AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Bennedsen, M. (2017). A rough multi-factor model of electricity spot prices. *Energy Economics*, 63, 301-313.
<https://doi.org/10.1016/j.eneco.2017.02.007>

Publication metadata

Title:	A rough multi-factor model of electricity spot prices
Author(s):	Bennedsen, M.
Journal:	<i>Energy Economics</i>
DOI/Link:	10.1016/j.eneco.2017.02.007
Document version:	Accepted manuscript (post-print)

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 recognize 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

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

A rough multi-factor model of electricity spot prices

Mikkel Bennedsen

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

Abstract

We introduce a new continuous-time 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, roughness, 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 rough model. Further, through the multi-factor approach we account for seasonality and spikes before estimating – and making inference on – the degree of roughness. This is novel in the literature and we present simulation evidence showing that these precautions are crucial for accurate estimation. Lastly, we estimate our model on recent data from six European energy exchanges and find statistical evidence of roughness in five out of six markets. As an application of our model, we show how, in these five markets, a rough component improves short term forecasting of the prices.

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

JEL: 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. 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 of very inelastic demand, prices 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 Przybyłowicz \(2000\)](#), [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 OU 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) rough characteristics (often termed *antipersistence* in the literature). Specifically, a number of studies show that the *roughness index* 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 the for Czech market in [Kristoufek and Lunackova \(2013\)](#). Although roughness/antipersistence 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 rough model for the spot prices. A notable exception is [Rypdal and Løvsletten \(2013\)](#) who compared a rough model based on a multifractal random walk with a (non-rough) model of the 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 rough approach. In short, we propose a rough multi-factor model of electricity spot prices, i.e. a multi-factor model where the base signal, X , is modelled by a rough 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 facts (i)-(iv)) or a rough/antipersistent model (accounting for stylized fact (v) only). What is more, our setup accomodates *testing*, in the rigorous statistical sense, for the presence of roughness in the price series; this is a novel contribution in the literature, as previous studies have focused on point estimates of the roughness 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 rough process (the alternative). Additionally, the multi-factor setup allows us to account for spikes *before* estimating and making inference on the roughness index; this is crucial because spikes can be detrimental to the estimation, a fact we will illustrate in the Appendix of the present paper.

As an application of the model we consider a simple forecasting exercise, showing that including a rough 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 rough 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 details the concept of roughness as it relates 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 roughness 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 roughness 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.¹ 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 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 selection². Both tests reject the null of a unit root in the raw prices series as well as in the log-prices, at all reasonable significance levels. That is, we find evidence

¹Unlike 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.

²As implemented in the MFE Toolbox of Kevin Sheppard, http://www.kevinsheppard.com/MFE_Toolbox.

Table 1: *Data summary*

Abbreviation	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

of stationarity in the de-seasonalized prices and we will therefore proceed to model all price series as stationary processes in the following.

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
nObs	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
Skew	0.558	2.569	0.725	−0.193	0.218	0.018	0.052	0.747	0.349	0.060	0.101	0.421
Kurt	12.549	78.412	616.925	653.490	21.017	9.474	23.412	15.590	57.139	115.594	6.610	10.128
JB	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). “nObs” denotes the number of observations, while “JB” is the Jarque-Bera test statistic.

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. Roughness in electricity prices

In this section, we explain how roughness relates to one-dimensional time series and hence to electricity prices. Let X be a zero-mean, real-valued time series with stationary increments, and define its (second order) *variogram*:

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

The concept of roughness refers to a particular behavior of the variogram for $h \approx 0$. That is, if it holds that

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

for a constant $c_2 > 0$ and some $\alpha \in (-1/2, 0)$, then we say that X is *rough*. In the following, we will therefore call α the *roughness index* or *roughness parameter*. The reason for this terminology is that when $\alpha < 0$, X will have sample paths which are rougher than the sample paths of the Brownian motion (the Brownian motion has $\alpha = 0$). What is meant by this is that, under quite general assumptions, a process X with roughness parameter α will be Hölder continuous of order γ for all $\gamma \in (0, \alpha + 1/2)$.³ It should be noted, that one usually allows for α to take values in $(-1/2, 1/2)$; the values where $\alpha > 0$ will correspond to processes with trajectories *smoother* than those of the Brownian motion. However, as we will see, the electricity data considered in this paper will always give rise to estimates of α which are negative and our focus will therefore be on the rough case, i.e. on $\alpha < 0$. A last thing to note is the well known fact that processes with roughness index $\alpha \neq 0$ are necessarily *not semimartingales*; see Proposition 2.2. in [Bennedsen \(2016\)](#) for a proof of this statement and Section 5 below for a discussion of the implications of using a rough statistical model for electricity prices.

