

Stochastic Behaviour of the Electricity Bid Stack: from Fundamental Drivers to Power Prices

Michael Coulon

*Mathematical Institute, University of Oxford,
24-29 St. Giles, Oxford OX1 3LB, UK
coulon@maths.ox.ac.uk (01865 280613)*

Sam Howison

*Oxford-Man Institute, University of Oxford,
Blue Boar Court, 9 Alfred Street, Oxford OX1 4EH, UK
& Mathematical Institute, University of Oxford
howison@maths.ox.ac.uk (01865 270500)*

23 October 2008

Abstract

We develop a fundamental model for spot electricity prices, based on stochastic processes for underlying factors (fuel prices, power demand and generation capacity availability), as well as a parametric form for the bid stack function which maps these price drivers to the power price. Using observed bid data, we find high correlations between the movements of bids and the corresponding fuel prices. We fit the model to the PJM and New England markets in the US, and discuss its performance, in terms of capturing key properties of simulated price trajectories, as well as comparing implied forward prices with observed data.

Keywords: electricity, bid stack, fundamental, margin, demand, natural gas

1 Introduction

The quest to model the unusual behaviour of electricity prices often leads to a choice between realistic, but complicated, models and simpler reduced-form models with convenient pricing formulas. In particular, due to the non-storability of electricity, there is a closer link between power and its fundamental underlying price drivers (in particular fuel prices, load and generating capacity) than in other markets. We propose a supply and demand based hybrid model, which exploits these links with underlying factors and allows flexibility and realism, while still retaining tractability for derivative pricing purposes.

Electricity spot prices typically exhibit periodicity (at annual, weekly and daily horizons), mean-reversion, very high volatility and sudden price spikes, all of which can be traced to supply and demand related causes. (see eg Burger *et al* (2007) and Eydeland and Wolyniec (2003)) Forward prices show none of this volatile behaviour, but instead suggest the need for multi-factor models to capture the movements of different parts of the forward curve; this is illustrated by the Principal Component Analysis of Koekebakker and Ollmar (2005) and by discussions of volatility term structure by Clewlow and Strickland (2000). Furthermore, both forwards and longer-term spot dynamics reveal close links with underlying drivers such as demand (or weather patterns) and especially fuel prices. Burger *et al* (2007) and Emery and Liu (2002) discuss the apparent cointegration, rather than the simple correlation of power and fuel prices, as long-term levels move together. All of these characteristics suggest the importance of looking beyond historic price series to better understand the dynamics of power prices in relation to fundamental drivers.

Nonetheless, the desire for analytic formulas and efficient pricing techniques has led to a large literature on direct spot or forward price modelling. Early spot price models by Lucia and Schwartz (2002), and Schwartz and Smith (2000) proposed a two-factor diffusion model to capture the different short and long-term dynamics of power prices. However, the importance of electricity spikes has led to the use of jump-diffusion processes by many authors, including Cartea and Figueroa (2005), and Kluge (2006). An affine jump diffusion framework leads to convenient formulas for derivative prices; see, for example Deng (1999), and Culot *et al* (2006), who also suggest regime-switching jumps to capture the behaviour of short-lived price spikes. Pure regime switching models have also been studied by De Jong and Huisman (2005) and Weron *et al* (2004), while an alternative approach is proposed by Geman and Roncoroni (2006), forcing jumps to be downwards when prices are above a certain threshold. While many of these models produce useful results and realistic price dynamics, they often face calibration challenges due either to unobservable factors, choices of probability measure, or to the complication of identifying historical spikes. In addition, they fail to capture the important correlations between power prices and other energy prices.

The alternative category of structural, fundamental, or supply and demand based models consists of a wide range of work including agent-based models of market power and bidding strategy (e.g. Ruibal and Mazumdar (2008), Supatgiat *et al* (2001)), production cost based equilibrium models (eg Bessembinder and Lemmon (2002)) and high-frequency analyses of the dominant spot price drivers (eg Karakatsani and Bunn (2008)). Many authors have studied the impact on power prices of movements in physical variables such as temperature, rainfall patterns and other demand-side factors (eg Huisman (2008), Vevhilainen and Pyykkonen (2004)), while some have focused primarily on specifying the correct shape of the electricity supply function (e.g. Kanamura and Ohashi (2004)). These models all attempt to bridge the gap between pure power price models and the complex equilibrium models used in industry to forecast prices based on detailed market knowledge such as specific generators' costs, schedules and constraints. Amongst the class of stochastic, econometric style models most useful for derivative pricing, Eydeland and Wolyniec (2003) discuss the role of hybrid models, though they still rely on fairly detailed local market knowledge and large simulations. At the other end of the spectrum, Barlow (2002) proposes a simple non-linear Ornstein-Uhlenbeck model consisting of a mean-reverting demand process combined with a supply curve. A similar approach is taken by Skantze *et al* (2000), Eydeland and Geman (1999), Villaplana (2004), and Cartea and Villaplana

(2007), with the inclusion of capacity as a stochastic process, incorporated through an exponentially-shaped supply curve. Burger *et al* (2004) consider a non-parametric supply curve, fitted using cubic splines, and assumed to be a function of demand over capacity. A slightly different perspective is gained by considering reserve margin (extra capacity available beyond demand) to be a key driver, particularly as a low level of margin corresponds to a period of market strain, and consequently a higher chance of a sudden price spike. This idea is exploited in a regime-switching framework by Mount *et al* (2006) and Anderson and Davison (2008), a non-parametric approach by Boogert and Dupont (2008) and a model which incorporates forward-looking margin information by Cartea *et al* (2008).

While the appropriate choice of factors may vary across different markets, the articles above focus mainly on load (demand) and capacity fluctuations, which are known to be the main short-term drivers of power prices. Moving to medium or long term dynamics, data suggests that it is very important to incorporate fuel price risks into the modelling framework. Pirrong and Jermakyan (2005,2005) propose a fairly simple and intuitive model based on the two risk factors of demand and natural gas price. As in our work, they study the PJM market and advocate the use of historical generator bid data to estimate the transformation from underlying factors to power prices, suggesting a simple non-parametric approximation to the bid curve, derived from the bid curve on the same date of a previous year and adjusting by the ratio of gas prices (or forward gas prices) on the two dates. We propose a more realistic, parametric approach to the bid stack function, allowing for the overlap of bids from generators of different fuel types. Thus, our model extends their approach by allowing multiple fuel prices as factors, as well as addressing capacity or margin issues such as outages. This allows us to capture the rather complex dependence structure of power and fuel prices, as needed to price a variety of cross-commodity spread options and other derivatives.

The remainder of the paper is organised as follows. Section 2 introduces the two North East US markets used in our analysis. Section 3 presents our approach to modelling the bid stack function, which links electricity spot prices to underlying drivers. We specify slightly different models for PJM and NEPOOL and demonstrate the strong fit to historical bid data. Section 4 completes our modelling framework by proposing stochastic processes for each of the underlying risk factors: fuel prices, demand or load, and available capacity or margin. Section 5 assesses the performance of the model through analysis of both simulated price paths and forward prices, in comparison to observed data. Section 6 concludes.

2 Electricity Market Data Set

We shall focus on two US electricity markets: the PJM market (Pennsylvania, New Jersey and Maryland, plus parts of nine other Eastern states) and the NEPOOL market (the New England region). While our methodology can be adapted to fit many different local characteristics, we choose these two US markets primarily for the availability of historical bid data (at a six-month lag), as well as other useful historical information also published on the PJM and NEPOOL websites.¹ The availability of bid data in a convenient form is still uncommon in power markets and is certainly central to the methodology described here, but it is possible to imagine approximations or variations of the model with parameters estimated for example from cost or heat rate data.

PJM is a large market currently serving over 50 million people, with a total capacity which grew from 50,000 MW to 160,000 MW between June 2000 and July 2007. This time period forms our dataset for parameter estimation, and we use PJM West price data throughout.² The PJM market

¹See <http://www.pjm.com/markets/energy-market/historical-bid-data.html> and <http://www.iso-ne.com/markets/hstdata/index.html> for bid data.

²Power prices in PJM are complicated by the existence of Locational Marginal Prices (LMPs), corresponding to the different regions of each market. Each LMP is determined by calculating the price required to deliver one extra unit of power to that point on the grid. Factors such as local transmission constraints can lead to significant variation across regions. In all our analysis, we use the PJM West region, as these prices are used to calculate the value of PJM futures

has an interesting mix of fuel types (power sources), with a significant proportion of capacity coming from coal, gas and oil, and nuclear. Table 1 illustrates how this fuel type breakdown has changed slightly from year to year over the period considered.³

Table 1: Fuel breakdown (percent) in the PJM market (2002-06), and NEPOOL (2006)

Fuel Type	Jan 02	Jan 03	Jan 04	Jan 05	Dec 05	Dec 06	NEPOOL
Nuclear	22	17.6	17.1	19.1	19.1	18.7	16
Hydro	5	5.3	5.4	3.6	3.6	4.6	12
Coal	34	37.9	36.2	42.1	41.2	40.7	9
Gas	4	7.5	6.8	16.2	15.6	16	16
Gas/Oil	14	15	18.9	9.6	10.5	10.5	16
Oil	18	15.5	14.3	8.5	8.4	8.9	12
Other	3	1.2	1.3	1	1	0.7	2

In contrast to PJM, New England has a much smaller and younger market with capacity fairly stable around 30,000 MW throughout the dataset, which covers the time period March 2003 to August 2007. NEPOOL has a simpler fuel mix than PJM, having little coal-powered generation. Furthermore, gas and oil together represent nearly half of the capacity in the market, while the remainder is primarily nuclear and hydro. Most nuclear and some hydro generators typically make bids of zero as they have little flexibility in terms of switching on or off in response to demand. As we shall see in Section 3, NEPOOL can therefore be treated in a simplified framework as a one-fuel market, whereas PJM requires at least two fuel types to capture the dynamics realistically.

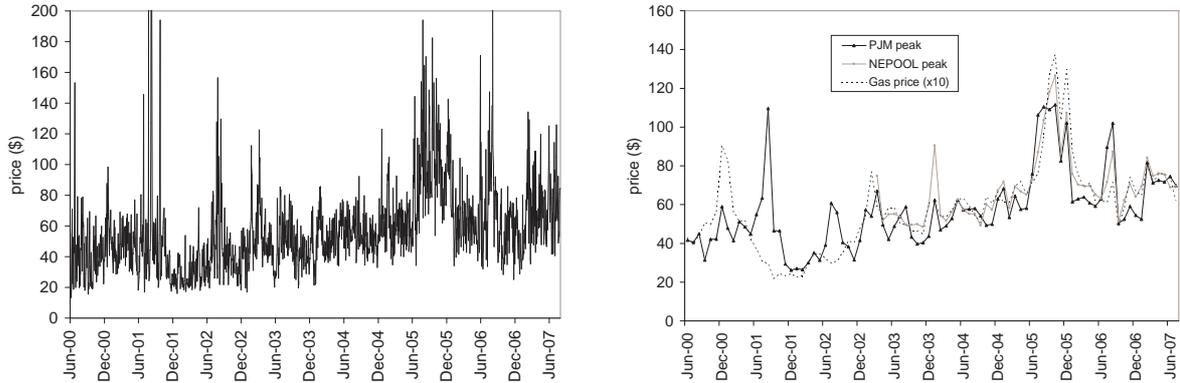


Figure 1: Daily average real-time peak prices for PJM (left graph); Monthly average real time peak prices for PJM and NEPOOL, as well as monthly average gas prices, multiplied by a factor of 10 for comparison purposes (right graph). (Note that the highest daily PJM price achieved in the left graph is \$637, and there are 8 data points beyond \$200, but only three since 2002, corresponding to the first three days of August 2006, with a maximum price of \$327.)

