

Rough electricity: a new fractal multi-factor model of electricity spot prices

Mikkel Bennedsen

CREATES Research Paper 2015-42

Rough electricity: a new fractal multi-factor model of electricity spot prices

Mikkel Bennedsen*

September 18, 2015

Abstract

We introduce a new mathematical model of electricity spot prices which accounts for the most important stylized facts of these time series: seasonality, spikes, stochastic volatility and mean reversion. Empirical studies have found a possible fifth stylized fact, fractality, and our approach explicitly incorporates this into the model of the prices. Our setup generalizes the popular Ornstein Uhlenbeck-based multi-factor framework of [Benth et al. \(2007\)](#) and allows us to perform statistical tests to distinguish between an Ornstein Uhlenbeck-based model and a fractal model. Further, through the multi-factor approach we account for seasonality and spikes before estimating – and making inference on – the degree of fractality. This is novel in the literature and we present simulation evidence showing that these precautions are crucial to accurate estimation. Lastly, we estimate our model on recent data from six European energy exchanges and we find statistical evidence of fractality in five out of six markets. As an application of our model, we show how, in these five markets, a fractal component improves short term forecasting of the prices.

Keywords: Energy markets; electricity prices; roughness; fractals; mean reversion; multi-factor modelling; forecasting.

JEL Classification: C22, C51, C52, C53, Q41

1 Introduction

This paper proposes a new mathematical model of electricity spot prices which parsimoniously accounts for the most important statistical properties of such prices. These *stylized facts* are empirical features consistently observed in the time series of electricity prices and at least four stylized facts are by now well-documented in the literature: (i) seasonality which is attributable to month-to-month variation in prices, mainly due to the seasons, and day-to-day variations, mainly due to weekly dependence of demand; (ii) spikes which are large up- or downwards movements followed by rapid reversion to previous levels often attributed to the non-storable nature of electricity combined with the fact that the energy grid needs to be in complete balance at all times. Because

*Department of Economics and Business Economics and CREATES, Aarhus University, Fuglesangs Allé 4, 8210 Aarhus V, Denmark. E-mail: mbennedsen@econ.au.dk

of very inelastic demand the price can surge (or drop) suddenly if the supply is limited (or in surplus) e.g. due to plant outages or unforeseen weather conditions. (iii) Extreme variability of prices as compared to traditional financial time series such as storable commodities or stocks; and (iv) mean-reversion of prices which has been found by several authors such as Weron and Przybylowicz (2000), Kaminski (2004) and Pilipovic (2007). The wish to incorporate these stylized facts into a mathematical model of electricity spot prices has motivated a large literature starting with the seminal work of Schwartz (1997) who introduced a model of commodity prices based on the Ornstein-Uhlenbeck (OU) process. Since Schwartz (1997), this model has been extended in several ways and applied to electricity price series, e.g. the jump-diffusion model of Cartea and Figueroa (2005) and the threshold model of Geman and Roncoroni (2006). Benth et al. (2007) generalized the Ornstein-Uhlenbeck framework to the so-called *multi-factor model*, where the (log-)price of electricity is modelled as

$$S_t = \Lambda_t + X_t + Y_t, \quad t \geq 0,$$

where Λ is a seasonal term, X is the *base signal* accounting for smaller day-to-day variations in the price, and Y is a *spike signal* accounting for large up- or down-movements and the ensuing rapid reversion to previous levels. A typical approach is to model Λ by a deterministic sinusoidal function, X as a Gaussian OU process and Y as a non-Gaussian (i.e. Lévy driven) OU process (e.g. Meyer-Brandis and Tankov, 2008; Benth et al., 2012). This model has been particularly useful as it elegantly accounts for the stylized facts (i)-(iv).

Meanwhile, a possible fifth stylized fact has been found in a different strand of literature, namely that the time series of electricity spot prices exhibit (v) fractal characteristics (often termed *antipersistence* in the literature). Specifically, a number of studies show that the *fractal index*¹ of electricity spot price time series is negative, meaning that the paths of the spot prices are very *rough* as compared to a model based on the Brownian motion such as the Brownian-driven OU process. This was for instance found for the Nordic market, Nord Pool, in Simonsen (2002) and Erzgräber et al. (2008), for the Spanish market in Norouzzadeh et al. (2007) and for the Czech market in Kristoufek and Lunackova (2013). Although fractality of the price series on several energy markets is by now well established, very few authors have utilized their findings to actually employ and estimate a fractal model for the spot prices. A notable exception is Rypdal and Løvsletten (2013) who compared a fractal model based on a multifractal random walk with a (non-fractal) model of Ornstein-Uhlenbeck type and concluded that the latter best described the Nord Pool spot prices from May 1992 to August 2011.

The main contribution of the present paper is to add to the above literature and merge the multi-factor approach with the fractal approach. In short, we propose a fractal multi-factor model of electricity spot prices, i.e. a multi-factor model where the base signal, X , is modelled by a fractal process. This will allow us to account for *all* the stylized facts (i)-(v), contrary to the above mentioned previous studies employing either a multi-factor model (accounting for stylized

¹The fractal index $\alpha \in (-\frac{1}{2}, \frac{1}{2})$ of a time series is a measure of the 'roughness' of the price curve. Negative values of α refer to curves more rough than those arising from the paths of a Brownian motion (which has $\alpha = 0$) while positive values of α refer to smooth paths, see also Section 3.

facts (i)-(iv)) or a fractal model (accounting for stylized fact (v) only). What is more, our setup accomodates *testing*, in a rigorous statistical sense, for the presence of fractality in the price series; this is a novel contribution in the literature, as previous studies have focused on point estimates of the fractal index in lieu of confidence intervals and hypothesis testing. In fact, as we shall see, our approach provides a direct statistical hypothesis test of whether the base signal is best described by a Brownian-driven OU process (the null) or of a fractal process (the alternative). Additionally, the multi-factor setup allows us to account for spikes *before* estimating and making inference on the fractal index; this is crucial because spikes can be detrimental to the estimation as shown in the Appendix.

As an application of the model we consider a simple forecasting exercise, showing that including a fractal component provides superior short-run forecasts of the base signal, as compared to three benchmark models. The exercise indicates that agents in the energy markets might improve their forecasts of prices by including a fractal component in their models.

The rest of the paper is structured as follows. In Section 2, we take a closer look at the six data sets we consider. Section 3 briefly explains the main ideas behind roughness and fractality as they relate to time series of electricity prices. Section 4 introduces the modelling framework we propose; this section also contains an estimation procedure for the model. Section 5 is concerned with the implications of our model as it relates to arbitrage and to forecasting. In Section 6, we present our empirical work where we find evidence of fractality in five out of the six markets considered. Section 7 discusses the findings and concludes. The Appendix contains an in-depth look at the estimation procedure by detailing the individual steps using the German EEX data set. Using simulations, the Appendix also briefly investigates possible sources of biases when estimating the fractal index of electricity prices and illustrates how the multi-factor framework helps in alleviating these.

2 The data

Electricity markets are organized as day-ahead markets where participants submit their bid/ask price for a given amount of electricity in a given hour of the day. The exchange then matches supply and demand, which yields 24 spot prices, one for each hour, for the following day.² In the following we define the 'spot price' as the *peak load price* which is computed as the mean of the spot prices from 9 a.m. to 8 p.m., the hours where load (demand) is highest. We consider data from six different European exchanges, as shown in Table 1 and plotted in Figure 1. All data series are excluding weekends as these tend not to add a lot of information not contained in the Friday price (Meyer-Brandis and Tankov, 2008). It is well known that electricity prices are subject to strong seasonal effects and before working with the data we therefore pre-process it by de-seasonalizing as explained in Section 4.1 below. Some empirical characteristics of these de-seasonalized prices and their increments are seen in Table 2. Looking at the table we see that the (log as well as raw) increments exhibit significant departures from normality. Indeed, the series display some skewness

²Unlike the five other markets, which records spot prices every hour, the UK market which we consider in this paper has prices for every 30 minutes, i.e. 48 prices per day.

Table 1: *Data summary*

Abbr.	Region	Start	End	Observations
EEX	Germany	Jan. 1, 2008	Jan. 23, 2015	1844
NP	Nordic	Jan. 1, 2008	Jan. 23, 2015	1844
PN	France	Jan. 1, 2008	Jan. 23, 2015	1844
GME	Italy	Jan. 1, 2008	Dec. 19, 2014	1819
UKAPX	United Kingdom	Jan. 1, 2008	Jan. 23, 2015	1844
NLAPX	Netherlands	Jan. 1, 2008	Jan. 23, 2015	1844

and a large amount of (excess) kurtosis ranging from 9.5 to 650 in the raw increment series and 6.6 to 116 in the log-increments. Naturally then, when conducting a formal test for normality (Jarque-Bera), we reject the null of Gaussianity of the increments in all time series considered.

The question of whether the *levels* of electricity prices constitute a stationary time series has been much debated in the literature. In Table 3 we therefore run two tests for a unit root in the data; the Phillips-Perron (PP) test and the augmented Dickey-Fuller (ADF) test with automatic lag selection³. Both tests reject the null of a unit root in the raw prices series as well as the log-prices at all reasonable significance levels. That is, we find evidence of stationarity in the de-seasonalized prices and we therefore proceed to model all price series as stationary in the following.

³As implemented in the MFE Toolbox of Kevin Sheppard, http://www.kevinsheppard.com/MFE_Toolbox.

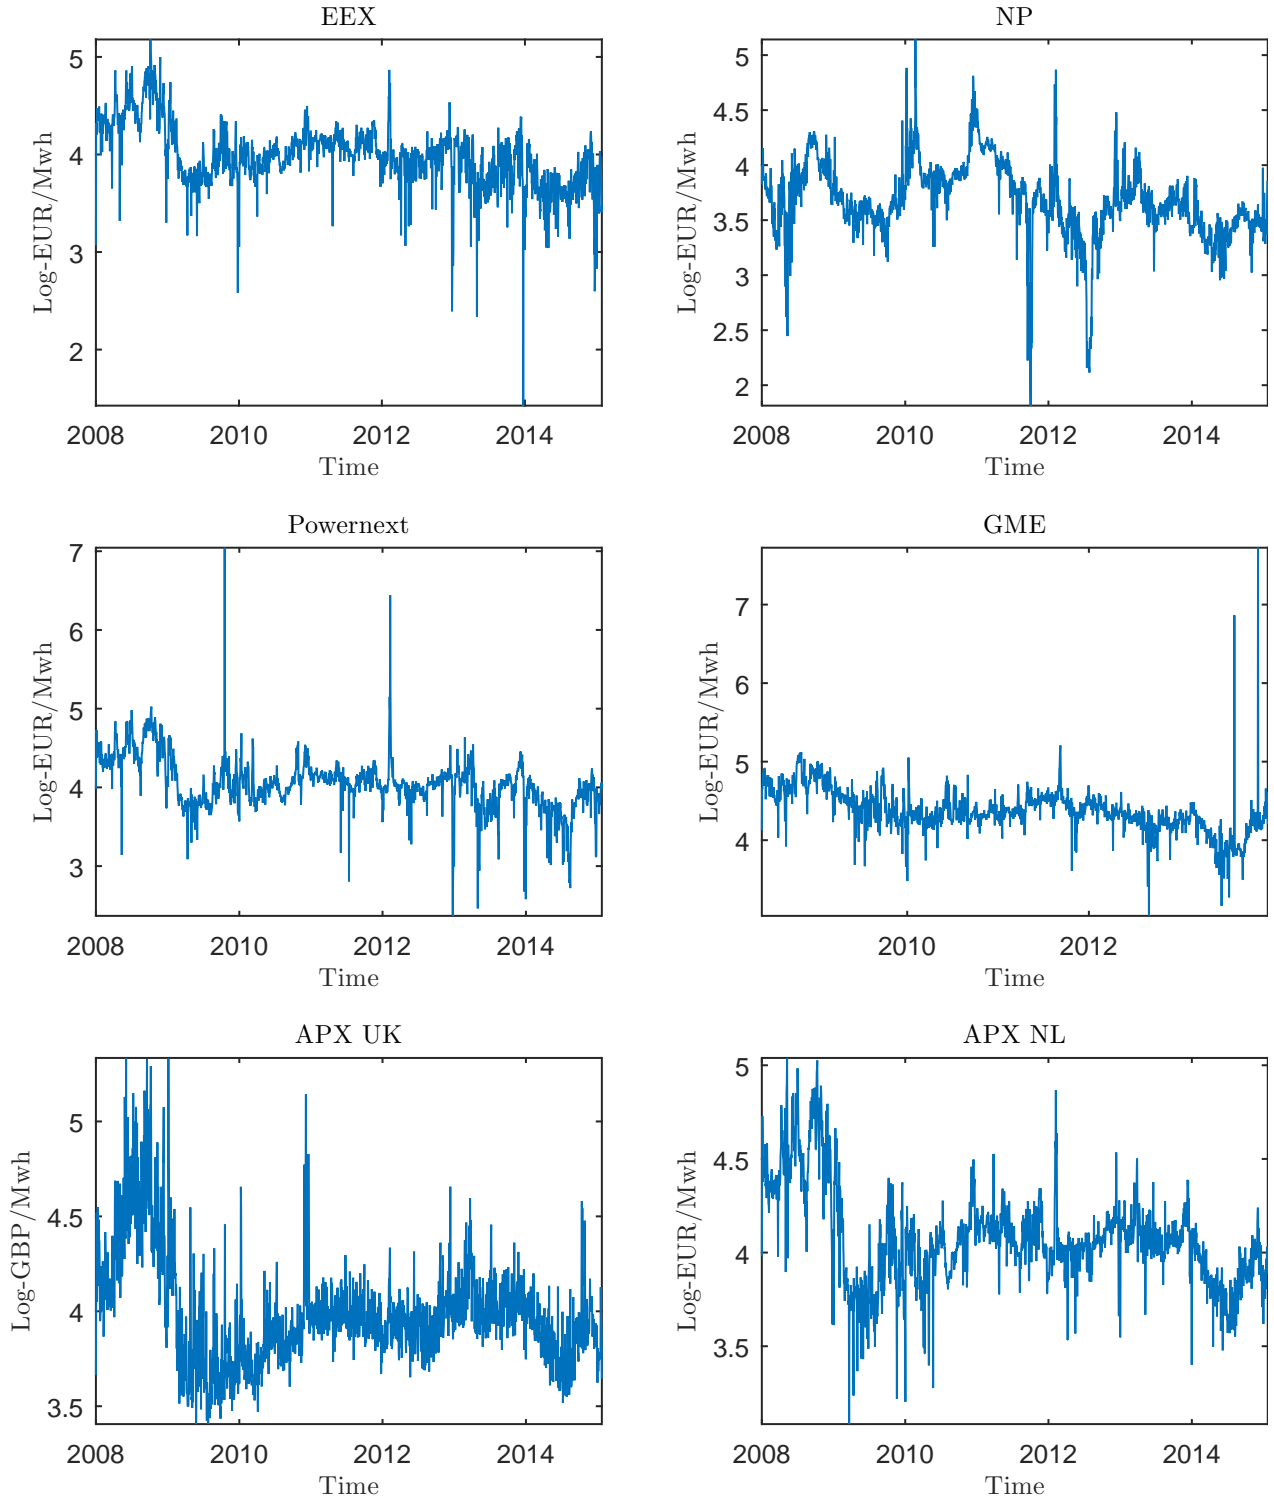


Figure 1: *Illustration of the six data sets: log-price evolution in log-Euros per Mwh (UK APX is in log-British Pounds).*

