

Modelling Electricity Prices with Forward Looking Capacity Constraints*

Álvaro Cartea,[†] Marcelo G. Figueroa,[‡] Hélyette Geman[§]

First version: August, 2006. This version: January 16, 2008

*We have benefited from the helpful comments of Gareth Davies, Murray Hartley, Andrew Shaw, Peter E. George, Roberto Renò, Jos van Bommel, Grainne O'Donnell, seminar attendants and conference participants at the University of Warwick, Birkbeck College, University of Toronto, University of Siena, University of Oxford, and TUM Munich. All remaining errors are our sole responsibility. We are also very grateful to Shanti Majithia and Richard Price at the National Grid Company and Oxera Consulting for providing data.

[†]Commodities Finance Centre, Birkbeck, University of London. a.cartea@bbk.ac.uk. Á. Cartea is thankful for the hospitality and generosity shown by the Finance Group at the Saïd Business School, Oxford, where part of this research was undertaken.

[‡]Commodities Finance Centre, Birkbeck, University of London and BP p.l.c. mfigueroa.keller@gmail.com

[§]Commodities Finance Centre, Birkbeck, University of London, and ESSEC, France. h.geman@bbk.ac.uk

Modelling Electricity Prices with Forward Looking Capacity Constraints

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 the electricity markets, improves the modeling of spikes (timing and magnitude) and the different speeds of mean reversion.

Keywords: capacity constraints, mean reversion, electricity indicated demand, electricity indicated generation, regime switching model.

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, for example the seminal work by 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) and Schwartz (1997)), or both stochastic convenience yield and interest rates, (Hilliard and Reis (1998) Miltersen and Schwartz (1998) and Casassus and Collin-Dufresne (2005)). Additional work based on the theory of storage includes that of Fama and French (1988) who look at business cycles and the behavior of metal prices. Eydeland and Geman (1999) argue that in the case of 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, Seppi, and Spatt (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, on 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 behavior 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 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 behavior of one player alone, for example a generating company, will be able to influence equilibrium spot and futures prices. For example in the UK there are around 50 companies that own generating capacity. A closer look at these companies reveals that 50% of production capacity is owned by only four companies, 90% of capacity is owned by 15 companies and over 99% is owned by 28 companies (see Table 7 in the appendix).

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 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 the UK is therefore dominated by the composition of its generation capacity, as shown in Table 1 below, and consequently seasonal levels and trajectory 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 behavior 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 behavior of players in terms of prices and quantities they are willing to buy or

¹We examine these forecasts in detail in Section 3.1.

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, Peña, and Villaplana (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, while Geman and Vu-Nhat (2005) model seasonality by sine and cosine functions.

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 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 in 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, Howison, and Kluge (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, Kallsen, and Meyer-Brandis (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 modeling 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, 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 modeling 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, labeled ‘high regime’, during which the probability of observing spikes is high, and a low-intensity level, labeled ‘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 to 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 processes $y(t)$

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

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

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

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

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

where the “high-regime” refers to periods where electricity prices may exhibit large spikes that are introduced in equation (2) through the stochastic shocks $\ln JdN(t)$. 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

²We can refer to equation (2) as representing an ‘Ornstein-Uhlenbeck type Lévy process’.

presence of huge spikes is unlikely.³

To solve (2), we first note that

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

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

$$\begin{aligned} y(t) &= e^{-\int_{t_0}^t \beta(u)du} y_{t_0} + \sigma \int_{t_0}^t e^{-\int_s^t \beta(u)du} dW(s) + \int_{t_0}^t e^{-\int_s^t \beta(u)du} \rho(s) \ln J dN(s) \\ &\quad + \int_{t_0}^t e^{-\int_s^t \beta(u)du} (1 - \rho(s)) dZ(s). \end{aligned} \quad (5)$$

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

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

³Below, in subsection 3.1, we provide a detailed explanation of how the deterministic function $\rho(t)$ is built.

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

$$\begin{aligned}
 dy(t) = & - [\alpha^H \rho(t) + \alpha^L (1 - \rho(t))] y(t) dt + \sigma dW(t) \\
 & + \rho(t) \ln J dN(t) + (1 - \rho(t)) dZ(t).
 \end{aligned}
 \tag{7}$$

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 $\rho(t)$

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

⁴This 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.

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 below.

