



UNIVERSITY OF LEEDS

This is a repository copy of *Bayesian calibration and number of jump components in electricity spot price models*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/94785/>

Version: Accepted Version

Article:

Gonzalez, J, Moriarty, J and Palczewski, J orcid.org/0000-0003-0235-8746 (2017)
Bayesian calibration and number of jump components in electricity spot price models.
Energy Economics, 65. pp. 375-388. ISSN 0140-9883

<https://doi.org/10.1016/j.eneco.2017.04.022>

© 2017 Elsevier B.V. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Accepted Manuscript

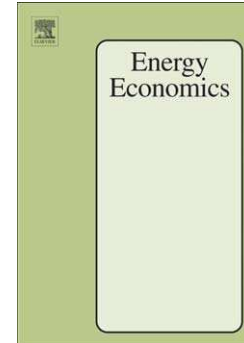
Bayesian calibration and number of jump components in electricity spot price models

Jhonny Gonzalez, John Moriarty, Jan Palczewski

PII: S0140-9883(17)30129-9
DOI: doi: [10.1016/j.eneco.2017.04.022](https://doi.org/10.1016/j.eneco.2017.04.022)
Reference: ENEECO 3622

To appear in: *Energy Economics*

Received date: 10 February 2016
Revised date: 20 April 2017
Accepted date: 28 April 2017



Please cite this article as: Gonzalez, Jhonny, Moriarty, John, Palczewski, Jan, Bayesian calibration and number of jump components in electricity spot price models, *Energy Economics* (2017), doi: [10.1016/j.eneco.2017.04.022](https://doi.org/10.1016/j.eneco.2017.04.022)

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

Bayesian calibration and number of jump components in electricity spot price models

Jhonny Gonzalez^a, John Moriarty^b, Jan Palczewski^{1c}

^a*School of Mathematics, University of Manchester, Manchester M13 9PL, UK*

^b*School of Mathematical Sciences, Queen Mary University of London, London E1 4NS, UK*

^c*School of Mathematics, University of Leeds, Leeds LS2 9JT, UK*

Abstract

We find empirical evidence that mean-reverting jump processes are not statistically adequate to model electricity spot price spikes but independent, signed sums of such processes are statistically adequate. Further we demonstrate a change in the composition of these sums after a major economic event. This is achieved by developing a Markov Chain Monte Carlo (MCMC) procedure for Bayesian model calibration and a Bayesian assessment of model adequacy (posterior predictive checking). In particular we determine the number of signed mean-reverting jump components required in the APXUK and EEX markets, in time periods both before and after the recent global financial crises. Statistically, consistent structural changes occur across both markets, with a reduction of the intensity and size, or the disappearance, of positive price spikes in the later period. All code and data are provided to enable replication of results.

Keywords: Multi-factor models, Bayesian calibration, Markov Chain Monte Carlo, Ornstein-Uhlenbeck process, Electricity spot price, Negative jumps

1. Introduction

Electricity spot markets have multiple fundamental drivers, for example baseload and renewable production (Würzburg et al., 2013). Disturbances in these drivers, such as plant outages and renewable gluts, can clearly have different dynamic characteristics and consequences. Since sharp disturbances create spikes in electricity spot prices (Seifert and Uhrig-Homburg, 2007) we may hypothesise that, over time, disturbances in different drivers give rise to spikes with statistically distinguishable directions, frequencies, height distributions and rates of decay. It has recently been demonstrated that electricity spot price formation can evolve over time (Brunner, 2014). Thus we may also hypothesise that the statistical characteristics of electricity price spikes

¹Corresponding author. *Email address:* J.Palczewski@leeds.ac.uk

will evolve in step with underlying economic events and developments, such as shifts in demand and increasing renewable penetrations.

In this paper we find empirical support for these two hypotheses. To this end we use *multi-factor* electricity spot price models, with multiple superposed mean-reverting components and a seasonal trend (Benth et al., 2007). This allows statistical patterns such as mean reversion, seasonality and spikes to be reproduced in modelling. Crucially for the present study, this approach also allows the statistical modelling of multiple spike components with differing frequencies, height distributions, decay rates, and directions (positive or negative). We demonstrate that in some electricity markets two types of positive spike are observed, while other markets require the inclusion of negative spikes. The modelling of negative spikes is an area of emerging interest (Fanone et al., 2013) as renewable penetrations, and hence gluts in renewable production, increase. Finally we document evolution of the statistical spike structure through periods of economic change by comparing two markets across two time periods, one before the recent global financial crises (2000-2007) and another afterwards (2011-2015) and reflect on possible interpretations of the results.

The calibration of multi-factor models is a highly challenging task and existing approaches typically involve making strong *a priori* assumptions, such as setting thresholds for jump sizes, which may mask the true statistical structure. Methodologically, we develop a Bayesian approach to calibration based on Markov Chain Monte Carlo (MCMC) methods. This goes beyond previous work by making minimal assumptions and enables us, for example, to estimate models with multiple spike components acting in the same direction, a feature which is confirmed empirically (in 2001–2006 data from the APXUK electricity spot market). In order to assess the number of mean-reverting jump components required we perform a Bayesian procedure of *posterior predictive checking*.

1.1. Background and related work

Econometric models of electricity spot prices have a number of important applications. They provide stochastic models which can be used by traders to analyse financial options on power (Benth et al., 2007), and by power system planners to conduct real options analyses for flexible physical assets such as storage and cogeneration (Moriarty and Palczewski, 2017; Kitapbayev et al., 2015). Further the pronounced price spikes which characterise spot electricity markets are of central interest to electricity market regulators who monitor and influence the economics of markets, aiming for example to prevent perceived abuses of market power (Stephenson and Paun, 2001).

The complexity of electricity spot price models, and multi-factor models in particular, makes their analysis statistically challenging and has given rise to a substantial literature. A single-factor model including the above stylised features was introduced by Clewlow and Strickland (2000). Through the use of a threshold, the single-factor model of Geman and Roncoroni (2006) incorporates two jump regimes: when the price is below the threshold jumps are positive, and when the price exceeds the threshold jumps are negative. Beginning with Lucia and Schwartz (2002) multi-factor models have expressed the price as a sum of unobservable or *latent* processes (*factors*) with distinct purposes, for example the modelling of short-term and long-term price variations respectively. Unlike many single factor models, multifactor models do not imply a perfect correlation between changes in spot, future and forward prices, which is consistent with the non-storability of electricity (Benth and Meyer-Brandis, 2009). The model of Lucia and Schwartz (2002) has two factors, namely a Gaussian mean-reverting process and an arithmetic Brownian motion (that is, a scaled Brownian motion with drift). Interestingly, while also developing a two-factor model, Seifert and Uhrig-Homburg (2007) explicitly refer to the physical origins of various types of jumps. Beyond two-factor models, a simple and flexible multi-factor model with jumps is given in Benth et al. (2007). Estimation procedures for this model are discussed in Meyer-Brandis and Tankov (2008), although the latter work adds strong assumptions in order to obtain tractable methods.

The interdependency between parameters in multi-factor models, in particular, is a challenge to calibration methods. A straightforward approach is to first separate the observed values into factors using signal filtering techniques, in order to subsequently employ classical maximum likelihood estimation. Such methods effectively assume that some of these interdependencies may be neglected, and this approach is taken for example in Meyer-Brandis and Tankov (2008) and Benth et al. (2012). An alternative is the joint estimation of latent factors, for which there are two leading methodologies in the literature: *expectation-maximisation (EM)* and *Markov Chain Monte Carlo (MCMC)* methods. While EM produces point estimates for parameters in either a Bayesian or frequentist framework² (see, for example, Rydén et al. (2008)), MCMC is able to generate samples from posterior parameter distributions. Particularly in models with multiple parameters and latent processes, these interdependencies may result in likelihood surfaces and

²Two possible approaches to the calibration of model parameters are commonly referred to as *frequentist* and *Bayesian*. In the frequentist approach one seeks to derive point estimates of 'true' parameter values from the data, for example by finding the maximiser of a likelihood function. An alternative viewpoint is taken in the Bayesian approach, where the unknown parameters are first assigned a probability *distribution* representing prior beliefs about their value. This prior distribution is combined with the observed data to produce an updated probability distribution representing the posterior beliefs about the parameters given both the prior and the data.

posterior distributions which are rather flat around their maxima. While EM suffers from Monte Carlo errors which amplify the usual difficulties in numerical optimisation for such problems, MCMC estimates the posterior distribution providing an analyst with a more complete picture of the interrelations between parameters.

In related contexts, MCMC has been applied to fit continuous-time stochastic volatility models to financial time series, where the price is a diffusion process whose volatility is a latent mean reverting jump process or the sum of a number of such processes (called a superposition model). In this line of research a missing data methodology is employed whereby the observed process is augmented with one or more latent marked Poisson processes and the MCMC procedure generates posterior samples in this high dimensional augmented state space. Examples include Roberts et al. (2004), Griffin and Steel (2006) and Frühwirth-Schnatter and Sögner (2009). Since energy prices additionally exhibit jumps directly in their paths, MCMC has been applied to extensions of these models in which a diffusion process with stochastic volatility is superposed with a jump process, see Green and Nossman (2008) in the context of electricity and Brix (2015) for gas prices. Technically the latter two papers estimate a discrete approximation of the models whereas in this study we pursue *exact* inference for continuous time models.

1.2. Contribution

From the modelling point of view a novelty of the present study is that the price is a superposition of more than one jump component, each with its own sign, frequency, size distribution and decay rate, along with a diffusion component. This approach acknowledges that the negative price spikes attributable to rapid wind power fluctuations may, for example, have quicker decay than the infrequent larger positive spikes due to major disturbances such as outages of a traditional generation plant. The inclusion of multiple jump components also addresses the following problem identified in Green and Nossman (2008) and Brix (2015). In two-factor models jumps of intermediate size must be accounted for either in the diffusion process (forcing unlikely spikes in the Brownian motion path) or the jump process (implying additional jumps). While the former can lead to an overestimation of volatility in the diffusion process, the latter may result in an overestimation of the intensity of the jump process, which is independent of the jump sizes. The inclusion of a second jump process with its own mean jump size and rate of mean reversion removes this dichotomy, offering an alternative to the inclusion of stochastic volatility in the diffusion process.

Our first methodological contribution is an MCMC algorithm for *exact* Bayesian inference on superposed OU models with diffusion and multiple jump components. We contrast exact inference with a commonly used estimation procedure using a discrete time model which is an

approximation to continuous dynamics. While this approximation is often used for practical reasons including simplified and/or tractable implementation, it is not possible to assess *a priori* the extent of the estimation error introduced by the approximation employed. Our MCMC procedure is not based on time discretisation of the model and the inference is therefore exact at the level of distributions. This is in contrast with the work in the aforementioned papers (Seifert and Uhrig-Homburg, 2007; Green and Nossman, 2008; Brix, 2015). Despite the relative simplicity of our multi-factor model, the MCMC procedure involves a number of challenging issues and in the electronic appendix we provide additional comments and details concerning efficient implementation.

