Modelling Electricity Prices with Forward Looking Capacity Constraints

— Source link ↗

Álvaro Cartea, Marcelo G. Figueroa, Hélyette Geman

Institutions: University of Oxford, Birkbeck, University of London, University of London

Related papers:

- Regime jumps in electricity prices

- A structural risk-neutral model of electricity prices

- Modelling the Structure of Long-Term Electricity Forward Prices at Nord Pool

- Random movements of power prices in competitive markets: a hybrid model approach

- Probabilistic Mid- and Long-Term Electricity Price Forecasting

Modelling Electricity Prices with Forward Looking Capacity Constraints

ÁLVARO CARTEA*, MARCELO G. FIGUEROA* & HÉLYETTE GEMAN*,**

*Commodities Finance Centre, Birkbeck College, University of London, London, UK, **ESSEC Business School, Cergy-Pontoise, France

ABSTRACT   We present a spot price model for wholesale electricity prices which incorporates forward looking information that is available to all market players. We focus on information that measures the extent to which the capacity of the England and Wales generation park will be constrained over the next 52 weeks. We propose a measure of 'tight market conditions', based on capacity constraints, which identifies the weeks of the year when price spikes are more likely to occur. We show that the incorporation of this type of forward looking information, not uncommon in electricity markets, improves the modelling of spikes (timing and magnitude) and the different speeds of mean reversion.

1. Introduction

The study of commodities and its role in the economy has a long tradition in financial economics. In theory, there are two key factors that link commodity spot and futures prices: storage costs and convenience yield. For example, early studies use the theory of storage to frame the relationship between spot and forward prices or between two forward prices of a commodity, such as the seminal work of Working (1949) and Brennan (1958). More recently, a number of models have assumed that these key factors are not deterministic and a series of alternative approaches have examined pricing applications assuming stochastic convenience yield (Gibson and Schwartz, 1990; Schwartz, 1997), or both stochastic convenience yield and interest rates (Casassus and Collin-Dufresne, 2005; Hilliard and Reis, 1998; Miltersen and Schwartz, 1998). Additional work based on the theory of storage includes that of Fama and French (1988), who look at business cycles and the behaviour of metal prices. Eydeland and Geman (1999) argue that, in the case of

Correspondence Address: Álvaro Cartea, Commodities Finance Centre, Birkbeck College, University of London, London, UK. Email: a.cartea@bbk.ac.uk

electricity, the existence of a convenience yield and the spot–forward relationship essentially collapses.

Another strand of the commodity literature has focused on understanding why futures are biased estimators of spot prices and what is the economic and financial rationale behind the emergence of situations of backwardation and contango in these markets. Studies along these lines include the work of Hirshleifer (1988), which looks at equilibrium models that explain how markets are organized and how market players hedge their positions in the forward markets. Further, the work of Fama and French (1987) aims at explaining the connection between spot and futures prices under storage considerations or by considering futures prices as expected spot prices plus a risk premium; Litzenberger and Rabinowitz (1995) examine the relation between spot and futures oil prices. Routledge *et al.* (2000) look at an equilibrium model of the term structure of forward prices for storable commodities. Bessembinder and Lemmon (2002), among other things, introduce two categories of players in electricity markets in order to derive an equilibrium approach in power markets.

Commodities have always been considered as an asset class very different from the more traditional ones such as equity, bonds, etc. However, although electricity is considered to be a member of the commodity class, it had not been the focus of attention in the literature as power markets were regulated around the world. It was not until the onset of deregulation in Western Europe, North America, New Zealand and Australia in the early 1990s, that the trading of power and related instruments became an interesting and attractive field of research for the academic community and industry participants.

One of the most distinctive features of power markets is that electricity cannot be stored, or it is too expensive to do so, and this characteristic alone is traditionally considered as the fundamental and unique reason that places electricity, within its asset class, in a league of its own. The inability to store electricity in a cost-effective manner renders useless traditional arguments of cash and carry, dynamic hedging and buy-and-hold strategies which are considered the cornerstone of no-arbitrage arguments employed to price the majority of traded instruments.

However, the inability to store electricity is not solely responsible for the exclusive idiosyncrasies of power markets. Price formation in these markets is intrinsically different because, unlike any other financial assets, there are further characteristics that contribute to the unique behaviour of prices, either on an equal footing or by magnifying the effects of non-storability. These characteristics include: number of players (generators and retailers), composition of the generation park and market design.

The number of players in this market is very small. On the generation side, there are few players per region or country as well as few retailers. Thus, the actions or behaviour of one player alone, for example a generating company, will be able to influence equilibrium spot and futures prices.

Further, the composition of the generation park also plays a crucial role in the determination of price dynamics. On the supply side, the 'supply stack' reflects the amount of electricity that generators are willing to produce at different prices. Therefore, the composition and heterogeneity of the generation park will determine the shape of the curve which usually exhibits a 'kink' to reflect a step change between

the low and high marginal production costs of the different plants. On the other hand, aggregate demand is hovering around this kink, making clearing prices very susceptible to abrupt changes in supply and/or demand. The supply stack present in any market is therefore dominated by the composition of its generation capacity and, consequently, seasonal levels and trajectile properties of prices will inherit the idiosyncrasies of the generation park.

Finally, the architecture and design of these markets clearly affect price levels and price dynamics. One clear example is how the behaviour of prices in England and Wales changed when the mechanism by which the market operated underwent a series of alterations as a result of the introduction of the New Electricity Trading Arrangements in March 2001. Moreover, the running of the electricity market in England and Wales, plus the recent addition of Scotland, is in the hands, by design, of the System Operator known as the National Grid Company (NGC).

One of the principal activities of the NGC is to organize the electricity market which requires, among other duties, that it provides a wealth of information to market players. There are two types of information that are publicly available: historical or out-turn data and forecasts. Historical data covers areas that include out-turn load (realized demand), out-turn capacity margins, fuel usage, etc. On the other hand, the NGC publishes forecasts on a range of crucial variables including demand, capacity margins, indicated demand and indicated generation.[1] These forecasts clearly affect the behaviour of players in terms of prices and quantities they are willing to buy or sell depending on market conditions. For example, if the NGC forecasts low levels of demand, some generators will see this as an indication that some or none of their capacity will be called on by the NGC to satisfy demand. Similarly, if the NGC forecasts of expected demand is high and capacity margin is low, this sends a clear signal to companies that own versatile generators, such as peaking plants, that they might be called on to generate power and that clearing prices will be high to meet the costs of the marginal plants coming online.

The unique characteristics of power markets therefore call for tailor-made models that can replicate the stylized facts of their price dynamics, but, ideally, must also take into account other information and characteristics of this peculiar market. The main contribution in this article is to argue for the first time that electricity spot models should incorporate forward looking data which are made available to all players by the System Operator, for demand and capacity forecasts are already incorporated in the information set upon which market players make buy and sell decisions.