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 in 2005 combined cycle gas turbine (CCGT) stations and conventional steam stations account 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




Figure 1: Fuel used in electricity generation in 1996 and 2004; data source: www.dti.gov.uk/energy/statistics.

gas prices to increase. This aspect of the markets has been studied in 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 gets tighter, and the volatility in the electricity market increases with the volatility in the gas market, creating a ‘contagion’ effect.

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

	2003		2004		2005	
Major power producers in 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	2,930	4.10%	2,930	4.10%	3,262	4.56%
Mixed or dual fired	6,413	8.97%	6,413	8.97%	6,403	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	1,537	2.15%	1,485	2.08%	1,346	1.88%
Hydro-electric stations:						
Natural flow	1,267	1.77%	1,270	1.78%	1,273	1.78%
Pumped storage	2,788	3.90%	2,788	3.90%	2,788	3.90%
Renewables other than hydro	117	0.16%	117	0.16%	117	0.16%

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

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

⁵In 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 $g(t)$.

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

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

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 $m, n \in \{1, 2, 3, \dots, 52\}$ denote weeks of the year. The notation t_n represents the time period in week n . For example

⁶Data can be accessed directly through www.bmreports.com.

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

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 4. 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 $\varrho(t_m, t_n)$ is, the tighter generation will become, leaving little manoeuvrability for the NGC to call on idle 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 did the market undergo 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

⁸Note 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} .

denotes the number of the spike, i.e. $i \in \{1, 2, \dots, N_J\}$, 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 Figure 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 3 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 $\varrho \in [0.90808, 0.92100)$.

bin	quantity of spikes	numb. 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

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

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

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 identified spikes. From Table 2 above we identify this threshold as $\delta = \mathbf{0.90808}$. We consequently define the high regime as the set $\Theta^{HR} : \{S(\tau_1), \dots, S(\tau_k)\}$ 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 $\Theta_{LR} : \{S(\tau_1), \dots, S(\tau_l)\}$ 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

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

In Figure 4 are represented 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

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 analyze 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

¹⁰An 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.

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

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 2 depicts the demand, the fitted seasonality and the residuals from the fit. The fitted seasonality takes the form

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

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, the result is observed in Figure 3 below.¹³ The significance of the coefficients is assessed below 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 volatility of $S(t)$ and the ratio. This confirms what had been previously assumed 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 .

¹²The 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.

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

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

Table 3: Parameter estimates

■




Figure 2: The dotted line represents the demand series, the dashed line the fitted seasonality and the solid line the residuals

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




Figure 3: Linear regression of deseasonalized ratio on the logarithm of the weekly volatility.

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 to 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, see for instance Geman and Roncoroni (2006), in the literature has been to adopt ‘Ornstein-Uhlenbeck-type Lévy’ models for price modeling. On one hand, this class of models is versatile at capturing most of the stylised features exhibited by electricity and other commodities. On the other hand, however, the task of calibration or




Figure 4: 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), $\rho(t) = 0$).

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, and the work of Schoutens, Tuerlinckx, and Valdivieso (2005).

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.

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

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

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 set of total 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 $\{q, p, \eta_1, \eta_2\}$.¹⁴

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

¹⁴The filter consists in recursively separating data points that are three standard deviations away from the mean.

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 $\Theta_2^s : \{\alpha_i^H, \alpha_j^L, \sigma_k, \sigma_z^l, \theta_m, \kappa_n\}$, with (i, j, k, l, m, n) taking values in $\{1, 2, 3\}$ since for each parameter we are starting from three different guesses.¹⁵

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 $(\bar{m}_1^s, \bar{m}_2^s, \bar{m}_3^s, \bar{m}_4^s)$ and $s \in \{1, 2, \dots, 729\}$. 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

$$\min_{s \in \{1, \dots, 729\}} \{(\bar{m}_1^s - m_1)^2 + (\bar{m}_2^s - m_2)^2 + (\bar{m}_3^s - m_3)^2 + (\bar{m}_4^s - m_4)^2\}, \quad (11)$$

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.¹⁶ The results are discussed in the following section and the parameter estimates are presented in Table 4.

¹⁵The total number of possible sets can be calculated by calculating first 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.

¹⁶It is important to note that we might not be finding a global but a local optimal set.

5 Results

When modeling 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 described 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 magnitudes in those regions where $\rho(t) = 0$.¹⁷ 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.

¹⁷Recall 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 to the periods where $\rho(t) = 1$.




Figure 5: Comparison of simulated and real price paths. The line with solid circle represents the spot price while the simulated price is represented by the line with hollow circle. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.