A typical example of a Gaussian process with non-trivial roughness index is the fractional Brownian motion (fBm) with Hurst parameter $H \in (0, 1)$, corresponding to $\alpha = H - \frac{1}{2}$. Here $\alpha = 0$ (so $H = \frac{1}{2}$) implies that the fBm is a standard Brownian motion. To illustrate the concept of roughness visually, several simulated fBm paths for different values of α are shown in Figure 2. The differing roughness of the paths, when α changes, is evident: we see how a low α corresponds to very rough paths while positive values results in increasingly smooth paths. Relating this 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.

4. A rough multi-factor modelling framework

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

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

³See, e.g., Proposition 2.1. in [Bennedsen \(2016\)](#) for a precise statement and proof of this result.

Figure 2: Simulated paths of an fBm for varying values of the roughness 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 .

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 the data we work with are peak load prices, which are averages 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 roughness index α) than what we will see in the following. That is, the conclusions drawn in this paper concerning roughness 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, most likely, be more realistic. However, since our focus is on the rough 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 mentioned in the introduction, several earlier studies have found evidence of rough characteristics in electricity price series; we therefore propose to model the base component X as a rough 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 Proposition 2.1. of Bennedsen et al. (2015) the roughness index of X , as defined in Section 3, is equal to α , and the estimators we consider later in this paper will therefore be valid for this model. 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 BSS process X are in order. First off, (4.3) 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 roughness 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 $\mathcal{BSS}$ 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 rough $\mathcal{BSS}$ model.

Secondly, (4.3) implies that the autocorrelations of the process decay asymptotically with exponential rate λ . The parameter λ is given in inverse units of time (days) and λ^{-1} is termed the *characteristic time scale* of the process (e.g., Rypdal and Løvsletten, 2013). An interesting question is to assess whether the electricity price data actually display a characteristic time scale; following the suggestion of an anonymous referee, we explore this question further in Section 4.2.1 below. In particular, we propose a simulation-based test for whether or not a characteristic time scale exists in the data.

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 6.3 we also find evidence of non-Gaussianity. The Section 4.2.2 below details how the $\mathcal{BSS}$ framework can accomodate this through the stochastic volatility component σ .

4.2.1. Models with and without characterisic time scale in their autocorrelations

In the literature on statistical properties of electricity prices, there are opposing views on which mathematical frameworks are adequate for modelling the data. On the one hand, authors such as Weron and Przybyłowicz (2000) and Simonsen (2002) argue for a rough (“antipersistent”) model without characteristic time scale in its autocorrelation function, such as the fractional Brownian motion. On the other hand, the OU-based models (e.g., Schwartz, 1997; Cartea and Figueroa, 2005; Meyer-Brandis and Tankov, 2008; Benth et al., 2012) implicitly assumes the presence of a characteristic time scale; however, very few studies have presented explicit evidence that such a time scale should actually be included in the model. A notable exception is Rypdal and Løvsletten (2013), where the authors find evidence that the Nordic energy market do display a characteristic time scale; however, Rypdal and Løvsletten (2013) are unable to conclude whether this finding is simply caused by seasonal effects in the data. The multi-factor setup allows us to bypass such doubts; since we filter out the seasonality (and spikes) in a preliminary step, any characteristic time scale found in the data will be in the base signal, i.e. it will be an inherent feature of the underlying stochastic process driving the prices.

One of the advantages of the $\mathcal{BSS}$ framework is that it conveniently and parsimoniously allows for *both* roughness (antipersistence) *and* a characteristic time scale in the autocorrelations of the model. The former is captured by the parameter α and the latter by λ . This is a flexible and powerful approach to decoupling the small-scale (roughness) behavior of the process from the long-term (memory) properties; see Bennedsen et al. (2016) for a further discussion of this decoupling and how it relates to financial time series. Loosely speaking, in the context of our proposed $\mathcal{BSS}$ model (4.3), $\lambda = 0$ corresponds to no characteristic time

scale in the model, while $\lambda > 0$ corresponds to a characteristic time scale of λ^{-1} days.⁴

As argued above, the $\mathcal{BSS}$ model allows us to test for the presence of roughness (antipersistence) by testing the null hypothesis $H_0 : \alpha = 0$ and we will do so in the empirical section below. However, given the discussion of this section, another interesting question is to investigate whether the electricity prices display a characteristic time scale in their autocorrelations as well, such as was found in Rypdal and Løvsetten (2013), or whether a rough model without such a characteristic scale is more adequate.

We would thus like to test the null hypothesis $\lambda = 0$, but as this value is not permitted in the model (4.3), we consider an alternative way to test whether the data display a characteristic time scale.⁵ Specifically, we propose a Monte Carlo approach where we assess the distribution of our proposed estimator of λ , when the underlying process does *not* possess a characteristic time scale in its autocorrelations. We then compare the estimated value of λ , obtained from electricity data, to the Monte Carlo distribution of $\hat{\lambda}$ obtained under the null of no characteristic time scale. In this way, the test will gauge whether an estimate of λ obtained from the data really corresponds to a characteristic scale in the autocorrelations, or whether the observed scale is “spurious”, i.e. whether it might as well, by chance, have arisen from a process without any characteristic scale. To be precise, our proposed test proceeds as follows.