The rest of this article is organized as follows. Section 2 reviews existing models and approaches proposed in the literature. In Section 3 we present the model and discuss in detail the inclusion of forward looking data: demand and capacity forecasts. Section 4 discusses the calibration of the model. Section 5 shows the results and Section 6 concludes.

2. Existing Models and Approaches in Power Markets

The most popular approach to model power prices adopts a reduced-form approach. Models in this category are mainly concerned with capturing the three main characteristics of its price dynamics (seasonality, large jumps and mean reversion)

with a view to pricing forwards, options and other derivatives such as interruptible contracts.

Recent models agree on the importance of including a deterministic seasonal trend which is estimated from historical data. For example, in Lucía and Schwartz (2002) the deterministic seasonal component is responsible for capturing any relevant predictable component of electricity prices. They include a constant and two terms that are responsible for the variation in price levels between working and non-working days and the seasonal evolution of prices throughout the year. Similarly, Escribano *et al.* (2002) take into account weekly and monthly seasonality by means of daily and monthly dummies and sinusoidal functions. Finally, Cartea and Figueroa (2005) construct an historical deterministic function by matching a Fourier series to the averages of the historical months.

It is important to note that although the inclusion of a deterministic seasonal component, based on historical data, is widespread in the literature, the question of how stable the seasonality function is must not be overlooked. For instance, the seasonal pattern in the Nord Pool market, where over 50% of the generation capacity is based on hydro, is linked to the amount of snow which has recently fallen, how full the reservoirs are and when and where will snow melt. Consequently, average prices in any given year in the Scandinavian market will reflect whether the previous year was wet, normal or dry. This raises concerns concerning the ability of an historical seasonal trend to capture average levels in markets like Nord Pool that depend so heavily on hydro-power. Similarly, in the UK, the principal sources of generation are coal and gas plants, thus future generation costs will also depend on how these markets are developing. Therefore, the composition of the generation park is one of the driving forces behind average power prices reflected in seasonal trends. The use of estimates for seasonal trends that are solely based on historical data will fail to reflect situations where fuel prices or reservoir levels have diverged from historical figures.

Another key component of spot-based models in electricity markets is the inclusion of jumps. Different versions of mean reverting jump-diffusion models applied to electricity markets have been considered. For instance, in Cartea and Figueroa (2005) and Hambly *et al.* (2007), large jumps or spikes occur according to a homogeneous Poisson process where the intensity parameter of the counting process is given by the historical average number of jumps. Geman and Roncoroni (2006) and Benth *et al.* (2007) improve on this by considering an inhomogeneous Poisson process where the jump intensity is a deterministic function, also based on historical data.

However sophisticated and accurate the modelling of the stochastic arrival of jumps is, the use of historical data to estimate or calibrate its parameters will only provide a broad pattern for the jump intensity that triggers large movements in electricity prices. Although this approach might work on average, it will clearly miss the timing of crucial events produced by constraints in the system such as plant outages, surges in demand, and periods where idle capacity is at low levels, some of which will not come as a surprise to market players.

In this article we focus on the England and Wales wholesale electricity market. To model price dynamics, rather than directly modelling the intensity parameter of jumps, we consider a regime-switching model where regime changes are governed by

a deterministic parameter alternating between a high-intensity level, labelled 'high regime', during which the probability of observing spikes is high, and a low-intensity level, labelled 'low regime', where prices may exhibit moderate jumps, if any. Central to our model is the way in which the spot dynamics switch between these two regimes. Changes between these two states of activity will be determined by an exogenous switching variable that is linked to publicly available forecasts, as will be explained in detail in the following section. Moreover, we propose the use of a seasonal component based on gas forward prices instead of the customary approach based on historical spot prices.

3. The Model

A simple inspection of market data reveals that mean reversion is present in wholesale electricity prices and the speed at which prices mean revert is not constant. Large positive spikes in the market are associated with unexpected shocks to the demand or supply side. For example, unanticipated surges in demand, that can only be met by calling on plants with high marginal production costs, result in short-term high prices. Similarly, shocks to the supply side come in the form of unforeseen plant failures that must be resolved by calling on other sources of generation with higher marginal costs. Generally, these high-impact deviations are short-lived and markets are back to normal within a few days. Furthermore, other price shocks not considered to be spikes, although being the result of unexpected variations in both demand and supply that are also short-lived, exhibit a mean reversion which is considerably slower than the speed at which large spikes die out.

Hence, we propose a model that captures, among other key elements such as seasonality, both large and normal deviations with different speeds of mean reversion. We assume that the log-price process is the result of a deterministic seasonal component $g(t)$ and a stochastic process $y(t)$

$$\ln S(t) = g(t) + y(t), \tag{1}$$

where $y(t)$, driven by three independent processes, satisfies the stochastic differential equation (SDE)

$$dy(t) = -\beta(t)y(t)dt + \sigma\,dW(t) + \rho(t)\ln J\,dN(t) + (1 - \rho(t))dZ(t). \tag{2}$$

Here, $\beta(t)$ is a time-dependent speed of mean reversion, $W(t)$ is a standard Brownian motion, $N(t)$ is a Poisson process with intensity ℓ, J are iid shocks (responsible for the large spikes) and $dZ(t)$ are the increments of a Lévy process with triplet (σ_Z, θ, M).[2] Moreover, $\rho(t)$ is an exogenous switching parameter such that

$$\rho(t) = \begin{cases} 1, & \text{for the high regime,} \\ 0, & \text{for the low regime,} \end{cases} \tag{3}$$

where the 'high regime' refers to periods where electricity prices may exhibit large spikes that are introduced in Equation (2) through the stochastic shocks $\ln J\,dN(t)$. And, on the other hand, the 'low regime' captures periods where price variations are present in the form of increments of the process $Z(t)$, where we assume that the presence of huge spikes is unlikely.[3]

To solve (2), we first note that

$$d\left(y(t)e^{\int_0^t \beta(u)\,du}\right) = e^{\int_0^t \beta(u)\,du}\,dy + y(t)\beta(t)e^{\int_0^t \beta(u)\,du}\,dt. \tag{4}$$