Figure 1 illustrates the dynamics of the real-time daily average peak price (average of hours 8-23, weekdays only) for PJM. The dynamics for NEPOOL are slightly less volatile but generally similar in appearance. Both real-time (RT) and day-ahead (DA) prices exist for both markets, with real-time prices typically more volatile. We shall only consider peak prices, as the modelling methodology is

contracts traded on NYMEX.

³PJM data has been taken from annual generating capacity reports (<http://www.pjm.com/services/system-performance/operations-analysis.html>) while NEPOOL data from Eydeland and Wolyniec (2003). Note that gas generators that also have oil-based generation capability are listed separately. So if all of these choose to use gas, the proportion of gas generators in the market is approximately 26% and 32% for PJM and NEPOOL respectively.

less well suited to describing off-peak prices, and thus a daily price will refer to a daily peak average. Although daily price series clearly reveal occasional price spikes, the magnitude and frequency of spikes is much higher at an hourly level (i.e., before averaging), as a result of brief outages or transmission constraints. The price series show high correlations with natural gas prices, also as illustrated in Figure 1. Here we have removed the noise by considering only monthly average prices, and obtain a remarkable visual correlation with monthly average Henry Hub gas prices. This link between gas and power prices is extremely strong for the gas and oil-dominated New England market, but also remarkably strong for PJM, as the market clearing price in peak hours is often set by the bids from gas generators. Capturing this relationship accurately is one of the primary advantages of our modelling approach.

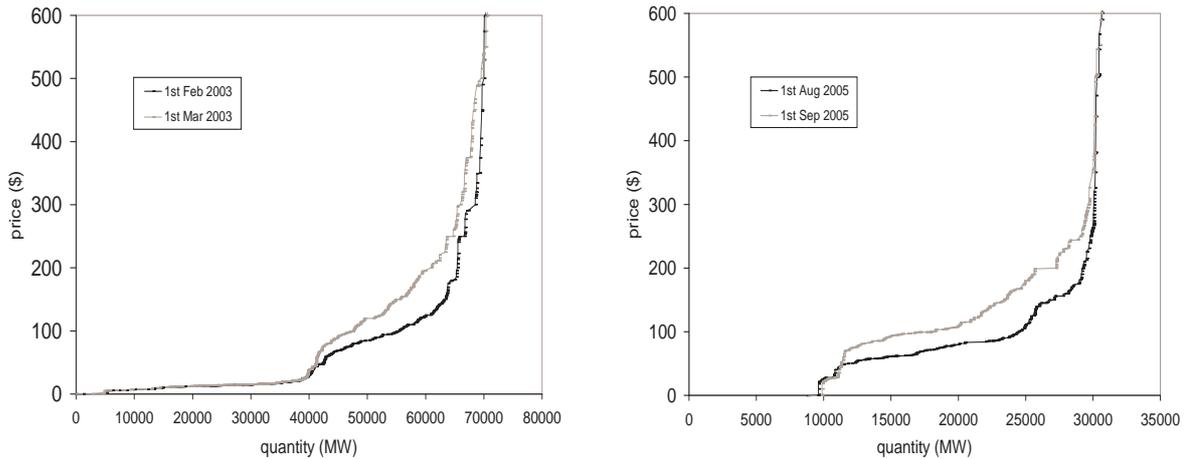


Figure 2: Sample bid stacks for PJM (left) and New England (right), showing movement as gas prices change. (The region \$600 to \$1000 is not shown as the bid stack is almost vertical for this range.)

In Section 3, we introduce our model for the bid stack, which can be understood as the map from the underlying random factors to the spot electricity price, and is thus the key component of the power price model. We use daily bid data, consisting of day-ahead bids from all available generators (close to 1000 for PJM, close to 300 for NEPOOL), for calibration.⁴ These bids describe the prices at which the generator is willing to sell varying amounts of electricity. Thus each generator submits a non-decreasing step function with a maximum of 10 steps, or price and quantity pairs. For example, a generating unit which bids (200MW, \$40), (300MW, \$50), and (350MW, \$80) is willing to sell its first 200MW of power at \$40, its next 100MW at \$50, and its final 50MW at \$80.⁵ By stacking all bids from all generators in order from lowest to highest price, we can create the market bid stack, typical examples of which are provided in Figure 2. The market administrator then determines the hourly spot price by matching with the total demand for power. Note that some generators submit ‘must-run’ bids, which we treat as bids of \$0, as they mean that the generator must sell its power no matter what the price is.⁶ As peak prices typically stay above \$30 but below \$150 (for over 90% of the hours observed during the most recent three years) the middle section of the bid stack is most relevant in determining prices, though the right hand side becomes important in the event of spikes. Figure 2 also shows that significant movement can occur from month to month, especially during times of large gas price increases, as was the case in both February 2003 and August 2005. We observe that

⁴Note that only daily observations of the bid stack are available for PJM, whereas hourly stacks are available for NEPOOL, though intra-day variation is low. We create an average bid stack for each day before performing the maximum likelihood estimation. Generator-specific issues such as start-up times and maximum run times per day are ignored for simplicity. Details of the name of or type of generator making each bid are not revealed.

⁵In PJM, generators have the alternative of connecting these bid points linearly instead of using a step function.

⁶In fact, this requirement of some generators can very occasionally produce negative prices in electricity markets, but this is only realistic during off-peak hours, which we do not model here.

during these months the majority of NEPOOL’s bid stack shifted upwards, while only the right hand half of PJM’s was affected.

Clearly power generators adjust their bids according to changes in their generation costs. We therefore expect a strong correlation between bid stack movements and fuel price changes, though the possible impact of other factors such as the exercise of market power or strategic bidding should be acknowledged. In order to understand this relationship, it is useful to consider the bid stack as a histogram of bids, as shown in Figure 3. We simply add up the total amount of capacity in MW that has been bid within each price bin.⁷ This provides an alternative perspective on the same data shown in Figure 2, and interestingly reveals one main cluster of bids for NEPOOL, but a pair of clusters for PJM separated by a region of few bids. We expect bids to be ordered roughly by fuel type corresponding to the merit order for each market (Table 1). This suggests that most bids in the left cluster of PJM’s histogram correspond to nuclear or coal generators, while the right cluster is primarily gas and oil, although low gas prices in particular can cause these clusters to merge somewhat. While this clustering provides our primary motivation for studying the bid stack, it is also important to discuss the far left and far right of the stack. The far left is less important as it consists of zero bids (including ‘must-run’) or very low bids, both corresponding primarily to nuclear power generators. This first 20-30% of capacity almost never determines the market clearing price during peak hours. On the other hand, the far right of the stack typically consists of a scattering of bids between about \$250 and \$1000 and therefore sets the price only during times of strain on the market. As we have seen, these instances occur fairly frequently during peak hours, as they correspond to the distinctive spikes visible in power prices. Our approach to modelling the entire bid stack focuses on the movement of these clusters of bids as fuel prices change.

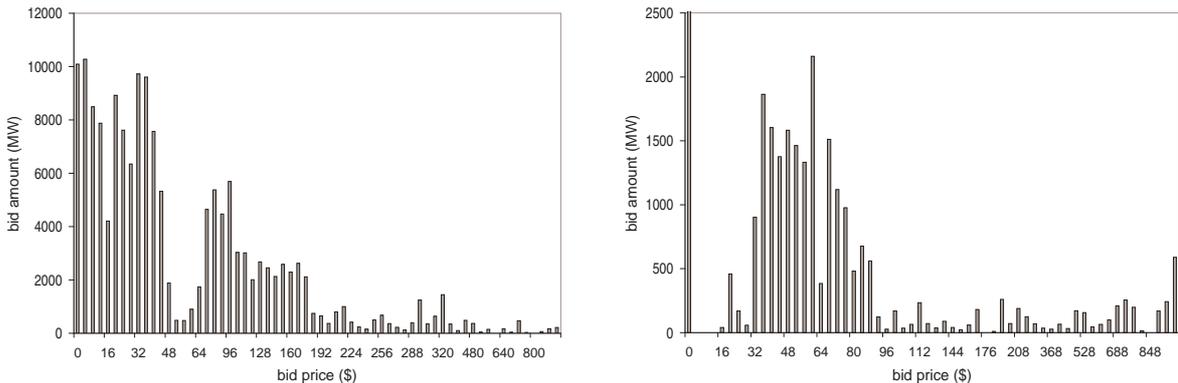


Figure 3: Sample histograms of bids for PJM (left) and New England (right). The quantity of bids at zero (including ‘must-run’ bids) for NEPOOL is far beyond the scale of the graph, at 8580MW.

3 The Bid Stack Model

Let S_t represent the spot price at time t , and $D_t \in [0, 1]$ the demand at time t , assumed to be inelastic with respect to price, as is often the case for electricity. We model demand not in terms of megawatt-hours but rather as a proportion of total market capacity, simplifying notation and allowing for growth in the market size. For generator i of n , let x_i be quantity supplied (again, normalised by total capacity) and $b_t^{(i)}(x_i)$ be the bid curve at time t , where $b_t^{(i)} : [0, c_i^{max}/c^{max}] \rightarrow [0, p^{max}]$.

⁷In order to plot the entire bid stack, we have chosen to vary the histogram bin size for different parts of the stack. For PJM, the bins covering the region $[\$0, \$64]$ have width $\$4$, while those covering $[\$64, \$320]$ have width $\$8$, and finally those in $[\$320, \$1000]$ have width $\$40$. Similarly, for New England, the bins covering the region $[\$0, \$112]$ have width $\$4$, while those covering $[\$112, \$208]$ have width $\$8$, and finally those in $[\$208, \$1000]$ have width $\$40$.