We simulate under the null of the process having no characteristic time scale (i.e. “ $\lambda = 0$ ”). Since the fBm is the “canonical” rough model without such a time scale (it is *self-similar*), we choose this model as our data generating process under the null. Then we estimate λ using a method-of-moments procedure, which we explain in Section 4.4.2 below. By performing these Monte Carlo simulations a large number of times, we can infer the distribution of the estimator of λ under the null that the process has no characteristic autocorrelation scale (here taking to mean that X is an fBm). Then, when we estimate λ from electricity price series we can perform inference by comparing the estimated value to the (Monte Carlo) distribution of $\hat{\lambda}$ under the null. For instance, if the estimate obtained from the data is larger than 95% of the Monte Carlo estimates from the simulations, we reject the null at a 5% significance level, and conclude that the data indeed appear to possess a characteristic time scale in its autocorrelations. Similarly, a P-value can be calculated from

⁴Strictly speaking, the $\mathcal{BSS}$ process X in (4.3) does not exist for $\lambda = 0$, as the integrand should be square integrable for the integral to be defined in the Itô sense. However, if one considers the *modified* $\mathcal{BSS}$ process

$$X_t^\# := \int_{-\infty}^t \left((t-s)^\alpha e^{-\lambda(t-s)} - (-s)_+^\alpha e^{-\lambda(-s)} \right) \sigma_s dW_s, \quad t \geq 0,$$

where $y_+ := \max(0, y)$ denotes the positive part of $y \in \mathbb{R}$, this process does exist for $\lambda = 0$. Indeed, this process reduces (up to a constant factor) to a *Type I fractional Brownian motion* (Marinucci and Robinson, 1999) when $\lambda = 0$. Note further, that the *increments* of processes X and $X^\#$ have similar behavior; in particular their variograms (B.1) coincide. When $\lambda > 0$ it is also the case that $\{X_t^\#\}_{t \geq 0} = \{X_t - X_0\}_{t \geq 0}$. Although the modified $\mathcal{BSS}$ process $X^\#$ is non-stationary, and therefore not suited for modelling of the base signal, cf. the discussion in Section 2, it does highlight the close connection between the fractional Brownian motion and the $\mathcal{BSS}$ process when $\lambda \approx 0$.

⁵Although $\lambda = 0$ is permitted for the modified $\mathcal{BSS}$ process $X^\#$, this value for λ is still on the boundary of the parameter space, making testing non-standard. We therefore proceed with the simulation-based test as described in the following.

the relative position of the estimated value of λ from the data, compared to the Monte Carlo estimates; for instance, if the estimated value from the data is greater than 98% of the Monte Carlo estimates, but smaller than the remaining 2%, the P-value of the null will be 2%.

4.2.2. Stochastic volatility: marginal distribution of the base component

The stochastic volatility (SV) process, σ , 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 6.3) 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 exisistence 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 ([Barndorff-Nielsen and Halgreen, 1977](#)). In Section 6.3 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 .

4.2.3. An alternative rough component

A perhaps more well-known choice for a stationary rough 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 roughness 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 Løvsetten (2013). However, we wish to employ a more flexible rough 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 autocorrelation function is available in closed form and given below in Equation (4.9). In Section 6.4 we will compare the two models when investigating if a rough 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 with 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 rough 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 rough framework to future studies.

For simplicity, we assume that X and Y are independent; this implies that the de-seasonalized price, $S - \Lambda = X + Y$, is (strictly) stationary.

4.4. Estimation of the rough 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 $\{S_t - \Lambda_t\}_{t \geq 0} = \{X_t + Y_t\}_{t \geq 0}$. 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 Appendix A.2 where we study the EEX data set in detail. Finally, after subtracting the jump process we are left with the base signal $\{X_t\}_{t \geq 0} = \{S_t - \Lambda_t - Y_t\}_{t \geq 0}$, which we proceed to model as the $\mathcal{BSS}$ process (4.3).

This procedure of estimating the different components separately is of course not optimal; at a minimum one loses efficiency and estimates can plausibly become biased as well. An alternative approach is to try to estimate the model in one go; this was recently done using particle Markov chain Monte Carlo methods in Brix et al. (2015). However, these methods are not straightforwardly applicable to the rough model of this paper; further, in simulations (not reported here for brevity) we found good performance of the piece-wise estimation procedure. For these reasons, we prefer the simple approach of estimating the signals separately and will do so in the following.