Rearranging Equation (2) and multiplying it through by $e^{\int_0^t \beta(u)\,du}$, we obtain

$$d\left(y(t)e^{\int_0^t \beta(u)\,du}\right) = e^{\int_0^t \beta(u)\,du}\,\sigma\,dW(t) + e^{\int_0^t \beta(u)\,du}\,\rho(t)\ln J\,dN(t)$$

$$+ e^{\int_0^t \beta(u)\,du}\,(1-\rho(t))dZ(t).$$

Finally, integrating between t_0 and t and re-arranging, we obtain

$$y(t) = e^{-\int_{t_0}^t \beta(u)\,du}\,y_{t_0} + \sigma\int_{t_0}^t e^{-\int_s^t \beta(u)\,du}\,dW(s)$$

$$+ \int_{t_0}^t e^{-\int_s^t \beta(u)\,du}\,\rho(s)\ln J\,dN(s)$$

$$+ \int_{t_0}^t e^{-\int_s^t \beta(u)\,du}\,(1-\rho(s))dZ(s). \tag{5}$$

It is clear that the model described in Equation (2), together with the specification of $\rho(t)$, alternates between periods when spikes are likely to occur ($\rho(t)=1$), and periods where prices also exhibit a great deal of variability, but large spikes are less likely to take place ($\rho(t)=0$). However, for the model (2) to capture the crucial feature that the speed at which deviations fade away will depend on the magnitude of the unexpected movements, we need to assume that $\beta(t)$ is also a function of the deterministic function $\rho(t)$. Consequently, we define

$$\beta(t) = \alpha^H\rho(t) + \alpha^L(1-\rho(t)), \tag{6}$$

with $\rho(t)$ being defined in (3) and $\alpha^H > \alpha^L > 0$. Then the stochastic differential Equation (2) may be written as

$$dy(t) = -\left[\alpha^H\rho(t) + \alpha^L(1-\rho(t))\right]y(t)\,dt + \sigma\,dW(t)$$

$$+ \rho(t)\ln J\,dN(t) + (1-\rho(t))\,dZ(t). \tag{7}$$

We note that another approach would be to extend the class of one factor mean reverting jump diffusion models by adding a second stochastic process and a switching parameter in order to alternate between processes. This other approach, although appealing from a mathematical standpoint, would be extremely difficult to calibrate or estimate since one cannot discern whether the shocks to the spot dynamics come from a high or low regime.

3.1 The Deterministic Switching Component ρ(t)

To motivate our choice for an appropriate forward-looking deterministic function $\rho(t)$, with the functional form (3), we will discuss what are the principal sources of generation in England and Wales and what is the production capacity of power producers.[4] We also discuss what type of information on generation capacity is publicly available.

The composition or structure of the generation park plays a crucial role in the determination of the level and volatility of power prices. Different power plants come on line at different price levels and the level of entry is determined by marginal generation costs which, in this particular market, increase at an increasing rate with output. Therefore, the shape of the power supply curve or supply stack reflects the degree of heterogeneity in the generation capabilities and marginal costs, which for most, if not all, power markets becomes very steep once expensive plants come on line.

In England and Wales the main fuel sources for generation have been coal, nuclear and gas. In recent years there has been a steady increase in the use of gas as a main fuel source until becoming the principal source of generation in 2004. This has been accompanied by a slight decrease in the dependence on coal, oil and nuclear and a very small increase in renewable sources. The latter can be explained in part by a shift towards cleaner sources of energy in order to meet targets set by the government (for instance, the Climate Change and Sustainable Energy Act 2006). In the case of nuclear energy, the decrease is explained by the fact that the existing plants are approaching the end of their expected life-cycles. Data reflecting these issues are summarized in Figure 1.

Table 1 shows in greater detail the production capacity, by type of generation, of power producers in England and Wales over the period 2003–2005. We observe that,
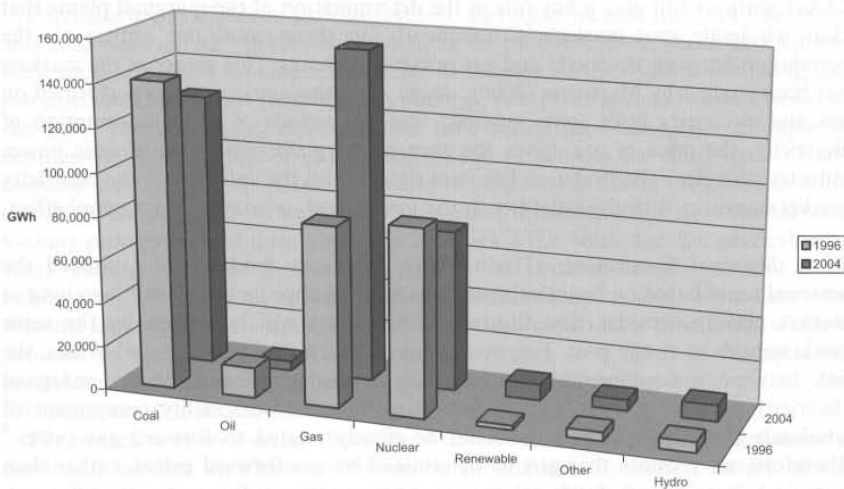
Figure 1. Fuel used in electricity generation in 1996 and 2004; data source: www.dti.gov.uk/energy/statistics.

Table 1. Type of plant and capacity of major power producers in the UK.

	2003		2004		2005	
Major power producers in the UK						
Total declared net capability (MWh)	71,465		73,277		74,041	
Conventional steam stations						
Coal fired	22,524	31.52%	22,639	31.68%	22,627	31.66%
Oil fired	2930	4.10%	2930	4.10%	3262	4.56%
Mixed or dual fired	6413	8.97%	6413	8.97%	6403	8.96%
Combined cycle gas turbine stations	22,037	30.84%	23,783	33.28%	24,373	34.10%
Nuclear stations	11,852	16.58%	11,852	16.58%	11,852	16.58%
Gas turbines and oil engines	1537	2.15%	1485	2.08%	1346	1.88%
Hydro-electric stations						
Natural flow	1267	1.77%	1270	1.78%	1273	1.78%
Pumped storage	2788	3.90%	2788	3.90%	2788	3.90%
Renewables other than hydro	117	0.16%	117	0.16%	117	0.16%

in 2005, combined cycle gas turbine (CCGT) stations and conventional steam stations accounted for 80% of the market's net capability of power production. Although coal-fired units and CCGT represent an almost similar share of the total capacity, the main difference between them stems from the fact that coal-fired units have lower marginal generation costs, but are less flexible in their response to address sudden demand fluctuations. One expects this flexibility in generation provided by CCGT stations to be more conspicuous during periods of high volatility. Economic rationale indicates that during periods of high uncertainty, CCGT stations will play a key role in the determination of the marginal plants that clear wholesale spot markets. Consequently, in these cases, one anticipates the correlation between electricity and gas prices to increase. This aspect of the markets has been studied by Martijena (2006), where an econometric analysis performed on gas and electricity price series suggests that, in periods of high consumption of electricity, the price of gas drives the average marginal cost in the electric power industry, therefore their relation becomes tighter, and the volatility in the electricity market increases with the volatility in the gas market, creating a 'contagion' effect.

3.1.1 Seasonal Component. Traditionally, electricity models had estimated the seasonal trend based on historical spot data, yet this approach is plausible so long as market players consider that future average prices will be driven by the same fundamentals as in the past. For example, in markets like England and Wales, the link between a fundamental variable such as gas prices and average marginal electricity costs is present, clearly indicating that the seasonality component of wholesale electricity prices $g(t)$ should be closely related to forward gas prices.[5] Therefore, we propose that $g(t)$ be determined by gas forward prices, rather than historical electricity wholesale prices, since at time t the gas forward curve is known, and it already reflects market expectations of the trend levels of average marginal electricity costs.