Generator i 's maximum capacity is c_i^{max} , total market capacity is $c^{max} = \sum_i c_i^{max}$, and p^{max} is the maximum bid allowed in the market (e.g. \$1000 for PJM). Then the market clearing price which allows supply and demand to match is given by

$$S_t = \max_{1 \leq i \leq n} \left\{ b_t^{(i)}(x_i) \right\}, \quad \text{where } \{x_1 \dots x_n\} = \operatorname{argmin}_{x_1, \dots, x_n} \left\{ \max_{1 \leq i \leq n} \left(b_t^{(i)}(x_i) \right) : \sum_{i=1}^n x_i = D_t \right\}.$$

Combining bid curves from different generators intuitively means stacking their component bids in order from lowest to highest. If we assume that all bid curves are strictly increasing (and step functions can be approximated by strictly increasing functions), then this corresponds to inverting each bid curve, adding the inverses, and inverting the sum. Letting $B_t^{obs}(\cdot)$ denote the exact bid stack observed in the market at time t , we can write

$$S_t = B_t^{obs}(D_t), \quad \text{where } B_t^{obs}(x) = I_t^{-1}(x), \quad \text{and } I_t(x) = \sum_{i=1}^n \left(b_t^{(i)} \right)^{-1}(x).$$

These equations provide a simplified description of an electricity market's structure, but rarely hold in practice. This is due to a variety of complications including generator outages, transmission constraints, imports or exports, variations in geographical distribution of demand, possible demand elasticity, and other rebalancing effects, especially for real-time prices. In order to capture these effects yet retain tractability, we introduce a process C_t for capacity available at time t (again normalised with c^{max}), and assume that now the bid stack is a function of D_t/C_t . In other words, any loss of supply is assumed to be equally spread throughout the stack, an approach also taken by Burger *et al* (2004). A fundamental requirement of this framework is that demand never exceeds capacity, so $0 < D_t/C_t < 1$. Typically we also observe $0 < D_t < C_t < 1$, though sometimes $C_t > 1$ (for example through imports of capacity). The spot price is now given by⁸

$$S_t = B_t \left(\frac{D_t}{C_t} \right), \quad \text{for } 0 < \frac{D_t}{C_t} < 1, \quad (1)$$

where the time dependence of the function $B_t(\cdot)$ is in fact a dependence on fuel prices, as these are the primary drivers of generators' bids. While this basic framework is only an approximation to the complexities of electricity markets, it allows us to retain the direct link to supply and demand factors, the flexibility to adapt to different markets, and the ability to price derivative products fairly easily. We now introduce our model for this function $B_t(\cdot)$, which we estimate directly from available bid data for PJM and NEPOOL.

3.1 General Case - fitting distributions to bids

As illustrated in Figure 3, a histogram of bids provides a useful alternative to simply observing the bid stack directly in Figure 2, and motivates fitting a density function to these histograms. Bids from generators with different fuel types are driven by different costs, leading to a mix of distributions in the overall market, with weights corresponding to the breakdown of fuel types in the market. With this new approach, the spot price $S_t = B_t(x)$ can be reinterpreted as the x -quantile of our bid distribution. Thus we fit a function to the density of bids and then deduce the quantile function (inverse cumulative distribution function), as opposed to fitting the bid stack (or quantile function) directly. One advantage is the wide range of well-known distributions that we can test. Furthermore, we can link distributions' parameters to the underlying fuel prices in an intuitive manner.

⁸We can interpret this spot price as being either a day-ahead or real-time price and either an hourly or daily average price, depending on what we are interested in modelling. The framework of the model remains the same. In our analysis, S_t is the hourly peak price process.

In the most general case, let $F_1(x), \dots, F_N(x)$ ($F_i(x) : \mathbb{R} \rightarrow [0, 1]$) be the proportion of bids below $\$x$ for generators of fuel type $i = 1, \dots, N$, with weights w_1, \dots, w_N , summing to unity. Then the spot power price S_t solves

$$F(S_t) = \sum_{i=1}^N w_i F_i(S_t) = \frac{D_t}{C_t},$$

and the bid stack is simply the inverse of the cdf, $F(x)$, of the mixture distribution.

To improve the fit in the most relevant region of the bid stack, we may wish to truncate the domain of D_t/C_t from $(0, 1)$ to (b_L, b_U) and ignore the tails of the bid distribution. This is only appropriate if $P[b_L < D_t/C_t < b_U] = 1$ (typically set $b_L = 0.2$ or 0.3 and $b_U = 0.9$ or 0.95). We then have

$$F(S_t) = \sum_{i=1}^N w_i F_i(S_t) = \frac{1}{b_U - b_L} \left(\frac{D_t}{C_t} - b_L \right).$$

The requirement that $D_t/C_t \in (b_L, b_U)$ poses problems from a modelling perspective. Therefore, we suggest an alternative approach of simply linearly rescaling both D_t and D_t/C_t such that for new variables \tilde{D}_t and \tilde{C}_t , we require $\tilde{D}_t/\tilde{C}_t \in (0, 1)$, just as in the untruncated case. In practice this means that the lowest portion of both demand and capacity is fixed, and changes in both demand and capacity only occur beyond this point in the stack. Moreover, any drop in available capacity is now assumed to affect only the region (b_L, b_U) of the bid stack. We then have

$$F(S_t) = \sum_{i=1}^N w_i F_i(S_t) = \frac{\tilde{D}_t}{\tilde{C}_t}, \quad \text{where} \quad \tilde{D}_t = \frac{1}{b_U - b_L} (D_t - b_L) \quad \text{and} \quad \frac{\tilde{D}_t}{\tilde{C}_t} = \frac{1}{b_U - b_L} \left(\frac{D_t}{C_t} - b_L \right) \quad (2)$$

Note that $\tilde{C}_t \neq C_t$, as \tilde{C}_t represents the percentage availability of capacity in the relevant region of the stack, not in the market in total.

We fit distributions to the bid data by maximum likelihood estimation, where a bid of q megawatts at price p is treated as q separate observations of a bid at p . Let (p_j, q_j) , $j = 1, \dots, M$ represent all the price quantity pairs that make up the portion $[b_L, b_U]$ of bid stack. So $\sum_{j=1}^M q_j$ is $(b_U - b_L)$ times the total capacity of the market. Consider the general case of fitting a mix of N distributions, with weights w_i (where $w_1 + \dots + w_N = 1$), and two-parameter density functions $f_i(x; \alpha_i, \beta_i)$, for $i = 1, \dots, N$. Then the log-likelihood function (for a given day) is

$$\begin{aligned} L(w_1, \dots, w_N, \alpha_1, \dots, \alpha_N, \beta_1, \dots, \beta_N) &= \log \left[\prod_{j=1}^M \left\{ \sum_{i=1}^N w_i f(p_j; \alpha_i, \beta_i) \right\}^{q_j} \right] \\ &= \sum_{j=1}^M q_j \log \left[\sum_{i=1}^N w_i f(p_j; \alpha_i, \beta_i) \right]. \end{aligned}$$

We have compared results using the following distributions: Gaussian, logistic, Cauchy, and Weibull. These all have appropriate humped shapes and only two parameters, one corresponding to at least roughly to the mean, and the other roughly to the standard deviation or shape. Hence we use the notation m_i, s_i , $i = 1 \dots N$ for these parameters. The performance of the four distributions is fairly similar in terms of both likelihood and capturing fuel price correlations, though thicker-tailed distributions dominate for higher choices of the cutoff point b_U , where the thin-tailed Gaussian performs erratically. Ultimately, we advocate the logistic distribution (with mean m_i and scale parameter s_i equal to $\sqrt{3}/\pi$ standard deviations) as the best choice, since it performs consistently for both markets and leads to the simplest mathematical expressions. For a mix of N logistic distributions, the log-likelihood function is given by

$$L(w_1, \dots, w_N, m_1, \dots, m_N, s_1, \dots, s_N) = \sum_{j=1}^M q_j \log \left\{ \sum_{i=1}^N \frac{w_i}{4s_i} \operatorname{sech}^2 \left(\frac{p_j - m_i}{2s_i} \right) \right\}. \quad (3)$$

3.2 One Fuel Case ($N = 1$) - NEPOOL

Beginning with the simpler NEPOOL case, the bids of generators can be split into bids of zero (roughly 30%) by nuclear and some hydro producers, and a cluster of bids primarily from oil and gas generators. Removing the lowest bids and considering the close relationship between gas and oil prices, a single fuel model is reasonable. We therefore estimate the bid stack parameters for each historical date as follows. Firstly, we ignore bids of below \$10 and also the highest 5%, 10% or 15% of bids, so $[b_L, b_U] \approx [0.3, 0.95]$, $[0.3, 0.9]$ or $[0.3, 0.85]$.⁹ Then (for the logistic distribution), we maximise with respect to m_1 and s_1 the likelihood function (3) with $N = 1$. We use a standard numerical optimisation scheme in MATLAB, while noting that closed form expressions exist for \hat{m}_1 and \hat{s}_1 only in the Gaussian case.

Figure 4 illustrates the MLE results for \hat{m}_1 and \hat{s}_1 with $b_U = 0.9$, also showing Henry Hub natural gas prices over the corresponding time period. Apart from two surprising spikes in bid levels in January 2004 and 2005, the correlation with gas prices is very high, and as much as 95% in the more recent data.¹⁰ The two spikes could perhaps have been caused by strategic bidding, though no particular information has been found to explain these events.¹¹ Nonetheless, the results strongly suggest assuming a linear dependence structure between m_1 and s_1 and the natural gas price G_t :

$$m_1 = \alpha_0 + \alpha_1 G_t, \quad s_1 = \beta_0 + \beta_1 G_t.$$

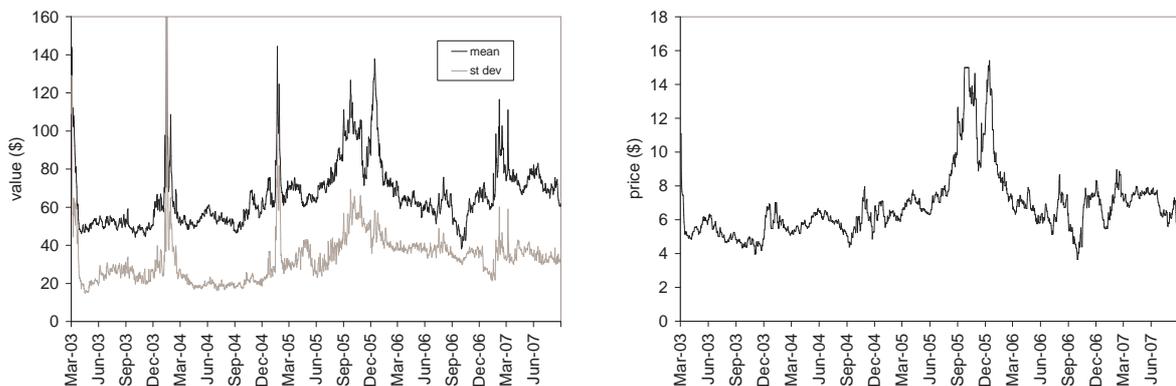


Figure 4: NEPOOL estimation results (left graph) for the mean \hat{m}_1 and standard deviation $\left(\frac{\pi}{\sqrt{3}}\right) \hat{s}_1$, compared with Henry Hub natural gas prices (right graph).

