics*, **51**, 31–44.
- Barndorff-Nielsen, O. E. (1997a). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics*, **24**(1), 1–13.
- Barndorff-Nielsen, O. E. (1997b). Processes of normal inverse Gaussian type. *Finance and Stochastics*, **2**(1), 41–68.
- Benth, F. E. and Šaltytė Benth, J. (2004). The normal inverse Gaussian distribution and spot price modelling in energy markets. *International Journal of Theoretical and Applied Finance*, **07**(02), 177.

- Benth, F. E. and Šaltytė Benth, J. (2011). Weather derivatives and stochastic modelling of temperature. *International Journal of Stochastic Analysis*, Vol. 2011, Article ID 576791, 21 pages.
- Benth, F.E. and Kettler, P. C. (2011). Dynamic copula models for the spark spread. *Quantitative Finance*, **3**, 407–421.
- Benth, F. E. and Pircalabu, A. (2018). A non-Gaussian Ornstein-Uhlenbeck model for pricing wind power futures. *Applied Mathematical Finance*, **25**(1), 36–65.
- Berg, D. (2009). Copula goodness-of-fit testing: An overview and power comparison. *European Journal of Finance*, **15**(7–8), 675–701.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, **31**(3), 307–327.
- Börger, R., Cartea, A., Kiesel, R. and Schindlmayr, G. (2009). Cross-commodity analysis and applications to risk management. *Journal of Futures Markets*, **29**(3), 197–217.
- Charalampous, G. and Madlener, R. (2016). Risk management and portfolio optimization for gas- and coal-fired power plants in Germany: A multivariate GARCH approach. *Journal of Energy Markets*, **9**(2), 69–94.
- Creal, D., Koopman, S.J., and Lucas, A. (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics*, **28**, 777–795.
- Dias, A. and Embrechts, P. (2009). Testing for structural changes in exchange rates dependence beyond linear correlation. *The European Journal of Finance*, **15**(7–8), 619–637.
- Elberg, C. and Hagspiel, S. (2015). Spatial dependencies of wind power and interrelations with spot price dynamics. *European Journal of Operational Research*, **241**, 260–272.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, **50**(4), 987–1007.
- European Commission, The EU Emissions Trading System, 2017. https://ec.europa.eu/clima/policies/ets_en (accessed 29 August 2017).
- Genest, C., Remillard, B., and Beaudoin, D. (2009). Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, **44**(2), 199–213.
- Gersema, G. and Wozabal, D. (2017). An equilibrium pricing model for wind power futures. *Energy Economics*, **65**, 64–74.
- Grothe, O. and Schneiders, J. (2011). Spatial dependence in wind and optimal wind power allocation: A copula-based analysis. *Energy Policy*, **39**, 4742–4754.
- Haldrup, N. and Nielsen, M. Ø., 2006. A regime switching long memory model for electricity prices. *Journal of Econometrics*, **135**(1–2), 349–376.
- Härdle, W.K. and López Cabrera, B. (2012). Implied market price of weather risk. *Applied Mathematical Finance*, **19**(1), 59–95.
- ICIS (2016). European Daily Electricity Markets Methodology.
- Jensen, M. B. and Lunde, A. (2001). The NIG-S&ARCH model: A fat-tailed, stochastic, and autoregressive conditional heteroskedastic volatility model. *The Econometrics Journal*, **4**(2), 319–342.
- Joe, H. (2005). Asymptotic efficiency of the two-stage estimation method for copula-based models. *Journal of Multivariate Analysis*, **94**(2), 401–419.
- McNeil, A. J., Frey, R. and Embrechts, P. (2005). Quantitative risk management. *Princeton Series in Finance*.
- Nelsen, R.B. (2006). *An Introduction to Copulas*, Springer.
- Patton, A.J. (2006(a)). Modelling asymmetric exchange rate dependence. *International Economic Review*, **47**(2), 527–556.
- Patton, A.J. (2006(b)). Estimation of multivariate models for time series of possibly different lengths. *Journal of Applied Econometrics*, **21**(2), 147–173.
- Patton, A.J. (2013). Copula methods for forecasting multivariate time series. In Graham Elliott and Allan Timmermann, editors, *Handbook of Economic Forecasting*, volume 2B, pages 899–960. Elsevier B.V.
- Pircalabu, A. and Jung, J. (2017). A mixed C-vine copula model for hedging price and volumetric risk in wind power trading. *Quantitative Finance*, **17**(10), 1583–1600.
- Pircalabu, A., Hvolby, T., Jung, J. and Høg, E. (2017). Joint price and volumetric risk in wind power trading: A copula approach. *Energy Economics*, **62**, 139–154.
- Rodriguez, J. C. (2007). Measuring financial contagion: A copula approach. *Journal of Empirical Finance*, **14**(3), 401–423.
- Rydberg, T. H. (1997). The normal inverse Gaussian Lévy process: Simulation and approximation. *Communications in Statistics. Stochastic Models*, **13**(4), 887–910.
- Shumway, R. H. and Stoffer, D. S. (2011). Time Series Analysis and Its Applications. *Springer Texts in Statistics*.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de L'Université de Paris*, **8**, 229–231.

Highlights

- The hedging potential of wind power futures for gas-fired power plants is studied.
- We consider the joint behavior of the clean spark spread and wind index.
- A seasonal copula mixture model is proposed to facilitate hedging decisions.
- We find that wind power futures are suitable for hedging the clean spark spread.
- Comparison studies are performed to highlight the benefits of the proposed model.