Table 2: *Descriptive statistics of electricity price increments*

	Panel A: <i>Raw increments</i>						Panel B: <i>Log-increments</i>					
	EEX	NP	Powernext	GME	APX UK	APX NL	EEX	NP	Powernext	GME	APX UK	APX NL
Observations	1848	1844	1844	1819	1844	1844	1848	1844	1844	1819	1844	1844
Mean	0.028	0.011	0.023	0.024	0.013	0.015	0.001	0.000	0.000	0.000	0.000	0.000
Min	-61.273	-60.748	-1064.780	-2195.649	-115.513	-50.791	-2.191	-0.573	-2.669	-3.518	-0.860	-0.708
Max	99.152	97.258	1080.133	2182.405	150.305	41.812	2.053	0.884	2.830	3.345	1.292	0.917
SD	9.710	5.790	39.397	79.447	14.430	7.830	0.203	0.099	0.194	0.215	0.184	0.127
Skewness	0.558	2.569	0.725	-0.193	0.218	0.018	0.052	0.747	0.349	0.060	0.101	0.421
Kurtosis	12.549	78.412	616.925	653.490	21.017	9.474	23.412	15.590	57.139	115.594	6.610	10.128
Jarque-Bera	7114	438737	28943260	32052661	24942	3219	32066	12343	225113	960318	1004	3956
P-value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001

Descriptive statistics for the increments of the de-seasonalized (see Section 4.1) electricity prices for the various power exchanges. Weekends are excluded. Panel A (left) is for the raw increments, while Panel B (right) is for the log-increments (log-returns).

Table 3: *Testing for unit roots in electricity price series*

	Panel A: <i>Raw prices</i>						Panel B: <i>Log-prices</i>					
	EEX	NP	Powernext	GME	APX UK	APX NL	EEX	NP	Powernext	GME	APX UK	APX NL
PP	-14.486	-9.594	-32.437	-40.137	-17.399	-11.086	-18.391	-7.448	-16.002	-23.681	-15.453	-12.147
P-value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
ADF	-4.136	-4.481	-10.276	-19.057	-2.696	-3.365	-5.832	-4.830	-6.206	-3.990	-3.046	-3.320
P-value	0.000	0.000	0.000	0.000	0.007	0.000	0.000	0.000	0.000	0.000	0.003	0.000

Unit root tests of the de-seasonalized (see Section 4.1) electricity prices for the various power exchanges. Weekends are excluded. Panel A (left) is for the raw prices, while Panel B (right) is for the log-prices. PP stands for the Phillips-Perron test for a unit root, while ADF is the augmented Dickey-Fuller test with automatic lag selection.

3 Fractality and roughness in electricity prices

In this section, we explain how fractals and roughness relate to one-dimensional time series and hence to electricity prices. In particular, we consider a time series $X_t \in \mathbb{R}$ and its graph

$$X = \{(t, X_t) \in \mathbb{R} \times \mathbb{R}, t \in \mathcal{T} \subset \mathbb{R}\} \subset \mathbb{R}^2,$$

where \mathcal{T} is a finite set of times where the process X is observed. The fractal dimension of the time series refers to the one-dimensional curve that arises in the (continuum) limit as the data gets observed at an infinitesimally dense subset of \mathbb{R} (e.g. $[0, 1]$). This approach of observing a time series at increasingly finer scales, is often termed infill asymptotics in the literature. If the limit curve is smooth and differentiable, then its fractal dimension, D , equals the topological dimension, i.e. $D = 1$ in the one-dimensional time series case. Conversely, for a rough, non-differentiable curve the fractal dimension may exceed the topological dimension. For instance, consider a Gaussian process $X = \{X_t\}_{t \in \mathbb{R}}$ with stationary increments with variogram

$$\gamma_2(h) = \frac{1}{2} \mathbb{E}[|X_{t+h} - X_t|^2], \quad h \geq 0,$$

satisfying

$$\gamma_2(h) = c_2 |h|^{2\alpha+1} + o(|h|^{2\alpha+1+\beta}), \quad \text{as } h \downarrow 0, \quad (3.1)$$

where $\alpha \in (-\frac{1}{2}, \frac{1}{2})$, $\beta \geq 0$ and $c_2 > 0$.⁴ The graph of a sample path of X now has fractal dimension

$$D = \frac{3}{2} - \alpha,$$

almost surely. We will refer to α as the *fractal index* of the process X : positive values of α – implying a fractal dimension close to the topological dimension, i.e. close to dimension 1 – corresponds to a smooth process (with $\alpha = \frac{1}{2}$ corresponding to a.s. differentiable sample path of X), while negative values of α – implying a fractal dimension close to one plus the topological dimension, i.e. close to dimension 2 – corresponds to very rough paths. A typical example of a Gaussian process with non-trivial fractal dimension is the fractional Brownian motion (fBm) with Hurst parameter $H \in (0, 1]$, corresponding to a fractal index $\alpha = H - \frac{1}{2}$ and fractal dimension $D = 2 - H$. Here $\alpha = 0$ (so $H = \frac{1}{2}$) implies that the fBm is a standard Brownian motion. Several simulated fBm paths for different values of α are shown in Figure 2. The differing roughness of the paths, when the fractal parameter changes, is evident: we see how a low α corresponds to a high fractal dimension (rough paths) and vice versa. The rougher the curve, the more it fills of the two-dimensional space: as α decreases, the fractal dimension increases and the one-dimensional curve becomes "closer to 2-dimensional". Relating this discussion back to time series of electricity prices, other authors have often found $\alpha < 0$ for these data sets, i.e. that the paths of the spot prices are rougher than what a Brownian motion-based model such as the Brownian-driven OU model would suggest. The next section introduces a model for electricity prices which is compatible with this observation.

⁴Note the non-standard parametrization as compared to e.g. [Gneiting et al. \(2012\)](#); their definition has the fractal index in $(0, 2]$ and corresponds to $2\alpha + 1$ in our parametrization. We find our choice more convenient for the present purposes as $\alpha = 0$ plays an important role here: this is where the process has sample paths of the same fractal dimension as a Brownian motion and hence also the Brownian-driven OU process.

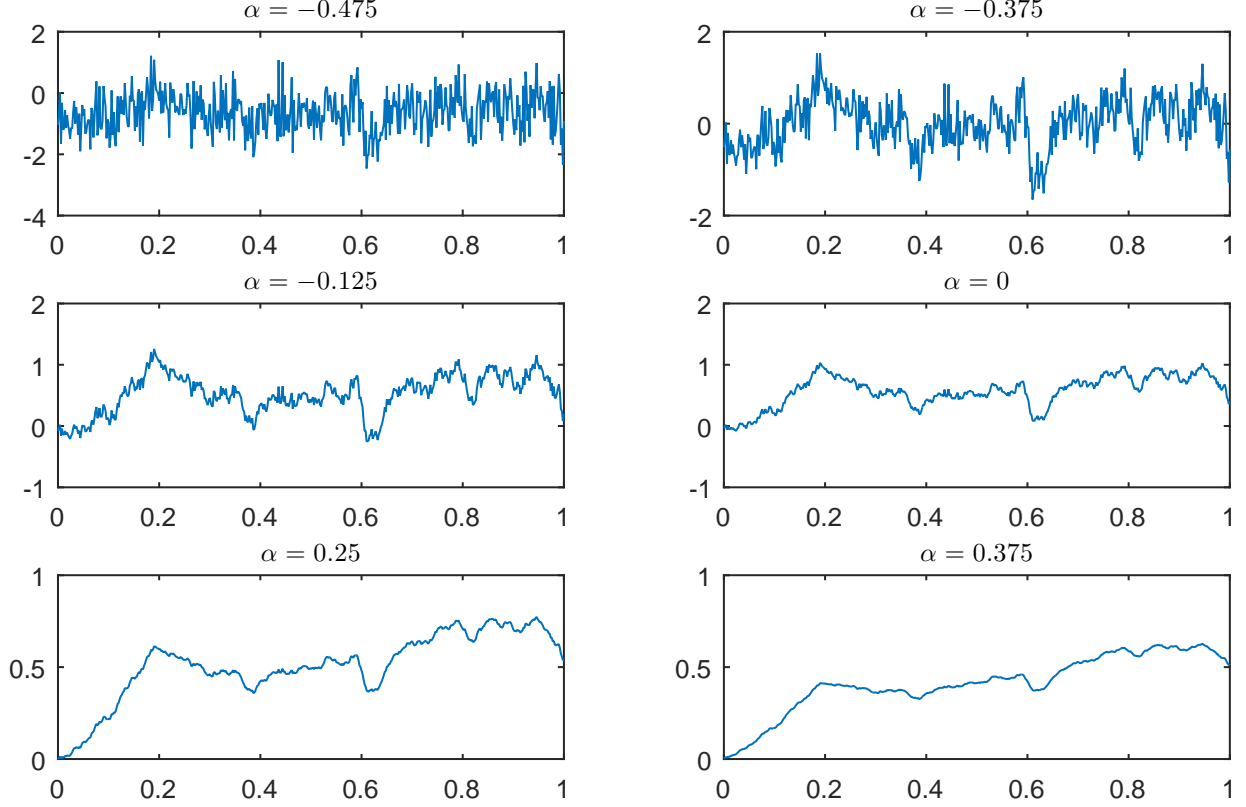


Figure 2: *Simulated paths of an fBm for varying values of the fractal index α . Recall that the Hurst index is $H = \alpha + \frac{1}{2}$, so that $\alpha = 0$ corresponds to the process being a standard Brownian motion. The same random numbers were used in the six simulation, hence the only difference is the roughness of the simulated paths as dictated by the value of α .*

4 A fractal multi-factor modelling framework

Consider the *multi-factor decomposition* of the electricity spot price

$$S_t = \Lambda_t + X_t + Y_t, \quad t \geq 0, \quad (4.1)$$

where $\Lambda = \{\Lambda_t\}_{t \geq 0}$, $X = \{X_t\}_{t \geq 0}$, and $Y = \{Y_t\}_{t \geq 0}$ are the seasonal component, the base signal and the spike signal respectively. This decomposition is the basis of the multi-factor model and has been extensively studied and successfully employed to electricity prices in several studies, such as Benth et al. (2007), Meyer-Brandis and Tankov (2008) and Hayfavi and Talasli (2014). The decomposition (4.1) can be seen as both an arithmetic and geometric model of the spot price by letting $S_t = P_t$ or $S_t = \log P_t$ accordingly, where $P = \{P_t\}_{t \geq 0}$ denotes the spot price of electricity. Note, that only the arithmetic model will allow for the possibility of negative spot prices, which has recently been observed in some markets. However, as we work with the peak load price, which is an average of the prices in 12 consecutive hours, and as negative prices do not tend to persist for many hours at the time, we do not observe any negative (or zero) prices in any of the time series under consideration. Hence, we proceed to work with the geometric model, $S_t = \log P_t$ in

the following. We want to note and stress that we also performed the analysis using the arithmetic model (i.e. using non-log prices), which did not yield qualitative differences and actually resulted in even rougher paths (i.e. a more negative fractal index α) than what we will see in the following. That is, the conclusions drawn in this paper concerning fractality of electricity prices hold a fortiori when considering an arithmetic model. As mentioned above, in the multi-factor framework the processes X and Y in (4.1) are usually modelled as (possibly non-Gaussian) OU processes but here we take a more general approach when modelling X , as explained in Section 4.2 below.

4.1 Modelling the seasonal component

Following Meyer-Brandis and Tankov (2008) we will take the seasonal component, Λ , to be a deterministic sinusoidal function plus an intercept and a trend. In full generality one could consider a more general stochastic process as a model for the seasonal component, which would, probably, be more realistic. However, since our focus is on the fractal nature of the prices, we will not pursue this here. We thus model the seasonal component as

$$\Lambda_t = c_1 + c_2 t + a_1 \sin\left(\frac{2\pi t}{5}\right) + a_2 \cos\left(\frac{2\pi t}{5}\right) + a_3 \sin\left(\frac{2\pi t}{260}\right) + a_4 \cos\left(\frac{2\pi t}{260}\right), \quad t \geq 0, \quad (4.2)$$

where $c_1, c_2, a_i \in \mathbb{R}$, $i = 1, \dots, 4$, are unknown coefficients to be estimated. The first two terms in (4.2) model the overall level and trend, while the following two takes care of the weekly seasonal component and the last two are for the yearly seasonal pattern (recall that weekends have been removed).