Estimating the parameters $\{\alpha_0, \alpha_1, \beta_0, \beta_1\}$ by regression produces slightly different results depending on the choice of distribution, upper cutoff of bids, b_U , and time period considered. The upper section of Table 2 shows the results for $b_U = 0.9$, which appears to be a reasonable choice, reducing the influence of the far right tail of the bid stack while keeping enough of the relevant region.¹² The results are very encouraging, showing high values of R^2 , particularly for \hat{m}_2 , and particularly over recent years, avoiding the two spikes described above.

While all four distributions studied share the useful property of having a fairly simple explicit inverse cumulative distribution function, this is particularly true for the logistic case. As a result,

⁹We first remove all bids between \$994 and \$1000, as there are sometimes large clusters of irrelevant bids at these levels which complicate matters if they are considered to be part of total capacity, c^{max} .

¹⁰Note also that correlations are higher if the gas price series has a lag of one day with respect to the bid stack parameters, as we would expect for day-ahead bidding. Thus we use a one day lag in our regression as well.

¹¹The PJM bid stack dynamics also show spikes in these months, but less dramatic ones. It is worth mentioning that no spike was observed in January 06 so there is no reason to expect this behaviour every January.

¹²Tests reveal that as we increase b_U from 0.8 to 1, values of R^2 in the regressions remain stable before falling off sharply after 0.9, in particular for \hat{s}_1 . For PJM, they increase gradually until about 0.95 before falling off, particularly for \hat{s}_2 .

the one-fuel model for NEPOOL leads to a convenient equation for the bid stack, and hence the spot electricity price. Under the assumptions introduced above, (2) can be written as follows:

$$S_t = \alpha_0 + \alpha_1 G_t + (\beta_0 + \beta_1 G_t) \left(\log(\tilde{D}_t) - \log(\tilde{C}_t - \tilde{D}_t) \right). \quad (4)$$

Thus we obtain a spot electricity price S_t which is linear in the natural gas price G_t , similar to the model for PJM prices by Pirrong and Jermakyan (2005). However, this only occurs in the one-fuel case in our model, so for NEPOOL, but not PJM. Note that the fairly simple form of (4) is very appealing, particularly for the pricing of forwards presented in Section 4.

3.3 Two Fuel Case ($N = 2$) - PJM

As discussed briefly in Section 2 and illustrated in Figure 3, the variety of fuel types in the PJM market suggests the use of at least two distributions to capture the behaviour of the bid stack. We choose a pair of distributions to roughly represent the coal and gas portions of the market, and estimate the bid stack parameters for each historical date as follows. Firstly set $[b_L, b_U] = [0.2, 0.85]$, $[0.2, 0.9]$ or $[0.2, 0.95]$, such that we ignore the highest 5, 10 or 15% and lowest 20% of bids for each date. The low bids in particular correspond primarily to the nuclear generators and hence do not move in the same manner as the neighbouring coal bids.¹³ Next, calculate fixed weights w_1 and $w_2 = 1 - w_1$ based on the split of coal versus gas and oil in Table 1. These weights change only at a few discrete points in time to reflect market changes. Finally (for the logistic distribution), maximise the likelihood function (3) for $N = 2$ with respect to $\{m_1, s_1, m_2, s_2\}$.

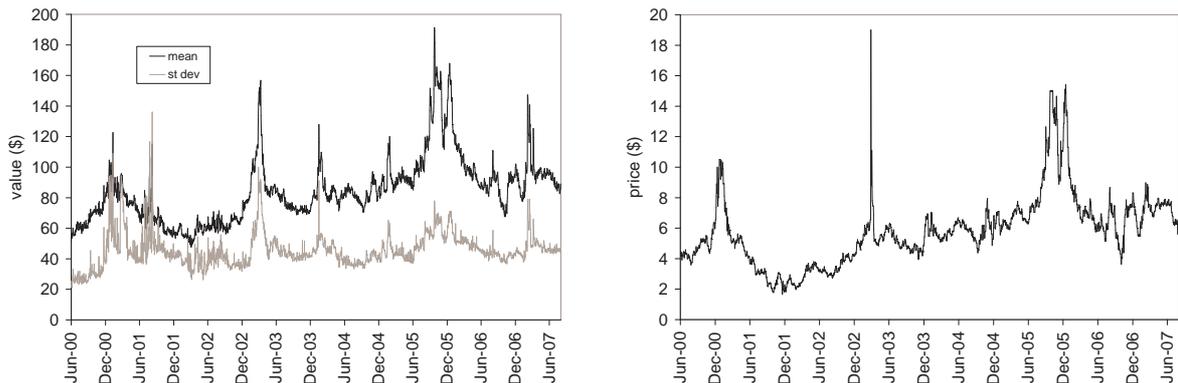


Figure 5: PJM estimation results (left graph) for the second distribution's mean \hat{m}_2 and standard deviation $\left(\frac{\pi}{\sqrt{3}}\right) \hat{s}_2$, compared with Henry Hub natural gas prices (right graph).

Figure 5 shows the gas distribution's estimated mean \hat{m}_2 and standard deviation \hat{s}_2 plotted against time in the logistic case with $b_U = 0.95$. As expected, both \hat{m}_2 and \hat{s}_2 show a strong correlation with the Henry Hub natural gas price, plotted again in Figure 5 over the corresponding time period. The results are particularly encouraging using the more recent data, with correlation as high as 96%. The results for \hat{m}_1 and \hat{s}_1 are plotted in Figure 6, along with the changes in Appalachian coal prices over the same period. Though not as striking as the gas correlation, some correlation is visible, particularly with the period of significant increase during the year 2004.

As for New England, these results for PJM suggest a linear dependence structure of the form,

$$m_1 = \tilde{\alpha}_0 + \tilde{\alpha}_1 P_t, \quad s_1 = \tilde{\beta}_0 + \tilde{\beta}_1 P_t, \quad m_2 = \alpha_0 + \alpha_1 G_t, \quad s_2 = \beta_0 + \beta_1 G_t,$$

¹³Tests reveal that including these lowest bids produces worse results by distorting the trends in coal parameters \hat{m}_1 and \hat{s}_1 . Tests with more than two distributions also do not appear to improve the results, and in fact tend to reduce the stability of the parameters \hat{m}_i and \hat{s}_i over time.

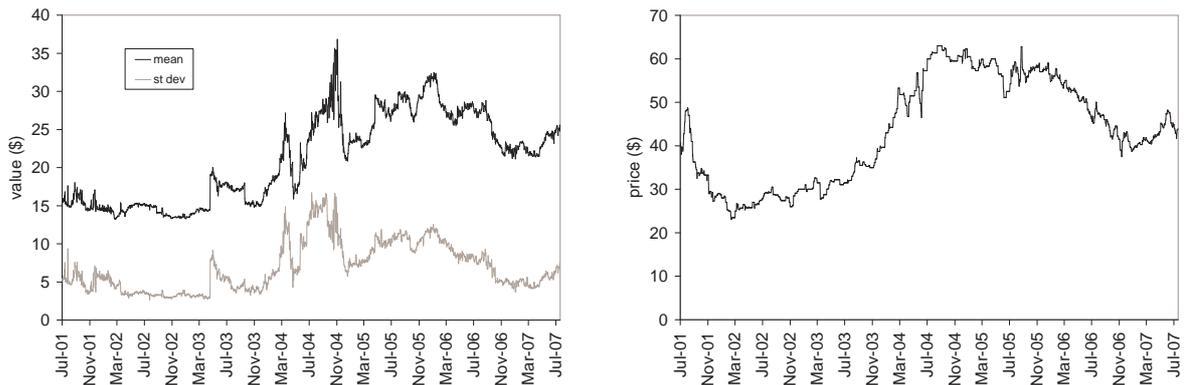


Figure 6: PJM estimation results (left graph) for the first distribution's mean \hat{m}_1 and standard deviation $\left(\frac{\pi}{\sqrt{3}}\right) \hat{s}_1$, compared with Appalachian coal prices (right graph).

where P_t and G_t are stochastic process for coal and gas prices respectively. Regression can again be used to estimate the parameters $\{\tilde{\alpha}_0, \tilde{\alpha}_1, \tilde{\beta}_0, \tilde{\beta}_1, \alpha_0, \alpha_1, \beta_0, \beta_1\}$. Results are shown in the middle and lower sections of Table 2, and are very positive, particularly for recent years. The R^2 values for \hat{m}_2 are remarkably high, and explain the source of the high correlation observed earlier in the monthly power and gas price series plotted in Figure 1. As before for NEPOOL, the standard deviation or scale parameter (here \hat{s}_2) shows a weaker relationship with gas prices than the mean, due in part to the method of bid data truncation at b_U .

Under the assumptions introduced above, our modelling framework from (2), can be written for PJM in the logistic case as follows:

$$S_t = x \quad \text{such that} \quad B_t^{-1}(x) = \frac{\tilde{D}_t}{\tilde{C}_t},$$

where

$$B_t^{-1}(x) = \frac{1}{2} + \frac{w_1}{2} \tanh\left(\frac{x - (\tilde{\alpha}_0 + \tilde{\alpha}_1 P_t)}{2(\tilde{\beta}_0 + \tilde{\beta}_1 P_t)}\right) + \frac{1 - w_1}{2} \tanh\left(\frac{x - (\alpha_0 + \alpha_1 G_t)}{2(\beta_0 + \beta_1 G_t)}\right) \quad (5)$$

As this function is not invertible explicitly, we cannot write down the bid stack function, but can easily solve for S_t numerically.

Table 2: Regression results for $\hat{m}_1, \hat{s}_1, \hat{m}_2, \hat{s}_2$ versus fuel prices, for NEPOOL ($b_U = 0.9$) and PJM ($b_U = 0.95$)

	Date range	\hat{m}_1 or \hat{m}_2			\hat{s}_1 or \hat{s}_2		
		inter	slope	R^2	inter	slope	R^2
NE (Gas)	Mar03-Aug07	17.35	7.67	0.701	7.36	1.29	0.168
	Mar05-Aug07	27.36	6.58	0.908	8.63	1.11	0.557
PJM (Coal)	Jun00-Jul07	3.38	0.408	0.727	-1.57	0.123	0.703
	Jun03-Jul07	6.43	0.355	0.487	-3.07	0.152	0.651
	Jun05-Jul07	7.02	0.390	0.749	-5.32	0.198	0.869
PJM (Gas)	Jun00-Jul07	35.15	8.51	0.833	17.82	1.25	0.233
	Jun03-Jul07	30.35	9.23	0.899	15.33	1.49	0.549
	Jun05-Jul07	31.03	9.20	0.927	15.23	1.53	0.674