In addition we demonstrate that model *adequacy* may also be addressed by our MCMC method. The complexity of electricity spot price models naturally gives rise to parsimony considerations. While multi-factor models (potentially also including latent volatility processes) offer great flexibility, the potential statistical pitfalls of overly flexible models, for example relating to issues of identifiability and out-of-sample prediction, are well known. In this context, the ability of MCMC to sample whole trajectories from the posterior distribution of the jump processes means in particular that the adequacy of latent variable models may be addressed. We exploit this fact by using the MCMC procedure to perform *posterior predictive checks* in the sense of Rubin (1984). For two different electricity spot markets, over two different periods of time, we determine in this way the minimum number of superposed processes required in the model. We find that two or three factors are sufficient in each case. We also show that taking either constant or periodic deterministic jump intensity rates can provide a relatively simple but sufficiently flexible modelling palette. Since the jump processes influence the spot price directly (additively) and have their own proper dynamics such models are also rather interpretable. More generally, since multi-factor models have been considered for a range of commodities including oil and gas (Schwartz and Smith, 2000; Brix, 2015) our algorithm is also potentially applicable in these contexts although this is outside the scope of the present paper (see Gonzalez (2015, Chapter 5) for an application to gas prices).

Section 2 describes the model and the data which animates our study, while Section 3 presents our MCMC algorithm including the approach to assessing model adequacy through posterior predictive checking. The data is analysed in Section 4. Section 5 contains discussion of results and Section 6 concludes. Notes on the efficient implementation of the algorithm are provided in the electronic appendix. Throughout we denote probability distributions as follows: $N(a, b)$ denotes the Normal distribution with mean a and variance b , $Ga(a, b)$ the Gamma distribution with mean a/b , $IG(a, b)$ the Inverse-Gamma distribution with mean $b(a - 1)^{-1}$ for $a > 1$, $Ex(a)$

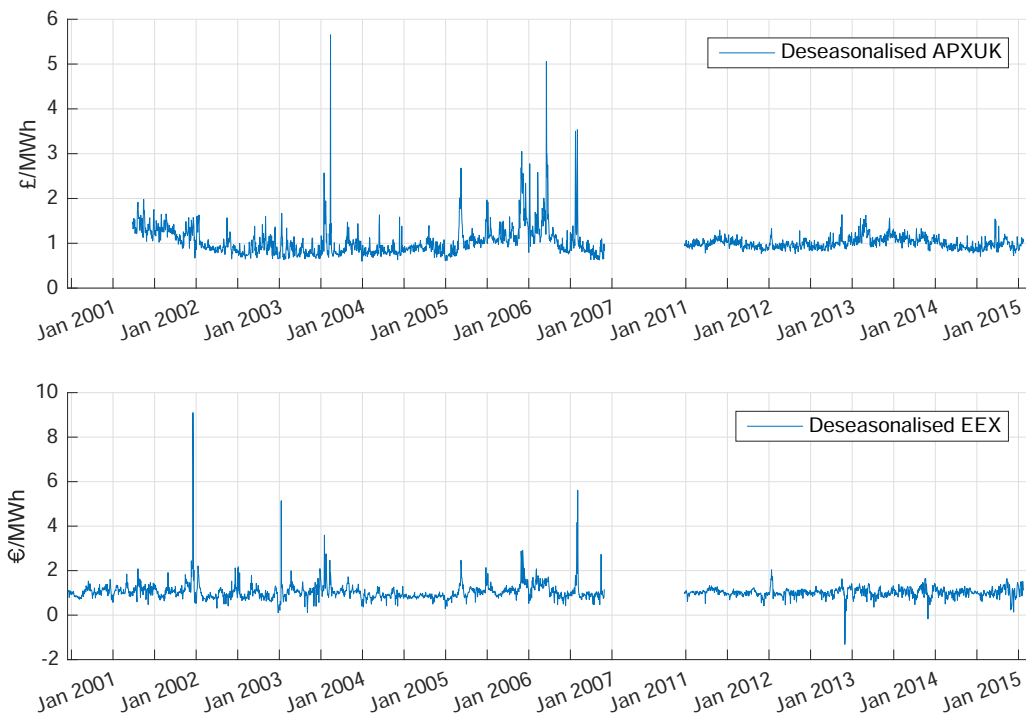


Figure 1: Deseasonalised APXUK (top panel) and EEX (bottom panel) daily average prices (excluding weekends) over two periods. The first period starts on March 27, 2001 for APXUK and on June 16, 2000 for EEX and finishes on November 21, 2006 for both markets. The second period is January 24, 2011 to February 16, 2015 for both series. Details of the deseasonalisation procedure are given in the electronic appendix.

the Exponential distribution with mean a and $U(a, b)$ the Uniform distribution on the interval (a, b) .

2. Model

2.1. Motivation

Figure 1 illustrates two electricity spot markets, the United Kingdom APXUK and European EEX, with weekend prices excluded. The left side of the figure plots daily average prices for the APXUK (March 2001 to November 2006) and the EEX (June 2000 to November 2006). This period was one of general growth in Europe, both economically and in electricity demand, and spot prices from this time have been studied by a number of authors including Green and Nossman (2008), Meyer-Brandis and Tankov (2008) and Benth et al. (2012). The right hand side of the figure plots daily average price data between 2011 and 2015, a period of decline in UK electricity

demand, although the picture across Europe was mixed.³ In order to reveal the structure of these price series more clearly the four time series have been separately deseasonalised (for details see the electronic appendix).

Reversion to a constant level is strongly suggested in the EEX data (bottom panel) and also, to a slightly lesser extent, in the APXUK series (top panel). Taking first the 2001–2006 APXUK data, the presence of significant positive price spikes is clear. However visual inspection also suggests that while some spikes decayed very quickly, a significant number showed more gradual decay. In contrast the positive spikes in the 2000–2006 EEX data appear uniformly to decay quickly and, in addition, the presence of smaller but rather frequent negative spikes is suggested.

While the 2011–2015 APXUK data also suggests regular positive spikes, their heights are significantly smaller than those observed in 2001–2006. In the 2011–2015 EEX data the presence of negative spikes is suggested perhaps more strongly than in 2000–2006. Further, once these negative spikes are taken into account, visual inspection reveals apparently little evidence of positive spikes.

For each series we apply the MCMC procedure described in Section 3 to verify the conclusions of our visual analysis and to establish the smallest number of signed jump components for which the posterior predictive check is favourable, in a sense made precise in Section 4.

In the following subsections we present in detail the class of spot price models to be calibrated.

2.2. Ornstein-Uhlenbeck processes

A process $Y(t)$, $0 \leq t \leq T$ which is right continuous with left limits is called an Ornstein-Uhlenbeck (OU) process if it is the unique strong solution to the stochastic differential equation (SDE)

$$dY(t) = \lambda^{-1}(\mu - Y(t))dt + \sigma dL(t), \quad Y(0-) = y_0, \quad (1)$$

where $L(t)$ is a driving noise process with independent increments, ie., a Lévy process. The initial state of the process Y is defined as the value at the left-hand limit $Y(0-)$ due to the possibility of a jump at time 0. In equation (1), $\mu \in \mathbb{R}$ is the mean level to which the process tends to revert, $\lambda^{-1} > 0$ denotes the speed of mean reversion and $\sigma > 0$ is the volatility of the

³Sources: European Commission Eurostat service; Digest of United Kingdom Energy Statistics (DUKES) 2015, UK Department of Energy & Climate Change.

OU process. The unique strong solution to the SDE (1) is given by

$$Y(t) = \mu + (y_0 - \mu)e^{-\lambda^{-1}t} + \int_0^t \sigma e^{-\lambda^{-1}(t-s)} dL(s). \quad (2)$$

We consider two different specifications for the Lévy process $L(t)$ driving $Y(t)$. On the one hand we consider an OU process where $L(t) = W(t)$ is a standard Wiener process. In this case the conditional distribution of $Y(t+s)$ given $Y(t)$, $t \in [0, T]$, $s \in [0, T-t]$, is Normal with the mean

$$\mathbb{E}[Y(t+s)|Y(t) = y] = \mu + (y - \mu)e^{-\lambda^{-1}s},$$

and the variance

$$\text{Var}[Y(t+s)|Y(t) = y] = \lambda\sigma^2(1 - e^{-2\lambda^{-1}s})/2.$$

Hence we call the process $Y(t)$ a Gaussian OU process. On the other hand we consider the case where $L(t)$ is a compound Poisson process, with the *interval representation*

$$L(t) = \sum_{j=1}^{\infty} \xi_j 1_{\{t \geq \tau_j\}}, \quad (3)$$

where the τ_j are the arrival times of a Poisson process and ξ_j represents the jump size at time τ_j (these jump sizes are independent and identically distributed (i.i.d.) random variables). The dynamics of $Y(t)$ are explicitly given by

$$Y(t+s) = \mu + (Y(t-) - \mu)e^{-\lambda^{-1}s} + \sum_{j:t \leq \tau_j \leq t+s} e^{-\lambda^{-1}(t+s-\tau_j)} \xi_j, \quad s \geq 0. \quad (4)$$

Below we shall model the stochastic part of energy spot prices by superimposing a number of OU processes.

2.3. A multi-factor model for energy spot prices

Let us denote by $X(t)$ the de-trended and deseasonalised spot price at time $t \geq 0$ (presentation of the relation between $X(t)$ and the electricity spot price $S(t)$ is deferred until the end of this section). We assume that the deseasonalised price $X(t)$ is a sum of $n+1$ OU processes

$$X(t) = \sum_{i=0}^n w_i Y_i(t), \quad (5)$$

where Y_0 is a Gaussian OU process

$$dY_0(t) = \lambda_0^{-1}(\mu - Y_0(t))dt + \sigma dW(t), \quad Y_0(0) = y_0, \quad (6)$$

and each $Y_i, i \geq 1$ is a jump OU process

$$dY_i(t) = -\lambda_i^{-1}Y_i(t)dt + dL_i(t), \quad Y_i(0-) = y_i, i = 1, \dots, n, \quad (7)$$

each L_i being a (possibly inhomogeneous) compound Poisson process with exponentially distributed jump sizes having mean β_i . We will refer to this as the $(n+1)$ -OU model. The constants $w_i \in \{1, -1\}$ are used to indicate whether positive or negative jumps are being modelled. Notice that each of the processes $Y_i(t), i \geq 1$, is non-negative since the L_i are increasing processes. Thus by setting $w_i = 1$, we employ $Y_i(t)$ to capture positive price spikes, whereas by setting $w_i = -1$, $Y_i(t)$ is assumed to model negative price spikes. Throughout we assume that w_0 is equal to 1.

For each compound Poisson process $L_i, i \geq 1$, we consider one of two specifications of the jump intensity rate. In the simpler specification we assume that the intensity rate is constant and equal to η_i , and hence that jump frequency is independent of time. In the alternative specification we take account of periodicity in the jump rate as in Geman and Roncoroni (2006) through the deterministic periodic intensity function

$$I_i(\eta_i, \theta_i, \delta_i, t) = \eta_i \left[\frac{2}{1 + |\sin(\pi(t - \theta_i)/k_i)|} - 1 \right]^{\delta_i}, \quad (8)$$

which has period k_i days, where $k_i \in (0, \infty)$ (see Figure 2 for a graph of the fitted intensity function I_1). The parameter $\eta_i \in (0, \infty)$ is the maximum jump rate whilst the exponent $\delta_i \in (0, \infty)$ controls the shape of the periodic function. In order to have a compact notation covering both the above model specifications, the intensity function parameter vector associated with the process Y_i will simply be denoted ϑ_i . In the constant intensity model specification we therefore understand that $\vartheta_i = \eta_i$, while in the periodic intensity model it is understood that $\vartheta_i = (\eta_i, \theta_i, \delta_i)$.