4.4.1. Estimating the roughness index of time series of electricity prices

Many estimators of the roughness index exist and the choice about which one to use is up for debate. Although most studies consider only one particular estimator, we recommend the use several of them to check if they agree. For this reason, Appendix B contains a brief review of a number of different estimators of α and in the empirical section below we will apply them all to the data. We do this for two reasons; firstly, we use them as a robustness check to see whether the 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 our purposes, we face two major obstacles when estimating the roughness index of electricity prices; (i) jumps (spikes) in the price process and (ii) relatively low frequency (daily) data. In Appendix C.1 and Appendix C.2 we will look closer at how the estimators behave under (i) and (ii), respectively. We find that the former causes a downward bias in the estimation of the roughness 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 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.

Let us briefly give some details on one of these estimators of the roughness index, namely the *change-of-frequency* (COF) estimator of Corcuera et al. (2013). The reasons for going into detail with this particular estimator, is that it comes with a central limit theorem (CLT),

so that inference can be conducted easily.⁶ Suppose we have $n \in \mathbb{N}$ equidistant observations of a process X observed in a fixed time interval $[0, 1]$, say. The COF statistic compares the *power variations* of the (second order) differences of the observations at two different frequencies:

$$\text{COF}_{p,n} = \frac{\sum_{k=5}^n |X_{k/n} - 2X_{(k-2)/n} + X_{(k-4)/n}|^p}{\sum_{k=3}^n |X_{k/n} - 2X_{(k-1)/n} + X_{(k-2)/n}|^p},$$

where $p > 0$ is a power parameter to be chosen by the researcher. [Corcuera et al. \(2013\)](#) show that when X is a rough process with fractal index $\alpha \in (-1/2, 1/2)$ in the sense of Section 3, then as $n \rightarrow \infty$

$$\text{COF}_{p,n} \rightarrow 2^{(2\alpha+1)p/2} \quad \text{in probability.} \quad (4.8)$$

This asymptotic regime, where the time period is fixed but the number of observations goes to infinity (so that the time between observations goes to zero), is often called *in-fill asymptotics*. Using (4.8), a straight forward estimator of α is

$$\hat{\alpha}_p := \frac{\log_2(\text{COF}_{p,n})}{p} - \frac{1}{2}, \quad p > 0,$$

which we will simply refer to as $\text{COF}(p)$ in the following. Both $p = 1$ and $p = 2$ can be found recommended in the literature; in what follows we will use both values. Note, however, that a feasible CLT is only available when $p = 2$, which will make this the most important value for us in the empirical section. We refer to [Corcuera et al. \(2013\)](#) for the details on the CLT.

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., [Barndorff-Nielsen, 2012](#))

$$\rho(h) = \text{Corr}(X_{t+h}, X_t) = \frac{2^{-\alpha+\frac{1}{2}}}{\Gamma(\alpha+1/2)} (\lambda h)^{\alpha+\frac{1}{2}} K_{\alpha+\frac{1}{2}}(\lambda h), \quad h \in \mathbb{R}, \quad (4.9)$$

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 employ a method-of-moments procedure by fitting (4.9), via non-linear least squares, to the empirically observed autocorrelations of the electricity (base signal) prices. In simulations, we found this two-step method of first estimating α and then plugging this estimate $\hat{\alpha}$ into (4.9) to estimate λ via method-of-moments, to work very well in a setup similar to our empirical applications below (simulation results not reported here, but available upon request).

⁶For a simulation-based, and therefore more computationally heavy, approach to inference on the roughness index see [Bennedsen \(2016\)](#).

5. Implications of a rough model of electricity prices

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

5.1. Arbitrage

It is well known that models with roughness index $\alpha \neq 0$, such as the fBm-OU model (4.7) with $H \neq 1/2$ and the rough $\mathcal{BSS}$ 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 rough 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 present.

5.2. Forecasting

Since roughness refers to the small scale behavior of the time series under consideration, it is plausible that including a rough 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 is the spike signal Y , given a model for this. However, as mentioned 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 6.4 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 open 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 (4.9). 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.)

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

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

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

To use this method in practice, we estimate the parameters of the $\mathcal{BSS}$ model, plug them into the functional form of the ACF in (4.9), and use this in the above to arrive at the best linear predictor.

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 , the following approximation is justified

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

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 6.4 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 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).

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.

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 roughness index α , which are negative across the board, indicating that the time series indeed do display a degree of roughness. 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 rough $\mathcal{BSS}$ 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 rough characteristics in their base signals.

To further check this fact, we employ nine different estimators of α , which are discussed in Appendix B; the results are seen in Table 6 which corroborates the findings of this section. 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. The DFA estimator has traditionally been popular in the literature; it is therefore evident that it is important to filter out spikes before estimation of the roughness index 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.