3.4 Discussion

The high fuel price correlations observed in both PJM and NEPOOL bids justify the methodology of the bid stack model, both in terms of the use of actual bid data and the representation of the bid stack as an inverse cdf function. Furthermore, the use of Gaussian, logistic, Cauchy or Weibull distributions with parameters linked to fuel prices can be understood intuitively in terms of the mix of different heat rates for different generating units in the market. For example, suppose that for a given fuel, H is a random variable describing the variety of heat rates existing among different generating units due to factors such as age and technology. Suppose also that constant fixed costs A exist for each generator per MWh of power generated. Finally suppose that generators make bids corresponding exactly to their costs. Then the random variable $X_t = A + HG_t$ describes the bid of a randomly chosen gas generating unit given some gas price G_t . If $H \sim N(\mu_H, \sigma_H^2)$, then $X \sim N(A + \mu_H G_t, \sigma_H^2 G_t^2)$. This gives us precisely the Gaussian version of the model above, with parameters $\alpha_0 = A$, $\alpha_1 = \mu_H$, $\beta_0 = 0$ and $\beta_1 = \sigma_H$. With fixed costs $A > 0$ we expect $\alpha_0 > 0$, as we observe in our regressions for both coal and gas. The observed mix of positive and negative values of β_0 could be reproduced by letting A instead have a Gaussian distribution correlated with H .¹⁴ Of course, this argument for heat rate distributions would not work so conveniently mathematically using other distributions, but the intuition still holds. Interestingly, the values we observe for α_1 (7.67 and 8.51 for NEPOOL and PJM respectively) also correspond closely to average PJM heat rates for gas generators, listed by PJM as approximately 7.3 MBtu/MWh in 2004 (PJM report (2005)). Similarly, $\tilde{\alpha}_1 = 0.408$ matches equally well with the coal heat rate of 0.378 t/MWh used by Fehr and Hinz (2006). This suggests the possible use of heat rate or cost data as an alternative to bid data, though it should be remembered that strategic bidding could also have an influence on the parameters. Interestingly, recent work by Hortacsu and Puller (2007) on strategic bidding suggests that while marginal cost curves often contain prominent flat or vertical sections, optimal bid curves are typically smoothed to more closely resemble the bid stacks above. Thus, even for a market with a very narrow range of heat rates, strategic bidding could result in the fairly wide bid distributions we observe.

The results presented above for both PJM and NEPOOL strongly support the overall framework of the bid stack model to connect electricity, gas and coal prices. However, there are several complications that should be handled carefully when estimating parameter values. Clearly, we need to find a compromise between striving for a perfect fit of the bid stack and a rough approximation which captures the basic relationships. The primary issues are the best choices for b_U , b_L and w_1 , and the corresponding impact on spot price behaviour. The value of w_1 is calculated from the percentages in Table 1, and changes at six different dates during the time period, corresponding to times of market expansion.¹⁵ The value of b_L is much less significant than b_U since the peak power price is much more likely to be set in the far right than far left of the stack. Tests of regressions using different values of b_U confirm the choices of 0.9 (NEPOOL) and 0.95 (PJM) as appropriate. Only a few hourly prices are set beyond these points, and Section 4.4 completes our methodology for capturing the far right tail of the price distribution, by compensating for errors produced by truncating the bid data at b_U .

4 Modelling the Price Drivers

The bid stack model is supplemented by stochastic processes for the primary risk factors which drive the spot power prices: gas prices G_t (and coal prices P_t), demand (or load) D_t , and capacity available C_t . An advantage of the supply and demand approach is that while choosing fairly simple processes for the underlying factors, we can still replicate the unusual features of power prices through the

¹⁴Of course the distributions for H and A may change over time for example as technology improves, which suggests that regressions over more recent data might be considered more useful.

¹⁵We make the assume that nuclear is grouped with coal and hydro with natural gas while ‘other’ is split equally between coal and gas, before removing $b_L = 0.2$ from coal/nuclear and rescaling. This process for w_1 is then tested by comparing with the value of w_1 for each date which minimises the sum of squared errors between the central portion (25%-75%) of the model bid stack and the observed stack. Results lend support the method of choosing w_1 as described.

choice of bid stack function.

4.1 Fuel Prices

For both the PJM and NEPOOL markets, the important fuel price to model is the US natural gas price, as discussed earlier and confirmed by the results of the bid stack fit. Coal prices are less important, even for PJM, because they are generally less volatile than gas, they drive a flatter and lower region of the bid stack and they are less significant in setting prices particularly during peak hours. Nevertheless, we propose a simple method for incorporating some coal price information without increasing the complexity of the model. In a more general setting, other fuel prices and even carbon emissions prices can also be included, as is clearly necessary for European power markets today.

4.1.1 Gas Prices

We choose a standard approach to modelling gas prices, by fitting log gas prices with the Schwartz 2-factor model described by Schwartz (1997),¹⁶ capturing the mean-reversion which is widely believed to exist in most commodity prices, as well as changes to the long term equilibrium level of prices. Let X_t^1 and X_t^2 be the two independent stochastic factors driving the spot gas price G_t , and $h(t)$ be a seasonal component. Dynamics under the risk-neutral measure \mathbb{Q} are given by

$$\begin{aligned} dX_t^1 &= \kappa(\mu_1 - X_t^1)dt + \sigma_1 dW_t, \\ dX_t^2 &= \mu_2 dt + \sigma_2 d\tilde{W}_t, \\ G_t &= \exp(h(t) + X_t^1 + X_t^2). \end{aligned} \tag{6}$$

Gas forward prices $F^G(t, T)$ for delivery at a discrete point in time T have value at time t given by $E_t^{\mathbb{Q}}[G_T]$, the conditional expectation of G_T under \mathbb{Q} . Hence we have:

$$\begin{aligned} \log(F^G(t, T)) &= \log\left(E_t^{\mathbb{Q}}\left[e^{h(T) + X_T^1 + X_T^2}\right]\right) \\ &= h(T) + X_t^1 e^{-\kappa(T-t)} + \mu_1\left(1 - e^{-\kappa(T-t)}\right) + X_t^2 + \mu_2(T-t) \dots \\ &\quad + \frac{\sigma_1^2}{4\kappa}\left(1 - e^{-2\kappa(T-t)}\right) + \frac{1}{2}\sigma_X^2(T-t) \end{aligned} \tag{7}$$

As X_t^1 and X_t^2 are unobservable factors, we use the Kalman Filter to calibrate our model to available Henry Hub gas data (see eg Schwartz (1997), Lucia and Schwartz (2002) and Culot *et al* (2006) in the electricity literature). The transition equation is easily written in vector form using the SDEs in (6). As liquid futures or forward prices are widely available for the US gas market, (7) forms the measurement equation, which importantly is linear in the state variables. NYMEX has provided us with daily historical forward curves from January 2000 through November 2006. We use forward prices for maturities of 1 month, 3 months, 6 months, 1 year, 2 years, 3 years (and when available 4, 5 and 6 years), thus keeping 6 to 9 forward prices for each historical date and reducing computation time. We also assume that the forwards mature at the middle of each month, though in fact there is a month-long delivery period.

The next step is to remove seasonality from the forward curves. Though the shape of the seasonal pattern appears fairly consistently throughout the historical data, the amplitude varies significantly over time for both the forward curve and log-forward curve. Thus, we cannot easily remove the seasonality using the same function $h(T)$ throughout. Instead, we deseasonalise each forward curve independently as follows. We firstly identify the linear trend in the curve beyond the one-year maturity point, thus avoiding short end deviations. We then let the current month equal the base month, and calculate for every other month the average difference between forwards with maturity in that

¹⁶This approach seems to give more reasonable parameter values than simply modelling G_t , particularly for recent data where the positive skew in gas prices is significant. It also ensures that gas prices remain positive.

month and the most recent base month forward. We account for linear trend in the process and average from the one-year point onwards. Finally, we deseasonalise the forward curve by adding or subtracting these average monthly differences as appropriate, relative to the mean of all months.

As is typical for the Kalman Filter, results are very sensitive to the choice of observation noise σ_3 . Hence we perform the optimisation in two stages, similarly to the method of Culot *et al* (2006). For fixed σ_3 , we implement the filtering with respect to $\{\kappa, \mu_1, \sigma_1, \mu_2, \sigma_2\}$, choosing initial parameters $X_0^1 = 0.7$, $X_0^2 = 0$ for Jan 2000 ($X_0^1 = 1.2$, $X_0^2 = 0.6$ for Jan 2003), and the initial value of the conditional covariance matrix of X_t to be $0.1I$. Thus we find the parameter set that maximises the total likelihood of observing the entire history of forward curves. We then vary σ_3 and rerun the filtering for each new value, in order to minimise the sum of squared errors of all natural gas option prices available on NYMEX. As these are options on forward gas contracts and forwards are lognormally distributed as in (7), option prices (with strike K , option maturity T_1 , forward maturity T_2 , constant interest rate r) are given in closed form by

$$\begin{aligned} V^G(t, T_1, T_2) &= e^{-r(T_1-t)} E_t [(F^G(T_1, T_2) - K)^+] \\ &= e^{-r(T_1-t)} \left[e^{\tilde{\mu} + \frac{1}{2}\tilde{\sigma}^2} \left\{ 1 - \Phi \left(\frac{\log(K) - \tilde{\mu} - \tilde{\sigma}^2}{\tilde{\sigma}} \right) \right\} - K \left\{ 1 - \Phi \left(\frac{\log(K) - \tilde{\mu}}{\tilde{\sigma}} \right) \right\} \right], \end{aligned}$$

where Φ is the standard Gaussian cdf,

$$\tilde{\mu} = h(T_2) + X_t^1 e^{-\kappa(T_2-t)} + \mu_1 \left(1 - e^{-\kappa(T_2-t)} \right) + X_t^2 + \mu_2(T_2-t) + \frac{1}{2}\sigma_X^2(T_2-T_1) + \frac{\sigma_1^2}{4\kappa} \left(1 - e^{-2\kappa(T_2-T_1)} \right),$$

and

$$\tilde{\sigma}^2 = \frac{\sigma_1^2 e^{-2\kappa(T_2-T_1)}}{2\kappa} \left(1 - e^{-2\kappa(T_1-t)} \right) + \sigma_2^2(T_1-t).$$

Table 3: Kalman Filter results for natural gas parameters

Date range	κ	μ_1	σ_1	μ_2	σ_2	σ_3
Jan00 - Mar06	0.869	1.004	0.631	-0.034	0.141	0.059
Jan03 - Mar06	0.684	1.171	0.597	-0.036	0.121	0.079
Jan00 - Nov06	1.143	1.034	0.700	-0.044	0.136	0.066
Jan03 - Nov06	1.580	1.330	0.836	-0.053	0.114	0.079

The results are shown in Table 3 for different date ranges. The parameters seem fairly stable over time, though the more recent data is characterised by a slightly higher volatility and faster speed of mean reversion for X^1 and larger negative drift for X^2 . Note that $\mu_2 < 0$ is necessary to account for the fact that the gas forward curve was in backwardation (downward sloping) in 82.5% of the observations in the dataset.

4.1.2 Coal Prices

As PJM is roughly 40% fuelled by coal, changes in coal prices can have a significant effect on the level of power prices. However, as we have seen (e.g. Figure 6), coal prices move only gradually, suggesting that volatility is low, and therefore that it is more important to capture the market's expectations of future trends in coal prices, than to capture the stochastic component. Hence we use NYMEX Appalachian coal futures curves, and take the very simple approach of assuming that the coal price at time T in the future will exactly equal the futures price $F^C(t, T)$ with maturity T . Thus we have effectively assumed a deterministic coal price model which is matched to the current forward curve. While this seems artificially simple, it provides satisfactory results in the model without introducing an additional stochastic factor which is likely to have little impact on power prices.