Inspecting the price series in Figure 1 and given that we reject the presence of a unit root in the series, there seems to be no evidence of any trend in the data; indeed, in Section 6 we will obtain estimates of c_2 which are practically zero. One could therefore argue that the linear trend term in the seasonal function is superfluous, but since the specification (4.2), including the linear trend component, is standard in the literature, we will keep it in the sequel.

4.2 Modelling the base component

As explained in the introduction, several earlier studies have found evidence of fractal characteristics of the prices and we therefore propose to model the base component X as a fractal process. A mathematical process which is particularly suited for our needs is the volatility modulated Brownian semistationary (BSS) process (Barndorff-Nielsen and Schmiegel, 2007, 2009),

$$X_t = \int_{-\infty}^t (t-s)^\alpha e^{-\lambda(t-s)} \sigma_s dW_s, \quad t \geq 0, \quad (4.3)$$

where $\lambda > 0$, $\alpha \in (-\frac{1}{2}, \frac{1}{2})$ and $\sigma = \{\sigma_t\}_{t \geq 0}$ is a stochastic volatility process, possibly correlated with the standard Brownian motion W . As shown in Bennedsen et al. (2015) the fractal index of X , as defined in Section 3, is equal to α almost surely. We only consider X_t for $t \geq 0$; the integration from minus infinity in (4.3) is purely for modelling reasons since it causes X to be a (strictly) stationary process as long as σ is stationary as well, which we assume henceforth. As discussed in Section 2 and evidenced in Table 2, stationarity is a desirable feature for our purposes,

as the (de-seasonalized) electricity prices we work with are found to be stationary. A few additional comments on the \mathcal{BSS} process X are in order. First off, (4.3) evidently extends the popular OU model, which is the go-to process in the multi-factor literature. To see this, consider (4.3) with $\alpha = 0$; now X indeed reduces to the (volatility modulated) OU process. This will provide the basis of our statistical test for fractality later: we will test the null hypothesis $H_0 : \alpha = 0$ against the alternative $H_a : \alpha < 0$, which gives rise to a statistically rigorous way to decide if the base signal is best described by a rough fractal \mathcal{BSS} process or whether one can do away with an OU process as is the standard choice in the multi-factor literature. In other words, this methodology provides a direct test comparing the OU model with the fractal \mathcal{BSS} model which improves on the comparison made in e.g. Rypdal and Løvsetten (2013) where the authors compared an OU process with a fractal model by comparing various statistical properties of the two processes "by eye". Lastly, the framework in equation (4.3) allows for non-Gaussianity through volatility modulation; in fact, in the multi-factor literature, the base signal is often modelled as a Gaussian OU process, despite studies such as Benth et al. (2012) finding evidence of non-Gaussianity. Indeed, in Section 6.2 we also find evidence of non-Gaussianity. The next section explains how the \mathcal{BSS} framework can accomodate this through the stochastic volatility component σ .

4.2.1 Stochastic volatility: marginal distribution of the base component

The stochastic volatility process allows for flexible modelling of the marginal distributon of the \mathcal{BSS} process X . When the SV process is constant, $\sigma(t) = 1$ for all t say, the resulting process X is Gaussian. When σ is stochastic, however, we get, conditionally on σ ,

$$X_t | \sigma \sim N(0, \xi_t^2), \quad t \geq 0, \quad (4.4)$$

where

$$\xi_t^2 := \int_0^\infty x^{2\alpha} e^{-2\lambda x} \sigma_{t-x}^2 dx, \quad t \geq 0.$$

This shows that the distribution of X_t for $t \geq 0$ is a mean-variance mixture and the form of σ will be determining the marginal distribution of the \mathcal{BSS} process. Specifying σ^2 as a *Lévy semistationary process* (\mathcal{LSS}) process, i.e.

$$\sigma_t^2 = \int_{-\infty}^t k(t-s) v_s dZ_s, \quad t \geq 0,$$

where k is a kernel function, v is volatility of volatility and Z is a subordinator (a non-decreasing Lévy process), provides a particularly powerful and convenient modelling framework.

As the marginal distribution of electricity spot price series has been found (e.g. Barndorff-Nielsen et al., 2013; Veraart and Veraart, 2014, and also Section 6.2) to be well described by the Normal Inverse Gaussian (NIG) distribution, it is an attractive feature of the \mathcal{BSS} framework that it accomodates this distribution for X_t . Indeed, as shown in Barndorff-Nielsen et al. (2013), when X is given as in (4.3) and when letting $k(x) = \frac{x^{-2\alpha-1} e^{-2\lambda x}}{\Gamma(2\alpha+1)\Gamma(-2\alpha)}$, $v_t := 1$ for all t and thus

$$\sigma_t^2 = \frac{1}{\Gamma(2\alpha+1)\Gamma(-2\alpha)} \int_{-\infty}^t (t-s)^{-2\alpha-1} e^{-2\lambda(t-s)} dZ_s, \quad t \geq 0, \quad (4.5)$$

with Γ being the gamma function, then, by the stochastic Fubini theorem,

$$\xi_t^2 = \int_{-\infty}^t e^{-2\lambda(t-s)} dZ_s, \quad t \geq 0. \quad (4.6)$$

In other words, $\xi^2 = \{\xi_t^2\}_{t \in \mathbb{R}}$ is a Lévy driven OU process with mean reversion parameter 2λ . As argued in [Barndorff-Nielsen et al. \(2013\)](#), (4.6) implies the existence of a Lévy process Z , such that ξ_t^2 has the Inverse Gaussian distribution; the upshot is that X , by (4.4), is NIG distributed, see e.g. [Barndorff-Nielsen and Halgreen \(1977\)](#). In Section 6.2 we illustrate how the NIG distribution fits the empirical distribution of the base signal much better than the Gaussian distribution does, strengthening our view that stochastic volatility is needed in the model. Note also, that the parameters in (4.5) are the same as those of the \mathcal{BSS} process (4.3). In other words, when modelling the SV component as in (4.5), and thus obtaining the NIG distribution as the marginal distribution of X , we need only estimate the parameters $\theta = (\alpha, \lambda)^T$ of the \mathcal{BSS} process to also estimate the parameters from the SV process σ .

4.2.2 An alternative fractal component

A perhaps more well-known choice for a stationary fractal component is the fBm-driven OU process, i.e.

$$X_t^H = \int_{-\infty}^t e^{-\lambda(t-s)} dB_s^H, \quad t \geq 0, \quad (4.7)$$

where $\lambda > 0$ and B^H is a fractional Brownian motion with Hurst index $H \in (0, 1)$. Recall, that in terms of the fractal index we have $\alpha = H - \frac{1}{2}$, so that $H = 1/2$ corresponds to B^H being a standard Brownian motion, while $H < 1/2$ indicates paths rougher than the Bm. This fBm-OU framework was indeed what was employed in [Rypdal and Løvstetten \(2013\)](#). However, we wish to employ a more flexible fractal model with a more tractable autocorrelation structure than (4.7) – which is not available in closed form – and we therefore prefer to use the \mathcal{BSS} process (4.3), which ACF is in closed form and given below in Equation 4.10. In Section 6.3 we will compare the two models when investigating if a fractal process can be utilized for increased forecast performance of prices, see also Section 5.2.

4.3 Modelling the spike component

When examining paths of electricity prices such as the ones seen in Figure 1, it seems reasonable to model the very extreme moves by a *spike process*, i.e. a process which suddenly jumps a large (positive or negative) value and rapidly reverts back to the previous level. The reasons for such moves are numerous and were briefly discussed in the introduction. A standard and very general approach to model the spikes in the multi-factor framework is using a non-Gaussian OU model for Y , such as was done in e.g. [Meyer-Brandis and Tankov \(2008\)](#) and [Benth et al. \(2012\)](#). That is, we take

$$Y_t = \int_{-\infty}^t e^{-\lambda_Y(t-s)} dL_s, \quad t \geq 0,$$

where $\lambda_Y > 0$ controls the speed of mean reversion and L is a pure jump Lévy process, e.g. a compound Poisson process. Typically, λ_Y will be quite large, implying fast decay of the spikes. Again, the integration from minus infinity is to ensure stationarity of Y . To focus on the fractal nature of the base signal we will be wholly agnostic about the distribution and general form of the spike component, Y . Examining the dynamics and distributional properties of the spike signal Y is an interesting subject in itself; we leave such investigations in the fractal-based framework to future studies.

For simplicity, we assume that X and Y are independent; this will cause the de-seasonalized prices, $S - \Lambda = X + Y$, to be (strictly) stationary.

4.4 Estimation of the fractal multi-factor model

The estimation procedure we propose and will use in the empirical study in Section 6 is straightforward given the above discussion. In a first step, we wish to estimate the seasonal component (4.2). However, the estimation can be very sensitive to outliers (spikes), and we therefore construct a smoothed prices series on which we estimate (4.2) by OLS. Specifically, the time series we use for this OLS regression is obtained by a block moving average of the log-prices, where the window length is 15 observations on each side of an individual observation. After having estimated the seasonal parameters we subtract the seasonal signal, (4.2), from the (original, non-smoothed) log-prices, leaving us with the time series $\{S_t - \Lambda_t\} = \{X_t + Y_t\}$. The next step is to separate the base signal X from the spike signal Y . We do this using the *hard thresholding* algorithm of Meyer-Brandis and Tankov (2008) – an example of such filtering is given in Section A.2 where we study the EEX data set in detail. Finally, after subtracting the jump process we are left with the base signal $\{X_t\} = \{S_t - \Lambda_t - Y_t\}$, which we proceed to model as the \mathcal{BSS} process (4.3).

4.4.1 Estimation the fractal index of time series of electricity prices

There are several estimators available in the literature with which to estimate the fractal index of a time series. In this section, we will briefly review some of them and explain how they fit in our setup; we refer to Gneiting et al. (2012) a more detailed exposition. For our purposes, we face two major obstacles when estimating the fractal index of electricity prices; (i) jumps (spikes) in the price process and (ii) relatively low frequency (daily) data. In Appendices B.1 and B.2 we will look closer at how the estimators behave under (i) and (ii), and we find that the former causes a downward bias in the estimation of the fractal index, while the latter is innocuous in our setup. The downward bias introduced by (i), i.e. the spikes, is the reason for us filtering out the spike signal Y , before estimating the fractal index α . This has not been done in previous studies, which leads us to conjecture that earlier studies might have found values of α which are more negative than they are in reality, cf. Section 6 and in particular Table 6.

The most straightforward estimation of the fractal index is to use relations similar to Equation (3.1). For a stationary Gaussian process X we actually have the more general (see Gneiting et al., 2012) result concerning the *variogram* of order $p > 0$,

$$\gamma_p(h) = \frac{1}{2} \mathbb{E}[|X_{t+h} - X_t|^p] = c_p |h|^{(2\alpha+1)p/2} + o\left(|h|^{(2\alpha+1+\beta)p/2}\right), \quad \text{as } h \downarrow 0, \quad (4.8)$$

where again α is the fractal index of the process, $\beta \geq 0$ and $c_p > 0$.⁵ Now, for a $p > 0$, an estimator of $a = (2\alpha + 1)p/2$ is obtained by regressing $\log |h|$ on an estimate of the log-variogram, i.e.

$$\log \hat{\gamma}_p(h) = C + a \log |h|, \quad (4.9)$$

where C is a constant. We then estimate $\hat{\alpha}(p) = \hat{a}/p - \frac{1}{2}$, where \hat{a} is obtained from the OLS regression in (4.9) with $\hat{\gamma}_p(h) \frac{1}{N-h} \sum_{i=1}^{N-h} |X_{i+h} - X_i|^p$. In practice it is prudent to use only very few lags of h in this regression, e.g. two or three lags. One also need to choose the power parameter p , a standard choice being $p = 2$ so that $\gamma_p(h) = \gamma_2(h)$ is the variogram of the process. However, $p = 1$ — which means that $\gamma_p(h) = \gamma_1(h)$ is the so-called *madogram* — has been found to be more robust to outliers (see Section B.1 and Gneiting et al. (2012)). In the following we will employ this method using both $p = 1$ and $p = 2$, and will term them "madogram" and "variogram" respectively.

Besides the straightforward method discussed above, there exist a wealth of more sophisticated estimators of the fractal index α (often in the form of an estimator of the Hurst index $H = \alpha + \frac{1}{2}$), see Taqqu et al. (1995) for a partial list along with a Monte Carlo study of their finite sample performances. We focus here on the one-dimensional, i.e. time series, case. Popular estimators used in the literature include Detrended Fluctuation Analysis (DFA) of Peng et al. (1994) and estimation of the generalized Hurst exponent (GEN), see e.g. Matteo (2007). Here we also consider the absolute value method (ABS) and the aggregated variance method (AGG), see Taqqu et al. (1995), as well as the Change-of-Frequency (COF) estimator of Barndorff-Nielsen et al. (2011). Both the GEN and the COF estimators rely on choosing a power parameter p , as was the case of the madogram/variogram in the regression (4.9). Thus, in the following we will write $GEN(p)$ and $COF(p)$ for the respective estimator using the power parameter p , typically $p = 1$ or $p = 2$. We note, that the $COF(2)$ estimator comes with a CLT (Barndorff-Nielsen et al., 2011), which makes this estimator particularly attractive for our purpose, since it will allow us to do inference on the fractal parameter and test whether it is statistically different from zero (recalling that zero corresponds to the \mathcal{BSS} process X being actually a OU process). Of the above estimators, DFA has arguably been the most popular one in the electricity price literature as it controls for the possible non-stationarity induced by seasonal effects. When estimating the fractal index in what follows, we will employ all the above mentioned estimators instead of relying on one single method, as is the usual practice in the electricity literature. We do this for two reasons; firstly, we use them as a robustness check to see whether all estimators agree. Secondly, presenting the estimations using all of the well-known (and some not so well-known) estimators should make it more comparable to earlier studies. For implementation of the respective estimators, we refer to the references above.