In Figure 3 we plot the estimates of α for all markets using the COF(2) estimator, as well as their 95% confidence bands. This is done to both the time series including spikes (blue squares) and the base signal (red crosses). A few conclusions emerge from this. First off, the estimates from two of the six markets (Powernext and GME) become noticeably less negative after filtering out the spikes, strengthening the claim that spikes might cause downward bias in the roughness estimates (see also Appendix C.1). Even more salient is the observation that the standard deviation of the estimator is generally huge when one does not filter out spikes before estimation, leading to too wide confidence intervals, and thus to

Figure 3: *Estimation of the roughness index of the spot prices from various energy exchanges. Estimation is done on the raw log-prices (blue squares) and on the base signal, i.e. after filtering out the seasonal component and the spike signal (red crosses). The solid lines correspond to 95% confidence intervals.*

(possibly) wrong inference. For instance, $\alpha = 0$ is contained in the confidence intervals for Nord Pool, Powernext, GME, and almost for EEX as well, when the analysis is done on the data including spikes. Conversely, when the analysis is done on the base signal, only the confidence bands from the Nord Pool market contain zero.

The above shows that accounting for spikes prior to estimation of α is crucial and the rough 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 rough/antipersistent 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*.

6.2. Characteristic time scale of electricity price series

Turning again to Table 5, we also see estimates of the mean reversion parameter λ from both the **BSS** model and the OU model. The value of $1/\hat{\lambda}$ is also reported (in days); this is the “characteristic time scale” of the process, i.e. a measure of how fast the autocorrelations decay (exponentially) to zero. In the OU model, $1/\lambda$ corresponds to the *half-life* of shocks to the process, i.e. the expected waiting time until the effect of a shock to the process is halved. For the **BSS** process with ACF (4.9) this interpretation is only valid approximately in finite samples. For the base signal, we see that the characteristic time scale for the **BSS** model varies from approximately 20 to 90 days. These numbers are significantly smaller for the OU model; the reason for this is, that the roughness in the **BSS** model 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 **BSS** model to better fit the ACF at the longer lags as compared to the OU model. In other words, not accounting for the roughness in electricity prices results in severe under estimation of the characteristic time scale, leading the researcher to conclude that shocks to the process die out more quickly than what is the case in actuality.

Table 5: *Estimating and testing for roughness*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
$\hat{\lambda}_{BSS}$	0.120	0.043	0.034	0.019	0.021	0.052
$1/\hat{\lambda}_{BSS}$ (days)	8.320	23.332	29.595	52.749	48.493	19.224
P-value	0	0	0.0002	0.0578	0.0116	0
$\hat{\lambda}_{OU}$	0.296	0.056	0.218	0.463	0.192	0.131
$1/\hat{\lambda}_{OU}$ (days)	3.376	17.768	4.581	2.159	5.216	7.636
P-value	0	0	0	0	0	0

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
$\hat{\lambda}_{BSS}$	0.048	0.026	0.033	0.031	0.011	0.019
$1/\hat{\lambda}_{BSS}$ (days)	20.960	38.723	30.553	32.067	88.295	51.584
P-value	0	0	0	0.0002	0.0832	0.0052
$\hat{\lambda}_{OU}$	0.149	0.037	0.105	0.144	0.144	0.074
$1/\hat{\lambda}_{OU}$ (days)	6.722	27.085	9.549	6.958	6.957	13.517
P-value	0	0	0	0	0	0

Estimating the roughness index, α , and testing $H_0 : \alpha = 0$ for various data sets. Panel A is using the standard approach of directly estimating the roughness 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 $H_a : \alpha < 0$. In the bottom half of the table the estimates of λ are reported from both the BSS model and the OU model. Here, the P-values refer to the test for a characteristic autocorrelation time scale, as explained in Section 4.2.1. For this test, 5 000 Monte Carlo replications of an fBm with Hurst index $H = \hat{\alpha} + 1/2$ where used to approximate the distribution of $\hat{\lambda}$ under the null.

Table 6: *Estimating the roughness 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 roughness index α for various data sets and various estimators. The estimators are reviewed in [Appendix B](#). 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.

Table 5 also contains P-values for the test of the presence of a characteristic time scale, detailed in Section 4.2.1. Looking at the base signal in Panel B, the P-values are all very small. In fact, for the OU model, the P-values are all equal to zero. All this tells us, however, is that a model without *both* roughness and a characteristic time scale in the autocorrelations is a poor model for the base signal (in essence, this test is whether a Brownian motion would be an adequate model for the base signal). More interestingly, the P-values are also very low when we consider the $\mathcal{BSS}$ model; only for one market (APX UK) is the P-value above 5%. This shows, that even after including a rough/antipersistent component, it is still important to also allow for a characteristic time scale in the model.

6.3. Marginal distribution of the base signal

As explained in Section 4.2.2, 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. Figure 4 shows 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.⁷ To examine the fit in the tails better, we present QQ-plots in the right column of Figure 4: here we again see that, for most markets, the NIG distribution is 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 4.2.2, the $\mathcal{BSS}$ framework is able to accurately capture the empirically observed marginal density of the extracted base signal of electricity prices.

6.4. Forecasting

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

⁷The NIG distribution were fitted using a maximum likelihood procedure.