Table 4: Results of fitting PJM and NEPOOL demand (Aug 04 - Jul 07)

Market	$\hat{\kappa}_Y$	$\hat{\mu}_Y$	$\hat{\sigma}_Y$
PJM (RT)	64.24	-0.001	1.393
PJM (DA)	52.311	-0.003	1.249
NEPOOL (RT)	81.69	-0.009	1.959
NEPOOL (DA)	132.06	-0.014	2.697

Market	a_1	a_2	a_3	a_4	a_5	a_6
PJM (RT)	-0.746	0	0.114	-0.862	0.186	-1.807
PJM (DA)	-0.776	0	0.106	-0.858	0.170	-1.633
NEPOOL (RT)	-0.894	0	-0.028	2.101	0.249	-1.826
NEPOOL (DA)	-1.078	0	-0.134	1.490	0.264	-1.958

4.2 Demand (Load)

Demand or load is easily observable in all electricity markets, and is typically characterised by multiple periodicities, at annual, weekly and intra-day levels. The lower lines in Figure 7 show average daily peak demand (real-time) in PJM and New England respectively. As we use daily data, we have averaged out intra-day effects and removed weekends and public holidays, so the primary seasonality remaining is annual. Peaks occur in both summer and winter corresponding to higher air conditioning and heating needs, with summer peaks larger than winter peaks, particularly for PJM. Initially we model this deterministic behaviour through a linear trend and a combination of two cosine functions with periods one year and six months. However, data shows that a linear trend is only necessary for the early years of PJM when the market expanded significantly in size. Thus we instead choose the three year period August 2004 to July 2007 and set $a_2 = 0$ for both markets, as there is no significant trend. Finally, the deseasonalised process is fitted by maximum likelihood estimation using an exponential OU process, as shocks to demand are considered to revert rapidly to the seasonal level. We work with rescaled demand \tilde{D}_t throughout, with $(b_L, b_U) = (0.2, 0.95)$ for PJM, and $(b_L, b_U) = (0.3, 0.9)$ for NEPOOL. Modelling $\log(\tilde{D}_t)$ ensures that \tilde{D}_t must remain positive, as required. Hence we have

$$\begin{aligned}
\log(\tilde{D}_t) &= f(t) + Y_t & (8) \\
f(t) &= a_1 + a_2 t + a_3 \cos(2\pi t + a_4) + a_5 \cos(4\pi t + a_6) \\
dY_t &= \kappa_Y (\mu_Y - Y_t) dt + \sigma_Y dB_t
\end{aligned}$$

where B_t is a Brownian Motion independent of W_t and \tilde{W}_t in the gas process. We assume throughout that fuel prices are independent of demand and capacity, which fluctuate on shorter time scales and are driven by more local conditions; Pirrong and Jermakyan (2005) also suggest this to be a reasonable assumption, though of course a prolonged cold spell over a large region is likely to impact both power demand and gas prices.

Table 4 lists the results for both PJM and NEPOOL, and both real-time (RT) and day-ahead (DA) demand, with $t = 0$ corresponding to June 1st 2000. The results show that mean-reversion rates κ_Y and volatility σ_Y are higher for NEPOOL than PJM over the chosen period. Although we are interested in hourly spot prices, we fit our demand model only to daily peak average demand as intra-day movements are fairly small and dominated by the intra-day periodicity. This can easily be incorporated into the model, for example with hourly indicator variables, but has very little effect. Intra-day demand movements are ultimately overshadowed by intra-day capacity jumps, as discussed below.

4.3 Capacity Available

As explained in Section 3, the process C_t (or \tilde{C}_t) captures a variety of supply-side information relating to outages, transmission constraints, exports, imports and other power delivery issues. Therefore it

is not easily observable or even intuitively understood, though we can think of it simply as the percentage of maximum capacity available. Since we observe hourly historical prices, demand and bid stacks, we can calculate the implied capacity available C_t^{imp} which allows (1) to hold as closely as possible.¹⁷ As $B^{obs}(\cdot)$ is non-decreasing, C_t^{imp} is uniquely defined by¹⁸

$$C_t^{imp} = \max \left\{ c \in \mathbb{R}^+ : B^{obs} \left(\frac{D_t}{c} \right) \geq S_t \right\}.$$

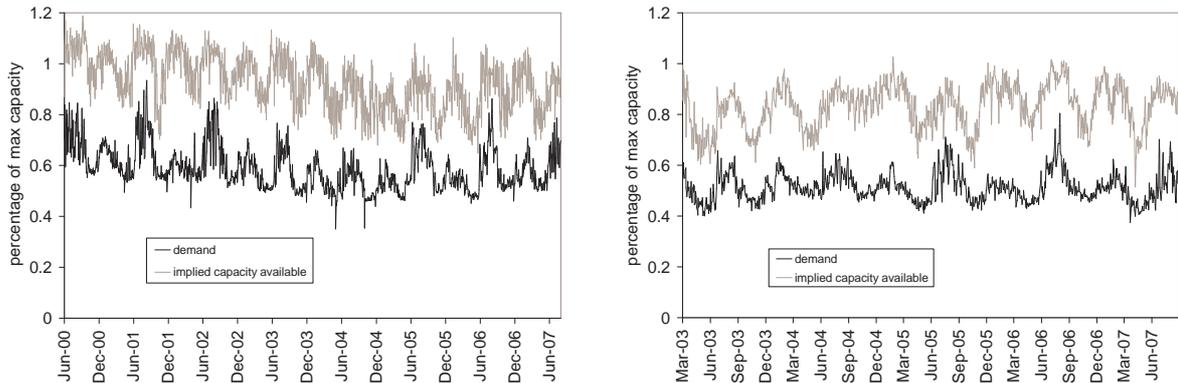


Figure 7: Daily average peak demand D_t (lower line) and implied capacity available C_t^{imp} (upper line), for PJM (left) and NEPOOL (right). (RT data is plotted for PJM, while DA is used for NEPOOL.)

The upper lines in Figure 7 show the historical evolution of the daily peak average implied capacity availability C_t^{imp} for both PJM and NEPOOL. It has clear seasonality matching roughly with demand seasonality, thus dampening the seasonality of prices. This is due to generators' maintenance schedules which are designed to avoid high demand periods.¹⁹ C_t^{imp} thus incorporates both expected and unexpected outages.

Using instead the truncated bid data and rescaled \tilde{D}_t and \tilde{C}_t involves solving for rescaled \tilde{C}_t^{imp} from (2) without breaking the condition $0 < \tilde{D}_t / \tilde{C}_t^{imp} < 1$. This holds as long as $B^{obs}(b_L) < S_t < B^{obs}(b_U)$ for all historical data, or equivalently, $b_L < D_t / C_t^{imp} < b_U$. However, for PJM, we find that $D_t / C_t^{imp} > 0.95$ for 0.41% of recent hourly data, while for NEPOOL $D_t / C_t^{imp} > 0.9$ for 0.6%. For these observations, we cannot define \tilde{C}_t^{imp} as above, without breaking the restriction $\tilde{D}_t < \tilde{C}_t^{imp}$. However, we can always define the (rescaled) model-implied capacity available \tilde{C}_t^{mod} as the unique solution to

$$S_t = B \left(\frac{\tilde{D}_t}{\tilde{C}_t^{mod}} \right),$$

where $B(\cdot)$ is our model bid stack, as defined in (4) or (5). In general, we expect \tilde{C}_t^{mod} to remain close to \tilde{C}_t^{imp} for the majority of data but to behave differently in the tails.

¹⁷Recall that B^{obs} is likely to be a step function, although some generators in PJM may bid continuous curves making it really a combination of step functions and piecewise linear functions.

¹⁸Occasionally, we observe $D_t > D_t / C_t^{imp}$, implying excess capacity available in the market relative to normal maximum capacity ($C_t^{imp} > 1$ in Figure 7). While this could realistically be caused by imports or slight demand elasticity to price, it can be a modelling concern particularly for low values of D_t . In some of these cases our assumption that the bid stack is a function of demand over capacity will lead to very high values of C_t^{imp} such as 1.5 or 2. In these cases, we adjust C_t^{imp} by assuming that the extra capacity enters the market only in the portion of the bid stack below where the price is set (as there would be no particular reason for the extra capacity to appear throughout the irrelevant portion of the stack). Hence we set $C_t^{imp} = D_t + (\text{extra capacity in the right of the stack}) = D_t + 1 - D_t / C_t^{imp}$.

¹⁹While it might be expected that bid data should not include generators known to be undergoing maintenance, historical bid data is in fact observed prior to the incorporation of scheduled outages.

4.4 Margin

Given our model for \tilde{D}_t in (8) above, modelling \tilde{C}_t separately makes it difficult to satisfy our fundamental requirement that demand does not exceed capacity. One approach is to model \tilde{D}_t/\tilde{C}_t directly, but this prevents us from disentangling demand and supply effects, and from using easily observable and well behaved demand data \tilde{D}_t . Furthermore, it still leaves us the problem of ensuring $\tilde{D}_t/\tilde{C}_t \in (0, 1)$. Instead we propose a stochastic process for the reserve margin $\tilde{M}_t = \tilde{C}_t - \tilde{D}_t$, representing the amount of extra capacity available in the market but not needed to match demand. By modelling both \tilde{D}_t and \tilde{M}_t as strictly positive processes, we automatically fulfil the required condition.

Hence we define (rescaled) implied margin \tilde{M}_t^{imp} and (rescaled) model-implied margin \tilde{M}_t^{mod} as:

$$\tilde{M}_t^{imp} = \tilde{C}_t^{imp} - \tilde{D}_t \quad \text{and} \quad \tilde{M}_t^{mod} = \tilde{C}_t^{mod} - \tilde{D}_t$$

Hourly historical data for \tilde{M}_t^{imp} suggests the need for a two-factor model for margin. The movement of daily peak averages over weeks or months shows both mean-reversion and some clear negative correlation with demand \tilde{D}_t , as one would expect. Upward shocks to demand often lead to downward shocks to margin, though market mechanisms such as imports, extra capacity reserves and transmission factors can dampen the effect and reduce the correlation. In addition, intra-day hourly margin reveals quite noisy behaviour with sudden and short-lived jumps due to a variety of short-term effects such as outages. We are interested less in describing the precise timing or autocorrelation of these spikes than in describing their magnitude and likelihood. Therefore, we propose a simple regime-switching model for $\log \tilde{M}_t$ consisting of an OU process for the ‘normal regime’ and an independent sample of a shifted exponential random variable for the ‘spike regime’:

$$\log(\tilde{M}_t) = \begin{cases} Z_t^{OU} & \text{with probability } 1 - p_i \\ Z_t^{SP} & \text{with probability } p_i \end{cases}$$

where the normal regime is given by

$$\begin{aligned} Z_t^{OU} &= \kappa_Z (\mu_Z - Z_t^{OU}) dt + \sigma_Z d\tilde{B}_t \\ dB_t d\tilde{B}_t &= \rho dt \end{aligned} \tag{9}$$

and the spike regime is given by

$$Z_t^{SP} = \alpha - J, \quad J \sim \text{Exp}(\lambda_i), \quad \text{for seasons } i = 1, 2, 3, 4.$$

Thus, each value of Z_t^{SP} (in practice sampled hourly) is independent of previous values, and the probability of being in the spike regime in any future hour is also independent of the current regime.²⁰ Data suggests that the left tail of the hourly margin distribution for PJM is significantly thicker in the summer months, suggesting a higher chance of outages. Therefore we fit seasonal spike parameters $p_i, \lambda_i : i = 1, 2, 3, 4$ where $i = 1$ corresponds to winter (Dec - Feb), $i = 2$ to spring (Mar - May), $i = 3$ to summer (Jun - Aug) and $i = 4$ to autumn (Sep - Nov).