4.4.2 Estimating the mean reversion parameter λ

As λ controls the mean reversion (at least on longer scales), we seek to estimate this parameter using the dependence structure of the base signal X . We do this using our estimate of α and the autocorrelation function (ACF) of the \mathcal{BSS} process (4.3), which can be shown to be equal to (e.g.

⁵Note, that even though this definition is stated for Gaussian processes, it equally well applies to volatility modulated \mathcal{BSS} (and OU) processes, see e.g. Bennedsen et al. (2015).

Barndorff-Nielsen, 2012)

$$\rho(h) = \text{Corr}(X_{t+h}, X_t) = \frac{2^{-\alpha+\frac{1}{2}}}{\Gamma(\alpha+1/2)} (\lambda h)^{\alpha+\frac{1}{2}} K_{\alpha+\frac{1}{2}}(\lambda h), \quad h \in \mathbb{R}, \quad (4.10)$$

where Γ is the gamma function and $K_\nu(x)$ is the modified Bessel function of the third kind with parameter ν , evaluated at x , see e.g. Abramowitz and Stegun (1972). To estimate λ , we perform a non-linear least squares estimation fitting (4.10) to the empirically observed autocorrelations.

5 Implications of a fractal model of electricity prices

Fractal models of financial prices are non-standard and will as such have different implications than a usual model based on a non-fractal process. We briefly consider some of these implications here.

5.1 Arbitrage

It is well known that fractal models such as the fBm-OU model (4.7) and the fractal \mathcal{BSS} process in (4.3) are, in general, *not* semimartingales. Classically, using such a process as the basis of the price of a financial asset would constitute an arbitrage opportunity (or, more precisely, a free lunch with vanishing risk), as was shown in Delbaen and Schachermayer (1994), see also Rogers (1997). One could fear that introducing a fractal model for the price of electricity would induce the same adverse effects. Fortunately, there are several reasons why this need not be. Most importantly, the electricity spot market is wholly different from the usual financial markets in that one can not buy electricity and hold it in a portfolio for a later date. For this reason, the usual 'buy-and-hold' arbitrage arguments break down and we can not rely on the classical no-arbitrage results such as the one in Delbaen and Schachermayer (1994). As an additional assurance, Pakkanen (2011) showed that \mathcal{BSS} processes such as the one in (4.3) will not be the source of arbitrage when transaction costs are considered.

5.2 Forecasting

Since fractality refers to the small scale behavior of the time series under consideration, it is plausible that including a fractal term will provide superior forecasts of prices *at least in the short run*. Since our three components in (4.1) are assumed independent, we will focus on forecasting the base signal only; the (deterministic) seasonal component is then easily adjusted for, as are the spike signal Y , given a model for this. However, as we explained in Section 4.3 we do not go into details with the spike part in this work. Hence, we will focus on predicting X in the following. The two forecasting methods we outline here will be empirically tested in Section 6.3 and visually illustrated for the EEX data set in Appendix A.4.

5.2.1 Forecasting the \mathcal{BSS} process

Ideally, for a forecast horizon $h > 0$ (days, say) we would like to compute

$$\hat{X}_{t+h} = \mathbb{E}[X_{t+h} | \mathcal{F}_t], \quad (5.1)$$

where $\mathcal{F}_t = \sigma\{X_s, s \in [0, t]\}$ is the filtration generated by the base signal X . Unfortunately, for the \mathcal{BSS} process, the expression for (5.1) is an unsolved problem. Instead, we will use a linear combination of past observed values to predict the future, i.e. we use the *best linear predictor*. The best linear predictor is given in terms of the autocorrelation function of the \mathcal{BSS} process, which we recall is given by Equation (4.10). The best linear predictor of X_{t+h} given X_1, X_2, \dots, X_t is (e.g. [Grimmett and Stirzaker, 2001](#), Section 9.2.)

$$\hat{X}_{t+h} = \sum_{i=1}^t a_i X_i, \quad h \geq 1,$$

where $a_1, a_2, \dots, a_t \in \mathbb{R}$ solves

$$\sum_{i=1}^t a_i \rho(|i-j|) = \rho(h+j), \quad 1 \leq j \leq t.$$

5.2.2 Forecasting the fBm-OU process

In the case of forecasting, there are advantages in choosing the fBm-based OU process in Equation (4.7) as a model of the base signal; the fBm has been extensively studied and in particular, we have a closed-form expression for the conditional expectation (5.1). To be precise, [Nuzman and Poor \(2000\)](#) show that for $H < \frac{1}{2}$

$$\mathbb{E}[B_{t+h}^H | \mathcal{F}_t^H] = \frac{\cos(H\pi)}{\pi} h^{H+1/2} \int_{-\infty}^t \frac{B_s^H}{(t-s+h)(t-s)^{H+1/2}} ds,$$

where $\mathcal{F}_t^H = \sigma\{B_s^H, s \leq t\}$ is the filtration generated by the fBm B^H during its entire history (i.e. from the infinite past until time t). Now, as argued in [Gatheral et al. \(2014\)](#), for reasonably small time scales h and for small values of the mean reversion parameter λ , we have the approximation

$$\mathbb{E}[X_{t+h}^H | \mathcal{F}_t^H] \approx \frac{\cos(H\pi)}{\pi} h^{H+1/2} \int_{-\infty}^t \frac{X_s^H}{(t-s+h)(t-s)^{H+1/2}} ds, \quad (5.2)$$

where X^H is the fBm-OU process in (4.7). We approximate this expression using the observed values of X^H together with a straightforward Riemann sum:

$$\hat{X}_{t+h}^H = \frac{\cos(H\pi)}{\pi} h^{H+1/2} \sum_{i=1}^{\lfloor t/\delta \rfloor} \frac{X_{i\delta}^H}{(t-(i-r)\delta+h)(t-(i-r)\delta)^{H+1/2}} \delta, \quad r \in (0, 1],$$

where $\delta > 0$ is the observation frequency (i.e. $\delta = 1$ day in our case) and $H = \alpha + \frac{1}{2}$ is the Hurst exponent. Note, that we face a rather arbitrary choice of $r \in (0, 1]$ when evaluating the integral at the points $s = (i-r)\delta$ for $i = 1, \dots, \lfloor t/\delta \rfloor$; in Section 6.3 we will simply choose r such that the forecast performance is optimized for the shortest forecast horizon $h = 1$ day.

6 Empirical analysis

In this section we apply the estimation procedure explained in Section 4.4 to the data in Table 1; in Appendix A we will illustrate the approach by considering the German EEX data in detail. The

Table 4: *Estimation of seasonality parameters*

	c_1	c_2	a_1	a_2	a_3	a_4
EEX	4.303 (0.011)	−0.000 (0.000)	0.030 (0.007)	−0.006 (0.007)	−0.079 (0.007)	0.071 (0.007)
NP	3.938 (0.013)	−0.000 (0.000)	0.017 (0.009)	0.000 (0.009)	0.010 (0.009)	0.101 (0.009)
Powernext	4.333 (0.011)	−0.000 (0.000)	0.024 (0.007)	−0.008 (0.007)	−0.060 (0.008)	0.118 (0.007)
GME	4.666 (0.009)	−0.000 (0.000)	0.011 (0.006)	−0.008 (0.006)	−0.064 (0.006)	0.041 (0.006)
APX UK	4.112 (0.012)	−0.000 (0.000)	0.005 (0.009)	0.001 (0.009)	−0.031 (0.009)	0.019 (0.009)
APX NL	4.270 (0.011)	−0.000 (0.000)	0.016 (0.007)	−0.011 (0.007)	−0.046 (0.008)	0.065 (0.007)

Estimates of the parameters of the seasonal function (4.2) with standard errors in parentheses.

estimation results for all six data sets are seen in Tables 4-6. Table 4 shows the estimates from the seasonal function in Equation (4.2): we see no noticeable trend (c_2), while there is a level effect (c_1) and some seasonality in all data series (a_1, a_2, a_3, a_4).

6.1 Roughness of electricity price series

Table 5 contains the most important information for our purposes. From the COF(2) estimator we have estimates of the fractal index α , which are negative across the board, indicating that the time series indeed do display a degree of fractality. Similarly, when conducting the formal hypothesis test of $H_0 : \alpha = 0$ against the alternative $H_a : \alpha < 0$, we reject H_0 for all markets under consideration, except Nord Pool, indicating that in these energy markets the \mathcal{BSS} process is to be preferred over the OU model as a model of the base signal. That is, we conclude that five out of the six markets studied here display fractal characteristics in their base signals. To further check this fact we employ all estimators of α discussed in Section 4.4.1; the results are seen in Table 6 which corroborates the findings. Indeed, *all* estimates of α are negative. Note that if one does not filter out the spikes (Panel A), the α estimates from Nord Pool using either of the GEN(1), GEN(2) and DFA estimators would – incorrectly – lead one to find $\alpha < 0$ in this market. Further, this is even the case after filtering out the spikes (Panel B) for the estimators ABS and AGG. Recall, that the DFA estimator has traditionally been popular in the literature. Therefore, it is evident that it is important to first filter out the spikes and, further, that which estimator one uses really does matter as ABS and AGG in general give more extreme results than do the other estimators. Additionally, it is not only in the Nord Pool case that some estimators become downward biased when one does not first filter out the spikes; to examine this closer we plot in Figure 3 the estimates of α for all markets using the madogram, variogram, COF(1) and COF(2) estimators. The first conclusion we can draw from this plot is that the four estimators agree more closely after filtering out the spikes, strengthening our conviction that some estimators are more

biased than others in the presence of outliers (see also Appendix B.1). The second conclusion is that the estimates coming from some markets become noticeably less negative, when first filtering out the jumps. This is most evident in the GME and Powernext data sets, but also to a lesser degree in EEX and APXUK. Again, this is evidence for the presence of biases in the estimation when one does not account for the spike signal separately. Indeed, the above shows that accounting for spikes prior to estimation of α is crucial and the fractal multi-factor framework is thus ideally suited for modelling electricity prices, since it allows us to do this in a straightforward manner. We note again that most authors in the fractal literature have generally *not* done any attempt to account for spikes, thus leading us to conjecture that earlier estimates in the literature might have been *too negative*.

Finally, turning again to Table 5, we also see estimates of the mean reversion parameter λ from both the \mathcal{BSS} model and the OU model. In all cases, the estimates coming from the \mathcal{BSS} model are smaller in magnitude as compared to the estimates coming from the OU model, implying slower (long term) mean reversion in the \mathcal{BSS} models. The reason for this is that the α parameter helps in fitting the autocorrelations in the short run: a negative value of α implies a sharp decrease in autocorrelation for the very short lags. This allows the λ parameter in the \mathcal{BSS} model to better fit the ACF at the longer lags. This effect is similar to what an OU two-factor model would achieve: here two different mean reversion parameters would also allow for both fast and slow decay of autocorrelations. The advantage of the \mathcal{BSS} model is that it does this while also allowing for fractality. Further, the \mathcal{BSS} model is more parsimonious as it only requires a single source of randomness, instead of two in the two-factor model.

Table 5: *Estimating and testing for fractality*Panel A: *Log-prices*

	EEX	NP	Powernext	GME	APX UK	APX NL
$\hat{\alpha}$	-0.204	-0.045	-0.300	-0.405	-0.306	-0.161
$\text{std}(\hat{\alpha})$	0.100	0.069	0.168	0.230	0.049	0.061
Test-stat	-2.044	-0.646	-1.792	-1.763	-6.283	-2.668
P-value	0.020	0.259	0.037	0.039	0.000	0.004
Reject at 5 perc	1	0	1	1	1	1
λ (\mathcal{BSS})	0.120	0.043	0.034	0.019	0.021	0.052
λ (OU)	0.296	0.056	0.218	0.463	0.192	0.131

Panel B: *Filtered log-prices (base signal)*

	EEX	NP	Powernext	GME	APX UK	APX NL
$\hat{\alpha}$	-0.211	-0.059	-0.198	-0.246	-0.308	-0.191
$\text{std}(\hat{\alpha})$	0.037	0.046	0.045	0.044	0.040	0.040
Test-stat	-5.683	-1.281	-4.410	-5.597	-7.664	-4.825
P-value	0.000	0.100	0.000	0.000	0.000	0.000
Reject at 5 perc	1	0	1	1	1	1
λ (\mathcal{BSS})	0.048	0.026	0.033	0.031	0.011	0.019
λ (OU)	0.149	0.037	0.105	0.144	0.144	0.074

Estimating the fractal index, α , and testing $H_0 : \alpha = 0$ for various data sets. Panel A is using the standard approach of directly estimating the fractal index of the de-seasonalized prices. In contrast, Panel B is after filtering the seasonal and spike signals from the data, i.e. the estimations are done on the base signal. $\hat{\alpha}$ is estimated using the COF(2) estimator. The test statistic and P-value refer to the hypothesis test $H_0 : \alpha = 0$ against the alternative $H_a : \alpha < 0$. The bottom two rows are estimates of the mean reversion parameter λ in the \mathcal{BSS} model as well as the OU model.