Figure 4: *Left: Empirical marginal distribution of the base signal (blue solid line) superimposed with a fitted Gaussian density (red solid line) and the corresponding fitted NIG density (green dashed line). Right: QQ-plots; the solid red curve denotes the Gaussian distribution and the green dashed line denotes the NIG distribution. If the data fits one of these distributions, the corresponding curve will lie on the black diagonal line.*

Table 7: *Forecasting the base signal*
RMSFE 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 forecast errors (RMSFE) relative to the $\mathcal{BSS}$ model. Values less than one favors the $\mathcal{BSS}$ model, while values greater than one favors the alternative method.

poorly in general. We believe that this is because the fBm lacks any characteristic time scale; in particular, the forecasts using (5.2) do not account for mean reversion. Secondly, the $\mathcal{BSS}$ forecast method generally fares better than the (non-rough) OU method, when the market in question display roughness, i.e. for all markets except Nord Pool (where we could not reject $H_0 : \alpha = 0$, see Table 5). Putting these two observations together, it seems to suggest that, at least for forecasting purposes, it is important to choose a model which is *both* rough *and* has a characteristic time scale in its autocorrelations. The $\mathcal{BSS}$ process is, in this regard, a very good candidate as a model for the base signal.

7. Conclusion

This article has conducted an in-depth investigation of the apparent roughness of time series of electricity spot prices. Similar to other authors we find evidence of rough (i.e. “antipersistent”) characteristics in electricity spot prices: in particular, in five out of the six markets studied, we reject the null hypothesis of a zero roughness 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 roughness index by filtering the prices for seasonality and spikes *before* estimating the degree of roughness 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 roughness 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 rough $\mathcal{BSS}$ model; a very interesting avenue of future research would be to include a rough component in a more realistic forecasting setup, such as the one considered in [Weron \(2014\)](#).

Although the OU model was found to be inadequate, the results of this paper show that it is important for a mathematical model of (the base signal of) electricity prices to contain a characteristic time scale in the autocorrelation function implied by the model. Indeed, when we performed the statistical tests for such a time scale, the P-values were all very low (only one market had a P-value above 5%, namely the APX UK market with a P-value of 8.3%). In particular, this finding rules out *self-similar* models, such as the fractional Brownian motion, as a model of the base signal. It also extends a finding in [Rypdal and Løvsletten \(2013\)](#), where the authors found a characteristic time scale in the Nord Pool market, but were unable to conclude whether this was simply due to seasonal effects in the data. In contrast, the multi-factor approach taken here, where we in an initial step filter out seasonality and spikes, allows us to conclude that even after controlling for jumps and seasonal effects, a characteristic time scale still exists in electricity price series. What is more, the $\mathcal{BSS}$ process seems to offer a particular convenient and tractable model which features both roughness and a characteristic time scale in its autocorrelations.

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 roughness hold a fortiori: we find even more extreme (negative) values of the roughness index α in this case.

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

Acknowledgements

I would like to thank the editor and two anonymous referees for valuable comments, which resulted in considerable improvements of the paper. I would also like to thank Professor 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.

- Abramowitz, M., Stegun, I. A., 1972. Handbook of mathematical functions with formulas, graphs, and mathematical tables, 10th Edition. Vol. 55. United States Department of Commerce.
- Barndorff-Nielsen, O. E., 2012. Notes on the gamma kernel. Thiele Research Reports (03).
- Barndorff-Nielsen, O. E., Benth, F. E., Veraart, A. E. D., 2013. Modelling energy spot prices by volatility modulated Lévy-driven Volterra processes. *Bernoulli* 19 (3), 803–845.
- Barndorff-Nielsen, O. E., Corcuera, J. M., Podolskij, M., 2011. Multipower variation for Brownian semistationary processes. *Bernoulli* 17 (4), 1159–1194.
- Barndorff-Nielsen, O. E., Halgreen, C., 1977. Infinite divisibility of the hyperbolic and generalized inverse gaussian distributions. *Z. Wahrscheinlichkeitstheorie verw. Gebiete* 38, 309–311.
- Barndorff-Nielsen, O. E., Schmiegel, J., 2007. Ambit processes: with applications to turbulence and tumour growth. In: *Stochastic analysis and applications*. Vol. 2 of Abel Symp. Springer, Berlin, pp. 93–124.
- Barndorff-Nielsen, O. E., Schmiegel, J., 2009. Brownian semistationary processes and volatility/intermittency. In: *Advanced financial modelling*. Vol. 8 of Radon Ser. Comput. Appl. Math. Walter de Gruyter, Berlin, pp. 1–25.
- Bennedsen, M., 2016. Semiparametric inference on the fractal index of Gaussian and conditionally Gaussian time series data, working paper available at arXiv:1608.01895.
- Bennedsen, M., Lunde, A., Pakkanen, M. S., 2015. Hybrid scheme for Brownian semistationary processes, working paper available at arXiv:1507.03004.
- Bennedsen, M., Lunde, A., Pakkanen, M. S., 2016. Decoupling the short- and long-term behavior of stochastic volatility, working paper available at arXiv:1610.00332.
- Benth, F. E., Kallsen, J., Meyer-Brandis, T., 2007. A non-Gaussian Ornstein–Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance* 14 (2).
- Benth, F. E., Kiesel, R., Nazarova, A., 2012. A critical empirical study of three electricity spot price models. *Energy Economics* 34 (5), 1589–1616.
- Brix, A. F., Lunde, A., Wei, W., 2015. A generalized Schwartz model for energy spot prices - estimation using a particle MCMC method, CREATES Working paper 2015-46.
- Cartea, A., Figueroa, M. G., 2005. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 12 (4), 313–335.
- Corcuera, J. M., Hedevarang, E., Pakkanen, M. S., Podolskij, M., 2013. Asymptotic theory for Brownian semistationary processes with application to turbulence. *Stochastic Process. Appl.* 123 (7), 2552–2574.