3.2 The 'High' and 'Low' Regimes

As mentioned in the introduction, we consider a regime-switching model where regime changes are governed by a forward looking deterministic parameter responsible for alternating between the high regime, where the probability of observing spikes is high, and the low regime, where prices may exhibit jumps, but huge spikes are unlikely. Changes between these two states of activity are determined in the model by the exogenous switching variable $\rho(t)$, which is calculated using publicly available forecasts.

To construct the deterministic function $\rho(t)$ we first look at the relation between National Demand Forecast $D(t_0, t_p)$ and forecasted Generation Capacity $C(t_0, t_p)$ which are both calculated at time t_0 for an upcoming period t_p. We define their ratio as

$$\varrho(t_0, t_p) = \frac{D(t_0, t_p)}{C(t_0, t_p)}. \tag{8}$$

These data are publicly available and supplied by the NGC and depending on the resolution of the data provided, the period t_p may be a half-hour slot, a day or a week.[6] Moreover, the NGC National Demand Forecast $D(t_0, t_p)$ is based on historically metered generation output for Great Britain; it takes into account transmission losses and includes station transformer load, pump storage demand and inter-connector demand. Similarly, the National Surplus forecast $NS(t_0, t_p)$ is based on forecasts of generator availability. Although unlikely, it is interesting to note that the ratio (8) may take values higher than unity, a situation that has occurred in the English and Wales market. In such circumstances, one would expect that, as time and other uncertainties unravel, market forces will act so that demand and supply meet at an equilibrium price.

Forecasts are available in various formats. Examples include the 2–14 day ahead and 2–52 week ahead.[7] In this work we focus on the NGC 2–52 weeks data set and draw on this forward looking information as follows. At every point in time t_m one can calculate the forecast ratio $\varrho(t_m, t_n)$ where $m, n \in \{1, 2, 3, ..., 52\}$ denote weeks of the year. The notation t_n represents the time period in week n. For example, t_2 denotes the second week of any calendar year, t_3 denotes the third week of any calendar year, and so on. Hence, for example, $\varrho(t_{36}, t_{38})$ is the forward looking ratio calculated during week 36 for week 38. Another example, $\varrho(t_{36}, t_4)$ is the forward looking ratio calculated during week 36 for week 4 (i.e. week 4 in the next calendar year). Furthermore, the 2–52 week forecasts are updated every week and made public every Thursday at around midday. Consequently, when we construct the forward looking ratio we update it every week and construct a time series which is depicted as 'crosses' in Figure 2. We point out that, as a result of the updating procedure, the ratio series will show the forecast made during week t_{n-1} for the following week t_n, $\varrho(t_{n-1}, t_n)$.[8]

The objective is to overlay this ratio with out-turn spot prices to determine a threshold, labelled δ, which enables us to differentiate for which ranges of the demand to capacity ratio (8) the market is more or less likely to exhibit large spikes. Intuitively, one would expect that the larger the ratio $\varrho(t_m, t_n)$, the tighter the generation will become, leaving little manoeuvrability for the NGC to call on idle
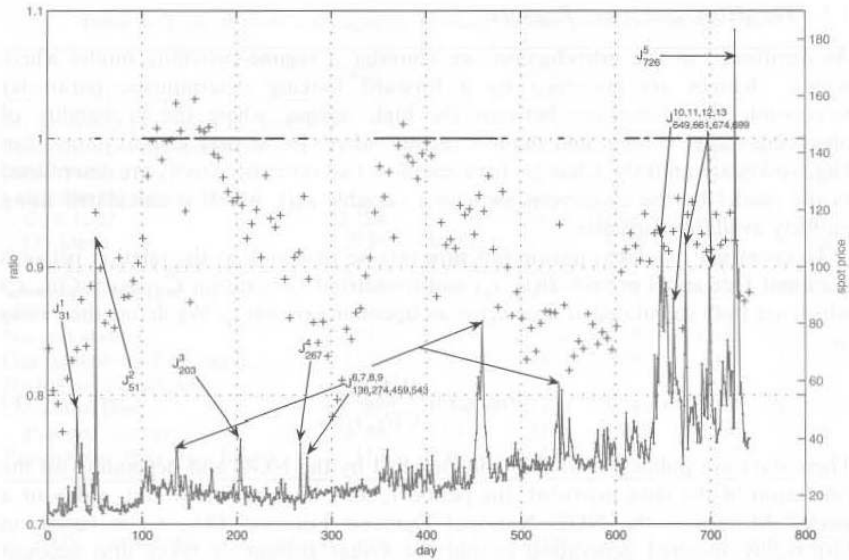
Figure 2. The spot price series is represented by the solid line and '•' marker. The series of the ratio, $\varrho(t_{n-1}, t_n)$, is represented with the '+' marker, and the deterministic function $\rho(t)$ is represented by a solid line. The identified spikes in the series are represented by J_j^i, where i denotes the spike name and j is the position on the series of the identified spike. (Note that the left axis is scaled between 0.7 and 1.1, hence $\rho(t)$ is only plotted when it takes the value 1; for all other cases, as defined by (9), $\rho(t)=0$.

capacity over the time period t_n in the event of a contingency. Similarly, for low values of $\varrho(t_m, t_n)$, one would anticipate unexpected events at time t_n not to have a large impact on equilibrium prices. Therefore, in order to determine δ for the England and Wales electricity market we proceed in three steps.

First, based on historical data of wholesale electricity prices $S(t)$ for the period June 03 through March 06, we establish, via a filter as in Cartea and Figueroa (2005), at which points in time the market underwent a price spike. The total number of spikes in our series is given by N_J and we index each spike with J_j^i, where i denotes the number of the spike, i.e. $i \in \{1, 2, ..., N_J\}$, and j indicates the position of the spike in the price series.[9] For example, J_{31}^1 denotes the first spike in our series and it occurred on the 31st point of our price series data, (see Equation (9)). Second, we create a histogram of the values taken by ϱ that links ranges of ϱ to the frequency of spikes. We show this in the first two columns of Table 2 where, for example, regardless of the time when these spikes occurred, the bin [0.92100, 0.93391) contains three spikes. The third column shows the number of weeks in which the ratio was in that bin. Finally, the fourth column in Table 2 identifies which bin each one of the spikes J_j^i belongs to. For example, there was a price spike in observation 267 of our price series and for this spike $\varrho \in [0.90808, 0.92100)$.

Third, the last step is to establish the threshold value δ according to the following criteria. We look for the first bin which signals the beginning of a sequence of

Table 2. The second column displays the number of spikes in each bin, the third column the number of observations, and the fourth the position of each of the identified spikes J_j^i in the time series.