Table 6: *Estimating the fractal index of all data sets using various estimators*
Panel A: *Log-prices*

	EEX	NP	Powernext	GME	APX UK	APX NL
<i>COF</i> (2)	−0.204	−0.045	−0.300	−0.405	−0.306	−0.161
<i>COF</i> (1)	−0.198	−0.022	−0.197	−0.260	−0.306	−0.184
Variogram	−0.270	−0.088	−0.294	−0.421	−0.342	−0.224
Madogram	−0.243	−0.052	−0.211	−0.290	−0.343	−0.201
<i>GEN</i> (2)	−0.340	−0.135	−0.286	−0.464	−0.365	−0.329
<i>GEN</i> (1)	−0.301	−0.112	−0.231	−0.341	−0.382	−0.296
ABS	−0.345	−0.278	−0.293	−0.350	−0.399	−0.273
DFA	−0.375	−0.169	−0.320	−0.454	−0.402	−0.335
AGG	−0.403	−0.350	−0.370	−0.431	−0.431	−0.307

Panel B: *Filtered log-prices (base signal)*

	EEX	NP	Powernext	GME	APX UK	APX NL
<i>COF</i> (2)	−0.211	−0.059	−0.198	−0.246	−0.308	−0.191
<i>COF</i> (1)	−0.230	−0.033	−0.194	−0.210	−0.314	−0.218
Variogram	−0.220	−0.073	−0.178	−0.253	−0.314	−0.198
Madogram	−0.232	−0.052	−0.175	−0.218	−0.314	−0.190
<i>GEN</i> (2)	−0.248	−0.082	−0.170	−0.267	−0.328	−0.222
<i>GEN</i> (1)	−0.242	−0.082	−0.176	−0.252	−0.349	−0.229
ABS	−0.323	−0.235	−0.294	−0.346	−0.446	−0.297
DFA	−0.309	−0.092	−0.243	−0.300	−0.369	−0.291
AGG	−0.334	−0.291	−0.341	−0.365	−0.460	−0.316

Estimating the fractal index for various data sets and with the estimators of Section 4.4.1. In Panel A estimation is done on the de-seasonalized log-prices. In contrast, Panel B (right) is after filtering the seasonal and spike signals from the data, i.e. the estimations are done on the base signal.

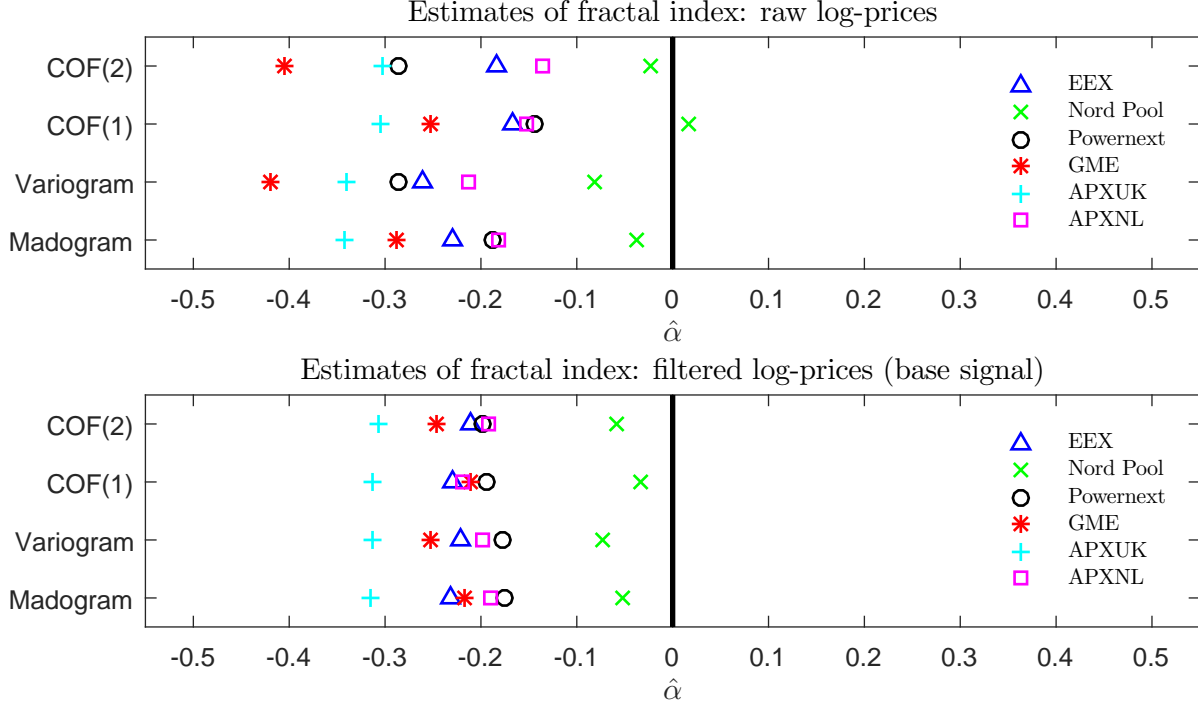


Figure 3: *Estimation of the fractal index of the spot prices from various energy exchanges. Estimation is done on the raw log-prices (top) and on base signal, i.e. after filtering out the seasonal component and the spike signal (bottom).*

6.2 Marginal distribution of the base signal

As explained in Section 4.2.1, the \mathcal{BSS} framework accomodates the NIG distribution as its marginals; we now take a closer look at how this distribution fits the empirical distribution of the base signal. We see below in Figure 4 that the NIG distribution indeed fits very well the empirically observed marginal distribution of the base signal, a fact that was also found in e.g. [Barndorff-Nielsen et al. \(2013\)](#) and [Veraart and Veraart \(2014\)](#). The plots in the left column of Figure 4 simply plots the empirical density of the base signal against a fitted Gaussian and a fitted NIG density: in most cases the NIG distribution fits the data much better than do the Gaussian distribution, especially around the center of the distribution. To examine the density in the tails better, we consider similars plot in the right column of Figure 4 but now using log-densities: here we again see that the NIG distribution is in general also better at capturing the tail behavior of the data, indicating fat tails of the marginal density of the time series under consideration. This is evidence of the inadequacy of the Gaussian OU process as a model of the base signal and thus indicates the presence of stochastic volatility. In summary, in view of the arguments in Section 4.2.1, the \mathcal{BSS} framework is able to accurately capture the empirically observed marginal density of the extracted base signal of electricity prices.

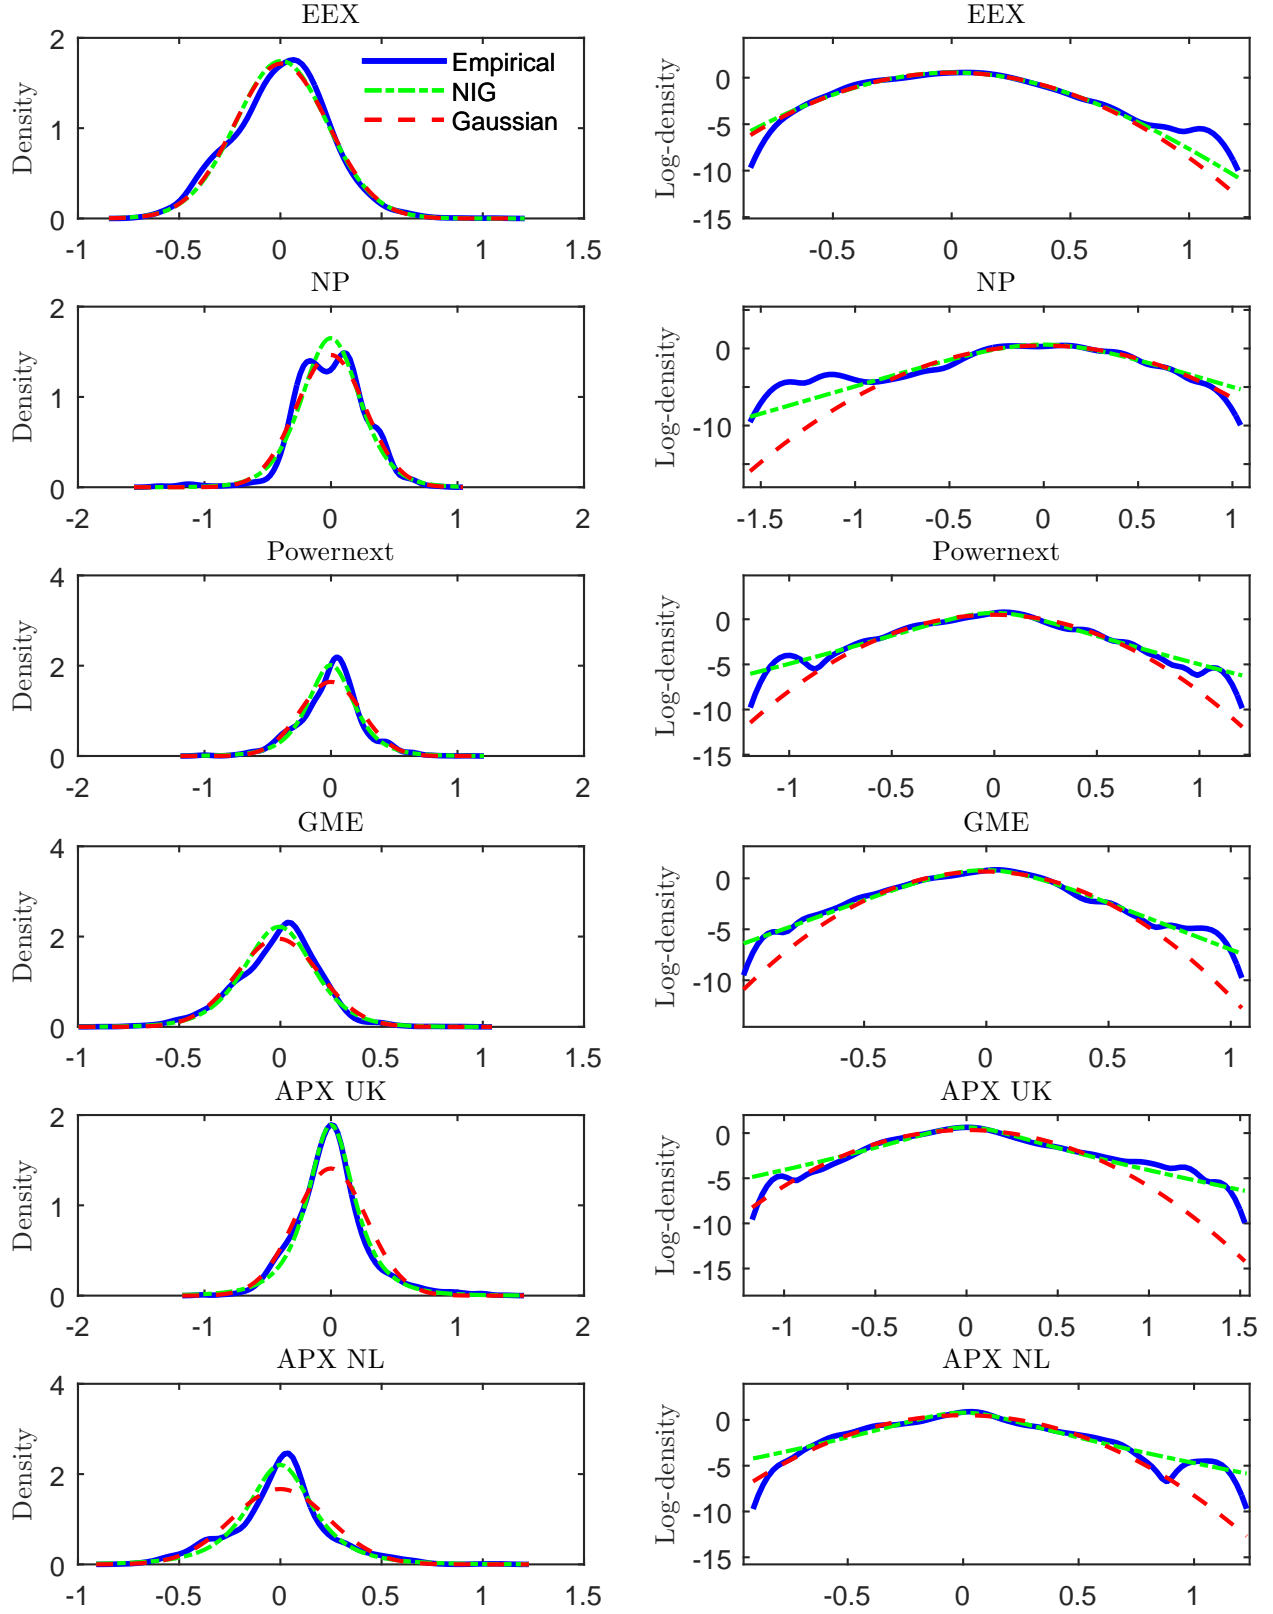


Figure 4: Empirical marginal distribution of the base signal (blue solid line) superimposed with a fitted NIG density (green half-broken line) and the corresponding fitted Gaussian density (red broken line). The plots in the right column is using log-densities.

6.3 Forecasting

Lastly, in Table 7 we present the results of a forecasting exercise. We took the extracted base signal from each of the six exchanges and applied the forecasting approaches of Section 5.2 using different values of forecast horizon h . We then compared the forecasts from the \mathcal{BSS} model with those from three benchmark models, the fBm model, the OU model, and a random walk (RW) model. The RW model is simply a "no-change" forecast, while the (discretized) OU model can be written as an AR(1) model and therefore easily forecast using standard methods. Table 7 shows root mean squared forecast errors (RMSFE) of the three approaches FBM, OU, and RW as compared to the \mathcal{BSS} approach. That is, the numbers reported are

$$\text{Relative RMSFE}(x) = \frac{\text{RMSFE}(\mathcal{BSS})}{\text{RMSFE}(x)}, \quad x = \text{FBM, OU, RW}.$$

Thus, values less than one indicate that the \mathcal{BSS} -based forecasts fare better than the alternative model. Some conclusions emerge: firstly, the fBm-based forecasts perform very poorly in general; this could be an indication that the model is not describing the data well, but more likely it is because the approximation in (5.2) is poor for our purpose. Note, that the autocorrelation of the OU-fBm model (4.7) is not available in closed form, hence making prediction using the best linear predictor, as was done for the \mathcal{BSS} process, difficult, at least without resorting to approximations or numerical calculations of the autocorrelation function. This makes fractal models with tractable autocorrelation structures, such as the \mathcal{BSS} process, preferable from a forecasting standpoint. Secondly, the \mathcal{BSS} forecast method generally fares better than the (non-fractal) OU method, when the market in question display fractality, i.e. for all markets except Nord Pool (where we could not reject $H_0 : \alpha = 0$, see Table 5). Lastly, the RW model performs poorly, especially for longer horizons, indicating that including some form of mean reversion in the model is indeed important for forecasting.