- Delbaen, F., Schachermayer, W., 1994. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen* 300 (1), 463–520.
- Erzgräber, H., Strozzi, F., Zaldívar, J., Touchette, H., Gutiérrez, E., Arrowsmith, D. K., 2008. Time series analysis and long range correlations of nordic spot electricity market data. *Physica A* 287, 6567–6574.
- Gatheral, J., Jaisson, T., Rosenbaum, M., 2014. Volatility is rough. Working paper.
- Geman, H., Roncoroni, A., 2006. Understanding the fine structure of electricity prices. *Journal of Business* 79 (3).
- Gneiting, T., Sevcikova, H., Percival, D. B., 2012. Estimators of Fractal Dimension: Assessing the Roughness of Time Series and Spatial Data. *Statistical Science* 27 (2), 247–277.
- Grimmett, G., Stirzaker, D., 2001. *Probability and Random Processes*, 3rd Edition. Oxford University Press.
- Hayfavi, A., Talasli, I., 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., Lunackova, P., 2013. Long-term memory in electricity prices: Czech market evidence. *Czech Journal of Economics and Finance* 63 (5), 407–424.
- Marinucci, D., Robinson, P. M., 1999. Alternative forms of fractional Brownian motion. *Journal of Statistical Planning and Inference* 80 (1–2), 111–122.
- Matteo, T. D., 2007. Multi-scaling in finance. *Quantitative Finance* 7 (1), 21–36.
- Meyer-Brandis, T., Tankov, P., 2008. Multi-factor jump-diffusion models of electricity prices. *International Journal of Theoretical and Applied Finance* 11 (5), 503–528.
- Norouzzadeh, P., Dullaert, W., Rahmani, B., 2007. Anti-correlation and multifractal features of Spain electricity spot market. *Physica A* 380, 333–342.
- Nuzman, C. J., Poor, V. H., 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., Buldyrev, S. V., Havlin, S., Simons, M., Stanley, H. E., Goldberger, A. L., 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., Løvsletten, O., 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., Teverovsky, V., Willinger, W., 1995. Estimators for long-range dependence: An empirical study. *Fractals* 3, 785–798.
- Veraart, A. E. D., Veraart, L. A. M., 2014. Modelling electricity day-ahead prices by multivariate Lévy semistationary processes. In: Benth, F. E., Kholodnyi, V. A., Laurence, P. (Eds.), *Quantitative Energy Finance*. Springer, New York, pp. 157–188.
- 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., Przybyłowicz, B., 2000. Hurst analysis of electricity price dynamics. *Physica A* 286, 462–468.

Appendix 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 price series was shown in the top left corner of Figure 1.

Appendix A.1. De-seasonalizing the data

We do OLS using the seasonal function (4.2) on the (locally smoothed) log-prices as explained in Section 4.1. The resulting fit is seen in Figure A.5 (top) where we also show the residuals $\{\log P_t - \Lambda_t\}_{t \geq 0} = \{X_t + Y_t\}_{t \geq 0}$ (bottom).

Figure A.5: Log-prices and fitted seasonality function (top) and residuals from the non-linear least squares regression (bottom).

Appendix 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 A.6.

Figure A.6: Result of hard thresholding where we see the initial signal (blue), the filtered base signal (red), and the filtered spike signal (green).

Figure A.7: Residuals of log-prices after filtering for seasonality and spikes. This signal will be modelled as a $\mathcal{BSS}$ process.