Bin	Quantity of spikes	Number of observations	J_j^i
[0.77893, 0.79185)	0	2	n.a.
[0.79185, 0.80476)	0	2	n.a.
[0.81768, 0.83059)	0	1	n.a.
[0.83059, 0.84351)	1	3	31
[0.84351, 0.85642)	0	9	n.a.
[0.85642, 0.86934)	0	7	n.a.
[0.86934, 0.88225)	0	13	n.a.
[0.88225, 0.89517)	0	12	n.a.
[0.89517, 0.90808)	0	7	n.a.
[0.90808, 0.92100)	3	7	267, 699, 726
[0.92100, 0.93391)	3	14	543, 649, 661
[0.93391, 0.94682)	3	15	51, 459, 674
[0.94682, 0.95974)	2	7	203, 274
[0.95974, 0.97265)	0	16	n.a.
[0.97265, 0.98557)	0	7	n.a.
[0.98557, 0.99848)	0	7	n.a.
[0.99848, 1.01140)	0	5	n.a.
[1.01140, 1.02430)	0	6	n.a.
[1.02430, 1.03720)	1	2	136

identified spikes. From Table 2 we identify this threshold as $\delta=0.90808$. We consequently define the high regime as the set $\Theta^{HR}:\{S(\tau_1), ..., S(\tau_k)\}$ where the times τ_i represent the corresponding points in the time series for which $\varrho(t_m, t_n) \geq \delta$ and $S(\tau_i)$ is the spot price evaluated at time $\tau_i \in t_n$. Similarly, the low regime is given by the set $\Theta_{LR}:\{S(\tau_1), ..., S(\tau_l)\}$ where the times τ_i represent the corresponding points in the time series for which $\varrho(t_m, t_n) < \delta$ and $S(\tau_i)$ is the spot price evaluated at time $\tau_i \in t_n$. Therefore, (3) becomes

$$\rho(t) = \begin{cases} 1, & \text{if } \varrho(t_{n-1}, t_n) \geq 0.90808, \quad \text{for } t \text{ in week } t_n, \\ 0, & \text{if } \varrho(t_{n-1}, t_n) < 0.90808, \quad \text{for } t \text{ in week } t_n. \end{cases} \quad (9)$$

Figure 2 presents the ratio time series $\varrho(t_{n-1}, t_n)$, the spot price series $S(t)$, and the deterministic function $\rho(t)$, together with the observed spikes throughout the historical sample considered.[10]

The evidence shown in Figure 2 and the fact that the information used to build the switching ratio undergoes weekly updates, suggest that an alternative way to model the dynamics of $\rho(t)$ is to assume that the ratio follows a stochastic process. For example, $\rho(t)$ could be composed of a deterministic trend plus noise. The crucial assumption, however, would be that the deterministic trend is calculated employing the forward looking information previously introduced. In our model we have assumed that the ratio is deterministic because more data would be necessary to model it stochastically and the calibration of the spot model would be far more complicated.

3.3   Dependence Between Ratio and Volatility

In order to justify the use of the ratio as an explanatory variable to determine the regime switching component of the model, we analyse the existence of a correlation or other dependence between the constructed ratio and the observed weekly volatility.[11]

We start by removing any seasonal component in the ratio which might obscure the relationship between the ratio and the volatility of $S(t)$. No significant seasonality was detected on the margin, defined as capacity minus demand, however a strong seasonality effect is observed on the demand $D(t_m, T_n)$, which in turn is used to construct the ratio as in (8).[12] Figure 3 depicts the demand, the fitted seasonality and the residuals from the fit. The fitted seasonality takes the form

$$f(\tau_D) = \sin\left(\frac{2\pi\tau_D}{52}\right) + \cos\left(\frac{2\pi\tau_D}{52}\right) \tag{10}$$

where τ_D is a dummy variable used to index the series in number of weeks and 52 is the annualization factor for a weekly-based estimate.

We then regress the residuals of the fitted demand and the weekly logarithm of the volatility, and the result is shown in Figure 4.[13] The significance of the coefficients is assessed in Table 3 by the p-values, which report the marginal significance level of the t-test. The test clearly rejects the null hypothesis of non-significant coefficients.

Further, to test the correlation between volatility and the ratio ϱ, we bootstrapped the data 5000 times. We find the median of the correlation is 19%, which indicates the presence of correlation. The most significant test is the one obtained by the significance of the beta coefficient that captures the relationship between the volatility of $S(t)$ and the ratio. This confirms what had been previously assumed
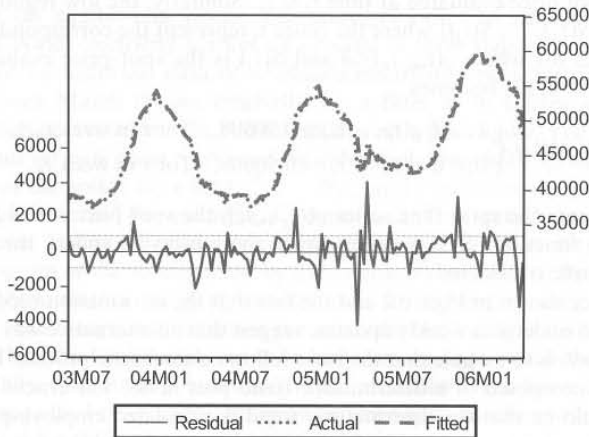
Figure 3. The dotted line represents the demand series, the dashed line the fitted seasonality and the solid line the residuals.
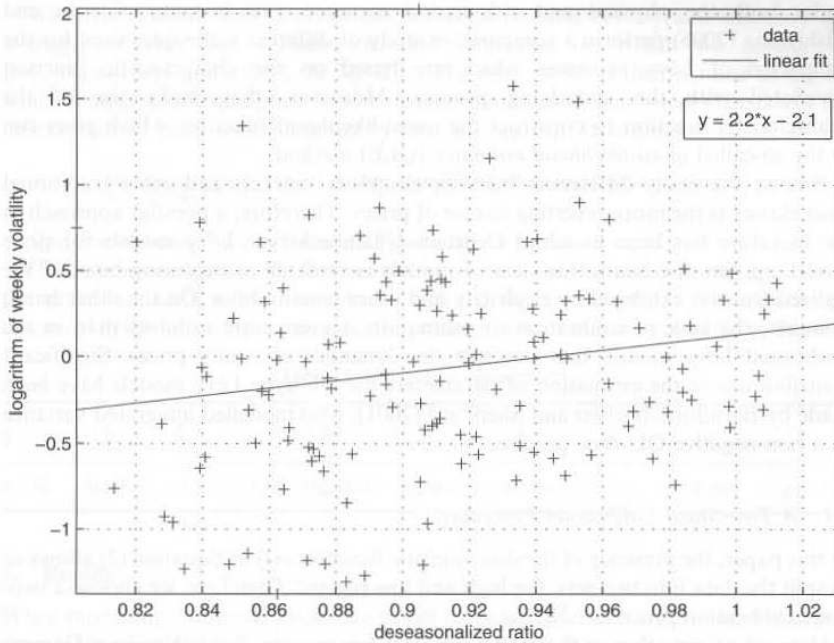
Figure 4. Linear regression of the deseasonalized ratio on the logarithm of the weekly volatility.

when considering the ratio as an indicator for the regime-switching component of the model. In other words, increases in the ratio $\varrho(t_{n-1}, t_n)$ are accompanied by an increase of volatility in week t_n.