Table 7: *Forecasting the base signal*
RMSE relative to \mathcal{BSS}

		$h = 1$	$h = 2$	$h = 4$	$h = 8$	$h = 12$	$h = 16$
EEX	RW	0.928	0.909	0.880	0.835	0.796	0.781
	OU	0.968	0.980	0.993	1.005	1.004	0.988
	FBM	0.821	0.846	0.861	0.856	0.842	0.844
NP	RW	1.007	0.990	0.978	0.945	0.914	0.882
	OU	1.015	1.005	1.006	1.001	1.000	0.996
	FBM	0.694	0.833	0.913	0.952	0.952	0.944
Powernext	RW	0.958	0.948	0.919	0.858	0.813	0.781
	OU	0.988	1.001	1.012	1.025	1.034	1.032
	FBM	0.866	0.888	0.890	0.867	0.845	0.830
GME	RW	0.921	0.904	0.900	0.885	0.884	0.873
	OU	0.958	0.954	0.943	0.910	0.901	0.907
	FBM	0.760	0.810	0.858	0.892	0.913	0.918
APX UK	RW	0.870	0.842	0.833	0.807	0.813	0.796
	OU	0.910	0.905	0.909	0.871	0.842	0.819
	FBM	0.680	0.747	0.814	0.858	0.887	0.895
APX NL	RW	0.942	0.924	0.902	0.874	0.853	0.835
	OU	0.964	0.963	0.967	0.973	0.973	0.964
	FBM	0.943	0.938	0.934	0.924	0.912	0.904

Forecasting the base signal using the four models, \mathcal{BSS} , FBM, OU and RW. The forecast horizon is h days and the numbers in the tables are root mean squared (RMSE) forecast errors relative to the \mathcal{BSS} model. Values less than one favors the \mathcal{BSS} model, while values greater than one favors the alternative method.

7 Conclusion

This article has conducted an in-depth investigation of the apparent fractality of time series of electricity spot prices. Similar to other authors we find evidence of fractal characteristics in electricity spot prices: in particular, in five out of the six markets studied, we reject the null hypothesis of a zero fractal index in favor of the alternative hypothesis of a negative index. Only in the Nord Pool market could we not reject the null $\alpha = 0$. Contrary to previous studies, we sought to avoid possible sources of biases in the estimators of the fractal index by filtering the prices for seasonality and spikes *before* estimating the degree of fractality in the data. We considered two pieces of evidence, which indicate that the \mathcal{BSS} model is indeed to be preferred over the OU model for most electricity price series. Firstly, we simply find statistically significant values of the fractal index and that these are negative. Secondly, the results of our forecasting exercise showed that there are gains to be had in this area by applying the fractal \mathcal{BSS} model; a very interesting avenue of future research would be to include a fractal component in a more realistic forecasting setup, such as the one considered in Weron (2014). Additionally, we showed that the \mathcal{BSS} framework admits the NIG distribution as the marginal distribution of electricity prices and that this is desirable since the empirical distribution of the base signals was well fit by this distribution.

The analysis in the present work was done on log-prices, but we stress that the approach would apply equally well to raw prices. In fact, we conducted the same analysis on these prices (not given here for brevity) and the results on fractality hold a fortiori: we find even more extreme (negative) values of the fractal parameter α in this case.

Finally, in the model presented here, we all but ignored a very important component of electricity prices, namely the spike process. Further research into a full-fledged fractal model and the implications for spikes, stochastic volatility, and stochastic seasonality is left for future research.

Acknowledgements

I would like to thank Asger Lunde for comments on an earlier draft of this paper as well as for many insightful discussions concerning statistical modelling of energy prices and Solveig Sørensen for competent and meticulous proof reading of the manuscript. The research has been supported by CREATES (DNRF78), funded by the Danish National Research Foundation, by Aarhus University Research Foundation (project “Stochastic and Econometric Analysis of Commodity Markets”), and by Aage and Ylva Nimbs Foundation.

References

- Abramowitz, M. and I. A. Stegun (1972). *Handbook of mathematical functions with formulas, graphs, and mathematical tables* (10th ed.), Volume 55. United States Department of Commerce.
- Barndorff-Nielsen, O. E. (2012). Notes on the gamma kernel. *Thiele Research Reports* (03).
- Barndorff-Nielsen, O. E., F. E. Benth, and A. E. D. Veraart (2013). Modelling energy spot prices by volatility modulated Lévy-driven Volterra processes. *Bernoulli* 19(3), 803–845.

- Barndorff-Nielsen, O. E., J. M. Corcuera, and M. Podolskij (2011). Multipower variation for Brownian semistationary processes. *Bernoulli* 17(4), 1159–1194.
- Barndorff-Nielsen, O. E. and C. Halgreen (1977). Infinite divisibility of the hyperbolic and generalized inverse gaussian distributions. *Z. Wahrscheinlichkeitstheorie verw. Gebiete* 38, 309–311.
- Barndorff-Nielsen, O. E. and J. Schmiegel (2007). Ambit processes: with applications to turbulence and tumour growth. In *Stochastic analysis and applications*, Volume 2 of *Abel Symp.*, pp. 93–124. Berlin: Springer.
- Barndorff-Nielsen, O. E. and J. Schmiegel (2009). Brownian semistationary processes and volatility/intermittency. In *Advanced financial modelling*, Volume 8 of *Radon Ser. Comput. Appl. Math.*, pp. 1–25. Berlin: Walter de Gruyter.
- Bennedsen, M., A. Lunde, and M. S. Pakkanen (2015). Hybrid scheme for Brownian semistationary processes. *Working paper*, *arXiv:1507.03004*.
- Benth, F. E., J. Kallsen, and T. Meyer-Brandis (2007). A non-Gaussian Ornstein–Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance* 14(2).
- Benth, F. E., R. Kiesel, and A. Nazarova (2012). A critical empirical study of three electricity spot price models. *Energy Economics* 34(5), 1589–1616.
- Cartea, A. and M. G. Figueroa (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 12(4), 313–335.
- Delbaen, F. and W. Schachermayer (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen* 300(1), 463–520.
- Erzgräber, H., F. Strozzi, J. Zaldívar, H. Touchette, E. Gutiérrez, and D. K. Arrowsmith (2008). Time series analysis and long range correlations of nordic spot electricity market data. *Physica A* 287, 6567–6574.
- Gatheral, J., T. Jaisson, and M. Rosenbaum (2014). Volatility is rough. *Working paper*.
- Geman, H. and A. Roncoroni (2006). Understanding the fine structure of electricity prices. *Journal of Business* 79(3).
- Gneiting, T., H. Sevcikova, and D. B. Percival (2012). Estimators of Fractal Dimension: Assessing the Roughness of Time Series and Spatial Data. *Statistical Science* 27(2), 247–277.
- Grimmett, G. and D. Stirzaker (2001). *Probability and Random Processes* (3rd ed.). Oxford University Press.
- Hayfavi, A. and I. Talasli (2014). Stochastic multifactor modeling of spot electricity prices. *Journal of Computational and Applied Mathematics* 259(B), 434–442.
- Kaminski, V. (2004). *Managing Energy Price Risk: Third Edition*. Risk Books.
- Kristoufek, L. and P. Lunackova (2013). Long-term memory in electricity prices: Czech market evidence. *Czech Journal of Economics and Finance* 63(5), 407–424.
- Matteo, T. D. (2007). Multi-scaling in finance. *Quantitative Finance* 7(1), 21–36.

- Meyer-Brandis, T. and P. Tankov (2008). Multi-factor jump-diffusion models of electricity prices. *International Journal of Theoretical and Applied Finance* 11(5), 503–528.
- Norouzzadeh, P., W. Dullaert, and B. Rahmani (2007). Anti-correlation and multifractal features of Spain electricity spot market. *Physica A* 380, 333–342.
- Nuzman, C. J. and V. H. Poor (2000). Linear estimation of self-similar processes via Lamberti’s transformation. *Journal of Applied Probability* 37(2), 429–452.
- Pakkanen, M. S. (2011). Brownian semistationary processes and conditional full support. *International Journal of Theoretical and Applied Finance* 14(4), 579–586.
- Peng, C.-K., S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger (1994). Mosaic organization of DNA nucleotides. *Physical Review E* 49(2), 1685–1689.
- Pilipovic, D. (2007). *Energy Risk: Valuing and Managing Energy Derivatives*. McGraw-Hill.
- Rogers, L. C. G. (1997). Arbitrage with fractional Brownian motion. *Mathematical Finance* 7, 95–105.
- Rypdal, M. and O. Løvstetten (2013). Modeling electricity spot prices using mean-reverting multifractal processes. *Physica A* 392(1), 194–207.
- Schwartz, E. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance* 52(3), 923–973.
- Simonsen, I. (2002). Measuring anti-correlations in the nordic electricity spot market by wavelets. *Physica A* 233, 597–606.
- Taqqu, M. S., V. Teverovsky, and W. Willinger (1995). Estimators for long-range dependence: An empirical study. *Fractals* 3, 785–798.
- Veraart, A. E. D. and L. A. M. Veraart (2014). Modelling electricity day-ahead prices by multivariate Lévy semistationary processes. In F. E. Benth, V. A. Kholodnyi, and P. Laurence (Eds.), *Quantitative Energy Finance*, pp. 157–188. New York: Springer.
- Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting* 30, 1030–1081.
- Weron, R. and B. Przybyłowicz (2000). Hurst analysis of electricity price dynamics. *Physica A* 286, 462–468.

A An empirical case study: The German EEX market

We analyze the German EEX spot market from January 1, 2008 to January 23, 2015 excluding weekends using the estimation approach outlined in the article. We work on the peak load price, i.e. the average price over Hour 9 to Hour 20 each day, yielding 1844 daily observations. The raw price (top) and log-price (bottom) are seen in Figure 5.

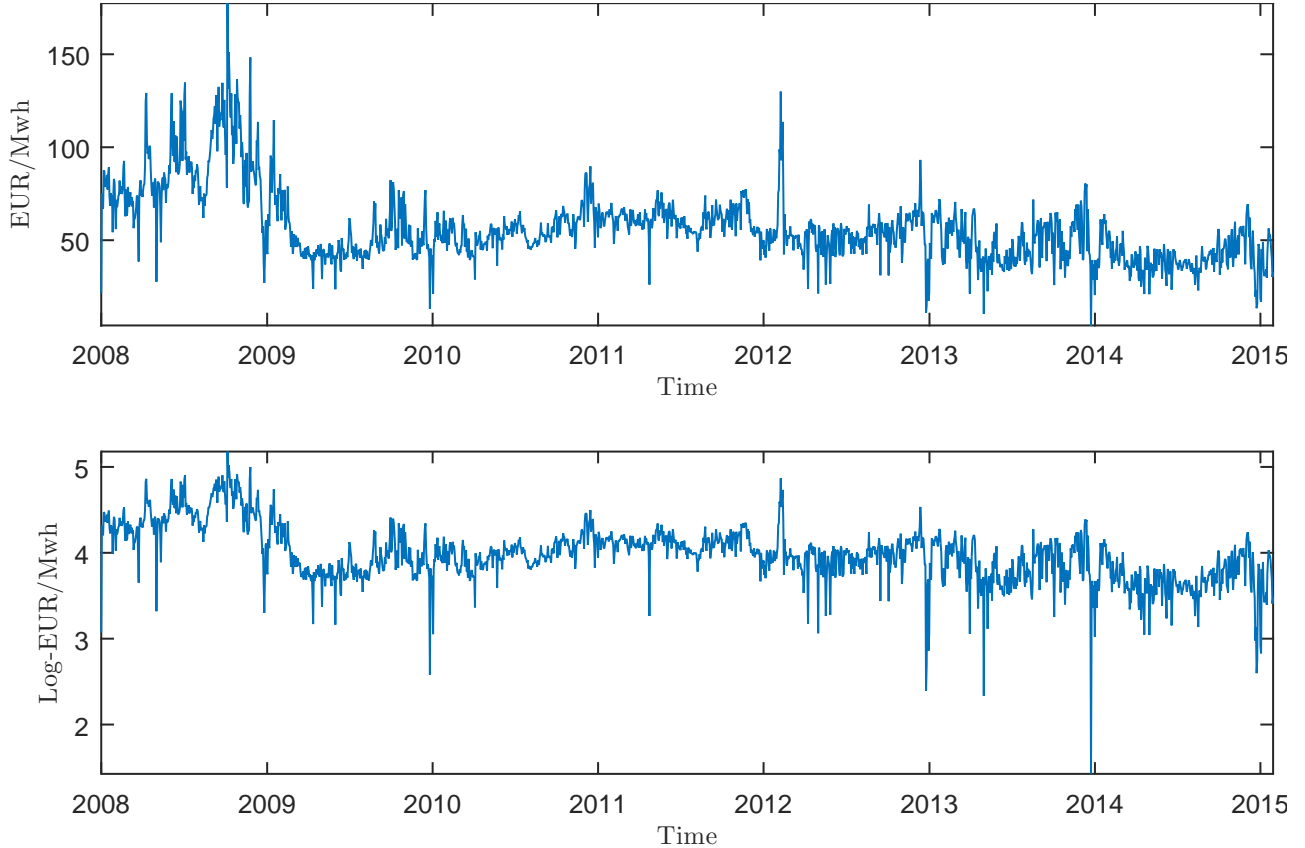


Figure 5: *Raw price (top) and log-price (bottom) of EEX peak load prices from January 1, 2008 to January 23, 2015.*

A.1 De-seasonalizing the data

We do OLS using the seasonal function (4.2) on the log-prices as explained in Section 4.1. The resulting fit is seen in Figure 6 (top) where we also show the residuals $\{\log P_t - \Lambda_t\} = \{X_t + Y_t\}$ (bottom).