We aim to show that using the sum of a number of such OU processes provides suitable flexibility for modelling electricity spot prices. A diffusive Gaussian component is used to model regular trading characterised by frequent small price variations. The jump components model the arrival of temporary system disturbances of various kinds causing imbalance between supply and demand. By specifying two jump components, say, such that $\lambda_1 > \lambda_2$, we can capture slowly and quickly decaying price spikes. As discussed in Seifert and Uhrig-Homburg (2007), different

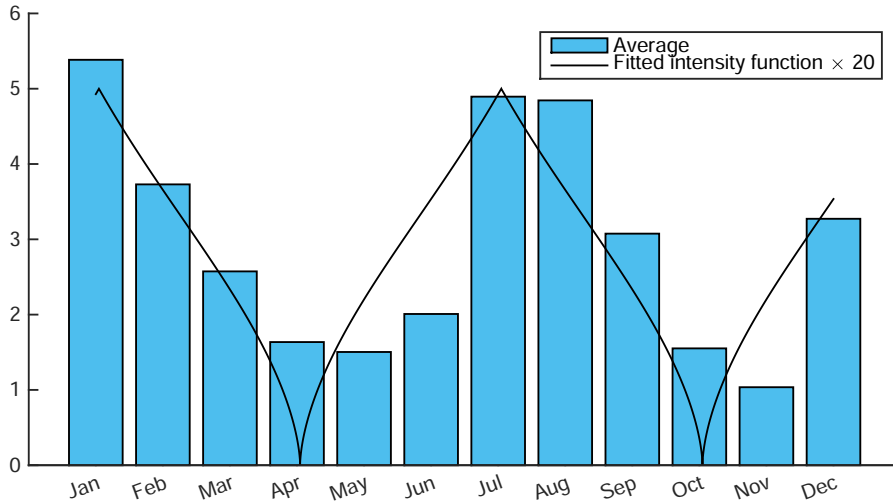


Figure 2: Monthly average number of positive jumps on the EEX market during 2000-6, together with the intensity function I_1 with parameter values taken from Table 4 for the 3-OU- I_1 model. As the jumps are not directly observable in the spot price series, we report the numbers inferred from the latent jump processes sampled in the MCMC procedure.

decay rates may correspond to different *physical causes* of spikes such as power plant outages or extreme changes in weather. The possibility of incorporating negative price spikes by taking $w_i = -1$ is explored in Section 4.

In the empirical studies of Section 4 we assume that the relation between the spot price $S(t)$ and the deseasonalised price $X(t)$ is of the following form:

$$S(t) = e^{f(t/260)} X(t), \quad (9)$$

where $f : [0, \infty) \rightarrow \mathbb{R}$ is a deterministic function that captures the long-term price trend and seasonality typically observed in energy spot prices. In our analysis we take a day as the unit of time and skip weekends due to their distinctly different price dynamics, resulting in a 260-day year. The function f is specified in terms of years to capture weather-induced market patterns linked to seasonal variations. The multiplicative seasonality in (9) is in line with the exponential price trends standard in mathematical economics. We take

$$f(\tau; a_1, \dots, a_6) = a_1 + a_2\tau + a_3 \sin(2\pi\tau) + a_4 \cos(2\pi\tau) + a_5 \sin(4\pi\tau) + a_6 \cos(4\pi\tau), \quad (10)$$

although our methodology applies to any other specification of seasonality provided its effect can be removed from the series of spot prices prior to statistical inference for $X(t)$.

3. Inference

In this section we present a Markov Chain Monte Carlo (MCMC) approach to Bayesian inference in the superposition model (5). We construct a Markov chain whose stationary distribution is the posterior distribution of the parameters in our model together with latent variables introduced to make the inference computationally tractable, see Section 3.1. The application of a Gibbs sampler allows single parameters or groups thereof to be updated conditioned on others being fixed – a standard MCMC approach which aids computational tractability. Central to the performance of MCMC and particularly the Gibbs sampler is the notion of *mixing* which is linked to the speed of convergence of the chain to its stationary distribution. Intuitively the better the mixing, the smaller the dependence between consecutive steps of the chain and, in effect, the less the chain gets blocked in small areas of the state space for long stretches of time. Mixing is negatively affected when the parameters which are updated at a given step of a Gibbs sampler depend on those upon which they are conditioned. This will be of particular importance in the choice of latent variables.

In Sections 3.1-3.3 we present techniques for Bayesian inference in the superposition model (5) in the case of one jump OU component ($n = 1$). This is then extended in Section 3.4 to the case of multiple jump components. For simplicity, when it does not lead to ambiguity, we drop the subscript in the jump process L_1 and its parameters, so that $L = L_1$, $\beta = \beta_1$ and $\vartheta = \vartheta_1$.

3.1. Data augmentation

Let $\mathcal{X} = \{x_0, \dots, x_N\}$ denote observations of the process (5) at times $0 = t_0, \dots, t_N = T$, and $\Delta_i = t_i - t_{i-1} > 0$, $i = 1, \dots, N$, the time increments between consecutive observations. The likelihood $\ell(\mathcal{X} \mid \mu, \lambda_0, \sigma, \lambda_1, \vartheta, \beta)$ of the data given parameters is neither analytically tractable nor amenable to numerical integration since it involves an infinite sum of integrals over high dimensional spaces. However by augmenting the state space with observations $\mathcal{Y}_1 = \{y_{1,0}, \dots, y_{1,N}\}$ of the process Y_1 at times t_i , the likelihood of \mathcal{X} given \mathcal{Y}_1 becomes independent of λ_1 , ϑ and β . Thanks to the explicit form of the transition density of a Gaussian OU process we have

$$\ell(\mathcal{X} \mid \mu, \lambda_0, \sigma, \mathcal{Y}_1) = \prod_{i=1}^N \frac{1}{\sqrt{2\pi}\Sigma_i} \exp \left\{ -\frac{1}{2\Sigma_i^2} \left(z_i - \mu - (z_{i-1} - \mu) e^{-\lambda_0^{-1}\Delta_i} \right)^2 \right\}, \quad (11)$$

where $\Sigma_i^2 = \lambda_0 \sigma^2 (1 - e^{-2\lambda_0^{-1}\Delta_i})/2$ and

$$z_i = x_i - y_{1,i}, \quad i = 0, \dots, N. \quad (12)$$

Space augmentation methods have been widely used in statistics to tackle computationally infeasible problems. However the choice of latent variables or processes has a profound influence on the properties of the resulting estimators, affecting in particular the mixing of a Markov chain approximating the posterior distribution. From a mathematical point of view, the body of work closest to the present inference problem is estimation in the context of stochastic volatility models, where the volatility process is driven by a jump process. Among these are the state space augmentations used in (Jacquier et al., 1994; Kim et al., 1998) and later criticised by Barndorff-Nielsen and Shephard (2001, p. 188) for high posterior correlation of the parameter λ_1 with the input trajectory of the process Y_1 . This correlation could lead to MCMC samplers based on this parametrisation performing poorly. Instead, Barndorff-Nielsen and Shephard (2001) propose an alternative data augmentation scheme based on a series representation of integrals with respect to a Poisson process $L(t)$, known as the Rosiński or Ferguson-Klass representation. Bayesian inference for a stochastic volatility model under this parametrisation was first explored in the discussion section of Barndorff-Nielsen and Shephard (2001) and further developed in Griffin and Steel (2006) and Frühwirth-Schnatter and Sögner (2009). In the present paper, however, we opt for the following more direct parametrisation independently suggested by several researchers (see, for example, Barndorff-Nielsen and Shephard (2001) and Roberts et al. (2004)).

Recall from Section 2.3 that the process $L(t)$ driving $Y_1(t)$ is a compound Poisson process with intensity function $I(\vartheta, t)$ and interval representation (3). Recall also that

$$Y_1(t+s) = Y_1(t-)e^{-\lambda_1^{-1}s} + \sum_{j:t \leq \tau_j \leq t+s} e^{-\lambda_1^{-1}(t+s-\tau_j)} \xi_j, \quad s \geq 0, \quad (13)$$

which motivates a data augmentation methodology where the set of pairs $\{(\tau_j, \xi_j)\}$, instead of the process Y , is treated as the missing data. This has the benefit of introducing independence between λ_1 and the latent variables, thus improving the mixing in Gibbs samplers.

Let us denote by Φ the marked Poisson process on $S = [0, T] \times (0, \infty)$ with locations τ_i on $[0, T]$ and marks ξ_i on $(0, \infty)$. The probability density of Φ is defined relative to a *dominating measure*, namely that of a Poisson process with unit intensity on $[0, T]$ and exponential jump sizes with parameter 1. Hence, thanks to the marking theorem (Kingman, 1992) and the likelihood ratio formula in Kutoyants (1998), the density of Φ with respect to this dominating measure is

$$\ell(\Phi | \vartheta, \beta) = L(\vartheta; \Phi) \cdot \beta^{-N_T} \exp \left\{ -(\beta^{-1} - 1) \sum_{j=1}^{N_T} \xi_j \right\}, \quad (14)$$

where N_T is the number of points in S . Here $L(\vartheta; \Phi)$ is the density, with respect to the Poisson

process with unit intensity, of the Poisson process with intensity $I(\vartheta, t)$:

$$L(\vartheta; \Phi) = \exp \left\{ \sum_{j=1}^{N_T} \log I(\vartheta, \tau_j) - \int_0^T I(\vartheta, t) dt + T \right\}. \quad (15)$$

When L is a homogeneous Poisson process with constant intensity η we obtain

$$L(\vartheta; \Phi) = \exp \{ -(\eta - 1)T \} \eta^{N_T}. \quad (16)$$

We will use a Gibbs sampler to simulate from the posterior distribution of the parameters and the missing data Φ given the observed data \mathcal{X} , using the factorisation

$$\pi(\mu, \lambda_0, \sigma, \lambda_1, \vartheta, \beta, \Phi | \mathcal{X}) \propto \ell(\mathcal{X} | \mu, \lambda_0, \sigma, \lambda_1, \Phi) \ell(\Phi | \vartheta, \beta) \pi(\mu, \lambda_0, \sigma, \lambda_1, \vartheta, \beta), \quad (17)$$

where $\pi(\mu, \lambda_0, \sigma, \lambda_1, \vartheta, \beta)$ is the joint prior density of the parameters.

3.2. Classes of prior distributions

To complete our Bayesian model we now specify classes of prior distributions for the parameters, which are assumed to be mutually independent⁴. For computational efficiency the classes chosen correspond to conjugate priors where possible. In the empirical studies presented in Section 4 the prior distributions are chosen with a large spread (for example variance, where this exists) in order to *let the data speak for itself*. Prior expectations are based on existing results in the literature, combined with further exploratory analysis of historical data as necessary. For details see the electronic appendix. Of course users of our methodology may also have prior beliefs about the model parameters, and in our Bayesian context the prior distributions may alternatively be chosen to reflect these beliefs where appropriate.

We specify a $N(a_\mu, b_\mu^2)$ prior distribution for the mean level μ of the Gaussian OU component, an $IG(a_\sigma, b_\sigma)$ for its volatility σ^2 , an $IG(a_\beta, b_\beta)$ for the jump size parameter β and an $IG(a_{\lambda_i}, b_{\lambda_i})$ for the mean reversion parameter λ_i , $i = 0, 1$. For the intensity function a $Ga(a_\eta, b_\eta)$ prior is chosen for η . Further when the intensity is periodic, a $Ga(a_\delta, b_\delta)$ prior is taken for δ and a $U(a_\theta, b_\theta)$ prior for θ (cf. (8)).

⁴In Section 3.4 and following, where more than one jump OU component is considered, the only statistical dependence we assume is a strict ordering of the mean reversion parameters λ_j , $j = 1, \dots, n$ when this is needed for identifiability.