In order to estimate the parameters above, we firstly use daily average implied margin \tilde{M}_t^{imp} , which averages over intra-day spikes, to help better identify the behaviour of Z_t^{OU} and especially its correlation with \tilde{D}_t . We estimate the parameters $\{\kappa_Z, \mu_Z, \sigma_Z, \rho\}$ by maximum likelihood, conditioning on the observed daily value of demand \tilde{D}_t .²¹ We use the same date range as for the MLE of the demand process. We then move to hourly data to fit the spike regime through a moment matching procedure. In order to estimate the point α beyond which the spike regime should operate, we exploit

²⁰Tests of historical data using a more formal continuous time Markov Chain with transition matrix lead to an expected duration of stay in the spike regime of approximately 2 hours for PJM and 3 hours for NEPOOL. Thus a more complete regime switching model for \tilde{M}_t leads to similar conclusions regarding the rapid speed of recovery from spikes.

²¹We cap $\tilde{D}_t/\tilde{C}_t^{imp}$ at 0.99 to avoid the rare cases of negative \tilde{M}_t^{imp} when $\tilde{D}_t/\tilde{C}_t^{imp} > b_U$ and reduce the largest downwards spikes in margin. This ultimately has little impact as the left tail of margin is captured in the second stage of estimation.

Table 5: Results of fitting PJM and NEPOOL margin (Aug 04 - Jul 07)

Market	$\hat{\kappa}_Z$	$\hat{\mu}_Z$	$\hat{\sigma}_Z$	$\hat{\rho}$	adjusted $\hat{\mu}_Z$	adjusted $\hat{\sigma}_Z$
PJM (RT)	133.59	-1.288	8.20	-0.358	-1.089	6.12
NEPOOL (DA)	76.00	-1.178	5.54	-0.125	-1.038	4.79

Market		winter	spring	summer	autumn
PJM (RT)	\hat{p}_i	0.090	0.103	0.174	0.109
	$\hat{\lambda}_i$	1.340	1.727	0.712	1.896
	$\hat{\alpha}$	-1.910	-1.910	-1.910	-1.910
	R^2 of fit	0.954	0.977	0.940	0.926
NEPOOL (DA)	\hat{p}_i	0.091	0.077	0.086	0.086
	$\hat{\lambda}_i$	1.501	1.481	1.405	1.487
	$\hat{\alpha}$	-1.830	-1.830	-1.830	-1.830
	R^2 of fit	0.949	0.937	0.983	0.900

the fact that Z_t^{OU} has a Gaussian distribution and hence a skew of zero. In contrast, the historical distribution of $\log \tilde{M}_t^{imp}$ has significant negative skew due to the thick left tail caused by outages. We remove as many data points as necessary to obtain a non-negative skew and choose the last point removed to be our parameter estimate for α . Since the remaining historical distribution (with points $\log \tilde{M}_t^{imp} < \alpha$ removed) is a more accurate representation of the invariant distribution of Z_t^{OU} , we re-estimate parameters for μ_Z and σ_Z to match the first two moments of this distribution. In other words, we equate $\hat{\mu}_Z$ to the mean of the truncated distribution, and $\hat{\sigma}_Z$ to its standard deviation multiplied by $\sqrt{2\hat{\kappa}_Z}$, with $\hat{\kappa}_Z$ as before. Essentially, we argue that parameter estimates $\hat{\kappa}_Z$ and $\hat{\rho}$ are well fitted by our initial procedure, while the initial estimators $\hat{\mu}_Z$ and $\hat{\sigma}_Z$ are distorted by the spikes and require adjustment via moment matching. Therefore,

$$\hat{\alpha} = \min \left\{ \alpha \in \mathbb{R} : \text{Skew} \left(\log \left(\tilde{M}^{imp} \right) 1_{\{\log(\tilde{M}^{imp}) \geq \alpha\}} \right) \geq 0 \right\},$$

$$\hat{\mu}_Z = E \left[\log \left(\tilde{M}^{imp} \right) 1_{\{\log(\tilde{M}^{imp}) \geq \hat{\alpha}\}} \right], \quad \hat{\sigma}_Z = \sqrt{2\hat{\kappa}_Z} \text{StDev} \left(\log \left(\tilde{M}^{imp} \right) 1_{\{\log(\tilde{M}^{imp}) \geq \hat{\alpha}\}} \right).$$

The final step of the parameter estimation is to find seasonal spike regime parameters p_i , and λ_i , for $i = 1, 2, 3, 4$. Here we switch from using the left tail of the distribution \tilde{M}^{imp} to that of \tilde{M}^{mod} . This key step allows us to compensate for any errors made in fitting the tail of the bid stack, in particular by not capturing the bids in the region $(b_U, 1)$. Though we may retain an artificially steep tail for the bid stack, we correct this by allowing the left tail of the margin distribution to be artificially stretched to produce the observed price spikes. Figure 8 shows log-histograms of $\log \tilde{M}^{mod}$ using the logistic distribution (and as usual $b_U = 0.9$ for NEPOOL and $b_U = 0.95$ for PJM). The observed linearity in the left tail justifies the use of an exponential distribution for the outage regime. Clearly for PJM the summer months require a different fit than other seasons, though for NEPOOL the difference between seasons is much less. The parameters p_i for each season are estimated simply by finding the proportion of observations below $\hat{\alpha}$, while λ_i is estimated as the slope of an ordinary least squares linear fit to the tail of the log-histograms.²² Table 5 lists all the estimated parameters for the margin process (with RT data for PJM, and DA data for NE), as well as the R^2 values for the linear fits to the tail, which are all above 0.9. Unlike demand, margin appears to be more volatile and ‘spikier’ for PJM than for NEPOOL, with higher values for κ_Z and σ_Z , as well as the spike regime probabilities p_i . Finally, PJM’s values of $p_3 = 0.174$ and $\lambda_3 = 0.712$ confirm that much larger and more frequent spikes occur in the summer, as also noted by Geman and Roncoroni (2006).

²²Firstly, a slight correction is required for p_i , since there exists a positive probability $q = \Phi \left(\frac{\hat{\alpha} - \hat{\mu}_Z}{\hat{\sigma}_Z} \right)$ of values below $\hat{\alpha}$ occurring in the non-spike regime. Hence, if \bar{p}_i is the percentage of observations below $\hat{\alpha}$, then $p_i = \frac{\bar{p}_i - q}{1 - q}$. Secondly, the regression to find λ_i has been performed over the ranges $[\hat{\alpha} - 3, \hat{\alpha}]$ and $[\hat{\alpha} - 1.5, \hat{\alpha}]$ for PJM and NEPOOL respectively, each split into six equal width probability bins.

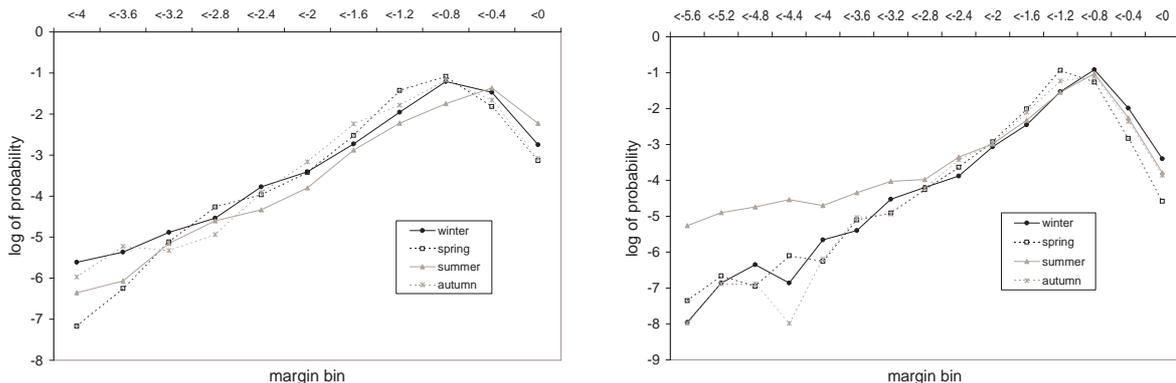


Figure 8: Log histograms of model implied log-margin $\log \tilde{M}_t$ for NEPOOL (left) and PJM (right), over the three year period from August 04 to July 07.

At this stage it is easy to demonstrate the additional convenience of choosing the logistic distribution for our bid stack model in Section 3. Firstly, for the one-fuel case of NEPOOL, our equation for electricity spot prices, (4), can be rewritten as follows:

$$S_t = \alpha_0 + \alpha_1 G_t + (\beta_0 + \beta_1 G_t) \left(\log \tilde{D}_t - \log \tilde{M}_t \right)$$

As S_t is linear in $\log \tilde{M}_t$, for a given demand and gas price, the exponentially distributed left tail of log-margin (from the spike regime) translates to exponentially distributed price spikes as well. Moreover, for a fixed demand, writing

$$\log \tilde{M}_t = \log \tilde{D}_t - \frac{S_t - m_1}{s_1},$$

we can translate the margin spike threshold α into a price spike threshold corresponding to a certain number of standard deviations in the bid distribution for gas.²³ Unlike typical regime-switching models for power prices (eg De Jong and Huisman (2005), Weron *et al* (2004)), we thus have an exponential ‘spike distribution’ for power prices which shifts over time as gas prices vary. For example, Figure 1 suggests that while an hourly price of \$150 could reasonably be considered a spike in early 2004, it could not be in late 2005 when high gas prices caused daily average peak prices to approach these levels.