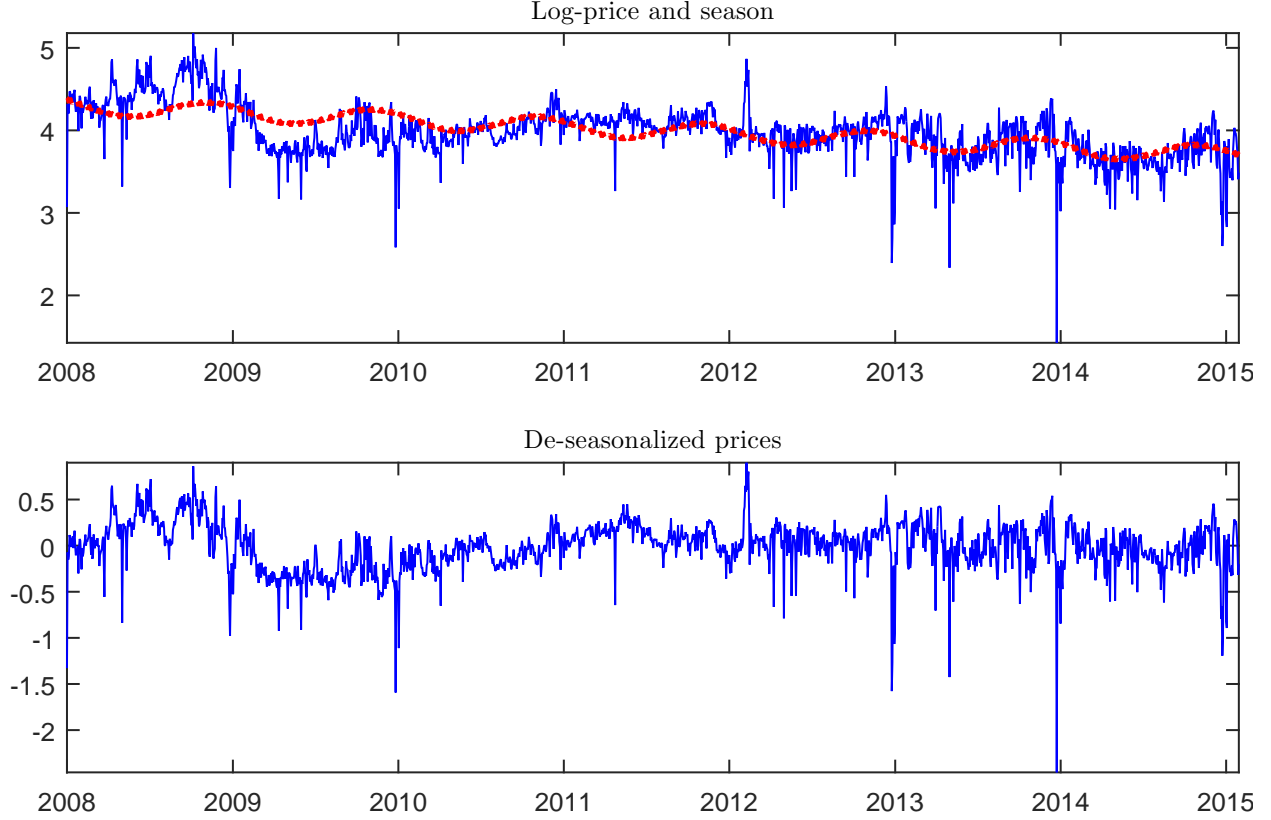


Figure 6: *Log-prices and fitted seasonality function (top) and residuals from the non-linear least squares regression (bottom).*

A.2 Filtering the jump process

We use the hard thresholding method described in [Meyer-Brandis and Tankov \(2008\)](#) and filter out the path of the spike process Y . One could estimate the mean reversion parameters and use them in the thresholding algorithm, but both [Meyer-Brandis and Tankov \(2008\)](#) and [Benth et al. \(2012\)](#) find that tuning them by eye works better. We follow their recommendations and choose $\lambda_1 = \frac{1}{115}$ and $\lambda_2 = \frac{1}{4}$. There is an additional tuning parameter, ϵ , controlling how many jumps to remove in an initial step. Specifically, we remove $\epsilon = 0.045 = 4.5\%$ of the largest (in absolute value) increments of the process before we calculate the (target) standard deviation of the resulting process. Then we insert back the ϵ values and then run the thresholding algorithm, removing spikes until the standard deviation of the base signal equals the target standard deviation. The result is seen in Figure 7.

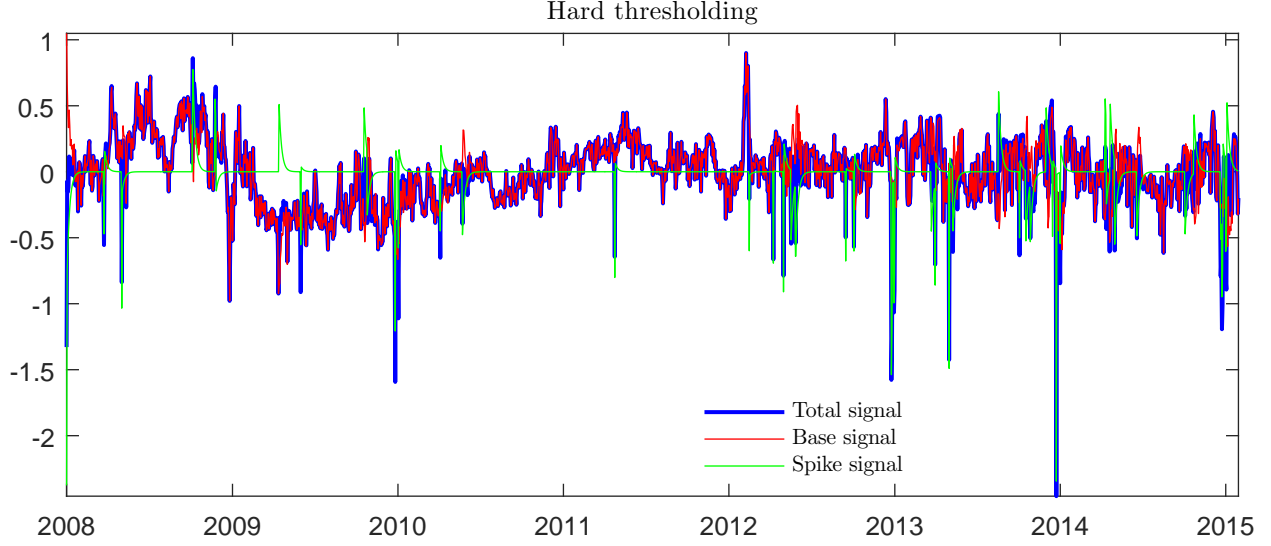


Figure 7: *Result of hard thresholding where we see the initial signal (blue), the filtered base signal (red), and the filtered spike signal (green).*

A.3 Analysis of the residual base signal

The filtered base signal from the hard thresholding algorithm is now assumed to be $\{X_t\} = \{\log P_t - \Lambda_t - Y_t\}$, where Y is the filtered spike signal; see Figure 8. We will model X as the \mathcal{BSS} process (4.3), which will be fit to the empirical base signal — specifically, we will estimate α using the $COF(2)$ estimator and then λ from the empirical autocorrelation (4.10). We then test the hypothesis $H_0 : \alpha = 0$ to see if we can do away with an OU model or if there is still the need for a fractional parameter in the base signal; the result of this test was seen in Table 5, Panel B, where the null hypothesis is rejected in favor of $H_a : \alpha < 0$.

We also estimate the mean-reversion parameter of an OU process, λ_{OU} , which is likewise fitted to the OU ACF, $\rho_{OU}(h) = e^{-\lambda_{OU}h}$. For the estimations involving the correlation function we use $\lfloor \sqrt{N} \rfloor + 1 = 43$ lags, where $N = 1844$ is the total number of observations of the base signal X . As discussed above, the fractal nature of the \mathcal{BSS} process causes the estimate of λ to be lower than in the OU model ($\hat{\lambda} = 0.048$ compared to $\hat{\lambda}_{OU} = 0.149$). This is evidence that the α parameter helps in the rapid mean reversion around the very short lags, allowing λ to take care of the longer lags and hence fit the ACF better at these lags.

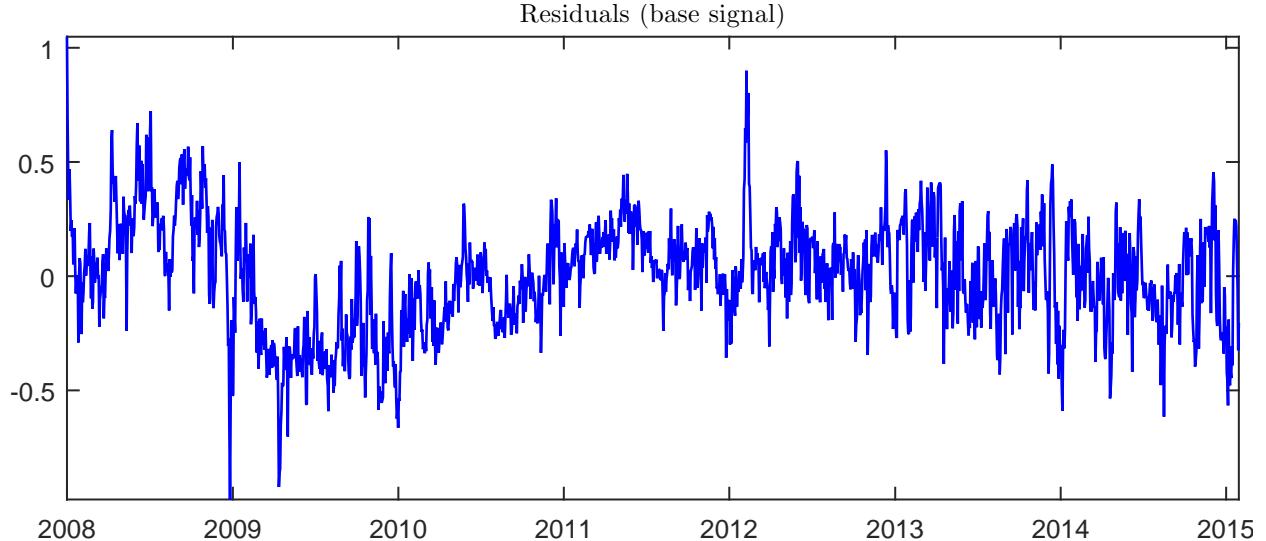


Figure 8: *Residuals of log-prices after filtering for seasonality and spikes. This signal will be modelled as a \mathcal{BSS} process.*

A.4 Forecasting the base signal

We here take a slightly closer look at the forecasting exercise of Section 6.3. Figure 9 plots the cumulative squared forecast errors (CSFE) for the three benchmark models as compared to the \mathcal{BSS} forecasts, when applied to the EEX data set. That is, we plot

$$\text{Excess CSFE}(x) = \text{CSFE}(x) - \text{CSFE}(\mathcal{BSS}), \quad x = \text{FBM}, \text{OU}, \text{RW},$$

so that positive values are in favor of the \mathcal{BSS} model. We use the first 400 observations for estimation and then forecast h periods into the future. After this we re-estimate the model the next day, forecast h periods into the future and repeat. The figure shows how the \mathcal{BSS} model fares best, but as the horizon h increases, the performance of the OU model gets closer to that of the \mathcal{BSS} . Both the fBm and RW models perform very poorly in general for the reasons discussed in Section 6.3.

B Estimating the degree of fractality in practice

B.1 Fractality estimation in the presence of spikes

To gauge the effect that possible jumps in the price process have on the estimators of the fractal index, we here conduct a small Monte Carlo study. Table 8 shows the bias of the various estimators, when outliers are introduced in a fractal process. We learn a few things from Table 8. Firstly, all estimators are virtually unbiased when no outliers are present. Secondly, when introducing outliers, all estimators become biased *downwards*; that is, in the presence of spikes, the estimate one obtains of the fractal index will indicate rougher paths than otherwise. Indeed, with only one outlier inserted, the estimators based on the power $p = 2$, i.e. the variogram, COF(2) and GEN(2),