Table 3. Parameter estimates.

β	se	p-value	R^2
$(-2.0787, 2.2042)$	$(0.9256, 1.0122)$	$(0.0263, 0.0311)$	0.0335

4. Calibration of Parameters

The calibration of models driven by Lévy processes is not a simple task. In most cases the marginal distributions of these models are not known in closed-form expression, while they are the cornerstone of robust estimating techniques such as the *maximum likelihood estimator* (MLE) or the *generalized method of moments* (GMM). Due to the popularity of Lévy-based models in finance, there has been a recent surge in the development of new methods devoted to parameter estimation

under both the physical and risk-neutral measure. For instance, Sueishi and Nishiyama (2006) perform a comparative study of different techniques used for the estimation of Lévy processes which are based on the characteristic function associated with the underlying process. Moreover, they make use of the characteristic function to construct the quasi-likelihood function, which gives rise to the so-called *quasi-likelihood estimator* (QLE) method.

Among the many differences between electricity markets and other traditional asset classes is the mean reverting nature of prices. Therefore, a popular approach in the literature has been to adopt Ornstein–Uhlenbeck-type Lévy models for price modelling. On one hand, this class of models is versatile at capturing most of the stylized features exhibited by electricity and other commodities. On the other hand, however, the task of calibration or estimation is even more arduous than in the traditional Lévy models that describe the dynamics of equity prices. Significant contributions in the estimation of parameters for OU-type Lévy models have been made by Barndorff-Nielsen and Shephard (2001), who modelled integrated variance as a non-negative OU-type process.

4.1 A Two-Stage Calibration Procedure

In this paper, the presence of the deterministic function $\rho(t)$ in Equation (2) allows us to split the data into two sets: the high and low regime. Therefore, we choose a two-stage calibration procedure.[14]

Here we assume that in the model (7) the Lévy process $Z(t)$ is Variance Gamma with parameters $(\sigma_z, \theta, \kappa)$. Moreover, we assume that the distribution of the spikes in the high regime, given by $U = \ln J$, is double exponential with density

$$f(u) = p\eta_1 \, e^{-\eta_1 u} 1_{u \geq 0} + q\eta_2 \, e^{\eta_2 u} 1_{u < 0},$$

where $\eta_1, \eta_2 > 0$, $p, q \geq 0$, s.t. $p + q = 1$. Further, p and q represent the probabilities of upward and downward jumps, respectively.

Therefore, the total set of parameters to estimate in the model specified by (7) is given by $\Theta : \{q, p, \eta_1, \eta_2, \alpha^H, \alpha^L, \sigma, \sigma_z, \theta, \kappa\}$, which we group as $\Theta : \{\Theta_1; \Theta_2\}$, with $\Theta_1 : \{q, p, \eta_1, \eta_2\}$ and $\Theta_2 : \{\alpha^H, \alpha^L, \sigma, \sigma_z, \theta, \kappa\}$. We group them in this manner to highlight the two-stage calibration process we perform.

Θ_1 contains the parameters responsible for the size and direction of spikes. Thus given the structure of our model we can separate the series, as discussed above, into high and low regimes. From the high-regime sub-sample we filter the spikes and proceed in a similar way as in Cartea and Figueroa (2005) to estimate the parameters $\{q, p, \eta_1, \eta_2\}$.[15]

In the second stage of the calibration, making use of the estimated parameters in Θ_1, we estimate the parameters in Θ_2 by matching the mean, variance, skewness and kurtosis of the deseasonalized returns of the spot price to those of simulated paths. We start by assuming three possible initial values for each one of the parameters in the set Θ_2. Hence, a possible set Θ_2^s, where superscript s indicates which combination out of the 729 possible ones we are looking at, will be denoted by $\Theta_2^s : \left\{ \alpha_i^H, \alpha_j^L, \sigma_k, \sigma_z^l, \theta_m, \kappa_n \right\}$, with (i, j, k, l, m, n) taking values in $\{1, 2, 3\}$ since, for each parameter, we are starting from three different guesses.[16]

We then perform 1000 simulations for each possible set Θ_2^s and calculate the averages of the four statistics we are interested in, which we denote by $(m_1^s, m_2^s, m_3^s, m_4^s)$ and $s \in \{1, 2, ..., 729\}$. The optimal set is then defined as the set which is closest, in a minimum square distance sense, to the empirical statistics of the deseasonalized price series denoted by (m_1, m_2, m_3, m_4). In other words, we solve

$$\min_{s \in \{1, ..., 729\}} \left\{ \left(m_1^s - m_1\right)^2 + \left(m_2^s - m_2\right)^2 + \left(m_3^s - m_3\right)^2 + \left(m_4^s - m_4\right)^2 \right\} \qquad (11)$$

to find the optimal $\tilde{s}$ given the initial guess of the parameters.

Once we have obtained the first set $\Theta_2^{\tilde{s}}$, we perturb the initial conditions and repeat the procedure until a local minimum, denoted $\hat{\Theta}_2$, is obtained.[17] The results are discussed in the following section and the parameter estimates are presented in Table 4.

Table 4. Calibrated parameters.

$\hat{q}$	$\hat{p}$	$\hat{\eta}_1$	$\hat{\eta}_2$	$\hat{\alpha}^H$	$\hat{\alpha}^L$	$\hat{\sigma}$	$\hat{\sigma}_z$	$\hat{\theta}$	$\hat{\kappa}$
8.570	0.600	5.150	2.130	182.000	52.000	0.750	0.250	14.000	0.025

5. Results

When modelling wholesale electricity prices there are two main criteria used to assess the ability of models to mimic price dynamics. First, the model has to be able to reproduce path properties, especially jumps and spikes that mean revert at speeds observed in the market. Second, the model should also be able to replicate statistical properties, understood as replicating the mean, variance, skewness and kurtosis of the returns series; in our case we expect the moment matching to be 'good' given the choice of the calibration procedure descried above.

Figures 5, 6 and 7 show sample paths produced by our model. It is clear that the model simulates the spikes in periods where the deterministic switching vector $\rho(t)=1$, as well as still allowing for jumps of lesser magnitude in those regions where $\rho(t)=0$.[18] Moreover, we can also observe that the mean reversion rates α^H and α^L in both regimes are such that temporary deviations revert back to the mean seasonal level at speeds observed in the market.