3.3. MCMC algorithm

In the algorithm below the Gibbs step for updating the λ_i employs a random-walk Metropolis-Hastings procedure. To ensure that the mixing is of the same order for small and large values of λ_i the step of the proposal should be state dependent; equivalently an appropriate transformation of λ_i may be applied. For computational convenience we opt for the latter, swapping λ_i in the inference procedure with $\rho_i = e^{-\lambda_i^{-1}}$.

After setting the initial state of the chain, the MCMC algorithm applied below cycles through the following steps:

MCMC algorithm for the 2-OU model

- Step 1: update $\mu \sim \pi(\mu \mid \rho_0, \sigma, \rho_1, \mathcal{X}, \Phi)$
- Step 2: update $\sigma^2 \sim \pi(\sigma^2 \mid \rho_0, \rho_1, \mathcal{X}, \Phi)$
- Step 3: update $\rho_0, \rho_1 \sim \pi(\rho_0, \rho_1 \mid \mu, \sigma, \mathcal{X}, \Phi)$
- Step 4: update $\vartheta \sim \pi(\vartheta \mid \Phi)$
- Step 5: update $\beta \sim \pi(\beta \mid \Phi)$
- Step 6: update $\Phi \sim \pi(\Phi \mid \mu, \rho_0, \sigma, \rho_1, \vartheta, \beta, \mathcal{X})$
- Step 7: Go to step 1.

Below we provide more details about each of these steps.

Step 1. Update μ

Recalling (11), the likelihood of the observed data conditional on the augmented state $(\mu, \lambda_0, \sigma, \rho_1, \Phi)$ is

$$\ell(\mathcal{X} \mid \mu, \lambda_0, \sigma, \rho_1, \Phi) \propto \frac{1}{\prod_{i=1}^N \Sigma_i} \exp \left\{ -\frac{1}{2} \sum_{i=1}^N \frac{1}{\Sigma_i^2} \left(z_i - z_{i-1} e^{-\lambda_0^{-1} \Delta_i} + \mu \left(e^{-\lambda_0^{-1} \Delta_i} - 1 \right) \right)^2 \right\},$$

where $\Sigma_i^2 = \lambda_0 \sigma^2 (1 - e^{-2\lambda_0^{-1} \Delta_i})/2$ and the z_i are computed as the difference between the observations of X and the trajectory of Y implied by the realisation Φ of the marked Poisson process. Using the conjugate prior for μ specified in the previous section it can be easily shown that the conditional distribution $\pi(\mu \mid \rho_0, \sigma, \rho_1, \mathcal{X}, \Phi)$ is

$$N \left(\frac{\sum_{i=1}^N (1 - e^{-\lambda_0^{-1} \Delta_i}) \Sigma_i^{-2} (z_i - z_{i-1} e^{-\lambda_0^{-1} \Delta_i}) + \frac{a_\mu}{\sigma_0^2}}{\sum_{i=1}^N (1 - e^{-\lambda_0^{-1} \Delta_i})^2 \Sigma_i^{-2} + \frac{1}{b_\mu^2}}, \frac{1}{\sum_{i=1}^N (1 - e^{-\lambda_0^{-1} \Delta_i})^2 \Sigma_i^{-2} + \frac{1}{b_\mu^2}} \right).$$

Step 2. Update σ^2

Due to the choice of prior, the conditional distribution $\pi(\sigma^2 \mid \rho_0, \rho_1, \mathcal{X}, \Phi)$ has the closed form

$$\text{IG} \left(\frac{N}{2} + a_\sigma, \frac{1}{\lambda_0} \sum_{i=1}^N \frac{s_i}{(1 - e^{-2\lambda_0^{-1}\Delta_i})} + b_\sigma \right),$$

where

$$s_i = \left(z_i - z_{i-1}e^{-\lambda_0^{-1}\Delta_i} + \mu \left(e^{-\lambda_0^{-1}\Delta_i} - 1 \right) \right)^2.$$

Step 3. Update ρ_0 and ρ_1

Explicit conditional distributions for ρ_0 and ρ_1 are not available and the density is only known up to a multiplicative constant:

$$\begin{aligned} \pi(\rho_0 \mid \mu, \sigma, \rho_1, \mathcal{X}, \Phi) &\propto \ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi)\pi(\rho_0), \\ \pi(\rho_1 \mid \mu, \sigma, \rho_0, \mathcal{X}, \Phi) &\propto \ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi)\pi(\rho_1). \end{aligned}$$

Hence we use a random-walk Metropolis-Hastings within Gibbs procedure to update ρ_0 and ρ_1 . The variance of the proposal distribution is tuned after pilot runs in order to achieve an acceptance rate between 20% and 50%.

Step 4. Update ϑ

In the case of constant intensity function, the conjugate prior for η yields an explicit conditional distribution

$$\eta \mid \Phi \sim \text{Ga}(a_\eta + N_T, T + b_\eta).$$

When the intensity function is time dependent we employ a random-walk Metropolis-Hastings within Gibbs procedure to update η, θ and δ :

$$\pi(\eta, \theta, \delta \mid \Phi) \propto \ell(\Phi \mid \eta, \theta, \delta)\pi(\eta)\pi(\theta)\pi(\delta).$$

Here the non-explicit function $L(\vartheta \mid \Phi)$ of (15) is numerically calculated by a quadrature method. The variance of the proposal distribution is tuned after pilot runs as above.

Step 5. Update β

The conditional distribution of β given Φ has the closed form

$$\beta \mid \Phi \sim \text{IG} \left(a_\beta + N_T, \sum_{i=1}^{N_T} \xi_i + b_\beta \right).$$

Step 6. Update the latent process Φ

The Metropolis-Hastings step we use to update the process Φ draws from the work of Geyer and Møller (1994), Roberts et al. (2004) and Frühwirth-Schnatter and Sögner (2009) on MCMC techniques for simulating point processes, extending it where appropriate to the case of inhomogeneous Poisson processes.

Let us assume that the current state of the Markov chain is

$$\Phi = \{(\tau_1, \xi_1), \dots, (\tau_{N_T}, \xi_{N_T})\},$$

that is, there are N_T points on the set S with jump times given by τ_j and the corresponding jump sizes by ξ_j . We choose randomly, with equal probability, one of the following three proposals.

Birth-and-death step

In the birth-and-death step we choose one of two moves. With probability $p \in (0, 1)$ we choose a birth move whereby a new-born point (τ, ξ) is added to the current configuration of Poisson points. The proposed new state is then $\Phi \cup \{(\tau, \xi)\}$. The point τ is drawn uniformly from $[0, T]$, whilst ξ is drawn from the jump size distribution $\text{Ex}(\beta)$. For this move the proposal transition kernel $q(\Phi, \Phi \cup \{(\tau, \xi)\})$ has the following density with respect to the product of Lebesgue measure on $[0, T]$ and $\text{Ex}(1)$ measure on $(0, \infty)$:

$$q(\Phi, \Phi \cup \{(\tau, \xi)\}) = \beta^{-1} \exp(-(\beta^{-1} - 1)\xi).$$

With probability $1 - p$ a death move is selected, a randomly chosen point (τ_i, ξ_i) being removed from Φ (provided that Φ is not empty). The proposal transition kernel (with respect to the counting measure) is

$$q(\Phi, \Phi \setminus \{(\tau_i, \xi_i)\}) = \frac{1}{N_T},$$

where N_T is the number of points in Φ before the death move. Then the Metropolis-Hastings acceptance ratio for a birth move from Φ to $\Phi \cup \{(\tau, \xi)\}$ is

$$\alpha(\Phi, \Phi \cup \{(\tau, \xi)\}) = \min \{1, r(\Phi, (\tau, \xi))\},$$

while the acceptance ratio for a death move from Φ to $\Phi \setminus \{(\tau_i, \xi_i)\}$ is

$$\alpha(\Phi, \Phi \setminus \{(\tau_i, \xi_i)\}) = \min \left\{ 1, \frac{1}{r(\Phi \setminus \{(\tau_i, \xi_i)\}, (\tau_i, \xi_i))} \right\},$$

where

$$\begin{aligned} r(\tilde{\Phi}, (\tau, \xi)) &= \frac{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \tilde{\Phi} \cup \{(\tau, \xi)\}) \pi(\tilde{\Phi} \cup \{(\tau, \xi)\} \mid \vartheta, \beta) \frac{1-p}{p}}{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \tilde{\Phi}) \pi(\tilde{\Phi} \mid \vartheta, \beta)} \\ &\quad \times \frac{1}{(N_T + 1)q(\Phi, \Phi \cup \{(\tau, \xi)\})} \\ &= \frac{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \tilde{\Phi} \cup \{(\tau, \xi)\}) \frac{1-p}{p} \frac{T}{\tilde{N}_T + 1} I(\vartheta, \tau)}{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \tilde{\Phi})} \end{aligned}$$

where \tilde{N}_T is the number of points of $\tilde{\Phi}$, cf. (14).

Local displacement move

Without loss of generality let us assume that the jump times of the Poisson process are ordered, so that $\tau_1 < \dots < \tau_{N_T}$. In the local displacement move we choose randomly one of the jump times, say τ_j , and generate a new jump time τ uniformly on $[\tau_{j-1}, \tau_{j+1}]$, putting $\tau_0 = 0$ and $\tau_{N_T+1} = T$. The point (τ_j, ξ_j) is then displaced and re-sized to (τ, ξ) , where $\xi = e^{-\lambda_1^{-1}(\tau-\tau_j)}\xi_j$. Formally we choose uniformly one of N_T transition kernels, with the j -th one preserving the conditional distribution $\pi(\tau, \xi \mid \mathcal{X}, \mu, \rho_0, \sigma, \rho_1, \vartheta, \beta, \Phi \setminus \{(\tau_j, \xi_j)\})$. The proposal for the j -th kernel has the Uniform distribution over (τ_{j-1}, τ_{j+1}) for the first variable with the second variable being then a deterministic transformation given by a 1-1 mapping $\mathcal{T}(\xi, \tau, \tau') = (\xi e^{-\lambda_1^{-1}(\tau'-\tau)}, \tau', \tau)$ such that $\mathcal{T} = \mathcal{T}^{-1}$. Following Tierney (1998, Section 2), the contribution of this deterministic transition to the Metropolis-Hastings acceptance ratio is $|\det \nabla \mathcal{T}(\xi_j, \tau_j, \tau)|$, where τ is the new proposed location of the jump. Hence the complete Metropolis-Hastings acceptance ratio is

$$\begin{aligned} r(\Phi, \Phi_{new}) &= \frac{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi_{new}) \pi(\tau, \xi \mid \vartheta, \beta) \tilde{q}(\tau, \tau_j)}{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi) \pi(\tau_j, \xi_j \mid \vartheta, \beta) \tilde{q}(\tau_j, \tau)} |\det \nabla \mathcal{T}(\xi_j, \tau_j, \tau)| \\ &= \frac{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi_{new}) I(\vartheta, \tau) e^{-\beta^{-1}\xi}}{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \Phi) I(\vartheta, \tau_j) e^{-\beta^{-1}\xi_j}} e^{-\lambda_1^{-1}(\tau-\tau_j)}, \end{aligned}$$

where $\tilde{q}(\tau, \tau') = (\tau_{j+1} - \tau_{j-1})^{-1}$ is the transition density for the jump location with respect to Lebesgue measure on (τ_{j-1}, τ_{j+1}) .