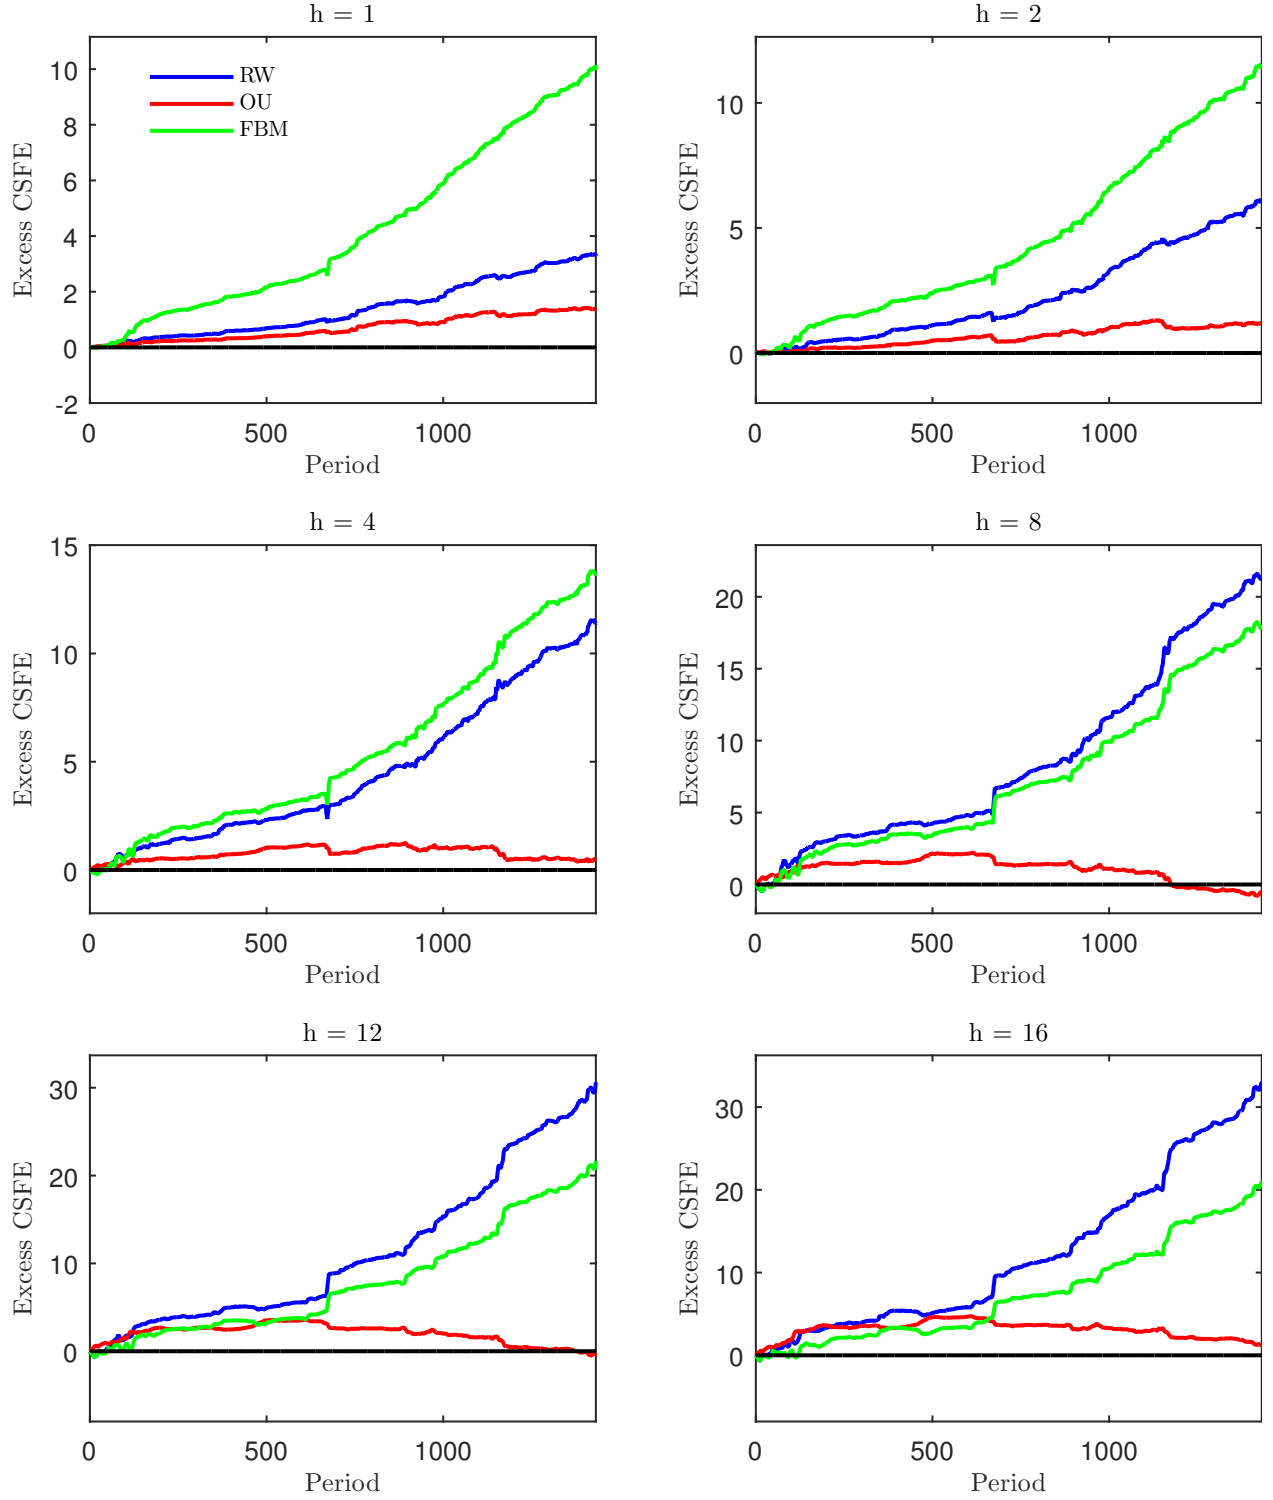


Figure 9: *Cumulated squared forecast errors for h periods ahead forecast using the three benchmark models, fBm, OU and a random walk (RW) as compared to the BSS model.*

Table 8: *Estimating the fractal index in the presence of outliers*

	Number of outliers									
	0	1	2	3	4	5	6	7	8	9
Madogram	-0.001	-0.020	-0.035	-0.049	-0.061	-0.072	-0.082	-0.091	-0.099	-0.106
Variogram	-0.001	-0.125	-0.165	-0.186	-0.199	-0.207	-0.213	-0.217	-0.221	-0.223
COF(1)	-0.008	-0.030	-0.048	-0.064	-0.078	-0.090	-0.100	-0.109	-0.118	-0.125
COF(2)	-0.005	-0.137	-0.175	-0.195	-0.206	-0.214	-0.219	-0.223	-0.226	-0.228
GEN(1)	-0.002	-0.018	-0.045	-0.069	-0.091	-0.087	-0.107	-0.109	-0.114	-0.124
GEN(2)	-0.004	-0.128	-0.195	-0.235	-0.266	-0.250	-0.264	-0.257	-0.249	-0.256
ABS	-0.015	0.005	0.005	0.005	-0.036	-0.075	-0.075	-0.087	-0.008	-0.030
DFA	0.013	-0.068	-0.101	-0.121	-0.136	-0.140	-0.147	-0.150	-0.154	-0.158
AGG	-0.009	-0.057	-0.057	-0.057	-0.110	-0.190	-0.190	-0.200	-0.078	-0.099

Bias of the fractal index estimators when outliers are placed in the data series. The simulated process, X , is a BSS process with $\alpha = -0.25$ and $\lambda = 1$ on $[0, 1]$ with $N = 500$ observations. See [Bennedsen et al. \(2015\)](#) for a method to simulate the BSS process. Outliers of size $\pm 5\text{std}(X_t)$ are then introduced equidistantly into the process. 5,000 Monte Carlo replications.

become noticeably biased. In contrast, the estimators based on $p = 1$, i.e. COF(1), GEN(1), ABS and the Madogram are significantly less biased, even for a moderate amount of spikes in the process. This robustness when $p = 1$ supports a similar finding in [Gneiting et al. \(2012\)](#). However, when spikes become more frequent, say more than 5 spikes (constituting 5% of the observations) even these estimators – except maybe ABS – also become biased. Given that most electricity price series arguably contain somewhat frequent spikes, this is a possible source of significant downwards bias. Indeed, this effect might explain the very low estimates of the fractal index found by previous authors, e.g. [Simonsen \(2002\)](#), [Norouzzadeh et al. \(2007\)](#) and [Erzgräber et al. \(2008\)](#).

B.2 Fractality estimation using low-frequency data

Besides spikes, another problem arises when we try to estimate the fractal index of (peak load) electricity prices. Electricity prices are recorded daily and are as such inherently sampled at a relatively low frequency. This is worrisome because the very definition of fractality is a local one; the fractal index is describing the very fine structure of the path of the stochastic process. This is also evident in the discussion in Section 3, where we saw that the definition of the fractal index is given in terms of in-fill asymptotics. In the same vein, the estimators such as the ones through the variogram in Equation (4.8) rely on the distance between observations, h , being very small. The question becomes whether it is possible at all to estimate the fractal index of a stochastic process observed so infrequently. Fortunately, the answer is in the affirmative. To see this, we conduct a Monte Carlo study where we simulate $N = 500$ observations of a fractal process on $[0, T]$ and estimate its fractal index. We do this with varying values of T and thus of the sampling frequency $dt = T/N$. Choosing $T = N$ and thinking of T as days, this mimicks daily data.

The results of the Monte Carlo study are seen in Table 9. Again a few conclusions emerge. As expected, the bias increases when we sample less frequently; further, this is a *downwards* bias, i.e. the paths of the price process are estimated to be *rougher* than what the underlying process truly

Table 9: *Estimating the fractal index with low frequency data*

	$\lambda = 0.1$			$\lambda = 0.5$			$\lambda = 1$		
$dt =$	0.1	0.5	1	0.1	0.5	1	0.1	0.5	1
Madogram	-0.001	-0.004	-0.010	-0.004	-0.034	-0.073	-0.011	-0.072	-0.137
Variogram	-0.001	-0.005	-0.010	-0.004	-0.034	-0.073	-0.011	-0.073	-0.137
COF(1)	-0.006	-0.007	-0.009	-0.006	-0.019	-0.050	-0.009	-0.050	-0.116
COF(2)	-0.004	-0.004	-0.005	-0.003	-0.017	-0.047	-0.006	-0.048	-0.113
GEN(1)	-0.003	-0.018	-0.039	-0.017	-0.097	-0.153	-0.040	-0.153	-0.203
GEN(2)	-0.005	-0.019	-0.040	-0.018	-0.098	-0.155	-0.041	-0.154	-0.204
ABS	-0.019	-0.073	-0.133	-0.075	-0.218	-0.252	-0.131	-0.253	-0.259
DFA	0.013	-0.003	-0.031	-0.002	-0.100	-0.152	-0.031	-0.152	-0.186
AGG	-0.013	-0.068	-0.129	-0.070	-0.214	-0.248	-0.127	-0.249	-0.256

Bias of the fractal index estimators when sampling infrequently and for different values of the mean-reversion parameter λ . The simulated process, X , is a BSS process with $\alpha = -0.25$ and $\lambda \in \{0.1, 0.5, 1, 2\}$ on $[0, T]$ with $N = 500$ observations and for $T \in \{50, 250, 500\}$. This corresponds to sampling the process at $dt \in \{0.1, 0.5, 1\}$ (days). 5,000 Monte Carlo replications.

is. This behavior is amplified when we increase the value of the mean reversion parameter λ , as also seen in the table. The reason for this is that strong mean reversion can have the same effect as a negative fractal index when the data is observed infrequently. Indeed, a negative fractal index causes the increments of the process to be negatively correlated and mean reversion mimicks this. The upshot is that when we sample infrequently (e.g. $dt = 1$), low fractal indices can simply be a product of strong mean reversion. As in the case of spikes, investigated in the previous section, this is a possible further source of (downward) bias in previous studies. Fortunately, when estimating the mean reversion parameter λ in the data, we often find it on the order $\hat{\lambda} \sim 0.01 - 0.10$; see e.g. [Benth et al. \(2012\)](#) and also Table 5 in Section 6 where the range $0.011 - 0.048$ is found for the data in Table 1. Looking again at Table 9 we notice that several of the estimators, such as $COF(1)$ and $COF(2)$, are not very biased in this situation, even for daily data ($dt = 1$). This makes us confident that it is, after all, feasible to estimate the fractal index, even using infrequently sampled data. This also holds for the DFA estimator, but not for the ABS and AGG estimators.

B.3 Lessons from the simulation experiment

The two previous sections above show that one should be careful when estimating the fractal index of electricity prices. In particular, spikes in the time series will cause the estimators to be downward biased. The same is true for data sampled with low frequency, especially when the mean reversion parameter, λ , is large in magnitude. We conclude that it is important to filter the prices for spikes before estimating the value of the fractal index. Further, it is preferable to use an estimator that is not too biased in the low frequency environment. Given the above and the fact that there is a CLT attached to it, we recommend using the $COF(2)$ estimator. However, it is prudent to use several estimators, including ones that are robust to outliers, such as the ones using $p = 1$; we recommend employing several estimators and only trusting the results when most of them agree.

Research Papers

2013



- 2015-25: Bent Jesper Christensen and Rasmus T. Varneskov: Medium Band Least Squares Estimation of Fractional Cointegration in the Presence of Low-Frequency Contamination
- 2015-26: Ulrich Hounyo and Rasmus T. Varneskov: A Local Stable Bootstrap for Power Variations of Pure-Jump Semimartingales and Activity Index Estimation
- 2015-27: Rasmus Søndergaard Pedersen and Anders Rahbek: Nonstationary ARCH and GARCH with t-distributed Innovations
- 2015-28: Tommaso Proietti and Eric Hillebrand: Seasonal Changes in Central England Temperatures
- 2015-29: Laurent Callot and Johannes Tang Kristensen: Regularized Estimation of Structural Instability in Factor Models: The US Macroeconomy and the Great Moderation
- 2015-30: Davide Delle Monache, Stefano Grassi and Paolo Santucci de Magistris: Testing for Level Shifts in Fractionally Integrated Processes: a State Space Approach
- 2015-31: Matias D. Cattaneo, Michael Jansson and Whitney K. Newey: Treatment Effects with Many Covariates and Heteroskedasticity
- 2015-32: Jean-Guy Simonato and Lars Stentoft: Which pricing approach for options under GARCH with non-normal innovations?
- 2015-33: Nina Munkholt Jakobsen and Michael Sørensen: Efficient Estimation for Diffusions Sampled at High Frequency Over a Fixed Time Interval
- 2015-34: Wei Wei and Denis Pelletier: A Jump-Diffusion Model with Stochastic Volatility and Durations
- 2015-35: Yunus Emre Ergemen and Carlos Velasco: Estimation of Fractionally Integrated Panels with Fixed Effects and Cross-Section Dependence
- 2015-36: Markku Lanne and Henri Nyberg: Nonlinear dynamic interrelationships between real activity and stock returns
- 2015-37: Markku Lanne and Jani Luoto: Estimation of DSGE Models under Diffuse Priors and Data-Driven Identification Constraints
- 2015-38: Lorenzo Boldrini and Eric Hillebrand: Supervision in Factor Models Using a Large Number of Predictors
- 2015-39: Lorenzo Boldrini and Eric Hillebrand: The Forecasting Power of the Yield Curve, a Supervised Factor Model Approach
- 2015-40: Lorenzo Boldrini: Forecasting the Global Mean Sea Level, a Continuous-Time State-Space Approach
- 2015-41: Yunus Emre Ergemen and Abderrahim Taamouti: Parametric Portfolio Policies with Common Volatility Dynamics
- 2015-42: Mikkel Bennedsen: Rough electricity: a new fractal multi-factor model of electricity spot prices