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 realised returns. Table 4 below presents first the values for the calibrated parameters in the model, followed by Table 5 where we present the simulated statistics from the model and the actual statistics from the actual distribution.

\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

Table 4: Calibrated parameters.

Finally, we compute the minimum square distance of the simulated and actual




Figure 6: Comparison of simulated and real price paths. The line with solid circle represents the spot price while the simulated price is represented by the line with hollow circle. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.

statistics using (11). We also compute this measure for the results of simulated and actual statistics reported in Geman and Roncoroni (2006) for comparative purposes. The results are summarized in Table 6 below.

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 to the results in Geman and Roncoroni (2006), at capturing higher order moments such as skewness and kurtosis. This should not come as a surprise since




Figure 7: Comparison of simulated and real price paths. The line with solid circle represents the spot price while the simulated price is represented by the line with hollow circle. Seasonality is represented by the thick solid line. On the second axis, the deterministic vector $\rho(t)$ is also shown.

the use of forward looking capacity constraints provides enough versatility for the model to switch between periods of large spikes with high mean reversion, ie 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

	emp.	sim.
mean	0.0023	0.0011
st. dev.	0.0315	0.0389
skewness	1.9943	1.1086
kurtosis	22.1243	21.48

Table 5: Comparison of simulated and empirical statistics.

	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	

Table 6: Comparison of empirical and simulated moments. The first three columns, labeled ‘ECAR’, ‘PJM’ and ‘COB’ refer to results from different American power markets, as reported in 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).

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.

Fourth, we have made use of exogenous observable variables in order to separate two distinct regimes where prices may jump and other 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 assess 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.

Appendix: Generator capacity in the UK

Power Stations in the United Kingdom			
	Company name	Share in Production	Cumulative share
1	British Energy	0.14903807	0.1490
2	RWE Npower Plc	0.12729716	0.2763
3	E.On UK	0.11811167	0.3944
4	Scottish & Southern Energy plc	0.11421521	0.5087
5	Scottish Power	0.07918333	0.5878
6	EDF Energy	0.06219059	0.6500
7	Drax Power Ltd	0.05090080	0.7009
8	Centrica	0.04437208	0.7453
9	International Power	0.03555960	0.7809
10	BNFL British Nuclear Group	0.02983084	0.8107
11	First Hydro Company	0.02694065	0.8376
12	Teesside Power Ltd	0.02419240	0.8618
13	Seabank Power Limited	0.01576699	0.8776
14	Barking Power	0.01290261	0.8905
15	Premier Power Ltd	0.01285100	0.9034
16	Spalding Energy Company Ltd	0.01165106	0.9150
17	Coryton Energy Company Ltd	0.00971567	0.9247
18	Rocksavage Power Co. Ltd	0.00967696	0.9344
19	Immingham CHP LLP	0.00956084	0.9440
20	AES	0.00851572	0.9525
21	Baglan Generation Ltd	0.00741900	0.9599
22	Alcan	0.00647066	0.9664
23	Coolkeeragh ESB Ltd	0.00541910	0.9718
24	Corby Power Ltd	0.00517395	0.9770
25	Uskmouth Power Company Ltd	0.00507073	0.9820
26	Beaufort Wind Ltd	0.00345790	0.9855
27	Derwent Cogeneration	0.00304502	0.9885
28	Gaz de France	0.00232247	0.9909
29	Fellside Heat and Power	0.00216764	0.9930
30	Paul's Hill WindLtd	0.00070964	0.9937
31	Roths Wind Ltd	0.00065803	0.9944
32	Crystal Rig Windfarm Ltd	0.00064513	0.9950
33	Airtricity	0.00056771	0.9956
34	Fenland Windfarms Ltd	0.00056771	0.9962
35	Fibrothetford	0.00050320	0.9967
36	EPR Ely Limited	0.00049030	0.9972
37	Combined Heat & Power Ltd	0.00041288	0.9976
38	Citigen (London) UK Ltd	0.00039998	0.9980
39	HG Capital	0.00027095	0.9982
40	Sita Tyre Recycling Ltd	0.00025805	0.9985
41	Western Power Generation	0.00020644	0.9987
42	Yorkshire Windpower Ltd	0.00020644	0.9989
43	Cemmaes Windfarm Ltd	0.00019354	0.9991
44	Fibrogen	0.00016773	0.9993
45	Fibropower Ltd	0.00016773	0.9994
46	K/S Winscales	0.00011612	0.9996
47	Llangwryfon Windfarm Ltd	0.00011612	0.9997
48	Cold Northcott Windfarm Ltd	9.0318E-05	0.9998
49	TPG Wind Ltd	9.0318E-05	0.9999
50	Blyth Offshore Wind Ltd	5.1610E-05	0.9999
51	Great Orton Windfarm Ltd	5.1610E-05	1.0000
52	Haverigg III Ltd	3.8708E-05	1.0000