Multiplicative jump size update

In this step the sizes of all jumps are independently updated. Specifically, for each jump (τ_j, ξ_j) we propose a new jump size $\xi'_j = \xi_j \phi_j$, where $\log(\phi_j) \sim N(0, c^2)$ are i.i.d. random variables. The variance c^2 is chosen inversely proportional to the current number of jumps, and the performance of this update step appears rather insensitive to the constant of proportionality. Denoting by Φ_{new} the Poisson point process with updated jump sizes, the Metropolis-Hastings acceptance ratio for this move is

$$\alpha(\Phi, \Phi_{new}) = \min \left\{ 1, \frac{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \vartheta, \beta, \Phi_{new})}{\ell(\mathcal{X} \mid \mu, \rho_0, \sigma, \rho_1, \vartheta, \beta, \Phi)} \exp \left\{ -(\beta^{-1} - 1) \sum_{i=1}^{N_T} (\xi'_i - \xi_i) \right\} \prod_{i=1}^{N_T} \frac{\xi'_i}{\xi_i} \right\}.$$

Note that the product $\prod_{i=1}^{N_T} \frac{\xi'_i}{\xi_i}$ is equivalent to $\prod_{i=1}^{N_T} \phi_i$.

3.4. Bayesian inference for a sum of three OU processes

As discussed in Section 2.3 we may believe *a priori* that the price spikes observed in the market are of a certain number of types, corresponding to their differing possible physical causes. Accordingly we now describe the extension of the 2-OU model by the addition of a further independent jump OU component $Y_2(t)$, henceforth referring to the first jump component as $Y_1(t)$ and using appropriate subscripts to distinguish their parameters; further jump components are incorporated similarly. The new jump component Y_2 may either have a positive contribution to the price process (ie. the sign w_2 in (5) is 1) with a rate of decay differing from that of the first component, or alternatively it may have a negative contribution (ie. $w_2 = -1$). For concreteness here we choose $w_1 = w_2 = 1$, so that:

$$X(t) = Y_0(t) + Y_1(t) + Y_2(t), \quad (18)$$

and we specify that $\lambda_1 > \lambda_2$ for identification purposes, ie. the jumps of $Y_1(t)$ have slower decay than those of $Y_2(t)$. When $w_1 = 1, w_2 = -1$ this constraint is not required.

We now have two marked Poisson processes Φ_1 and Φ_2 , corresponding to $L_1(t)$ and $L_2(t)$ respectively, which are conditionally independent given their parameters. The augmented likelihood $\ell(\mathcal{X} \mid \mu, \lambda_0, \sigma, \mathcal{Y}_1, \mathcal{Y}_2)$ is given by equation (11) with $z_j = x_j - y_{1,j} - y_{2,j}$. The likelihood of $\Phi = (\Phi_1, \Phi_2)$ with respect to the dominating measure of a pair of independent marked Poisson process with unit intensity and jump sizes with a $\text{Ex}(1)$ distribution is given by

$$\pi(\Phi \mid \vartheta_1, \vartheta_2, \beta_1, \beta_2) = \pi(\Phi_1 \mid \vartheta_1, \beta_1) \pi(\Phi_2 \mid \vartheta_2, \beta_2),$$

where for $i = 1, 2$,

$$\pi(\Phi_i | \vartheta_i, \beta_i) = \exp \left\{ \sum_{j=1}^{N_T^i} \log I_i(\vartheta_i, \tau_{i,j}) - \int_0^T I_i(\vartheta_i, t) dt + T \right\} \beta_i^{N_T^i} \exp \left\{ -(\beta_i^{-1} - 1) \sum_{j=1}^{N_T^i} \xi_{i,j} \right\}.$$

Here N_T^i denotes the number of points of Φ_i in the set S and

$$\Phi_i = \{(\tau_{i,1}, \xi_{i,1}), \dots, (\tau_{i,N_T^i}, \xi_{i,N_T^i})\}.$$

We impose the condition $\lambda_1 > \lambda_2$ by defining the prior distribution for ρ_2 as

$$\rho_2 | \rho_1 \sim \rho_1 \mathbf{U}(0, 1),$$

where $\rho_i = e^{-1/\lambda_i}$, $i \geq 1$, as before (see the electronic appendix for properties of ρ_2). Since all other parameters are *a priori* mutually independent, their update steps are identical to those for the 2-OU model.

Updates of ρ_1 and ρ_2 are made using a random-walk Metropolis-Hastings algorithm (with independent proposals for each variable). The proposal distribution is Normal with the variance tuned as before after pilot runs.

3.5. Posterior predictive check

Our approach to assessing model adequacy is the following. Since Y_0 is a Gaussian OU process, its transition density specified in Section 2.2 implies that ε_j , $j = 1, \dots, N$, defined implicitly by

$$Y_0(t_j) = \mu + (Y_0(t_{j-1}) - \mu)e^{-\lambda_0^{-1}\Delta_j} + \left(\frac{\sigma^2 \lambda_0}{2} (1 - e^{-2\lambda_0^{-1}\Delta_j}) \right)^{1/2} \varepsilon_j, \quad (19)$$

are independent and distributed as $\mathbf{N}(0, 1)$. Given observations of Y_0 at sampling times t_0, \dots, t_N the distribution of $\{\varepsilon_1, \dots, \varepsilon_N\}$ may then be tested, and this test may be repeated across MCMC iterations. At each iteration k of the MCMC algorithm (assuming the Markov chain has reached stationarity) we use the current state of parameters $\Theta^{(k)} = \{\mu^{(k)}, \lambda_i^{(k)}, \sigma^{(k)}, \vartheta_i^{(k)}, \beta_i^{(k)}\}$ and missing data $\Phi^{(k)}$ to recover the path of each jump process $y_{i,j}^{(k)}$, $j = 0, \dots, N$. The path of $Y_0^{(k)}$ at times t_0, \dots, t_N is then computed as

$$z_j^{(k)} = x_j - \sum_{i=1}^n w_i y_{i,j}^{(k)}, \quad j = 0, \dots, N, \quad (20)$$

where x_i is the deseasonalised price at time t_i and w_i is the sign of the i -th jump component. From this the noise data $\{\varepsilon_j^{(k)}\}_{j=1,\dots,N}$ for each MCMC iteration k is obtained and subjected to a Kolmogorov-Smirnov (KS) test for the standard Normal distribution, yielding a p -value $p^{(k)}$. Following Rubin (1984) we call the distribution of $p^{(k)}$ the posterior predictive check distribution. We refer to its mean as the *posterior predictive p -value* and interpret it accordingly, cf. Gelman (2003) and references therein.

Similarly we perform diagnostics on the jump processes sampled from our Markov chain. For each jump process L_i at iteration k of the Markov chain, the set of jump sizes and the set of inter-arrival times are both subjected to KS tests. The former set undergoes a KS test for the Exponential distribution with mean $\beta_i^{(k)}$. In the homogeneous Poisson process model, the latter set undergoes a KS test for the Exponential distribution with mean $(\eta_i^{(k)})^{-1}$; otherwise an independent sample is taken from the inter-arrival times of an inhomogeneous Poisson process with time varying intensity $I_i(\vartheta_i^{(k)}, t)$ and its distribution compared with the latter set in a two-sample KS test. Posterior predictive p -values are reported.

3.6. Implementation

The parameter-dependent balance between jumps and diffusion in the spot price model (5) raises a number of potential issues regarding the implementation of the MCMC procedure described above. Extensive testing was carried out with simulated data in order to probe these issues, and details are given in the electronic appendix.

Our MCMC algorithms were implemented by combining Matlab and C++ MEX code (provided as an electronic supplement), and were run on a 2.5GHz Intel Xeon E5 processor. For the 2-OU model with 1500 observations the computation time for completing 1000 MCMC iterations using a single core was approximately 1.9 seconds. The corresponding figure for the 3-OU model was about 3.6 seconds with a single update of the latent process Φ , and 8.8 seconds with 5 updates of Φ per MCMC iteration. In the numerical examples of Section 4, a burn-in period of 500 000 iterations was allocated. The following 1.5 million were thinned by taking one sample every 100 iterations and used to establish the posterior distribution. This corresponds to around 1 hour running time for the 2-OU model and below 5 hours for the 3-OU model with the fivefold update of the latent process Φ .

4. Case study application to the APXUK and EEX markets

In this section we apply the inference procedure described in Section 3 to daily average electricity prices corresponding to the APX Power UK spot base index (APXUK hereafter, quoted in

£/MWh)⁵ and the European Energy Exchange Phelix day base index (EEX, quoted in €/MWh)⁶. The data were retrieved from Thompson Reuters Datastream. Weekends (which tend to differ statistically from weekdays due to a reduction in trading) are removed in both cases. Hence we assume a calendar year of 260 days and take $\Delta_j = 1$ so that parameters are reported in daily units. We consider two time periods, namely 2000–2006 and 2011–2015. Figure 1 displays de-seasonalised versions of this data. The pattern of positive and negative spikes appears to differ across the two markets and time periods and correspondingly we perform four separate analyses.

A step-by-step guide to the practical application of our techniques is first presented in Section 4.1, followed in Section 4.2 by an economics-oriented discussion of the models fitted to the above datasets. The steps involved in fitting are collected in Sections 4.3–4.6, together with related econometric discussion.

4.1. Step-by-step guide

This section contains a subjective guide to estimation for the above multi-factor price model using discrete observations. The analytical procedure is summarised in an algorithm below and later exemplified in Sections 4.3–4.6.

1. Deseasonalise time series
2. Set $n = 1$
3. Fit all combinations of $(n + 1)$ -OU models (by a combination we mean the number of positive and negative jump components)
4. For each combination in step 3:
 - (a) compute posterior predictive p -values for the increments of the residual Gaussian OU process,
 - (b) for each latent Poisson process, compute posterior predictive p -values for the jump arrival rates and the jump sizes
5. Accept a model if all p -values are above a selected threshold (we have used 10%)
6. If no model has been accepted then, for each combination in step 3: introduce time-varying intensities for every subset of the jump components (see Eq.(8)); repeat model estimation, compute predictive p -values and accept / reject model as in step 5.
7. If no model has been accepted then set $n := n + 1$ and go back to step 3.

⁵<https://www.apxgroup.com/market-results/apx-power-uk/ukpx-rpd-index-methodology/>

⁶<http://www.epexspot.com/en/>

The above procedure suffers from the classical problem of multiple testing (Benjamini and Hochberg, 1995), ie., if sufficiently many models are tested one of them will prove statistically significant purely due to randomness. However, for the sake of simplicity we decided not to include this aspect of model selection in the above procedure. Instead we suggest verifying the selected model on a subset of the data, or on another dataset with similar characteristics, in order to confirm that the model choice is robust. This involves estimating the selected model on the new dataset and checking that all predictive p -values remain above the threshold applied.

The appropriate modelling of the seasonal component, in order to remove it from the data, depends heavily on the market under study (Weron, 2007). Electricity spot markets experience significant yearly variations and the specification (10) of the seasonality function takes these into account. We used a minimum least-squares fit for the log price, which corresponds to a linear regression of the natural logarithm of price on each of the terms of the seasonality function f .