To assess the statistical performance of the model we calculate the first four moments of 1000 simulated price paths and compare these moments with those obtained from the distribution of realized returns. Table 4 first presents the values for the calibrated parameters in the model, followed by Table 5 where we present the

Table 5. Comparison of simulated and empirical statistics.

|	Empirical	Simulated
Mean	0.0023	0.0011
St. dev.	0.0315	0.0389
Skewness	1.9943	1.1086
Kurtosis	22.1243	21.48
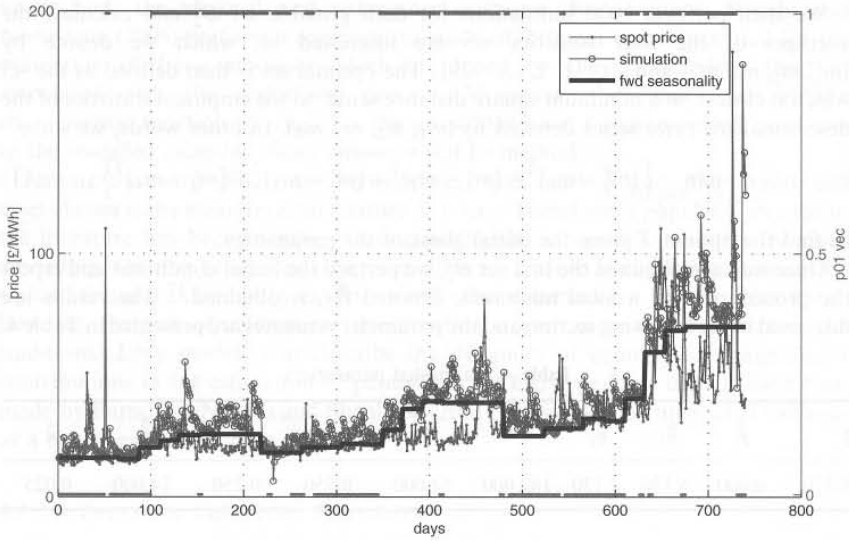
Figure 5. Comparison of simulated and real price paths. The line with solid circles represents the spot price while the simulated price is represented by the line with hollow circles. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.
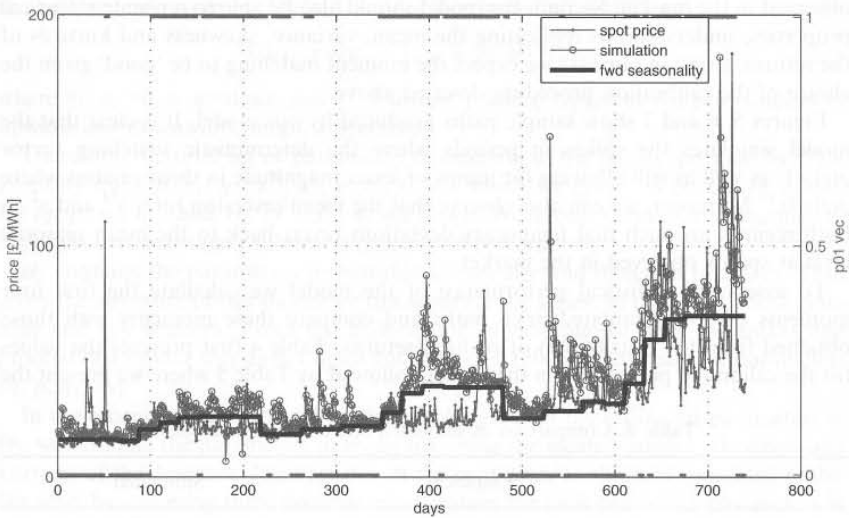
Figure 6. Comparison of simulated and real price paths. The line with solid circles represents the spot price while the simulated price is represented by the line with hollow circles. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.
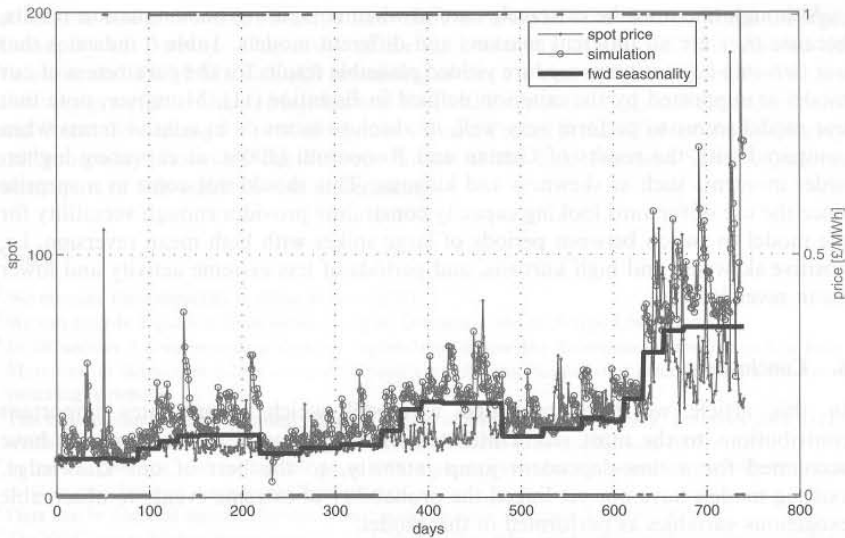
Figure 7. Comparison of simulated and real price paths. The line with solid circles represents the spot price while the simulated price is represented by the line with hollow circles. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.

simulated statistics from the model and the actual statistics from the actual distribution.

Finally, we compute the minimum square distance of the simulated and actual statistics using (11). We also compute this measure for the results of simulated and actual statistics reported by Geman and Roncoroni (2006) for comparative purposes. The results are summarized in Table 6.

Table 6. Comparison of empirical and simulated moments. The first three columns, labelled 'ECAR', 'PJM' and 'COB' refer to results from different American power markets, as reported by Geman and Roncoroni (2006), whereas the last column, 'UK', corresponds to the results obtained in this paper. Finally, d represents the minimum square distance defined in Equation (11).

	ECAR		PJM		COB		UK	
	Emp.	Sim.	Emp.	Sim.	Emp.	Sim.	Emp.	Sim.
Mean	−0.0002	−0.0001	−0.0006	0	0.0009	0.0006	0.0023	0.0011
St. dev.	0.3531	0.3382	0.2364	0.2305	0.1586	0.1121	0.0315	0.0389
Skewness	−0.5575	2.1686	0.3949	1.6536	0.1587	0.961	1.9943	1.1086
Kurtosis	21.6833	22.5825	13.1507	14.8429	6.7706	6.5402	22.1243	21.48
d	8.2404		4.4479		0.6989		1.1997	

Although one must be extremely careful when comparing our estimation results, because they are all different markets and different models, Table 6 indicates that our two-step estimation procedure yielded plausible results for the parameters of our model as supported by the criterion defined in Equation (11). Moreover, note that our model seems to perform very well, in absolute terms or in relative terms when compared with the results of Geman and Roncoroni (2006), at capturing higher-order moments such as skewness and kurtosis. This should not come as a surprise since the use of forward looking capacity constraints provides enough versatility for the model to switch between periods of large spikes with high mean reversion, i.e. positive skewness and high kurtosis, and periods of less extreme activity and lower mean reversion.

6. Conclusions

In this article we have presented a model which incorporates important contributions to the most recent literature. First, although different models have accounted for a time-dependent jump intensity, to the best of our knowledge, existing models have not yet linked the probability of extreme events to observable exogenous variables as performed in this model.