Appendix A.3. Analysis of the residual base signal

The filtered base signal from the hard thresholding algorithm is now assumed to be $\{X_t\}_{t \geq 0} = \{\log P_t - \Lambda_t - Y_t\}_{t \geq 0}$, where Y is the filtered spike signal; see Figure A.7. 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 (4.9). We then test the hypothesis $H_0 : \alpha = 0$ to see if an OU model is adequate to describe the data or if there is still the need for a roughness 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 autocorrelation 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 rough 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 A.8: *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 $\mathcal{BSS}$ model.*

Appendix A.4. Forecasting the base signal

We here take a slightly closer look at the forecasting exercise of Section 6.4. Figure A.8 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 6.4.

Appendix B. Estimators of the roughness index

In this section, we will briefly review some of the estimators of the roughness index α 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 roughness index of electricity prices; (i) jumps (spikes) in the price process and (ii) relatively low frequency (daily) data. In [Appendix C.1](#) and [Appendix C.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 roughness 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 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 roughness 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 $p > 0$,

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

where again α is the roughness index of the process, $\beta \geq 0$ and $c_p > 0$.⁸ 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.

$$\log \hat{\gamma}_p(h) = C + a \log |h| + \epsilon_h, \quad h = 1, 2, \dots, m, \quad (\text{B.2})$$

where C is a constant, m is a “bandwidth” parameter, and ϵ_h an error term. We then estimate $\hat{\alpha}(p) = \hat{a}/p - \frac{1}{2}$, where $\hat{a}$ is obtained from the OLS regression in [\(B.2\)](#) with

$$\hat{\gamma}_p(h) := \frac{1}{N-h} \sum_{i=1}^{N-h} |X_{i+h} - X_i|^p.$$

Intuitively, this estimator will work, since the second part of [\(B.1\)](#) will be asymptotically negligible as $h \rightarrow 0$.⁹ For this reason, it is prudent to use only very few lags of h in this regression, i.e. to set $m = 2$ or $m = 3$. In our application we set $m = 3$. For the power parameter p , a standard choice is $p = 2$ so that $\gamma_p(h) = \gamma_2(h)$ is the (second order) 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 [Appendix C.1](#) and [Gneiting et al., 2012](#); [Bennedsen, 2016](#)). 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 roughness index α (often in the form of an estimator of the Hurst

⁸Note, that even though this definition is stated for Gaussian processes, it equally well applies to volatility modulated $\mathcal{BSS}$ (and OU) processes as shown in [Proposition 2.1.](#) of [Bennedsen et al. \(2015\)](#).

⁹The properties of this estimator has been widely studied. See, e.g., [Bennedsen \(2016\)](#) [Proposition 3.1.](#) for a proof of consistency of the estimator.

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 (B.2). 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 roughness 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.

These estimators are applied to the data of Section 6 in Table 6.

Appendix C. Estimating the degree of roughness in practice

Appendix C.1. Roughness estimation in the presence of spikes

To gauge the effect that possible jumps in the price process have on the estimators of the roughness index, we here conduct a small Monte Carlo study. Table C.8 shows the bias of the various estimators, when outliers are introduced in a rough process. We learn a few things from Table C.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 roughness 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 $GEN(2)$, 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 findings in [Gneiting et al. \(2012\)](#) and [Bennedsen \(2016\)](#). However, when spikes become more frequent, say more than 5 spikes (constituting 1% 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 roughness index found by previous authors, e.g., [Simonsen \(2002\)](#), [Norouzzadeh et al. \(2007\)](#), and [Erzgräber et al. \(2008\)](#).

Appendix C.2. Roughness estimation using low-frequency data

Besides spikes, another problem arises when we try to estimate the roughness 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 roughness is a local one; the roughness 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

Table C.8: *Estimating the roughness index in the presence of outliers*
Number of outliers

	0	1	2	3	4	5	6	7	8	9
Madogr.	-0.001	-0.020	-0.035	-0.049	-0.061	-0.072	-0.082	-0.091	-0.099	-0.106
Variogr.	-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 roughness index estimators when outliers are placed in the data series. The simulated process, X , is a $\mathcal{BSS}$ 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 $\mathcal{BSS}$ process. Outliers of size $\pm 5\text{std}(X_t)$ are then introduced equidistantly into the process. 5 000 Monte Carlo replications.

the definition of the roughness index is given in terms of in-fill asymptotics. In the same vein, the estimators such as the ones through the variogram in Equation (B.1) rely on the distance between observations, h , being small. The question becomes whether it is possible at all to estimate the roughness 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 rough process on $[0, T]$ and estimate its roughness index α . 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 C.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 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 roughness index when the data is observed infrequently. Indeed, a negative roughness 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 roughness 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. However, 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 C.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 roughness index, even using infrequently sampled data. This also holds for the DFA estimator, but not for the ABS and AGG estimators.

Table C.9: *Estimating the roughness 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 roughness index estimators when sampling infrequently and for different values of the mean-reversion parameter λ . The simulated process, X , is a $\mathcal{BSS}$ 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.

Appendix C.3. Lessons from the simulation experiment

The two previous sections above show that one should be careful when estimating the roughness 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 roughness 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, in practice 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.