According to the algorithm above, the data is first deseasonalised. It is then checked whether the simplest model with just one jump component provides a suitable statistical description of the data, as follows. The MCMC procedure is applied both to the model with one positive jump component and to the model with one negative jump component. For each model assessed this generates a sequence of posterior samples of: the latent Poisson process driving the price spikes; the spike sizes; and the implied discrete increments of the latent Gaussian Ornstein-Uhlenbeck process. These are used to compute posterior predictive p -values for the jump times, jump sizes and the increments of the Gaussian OU process as in Section 3.5. If all computed p -values lie above a predetermined threshold (we have used 10%) the model may be deemed statistically valid. If neither of the simple models is statistically valid, we suggest relaxing the assumption of constant jump intensity and repeating the estimation procedure. If those models also fail the statistical validity test, the number of jump components may be increased by one and the above estimation and verification procedure repeated. Since the inclusion of further jump components clearly improves the model fit, this motivates the acceptance of the simplest model satisfying the above criterion.

4.2. Summary of empirical results

Our first hypothesis in this work is that, over time, disturbances in the different drivers involved in electricity spot price formation give rise to spikes with statistically distinguishable directions, frequencies, height distributions and rates of decay. Our results provide evidence for this hypothesis in data from the period 2000–2006. For the 2001–2006 APXUK data we find that the model with a single positive jump component has posterior predictive p -values which are too low to be judged adequate. In contrast when two independent positive jump components

are included in the model, the MCMC procedure is able to distinguish these two factors statistically. This is evidenced by posterior predictive p -values for each fitted component which are statistically acceptable.

Similarly in the 2000–2006 EEX data our calibration procedure is able to statistically distinguish two independent jump components. In this case it is the signed combination of one positive and one negative jump component which has empirical support (that is, acceptable posterior predictive p -values for each fitted component). Both the model with a single positive jump component, and the model specifying two positive jump components, have posterior predictive p -values which are too low to be judged adequate.

Our second hypothesis is that the composition of these statistical models evolves in parallel with changes in underlying economic factors. This hypothesis is supported by comparing the above models for 2000–2006 with models for the same markets over the period 2011–2015. In both cases statistical changes are detected and, further, the nature of the change is consistent across the two markets. In this later period a single positive jump component provides an adequate fit to the APXUK data. There is also an apparent decrease in the frequency of price spikes: on average the total number of jumps (of any size) per unit time is less than a third for the 2011–2015 series compared to 2001–2006. The 2011–2015 EEX data also has one fewer positive jump component relative to the period 2000–2006. Thus in the later period a single negative jump component is adequate.

4.3. Deseasonalisation of the time series

In this paper we perform inference on deseasonalised data, treating the seasonal trend function $f(t)$ as a known characteristic of the particular energy market under study. For the purposes of the numerical illustration in this section we assume the form (10) and solve a least squares problem

$$\sum_{i=0}^N (\log S_{obs}(t_i) - f(t_i/260))^2 \rightarrow \min,$$

where $S_{obs}(t)$ denotes the observed spot price at time t . Table A.3 in the electronic appendix presents estimated parameters for the trend functions and Figure 1 displays the resulting deseasonalised time series $X(t) = S_{obs}(t)e^{-f(t/260)}$.

4.4. 2001–2006 APXUK data

4.4.1. One jump component

We take the priors specified in Table 1 as input to the 2-OU model with a single, positive jump component and constant intensity rate. The Markov chain was initialised with the state

Parameter	Prior properties			Posterior properties	
	Prior	Mean	SD	Mean	SD
μ	$N(1, 20^2)$	1	20	0.9592	0.0308
σ^2	$IG(1.5, 0.005)$	0.01	--	0.0096	0.0008
e^{-1/λ_0}	$U(0, 1)$	0.5	0.2887	0.9170	0.0116
(λ_0)	$(IG(1,1))$	(- -)	(- -)	(11.7936)	(1.8298)
e^{-1/λ_1}	$U(0, 1)$	0.5	0.2887	0.1570	0.0189
(λ_1)	$(IG(1,1))$	(- -)	(- -)	(0.5403)	(0.0352)
η	$Ga(1, 10)$	0.1	0.1	0.2499	0.0297
β	$IG(1, 1)$	--	--	0.7159	0.0738

Table 1: Prior distributions and posterior moments obtained when calibrating the 2-OU model to the 2001–2006 APXUK data. The posteriors for the ‘indirect’ parameters λ_i were obtained by transformation of the parameters e^{-1/λ_i} of the Markov chain at each step and their entries are given in brackets.

$(\mu, \lambda_0, \sigma, \lambda_1, \eta, \beta, \Phi) = (1, 5, 0.1, 2, 0.1, 0.5, \mathbf{0})$ where $\mathbf{0}$ denotes the absence of jumps, and the birth-and-death parameter p was set equal to 0.5 (ie. equal probability for the birth or death of a jump). Table 1 presents summary statistics for the prior and posterior distributions of the univariate parameters. For those priors with a second moment, it may be observed that the standard deviation of the posterior is typically at least an order of magnitude smaller.

4.4.2. Two jump components

We also calibrate a 3-OU model with constant jump intensity and two positive jump components with the condition $\lambda_1 > \lambda_2$ introduced for identifiability (through the use of an appropriate prior as specified in Subsection 3.4), ie. the jumps of Y_2 decay faster than those of Y_1 . The initial state of the chain was set to

$$(\mu, \lambda_0, \sigma, \lambda_1, \eta_1, \beta_1, \lambda_2, \eta_2, \beta_2, \Phi) = (1, 5, 0.2, 5, 0.001, 0.5, 1, 0.001, 0.5, \mathbf{0}).$$

Summary statistics for the prior and posterior distributions of the univariate parameters are given in Table 2. For each of the jump component parameters λ_i, η_i and β_i the posterior distributions for the two jump components $i = 1, 2$ are well separated. In particular the ‘new’ jump component Y_2 suggests that the quickly decaying price shocks are both less frequent and larger on average than the more slowly decaying jumps given by Y_1 . As may be anticipated, the posterior distribution of the volatility σ of the diffusion component Y_0 is correspondingly shifted lower with the inclusion of Y_2 . However the posterior moments of the speed of mean reversion λ_0 remained virtually

unchanged.

Parameter	Prior properties			Posterior properties	
	Prior	Mean	SD	Mean	SD
μ	$N(1, 20^2)$	1	20	0.8693	0.0284
σ^2	$IG(1.5, 0.005)$	0.01	--	0.0057	0.0006
e^{-1/λ_0}	$U(0, 1)$	0.5	0.2887	0.9176	0.0128
(λ_0)	$(IG(1, 1))$	(--)	(--)	(11.9352)	(1.9892)
e^{-1/λ_1}	$U(0, 1)$	0.5	0.2887	0.6605	0.0342
(λ_1)	$(IG(1, 1))$	(--)	(--)	(2.4408)	(0.3075)
e^{-1/λ_2}	--	0.25	0.2205	0.0742	0.0148
(λ_2)	(--)	(--)	(--)	(0.3839)	(0.0297)
η_1	$Ga(1, 10)$	0.1	0.1	0.2412	0.0487
η_2	$Ga(1, 10)$	0.1	0.1	0.1698	0.0250
β_1	$IG(1, 1)$	--	--	0.2243	0.0291
β_2	$IG(1, 1)$	--	--	0.8544	0.1047

Table 2: Prior distributions and posterior properties obtained when calibrating the 3-OU model to the 2001–2006 APXUK data. The indirect parameters λ_i are treated as described in the caption to Table 1. The distribution and moments of $\rho_2 = e^{-1/\lambda_2}$ are calculated in the electronic appendix.

4.4.3. Augmentation of the state space

In order to illustrate the role of the latent variables introduced in the augmented state space of our Markov chain, and to show how these may vary between the 2-OU and 3-OU models, Figure 3 gives a representation of one posterior sample for each model via their respective latent processes Y_i . For clarity of the plot a restricted period of 200 days is shown, giving the paths of the jump processes, plus the deseasonalised APXUK price superimposed on the implied diffusion $Y_0 = X - \sum_{i=1}^n Y_i$. (For both models these samples are in fact the last state of the simulated Markov chain.)

From Table 1 the jumps of Y_1 are relatively large (their distributional mean size β has posterior expected value 0.72) and the decay rate λ_1 has posterior mean 0.54. The 3-OU model identifies both slowly decaying small jumps and rapidly decaying large jumps. Inspecting the plots in Figure 3 for the 2-OU model around day 600, it is therefore apparent that in this illustrative example runs of consecutive quickly decaying jumps combine to produce an apparently larger and more slowly decaying disturbance, while some single jumps such as that around day

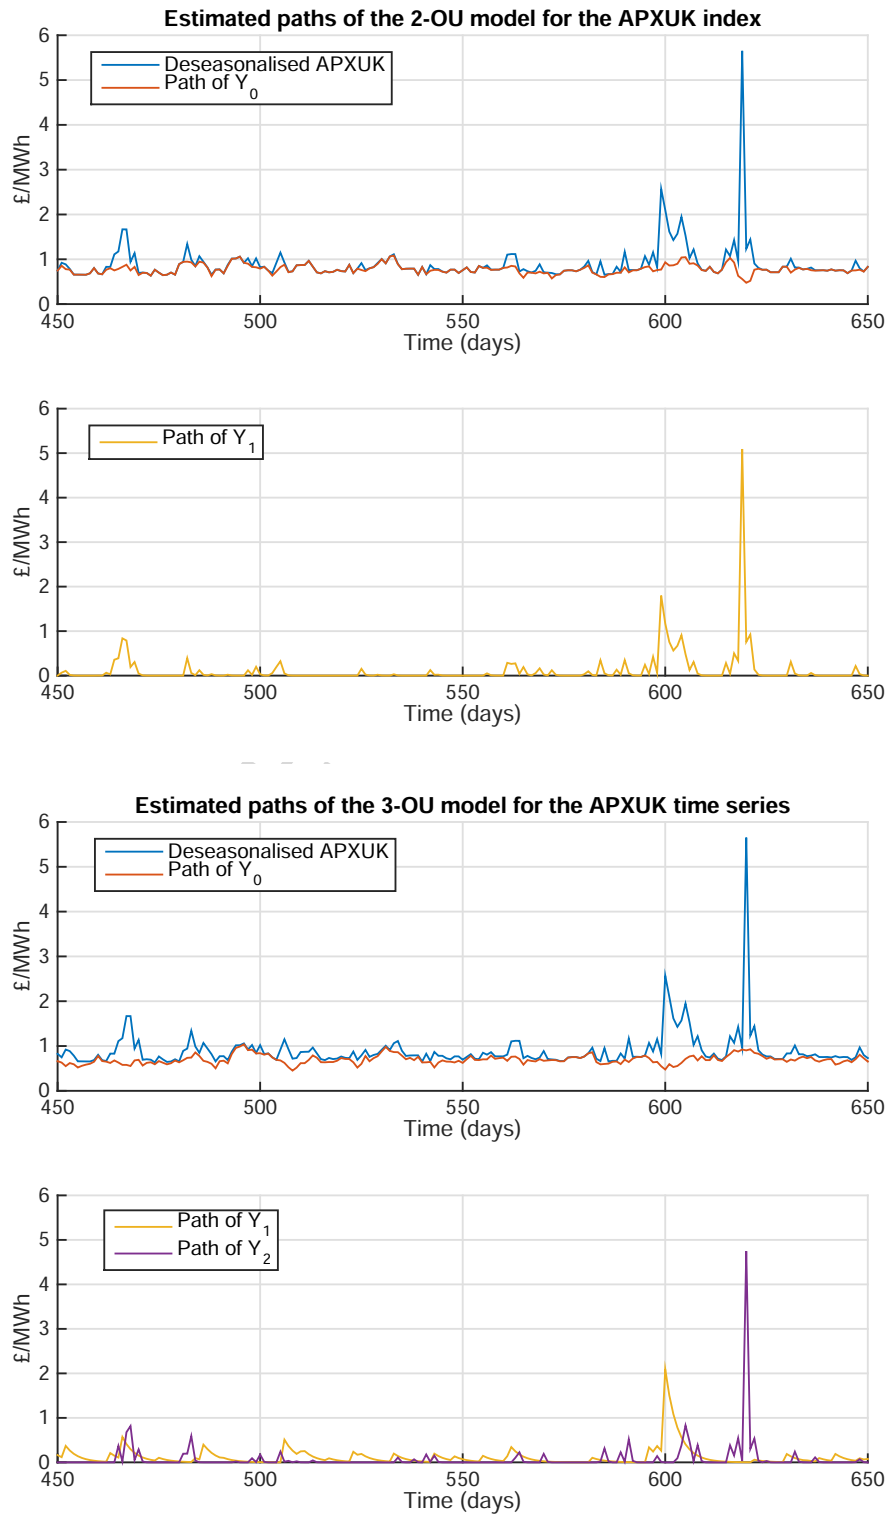


Figure 3: Samples from the final state of the Markov chain for the jump processes, plus a section of the deseasonalised 2001–2006 APXUK time series superimposed on the implied diffusion process Y_0 . Top: 2-OU model, bottom: 3-OU model.