Second, we have also extended the literature in the field by allowing for time-varying mean reverting processes. In particular, this is aimed at solving the critical problem encountered by one-factor mean reverting models regarding the flatness of the long-end of the forward curves. Although other models account for sums of OU–Lévy processes with different mean reversion rates, we believe this model allows for a time-varying structure while preserving simplicity.

Third, we have tackled a common drawback of this class of models through the incorporation of a forward-looking seasonality which enables the model to mean revert to more realistic scenarios that reflect changes in fuel prices.

Fourth, we have made use of exogenous observable variables in order to separate two distinct regimes where prices may jump and another where prices are allowed to spike. We believe the difference between jumps and spikes is critical in power markets and hence a model should be able to distinguish between both processes.

The results obtained through the simulations of the spot prices seem to be in reasonable qualitative accordance with observed historical price paths. More importantly, by measuring the square distance between the first four moments of the simulated paths and the actual distribution we are able to asses quantitatively the performance of the model. Comparing our results with those of other authors, we conclude that the model performs well. Indeed, our results outperform the benchmark used in two out of three cases.

Notes

[1] We examine these forecasts in detail in Section 3.1.

[2] We can refer to Equation (2) as representing an Ornstein–Uhlenbeck-type Lévy process.

[3] In Subsection 3.1 we provide a detailed explanation of how the deterministic function $\rho(t)$ is built. Moreover, in Subsection 3.2 we discuss the possibility of using, instead of a deterministic, a stochastic switching parameter.

[4] This information has been obtained from the 'Digest of United Kingdom Energy Statistics 2004, DTI', which is publicly available at www.dti.gov.uk/energy/statistics (accessed January 2007).

[5] In the UK presently, and the last few years, CCGT plants have been the plants 'on the margin', which explains our choice of the seasonal function $g(t)$.

[6] Data can be accessed directly through www.bmreports.com (accessed January 2007).

[7] The NGC also publishes shorter-term forecasts that include other information such as indicated demand and indicated generation for the day ahead market.

[8] Note that our updating procedure is equivalent to picking the first observation of the 2–52 forecast for every week. In other words, the forecast used for week t_n is the one which was made public as the first data point of the 2–52 forecasts in week t_{n-1}.

[9] The series considered in our study coincides with the spot series of 740 daily observations (excluding weekends) between 2/06/03 and 31/03/06.

[10] An important issue is the robustness of $\rho(t)$ if we calculate $\varrho(t_m, t_n)$ for other values of m. Hence, we repeated our study for different values of m and found that the threshold was stable. For instance, if the forward looking ratio is calculated using 13-week-ahead forecasts, i.e. $\varrho(t_{n-13}, t_n)$, our findings show that the threshold remains unchanged.

[11] Since the ratio is based on weekly observations we compare it with the weekly volatility.

[12] The capacity $C(t_m, t_n)$ is calculated as demand forecast $D(t_m, t_n)$ plus margin, which is also made public by the NGC.

[13] We regress the logarithm of the volatility in order to reduce the effect of apparent outliers in the ratio.

[14] Alternative approaches include the use of Markov regime switching models; see, for example, Hamilton (1994).

[15] The filter consists of recursively separating data points that are three standard deviations away from the mean.

[16] The total number of possible sets can be calculated by first calculating the combinatorial number nC_k, which gives the number of $k=6$ subsets possible out of a set of $n=18$ distinct items, and later by excluding those sets with more than one element of each sub-group.

[17] It is important to note that we might not be finding a global, but a local, optimal set.

[18] Recall that the fact that $\rho(t)=0$ does not preclude the model from exhibiting large spikes, it just signals that the probability of observing a spike is very low in comparison with the periods where $\rho(t)=1$.

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* B – Part 2, 63, pp. 167–241.

Benth, F. E., Kallsen, J. and Meyer-Brandis, T. (2007) A non-Gaussian Ornstein–Uhlenbeck process for electricity spot price modelling and derivatives pricing. *Applied Mathematical Finance*, 14(2), pp. 153–169.

Bessembinder, H. and Lemmon, M. L. (2002) Equilibrium pricing and optimal hedging in electricity forward markets. *The Journal of Finance*, LVII(3), pp. 1347–1382.

Brennan, M. (1958) The supply of storage. *The American Economic Review*, 48(1), pp. 50–72.

Cartea, Á. and Figueroa, M. G. (2005) Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4), pp. 313–335.

Casassus, J. and Collin-Dufresne, P. (2005) Stochastic convenience yield implied from commodity futures and interest rates. *The Journal of Finance*, LX(5), pp. 2283–2331.

Escribano, Á., Peña, J. I. and Villaplana, P. (2002) Modelling electricity prices: International evidence. Working paper 02-27, Universidad Carlos III de Madrid.

Eydeland, A. and Geman, H. (1999) Pricing power derivatives. *Risk*, pp. 71–74.

Fama, E. F. and French, K. R. (1987) Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *The Journal of Business*, 60(1), pp. 55–73.

Fama, E. F. and French, K. R. (1988) Business cycles and the behavior of metals prices. *The Journal of Finance*, 43(5), pp. 1075–1093.

Geman, H. and Roncoroni, A. (2006) Understanding the fine structure of electricity prices. *The Journal of Business*, 79(3), pp. 1225–1262.

Gibson, R. and Schwartz, E. S. (1990) Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance*, XLV(3), pp. 959–976.

Hambly, B., Howison, S. and Kluge, T. (2007) Modelling spikes and pricing swing options in electricity markets. Working paper, University of Oxford.

Hamilton, J. (1994) *Time Series Analysis* (Princeton, NJ: Princeton University Press).

Hilliard, J. E. and Reis, J. (1998) Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot. *The Journal of Financial and Quantitative Analysis*, 33(1), pp. 61–86.

Hirshleifer, D. (1988) Risk, futures pricing, and the organization of production in commodity markets. *Journal of Political Economy*, 96(6), pp. 1206–1220.

Litzenberger, R. H. and Rabinowitz, N. (1995) Backwardation in oil futures markets: Theory and empirical evidence. *The Journal of Finance*, 50(5), pp. 1517–1545.

Lucía, J. J. and Schwartz, E. S. (2002) Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research*, 5, pp. 5–50.

Martijena, M. F. (2006) Evaluating the relation between gas and electricity prices volatilities. Case study: United Kingdom. MSc. thesis, Birkbeck, University of London.

Miltersen, K. R. and Schwartz, E. S. (1998) Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates. *The Journal of Finance*, 33(1), pp. 33–59.

Routledge, B. R., Seppi, D. J. and Spatt, C. S. (2000) Equilibrium forward curves for commodities. *The Journal of Finance*, LV(3), pp. 1297–1338.

Schwartz, E. S. (1997) The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3), pp. 923–973.

Sueishi, N. and Nishiyama, Y. (2006) Estimation of Lévy processes in mathematical finance: A comparatative study. Graduate School of Economics, Kyoto University and Kyoto Institute of Economic Research, Kyoto University.

Working, H. (1949) The theory of price of storage. *American Economic Review*, 39(6), pp. 1254–1262.