Table 7: Companies with power stations operational at the end of May 2006 in the UK. Note that 4 companies hold 50% of output capacity; the first 15 companies already comprise over 90% of the output capacity; and the first 28 comprise over 99% of production capacity. (Department of Trade and Industry 2006).

References

- Barndorff-Nielsen, Ole E., and Neil Shephard, 2001, Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics, *Journal Royal Statistical Society B -Part 2*, 167–241.
- Benth, Fred Espen, Jan Kallsen, and Thilo Meyer-Brandis, 2007, A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modelling and derivatives pricing, *Applied Mathematical Finance* 14, 153–169.
- Bessembinder, Hendrik, and Michael L. Lemmon, 2002, Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets, *The Journal of Finance* LVII, 1347–1382.
- Brennan, M., 1958, The Supply of Storage, *American Economic Review* 48, 50–72.
- Cartea, Álvaro, and Marcelo G. Figueroa, 2005, Pricing in Electricity Markets: a Mean Reverting Jump Diffusion Model with Seasonality, *Applied Mathematical Finance* 12, 313–335.
- Casassus, J., and P. Collin-Dufresne, 2005, Stochastic Convenience Yield implied from Commodity Futures and Interest Rates, *The Journal of Finance* LX, 2283–2331.
- Escribano, Álvaro, Juan Ignacio Peña, and Pablo Villaplana, 2002, Modelling Electricity Prices: International Evidence, Universidad Carlos III de Madrid – working paper 02-27.
- Eydeland, Alexander, and Helyette Geman, 1999, Pricing Power Derivatives, *Risk* pp. 71–74.

- Fama, E.F., and K.R. French, 1987, Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage, *The Journal of Business* 60, 55–73.
- Fama, E.F., and K.R. French, 1988, Business Cycles and the Behavior of Metals Prices, *The Journal of Finance* 43, 1075–1093.
- Geman, Helyette, and Andrea Roncoroni, 2006, Understanding the Fine Structure of Electricity Prices, *The Journal of Business* 79.
- Geman, Helyette, and Nguyen Vu-Nhat, 2005, Soybean Inventory and Forward Curve Dynamics, *Management Science* 51, 1076–1091.
- Gibson, R., and E.S. Schwartz, 1990, Stochastic Convenience Yield and the pricing of Oil Contingent Claims, *The Journal of Finance* XLV, 959–976.
- Hambly, Ben, Sam Howison, and Tino Kluge, 2007, Modelling spikes and pricing swing options in electricity markets, Working paper, University of Oxford.
- Hilliard, J.E., and Jorge Reis, 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, 61–86.
- Hirshleifer, D., 1988, Risk, Futures Pricing, and the Organization of Production in Commodity Markets, *Journal of Political Economy* 96, 1206–1220.
- Litzenberger, R.H., and N. Rabinowitz, 1995, Backwardation in Oil Futures Markets: Theory and Empirical Evidence, *The Journal of Finance*, 50, 1517–1545.

- Lucía, Julio J., and Eduardo S. Schwartz, 2002, Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange, *Review of Derivatives Research* 5, 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 E.S. Schwartz, 1998, Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates, *The Journal of Finance* 33, 33–59.
- Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt, 2000, Equilibrium Forward Curves for Commodities, *The Journal of Finance* LV, 1297–1338.
- Schoutens, Wim, Francis Tuerlinckx, and Luis Valdivieso, 2005, Maximum likelihood estimation in processes of Ornstein-Uhlenbeck type, Universitair Centrum voor Statistiek, K.U Leuven.
- Schwartz, E.S., 1997, The Stochastic Behavior of commodity Prices: Implications for Valuation and Hedging, *The Journal of Finance* 52, 923–973.
- Sueishi, N., and Y Nishiyama, 2006, Estimation of Lévy Processes in Mathematical Finance: a Comparative Study, Graduate School of Economics, Kyoto University and Kyoto Institute of Economic Research, Kyoto University.
- Working, H., 1949, The Theory of Price of Storage, *The American Economic Review*, 39, 1254–1262.