630 yield quickly decaying large spikes. In contrast the runs of overlapping spikes are much reduced in the plots for the 3-OU model.

4.4.4. Diagnostics

The posterior predictive p -values for the diffusion process Y_0 are 0.0617 and 0.335 for the 2-OU and 3-OU model respectively. Taking a posterior predictive p -value in excess of 0.1 to be acceptable, two jump components are therefore required in order to give the diffusion process an acceptable fit on the basis of this diagnostic.

We also test the modelling assumption of Poisson jump arrivals with a constant intensity using the diagnostic described in Section 3.5. The constant jump intensity model is acceptable for the APXUK data with a posterior predictive p -value for the distribution of spike inter-arrival times of approximately 0.4, see Table 3. We note finally that the exponential model for jump sizes is acceptable in all cases (with the posterior predictive p -value exceeding 0.3).

Jump times of	APXUK (2001-6)		EEX (2000-6)		
	2-OU	3-OU	2-OU	3-OU	3-OU- I_1
Φ_1	0.0525	0.4003	0.0099	0.0185	0.1643
Φ_2	--	0.3089	--	0.4681	0.4738

Table 3: Posterior predictive p -values for the model of jump times for processes Φ_i .

4.5. 2000–2006 EEX data

We also calibrate 2-OU and 3-OU models to the 2000–2006 EEX data. Since exploratory analysis of the EEX dataset suggests the presence of frequent negative price spikes, our 3-OU model for the EEX series will differ from that for the APXUK dataset by specifying a negative sign for the second jump component Y_2 . Indeed, calibration of the 3-OU model with two positive jump components yields a posterior predictive p -value for the increments of the process Y_0 less than 0.005 and the posterior distributions for the parameters of the two positive jump components are not well separated (data not shown), indicating that the *positive* price spikes in the EEX market tend to be driven by one jump component. In the 3-OU model of (5) we therefore set $w_0 = w_1 = 1$ and $w_2 = -1$.

Further, taking into account the experience of past studies we consider both constant and periodic jump rates for the positive jump process $L_1(t)$, taking the periodic intensity function $I_1(\vartheta_1, t)$ given in (8) with $k = 130$ days (which corresponds to a period of one half-year). We refer to the latter model as the 3-OU- I_1 model. We take the same priors as in the APXUK studies above, now removing the restriction on the decay rates so that the prior for ρ_2 (or equivalently for

λ_2) is independent and distributed as that for ρ_1 (or λ_1), since with jumps of opposite direction there should be no issue of identifiability. Further, for the 3-OU- I_1 model the priors for η_1 and δ_1 are both $\text{Ga}(1, 10)$, while for θ_1 a $\text{U}(65, 195)$ prior is taken.

4.5.1. Number of jump components

Table 4 presents summary statistics for the posterior distributions of the parameters for both the 2-OU model and the 3-OU models applied to the 2000–2006 EEX data. The posterior diagnostic for the diffusion component Y_0 is acceptable for both three-component models while, as in the APXUK case, the two-component model does not appear to be satisfactory: the posterior p -value for Y_0 is equal to 0.0021 for the 2-OU model and equal to 0.192 and 0.26 for the 3-OU and 3-OU- I_1 models respectively. This is explained by the fact that negative price jumps are not accounted for with the 2-OU model, frequently resulting in large residuals for Y_0 .

There is agreement across the first two moments of the posterior distributions for all parameters common to the 3-OU and 3-OU- I_1 models. In contrast with the 2001–2006 APXUK dataset, however, the results for EEX support the presence of seasonality in the occurrence of price spikes, see Table 3. The constant jump intensity model appears to be unsatisfactory for the first jump component Φ_1 , with a corresponding p -value of 0.0185 in the 3-OU model, while the 3-OU- I_1 returns a p -value of 0.1643. Figure 2 displays the number of positive jumps on the EEX market by month, averaged over our posterior samples of the process Φ_1 in the 3-OU- I_1 model.

4.6. 2011 - 2015 data

Motivated by visual inspection of the price data in Figure 1, as discussed in Section 2, we wish to examine whether the statistical structure of the price data differs in periods before and after the global financial crises of 2007-8 and 2009. The models given by (5) were therefore calibrated to the APXUK and EEX indices over the sample period ranging from January 24, 2011 to February 16, 2015, and the simplest acceptable models were identified on the basis of posterior predictive p -values (again taking 0.1 as the minimum acceptable level). It may be seen from Table 5 that the 2011-2015 APXUK data supports the 2-OU model with one positive jump component. In order to discuss the statistics of the jump processes we will refer to the posterior mean values presented in Table 2 as ‘Old’ and in Table 6 as ‘New’. Although both the ‘New’ values lie between the corresponding ‘Old’ values, $\beta_1^{\text{Old}} < \beta_1^{\text{New}} < \beta_2^{\text{Old}}$ and $\lambda_1^{\text{Old}} < \lambda_1^{\text{New}} < \lambda_2^{\text{Old}}$, the new jump process cannot be interpreted as simply a statistical mixture of the two old jump processes since its intensity is lower than both of the old jump intensities.⁷ Indeed, on average

⁷Also, the sum of two jump OU processes with different mean reversion rates is statistically significantly different from one jump OU process and cannot, therefore, be successfully approximated by the latter.

Parameter	Prior properties			2-OU		2-OU- I_1		3-OU		3-OU- I_1	
	Prior	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
μ	N(1, 20 ²)	1	20	0.9978	0.0182	0.9954	(0.0184)	1.0171	0.0230	1.0146	0.0226
σ^2	IG(1.5, 0.005)	0.01	--	0.0271	0.0014	0.0269	(0.0014)	0.0122	0.0013	0.0119	0.0012
e^{-1/λ_0}	U(0, 1)	0.5	0.2887	0.7978	0.0166	0.7980	(0.0162)	0.8821	0.0151	0.8835	0.0150
(λ_0)	(IG(1, 1))	(- -)	(- -)	(4.4612)	(0.4203)	(4.4635)	(0.4071)	(8.1130)	(1.1438)	(8.2193)	(1.1698)
e^{-1/λ_1}	U(0, 1)	0.5	0.2887	0.1057	0.0267	0.1142	(0.0246)	0.1915	0.0313	0.1809	0.0344
(λ_1)	(IG(1, 1))	(- -)	(- -)	(0.4441)	(0.0512)	(0.4603)	(0.0465)	(0.6060)	(0.0601)	(0.5859)	(0.0657)
e^{-1/λ_2}	U(0, 1)	0.5	0.2887	--	--	--	--	0.2295	0.0348	0.2230	0.0384
(λ_2)	(IG(1, 1))	(- -)	(- -)	--	--	--	--	(0.6814)	(0.0709)	(0.6687)	(0.0775)
η_1	Ga(1, 10)	0.1	0.1	0.1049	0.0185	--	--	0.1274	0.0192	--	--
η_1^*	Ga(1, 10)	0.1	0.1	--	--	0.2881	(0.0619)	--	--	0.2515	0.0428
η_2	Ga(1, 10)	0.1	0.1	--	--	--	--	0.1438	0.0335	0.1422	0.0320
β_1	IG(1, 1)	--	--	1.0971	0.1435	1.0737	(0.1401)	0.9045	0.1088	0.8998	0.1052
β_2	IG(1, 1)	--	--	--	--	--	--	0.4176	0.0734	0.4308	0.0793
θ_1	U(65, 195)	130	37.5278	--	--	135.0048	(2.8474)	--	--	141.3725	3.5071
δ_1	Ga(1, 10)	0.1	0.1	--	--	0.6271	(0.1281)	--	--	0.3408	0.0884

Table 4: Prior distributions and posterior properties when fitting the 2- and 3-OU models to the 2000–2006 EEX data. * In the 2-OU- I_1 and 3-OU- I_1 models the parameter η_1 indicates the maximum jump rate of the periodic intensity function I_1 . The third OU component Y_2 of the models 3-OU and 3-OU- I_1 is negative. The indirect parameters λ_i are treated as described in the caption to Table 1.

	APXUK (2011-15)	EEX (2011-15)		
	2-OU	2-OU	2-OU- I_1	2-OU ⁻
Y_0	0.2167	0.0001	0.0012	0.1502
Jump times of Φ_1	0.3521	0.4219	0.3207	0.4452
Jump sizes of Φ_1	0.4574	0.5121	0.4554	0.4998

Table 5: Posterior predictive p -values for a range of models for the APXUK and EEX indices over the sample period January 24, 2011 to February 2, 2015.

the total number of jumps (of any size) per unit time is less than a third for the 2011-2015 data compared to 2001–2006.

From Table 5, the 2011-2015 EEX data in fact supports the 2-OU⁻ model which has a single, negative jump component (motivating the superscript minus in the notation). In this model the small number of larger upward price movements above the mean level in Figure 1 must therefore be accounted for by correspondingly large residuals in the diffusion component Y_0 , and this explains the relatively low predictive p -value (0.1502) for Y_0 . With regard to the statistics of the negative jump component, the values $\lambda_1, \eta_1, \beta_1$ in Table 6 should be compared to the values of $\lambda_2, \eta_2, \beta_2$ in Table 4. Since negative prices were introduced in this market on September 1, 2008 (Genoese et al., 2010), in general larger negative jumps were possible in the 2011-2015 data. Indeed the most significant downward jump in 2011-2015 was to a large negative price, and our finding $0.593 = \beta_1^{\text{New}} > \beta_2^{\text{Old}} = 0.4308$ is consistent with this change to the EEX market structure. For the diffusion component, the coefficient λ_0 decreases from 11.9 and 8.22 for the APXUK and EEX markets respectively to approximately 3.6, a value which happens to be consistent across both markets.

5. Discussion

5.1. Scope of contribution

We have shown that multiple components can be required to obtain statistically adequate mean-reverting models of electricity spot prices and, further, that the required combination of components can change over time. To clarify the value of this contribution we note that spot price models have two principal areas of application in the literature. Firstly, price forecasting is concerned with the prediction of prices over future time points or periods given the current and past values of relevant variables (see for example Weron (2007)). As such it is not concerned with the detailed statistical properties of spot price *trajectories* such as the long-run statistical patterns in price spikes, which are the object of our work. Instead studies (such as ours) of spot price dynamics are suitable both for derivative pricing (Hull, 2009) and operational analyses in

Parameter	APXUK (2011-15)		EEX (2011-15)	
	2-OU model		2-OU ⁻ model	
	Mean	SD	Mean	SD
μ	0.9865	0.0095	1.0480	0.0149
σ^2	0.0071	0.0006	0.0170	0.0016
e^{-1/λ_0}	0.7537	0.0232	0.7590	0.0251
(λ_0)	(3.5727)	(0.3976)	(3.6736)	(0.4526)
e^{-1/λ_1}	0.2104	0.0394	0.1941	0.0381
(λ_1)	(0.6435)	(0.0781)	(0.6115)	(0.0741)
η_1	0.1172	0.0324	0.1105	0.0310
β_1	0.3981	0.0921	0.5930	0.1325

Table 6: Posterior properties obtained when calibrating the 2-OU (one positive jump component) and 2-OU⁻ (one negative jump component) models to the APXUK and EEX datasets respectively. The sample period is January 24, 2011 to February 16, 2015.

the real options framework (see for example Kitapbayev et al. (2015); Moriarty and Palczewski (2017)). In derivative pricing a main goal is to combine the latter models with observed derivative prices to construct so-called *risk-neutral* or *martingale* probability measures. Thus while model parameters are inferred from derivative prices it is important to identify the right class of price models, and our work provides an approach to this question via posterior predictive checking. In contrast, in real options analyses the fact that real projects are not traded means that the *physical* or *historic* probability measure is often the one used. In this context our work provides an approach both to model specification and to the calibration of model parameters to historic data.

The methodological advantages of our approach to calibration, which aims to make minimal assumptions about the spike processes, are confirmed by the 2000–2006 APXUK analysis. In mean-reverting models jumps do not immediately vanish but instead decay over time. This means that jump components, particularly those having the same direction, can interact when superposed. Inference on the individual spike components is then more challenging and simple signal processing approaches (c.f. Meyer-Brandis and Tankov (2008)), such as the use of thresholds to identify jumps, are rendered unsuitable. Nevertheless we have shown that a statistically adequate model can be extracted. The ability of our MCMC procedure to distinguish spikes in the same direction is confirmed in the electronic appendix using simulated price data. There, Figure A.2 plots the simulated jumps (in red) and a visual representation of the posterior distri-

bution of the latent jump processes (blue, see Section A.3 for details) so that the agreement can be assessed.

5.2. Fundamental drivers

In this section we attempt to relate our empirical results to their underlying physical and economic drivers. Negative price spikes are associated with the priority given to wind energy in the spot market (Benth, 2013). A glut in wind power production can lead to a corresponding decrease in demand for other sources of generation. It can be impossible for conventional generators to reduce production sufficiently so they may temporarily accept low (or even negative) prices. In support of this analysis, we observed that the inclusion of a negative jump component was necessary for adequate modelling of the EEX data in both periods. Further the negative jumps had higher mean size in the later period, a change consistent with the increasing penetration of renewable generation.

While preprocessing the data we removed a deterministic seasonal component from the spot prices. However in some electricity markets (particularly in the US and Europe) seasonality has also been observed in the frequency of price spikes (Geman and Roncoroni, 2006; Benth et al., 2012). A priori this may be explained by greater levels of stress in the power system during the extremes of seasonal variation in weather due, for example, to heating load during cold snaps. The presence of jump seasonality was suggested in the 2000–2006 EEX data as illustrated in Figure 2, which displays the number of positive jumps on the EEX market by month, averaged over samples from our MCMC procedure. (It should be noted that Figure 2 is indicative and does not represent direct statistical estimates of jump frequency in spot prices.) Indeed for the latter series it was necessary to incorporate seasonality in the arrival rate of the positive jump component in order to obtain a statistically adequate model. In contrast seasonal jump components were not statistically necessary for the APXUK data in either period, which may be related to the less severe extremes of UK winter weather.

As presented in Section 4.2, there is statistical evidence for a reduction across both markets in both the number of positive jump components and the frequency of positive price spikes after the global financial crises of 2007–8 and 2009. It is true that in both the UK and Germany, electricity consumption generally increased in the period 2000–2006 and was generally level or decreased during 2011–2015.⁸ It follows that the power systems under study faced less stress from constraints in either production or transmission capacity during the latter period, and this is consistent with the observed reduction in positive price spikes. It must be noted however that

⁸<http://data.worldbank.org>

the outlook for the future is somewhat different, with changes on the supply side including the decommissioning of ageing and carbon-intensive conventional generation and increased reliance on intermittent generation (see, for example, National Grid plc (2016)), suggesting that positive spikes may return to the electricity spot market.

6. Conclusions

By modelling mean-reverting deseasonalised electricity spot prices as the sum of a diffusion process and multiple signed jump processes of deterministic intensity, and applying a Bayesian calibration procedure and posterior diagnostics, we have identified a class of multi-factor models suitable for modelling empirical prices across two different markets and two different time periods. In contrast with several recent studies using stochastic volatility models we have employed multiple signed jump components, albeit with simpler deterministic volatilities (either constant, or deterministic and periodic). This approach allows straightforward comparison of the statistical structure of prices across different markets and time periods: each model has a number of signed jump components with distinguished jump intensities, decay rates and size distributions. In both the APXUK and EEX markets it was found that the statistical structure of the price series differs before and after the period 2007-2010 and, in particular, that the number of positive jump components decreased (from 2 to 1 and 1 to 0 respectively), with the mean reversion speed of the diffusive price component increasing in both markets. Seasonality in the jump intensity was found to be necessary only in the earlier (2000–2006) EEX data and only for its positive jump component.

Acknowledgements

JM was supported by grants EP/K00557X/1 and EP/K00557X/2 and JG was supported by grant EP/I031650/1 from the UK Engineering and Physical Sciences Research Council. JP was supported in part by MNiSzW grant UMO-2012/07/B/ST1/03298.

References

Barndorff-Nielsen, O. E. and Shephard, N. (2001). Non-Gaussian Ornstein–Uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):167–241.

- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 57(1):289–300.
- Benth, F. E. (2013). Stochastic volatility and dependency in energy markets: Multi-factor modelling. In Henderson, V. and Sircar, R., editors, *Paris-Princeton Lectures on Mathematical Finance 2013*. Springer.
- Benth, F. E., Kallsen, J., and Meyer-Brandis, T. (2007). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, 14(2):153–169.
- Benth, F. E., Kiesel, R., and Nazarova, A. (2012). A critical empirical study of three electricity spot price models. *Energy Economics*, 34(5):1589–1616.
- Benth, F. E. and Meyer-Brandis, T. (2009). The information premium for non-storable commodities. *Journal of Energy Markets*, 2(3):111–140.
- Brix, A. F. (2015). *Estimation of Continuous Time Models Driven by Lévy Processes*. PhD thesis, Aarhus University, School of Business and Social Sciences.
- Brunner, C. (2014). Changes in electricity spot price formation in Germany caused by a high share of renewable energies. *Energy Systems*, 5(1):45–64.
- Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management*. Lacima Publications.
- Fanone, E., Gamba, A., and Prokopczuk, M. (2013). The case of negative day-ahead electricity prices. *Energy Economics*, 35:22–34.
- Frühwirth-Schnatter, S. and Sögner, L. (2009). Bayesian estimation of stochastic volatility models based on OU processes with marginal Gamma law. *Annals of the Institute of Statistical Mathematics*, 61(1):159–179.
- Gelman, A. (2003). A Bayesian formulation of exploratory data analysis and goodness-of-fit testing. *International Statistical Review*, 71(2):369–382.
- Geman, H. and Roncoroni, A. (2006). Understanding the fine structure of electricity prices. *The Journal of Business*, 79(3):1225–1261.

- Genoese, F., Genoese, M., and Wietschel, M. (2010). Occurrence of negative prices on the German spot market for electricity and their influence on balancing power markets. In *7th International Conference on the European Energy Market (EEM)*, pages 1–6. IEEE.
- Geyer, C. J. and Møller, J. (1994). Simulation procedures and likelihood inference for spatial point processes. *Scandinavian Journal of Statistics*, 21(4):359–373.
- Gonzalez, J. (2015). *Modelling and controlling risk in energy systems*. PhD thesis, The University of Manchester, Manchester, UK.
- Green, R. and Nossman, M. (2008). Markov chain Monte Carlo estimation of a multi-factor jump diffusion model for power prices. *The Journal of Energy Markets*, 1(4):65–90.
- Griffin, J. E. and Steel, M. F. J. (2006). Inference with non-Gaussian Ornstein-Uhlenbeck processes for stochastic volatility. *Journal of Econometrics*, 134(2):605–644.
- Hull, J. (2009). *Options, Futures and Other Derivatives*. Pearson/Prentice Hall, 7th edition.
- Jacquier, E., Polson, N. G., and Rossi, P. E. (1994). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business & Economic Statistics*, 12(4):371–389.
- Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3):361–393.
- Kingman, J. F. C. (1992). *Poisson processes*. Oxford University Press.
- Kitapbayev, Y., Moriarty, J., and Mancarella, P. (2015). Stochastic control and real options valuation of thermal storage-enabled demand response from flexible district energy systems. *Applied Energy*, 137:823–831.
- Kutoyants, Y. A. (1998). *Statistical Inference for Spatial Poisson Processes*. Lecture Notes in Statistics. Springer, New York.
- Lucia, J. J. and Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5(1):5–50.
- Meyer-Brandis, T. and Tankov, P. (2008). Multi-factor jump-diffusion models of electricity prices. *International Journal of Theoretical and Applied Finance*, 11(05):503–528.
- Moriarty, J. and Palczewski, J. (2017). Real option valuation for reserve capacity. *European Journal of Operational Research*, 257(1):251–260.

- National Grid plc (2016). Electricity ten year statement. <http://www2.nationalgrid.com/UK/Industry-information/Future-of-Energy/Electricity-Ten-Year-Statement/>. Accessed: 2017-04-10.
- Roberts, G. O., Papaspiliopoulos, O., and Dellaportas, P. (2004). Bayesian inference for non-Gaussian Ornstein-Uhlenbeck stochastic volatility processes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 66(2):369–393.
- Rubin, D. B. (1984). Bayesianly justifiable and relevant frequency calculations for the applied statistician. *The Annals of Statistics*, 12(4):1151–1172.
- Rydén, T. et al. (2008). EM versus Markov chain Monte Carlo for estimation of hidden Markov models: A computational perspective. *Bayesian Analysis*, 3(4):659–688.
- Schwartz, E. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911.
- Seifert, J. and Uhrig-Homburg, M. (2007). Modelling jumps in electricity prices: theory and empirical evidence. *Review of Derivatives Research*, 10(1):59–85.
- Stephenson, P. and Paun, M. (2001). Electricity market trading. *Power Engineering Journal*, 15(6):277–288.
- Tierney, L. (1998). A note on Metropolis-Hastings kernels for general state spaces. *The Annals of Applied Probability*, 8(1):1–9.
- Weron, R. (2007). *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*. John Wiley & Sons.
- Würzburg, K., Labandeira, X., and Linares, P. (2013). Renewable generation and electricity prices: Taking stock and new evidence for germany and austria. *Energy Economics*, 40:S159–S171.

Highlights

- We perform Bayesian calibration of multifactor electricity spot price models via MCMC.
- Two or three signed jump components are sufficient for EEX and APXUK markets.
- Fewer positive jump components are found in recent data.
- Seasonal jump intensity is needed only in earlier EEX data.