While formulas are not as simple in the two-fuel case for PJM, the linear relationship between S_t and $\log \tilde{M}_t$ still holds approximately for the tail, and hence the spike regime. This follows because the influence of the coal distribution is negligible in the far right of the bid stack, so we can think of using a one-fuel model as an approximation. Then

$$\log \tilde{M}_t \approx \log \tilde{D}_t - \frac{S_t - m_2}{s_2}, \quad (10)$$

where \tilde{M}_t and \tilde{D}_t (and \tilde{C}_t below) represent a second rescaling of demand and margin from (b_L, b_U) to the gas-dominated portion of the bid stack $(b_L + w_1(b_U - b_L), b_U)$. The rescaling technique introduced

²³Ignoring the role of demand in determining price spikes is an approximation, but demand varies much less than margin. As the bid stack is an increasing function of $\frac{\tilde{D}_t}{\tilde{D}_t + \tilde{M}_t}$, demand shocks combined with margin shocks will clearly produce higher spikes than margin shocks alone.

in (2) also implies

$$\begin{aligned}
\log(\tilde{M}_t) &= \log \left[\tilde{D}_t \left(\frac{1}{\tilde{D}_t/\tilde{C}_t} \right) - \tilde{D}_t \right] \\
&= \log(\tilde{D}_t) + \log \left[\frac{1 - w_1}{\tilde{D}_t/\tilde{C}_t - w_1} - 1 \right] \\
&= \log(\tilde{D}_t) + \log \left[\frac{1 - \frac{\tilde{D}_t}{\tilde{D}_t + \tilde{M}_t}}{\frac{\tilde{D}_t}{\tilde{D}_t + \tilde{M}_t} - w_1} \right] \\
&= \log(\tilde{D}_t) + \log(\tilde{M}_t) - \log \left((1 - w_1)\tilde{D}_t - w_1\tilde{M}_t \right) \\
&\approx \log(\tilde{D}_t) + \log(\tilde{M}_t) - \log(\tilde{D}_t) - \log(1 - w_1)
\end{aligned}$$

where the last line holds for small margin \tilde{M}_t . Thus, in the right tail of the bid stack, using (10),

$$\log \tilde{M}_t \approx \log \tilde{D}_t + \log(1 - w_1) - \frac{S_t - m_2}{s_2}.$$

The approximate linear relationship between S_t and $\log \tilde{M}_t$ again suggests that, holding gas and demand constant, price spikes caused by margin spikes are also exponentially distributed.

5 Empirical Results

The performance of the model can be evaluated according to several different criteria for each of the two markets considered. We firstly compare the properties of simulated electricity price paths with those observed in the market, as is suggested by Geman and Roncoroni (2006). Typical statistics such as mean, variance, skew and kurtosis are considered, but also other key features in power markets, such as correlation with fuel prices and the probability of spikes above certain threshold prices. We make comparisons of price series at various time-scales, since statistics such as correlation are most relevant in terms of daily, weekly or even monthly averages. Though most of our state variables are mean-reverting, the component X_t^2 of gas prices is simply a Brownian Motion with drift, implying that no invariant distribution for spot prices S_t exists. Hence some statistics are likely to be unstable over time and should be analysed with care. Our second tool for evaluating model performance is a comparison of model-implied forward prices and observed forward prices. In power markets, it is particularly important for a model to capture the forward curve accurately, a comparison we make for the sample dates 30 December 2005, 31 March 2006, and 29 September 2006. These form a fairly representative sample as they correspond to times of high, medium and low gas prices respectively, and thus will also provide three different starting points for the simulation analysis.

Throughout this section, we use the logistic distribution for the bid stack model ($b_U = 0.9$ for NEPOOL, $b_U = 0.95$ for PJM), with parameters given in Table 4. For the PJM gas distribution, we use the regression results for $\{\alpha_0, \alpha_1, \beta_0, \beta_1\}$ over the entire dataset June 2000 - July 2007. However for the coal distribution, this leads to a consistent underestimation of \hat{m}_1 and \hat{s}_1 during 2006. This could be due in part to estimation difficulties during periods of significant gas and coal distribution overlap (2004-05), or perhaps to a gradual change in the heat rates of coal generators over time. As a result we use instead the regression results for $\{\tilde{\alpha}_0, \tilde{\alpha}_1, \tilde{\beta}_0, \tilde{\beta}_1\}$ over the most recent period June 2005 to July 2007. For NEPOOL, we use the parameters $\{\alpha_0, \alpha_1, \beta_0, \beta_1\}$ from regression over the entire period March 2003 to August 2007. For the gas process G_t , we use parameters in Table 5 for January 2000 - November 2006, making use of the most data. Finally, for demand and margin, we use the parameters in Tables 6 and 7, noting that forward prices (taken from NYMEX) are settled using real-time prices for PJM, but day-ahead prices for NEPOOL.

5.1 Simulation Analysis

Tables 6 and 7 show a comparison of statistics for observed and simulated price data. We include two different observed time periods (Mar 03 to Aug 07 and Mar 05 to Aug 07 for NEPOOL, Jun 00 to Jul 07 and Aug 04 to Jul 07 for PJM) to illustrate the significant variation in all of these parameters over time. In particular, the estimates of skew, kurtosis and the probability of hourly prices greater than \$100 are very unstable over time, and other statistics to a lesser extent. As discussed, part of the problem lies in the the non-stationarity of gas prices, while other issues such as growth and development of the market may also play a role. For the simulated statistics, we average over 2000 two-year Monte Carlo simulations of hourly prices with three different starting values for natural gas prices, $G_0 = 4.11, 7.19$ and 10.04 .²⁴ Note that while \tilde{M}_t and \tilde{D}_t are simulated under the real-world measure \mathbb{P} , gas prices G_t are simulated under the risk-neutral measure \mathbb{Q} . Therefore, there is likely to be some difference in simulated statistics and observed statistics due to the market price of natural gas risk. In addition, comparing an average over 2000 gas trajectories with one observed historical trajectory will inevitably lead to significant differences. As a result, we include a further simulation which uses instead a fixed gas price trajectory matching the path of gas prices during the more recent observed period for each market (05-07 for NE, 04-07 for PJM).

Table 6: NEPOOL statistics of observed and simulated paths

Time scale	Statistic	Observed			Simulated			
		03-07 (DA)	05-07 (DA)	05-07 (RT)	Low G_0 (Sep06)	Mid G_0 (Mar06)	High G_0 (Dec05)	Fixed G (05-07)
hourly	mean	64.20	67.96	77.25	76.18	87.85	96.05	77.74
hourly	st dev	25.90	28.49	35.00	24.36	28.68	30.71	20.47
hourly	skew	2.72	0.61	12.18	0.83	0.88	0.86	1.47
hourly	kurtosis	25.89	0.14	291.47	4.12	4.17	4.11	6.09
hourly	prob>\$100	9.07%	13.68%	15.62%	18.06%	27.98%	37.58%	12.61%
hourly	prob>\$200	0.10%	0.00%	0.26%	0.82%	1.80%	2.47%	0.03%
hourly	prob>\$300	0.08%	0.00%	0.10%	0.07%	0.18%	0.22%	0.00%
daily	st dev	23.26	24.92	25.00	22.84	27.05	28.99	18.72
daily	skew	3.20	0.85	3.94	0.61	0.68	0.67	1.40
daily	kurtosis	28.40	0.34	36.79	3.16	3.26	3.27	5.24
daily	corr with G	0.674	0.793	0.645	0.916	0.926	0.932	0.908
weekly	corr with G	0.768	0.846	0.796	0.942	0.949	0.954	0.937
monthly	corr with G	0.804	0.882	0.919	0.945	0.962	0.949	0.956

Overall, the results are satisfactory considering the weaknesses of the approach. In addition to the non-stationarity of gas prices discussed above, statistics such as kurtosis are highly sensitive to very rare but large spikes, implying that a very long historical time series is required to have any confidence in our estimate. However, statistics such as correlation with natural gas prices can be considered more reliable, and for these the model appears to capture the dynamics of the market quite well. In particular, as expected, correlations increase as we move from daily to weekly to monthly averages, since the impact of the shorter-term factors (demand and margin) diminishes. Also, the one-fuel model for NEPOOL leads to higher correlations with gas than PJM's two-fuel model. However, the correlation is generally slightly overestimated, perhaps due to the grouping of oil with gas and also the small amounts of other fuels present. For PJM, correlation is slightly underestimated by the simulations except in the case of fixing the gas price path to match history. As expected, this case produces the best overall results among the simulations by eliminating differences caused by the gas price dynamics, and in particular shows a large improvement in the probabilities of hourly prices above \$100, \$200 or \$300. Finally, Tables 6 and 7 underscore the difficulty in comparing statistics

²⁴These values are taken from the gas prices of 29 September 2006, 31 March 2006, and 30 December 2005 respectively, and correspond to Kalman filter estimates of 0.67, 0.93 and 1.31 respectively for the mean-reverting component X_0^1 .

Table 7: PJM statistics of observed and simulated paths

Time scale	Statistic	Observed			Simulated			
		00–07 (RT)	04–07 (RT)	04–07 (DA)	Low G_0 (Sep06)	Mid G_0 (Mar06)	High G_0 (Dec05)	Fixed G (04–07)
hourly	mean	57.25	71.02	65.38	67.46	75.97	83.54	67.71
hourly	st dev	45.40	41.77	28.06	36.76	47.18	50.74	35.79
hourly	skew	8.57	3.81	2.23	1.76	1.77	1.74	2.15
hourly	kurtosis	143.69	37.59	9.97	9.09	9.77	10.22	11.32
hourly	prob>100	8.99%	16.57%	8.61%	16.49%	22.03%	28.30%	14.54%
hourly	prob>200	0.78%	1.28%	0.27%	1.25%	2.27%	3.10%	0.91%
hourly	prob>300	0.26%	0.23%	0.06%	0.16%	0.40%	0.56%	0.10%
daily	st dev	33.11	29.39	23.40	25.79	33.42	36.13	25.31
daily	skew	5.84	2.39	2.09	0.78	0.93	0.96	1.41
daily	kurtosis	74.31	11.87	8.86	3.35	4.23	4.40	5.30
daily	corr with G	0.432	0.497	0.622	0.462	0.450	0.478	0.641
weekly	corr with G	0.585	0.668	0.707	0.556	0.548	0.578	0.741
monthly	corr with G	0.708	0.735	0.742	0.626	0.620	0.653	0.810

such as moments of price distributions, as these tend to vary significantly for different observed time periods (and for real-time vs day-ahead), as well as for different starting values of simulations.

5.2 Forward Price Analysis

In the one-fuel logistic case of the bid stack model, forward electricity prices $F^P(t, T)$ at time t with maturity T can be found in closed form under the following assumptions. We first ignore delivery periods by choosing T to be the midpoint of the delivery period $[T_1, T_2]$. In the case of monthly forwards with maturity more than a month away, this should have little impact. Secondly, we assume that the market prices of demand risk and margin risk are both zero. In other words, (8) and (9) hold under both the real world measure \mathbb{P} and the risk-neutral measure \mathbb{Q} . (Note again that by using the Kalman Filter, the dynamics of gas prices G_t in (6) are already under \mathbb{Q} .) Thirdly, as mentioned before, G_t (driven by X_t^1 and X_t^2) is assumed to be independent of both \tilde{D}_t (driven by Y_t) and \tilde{M}_t (driven by Z_t^{OU} and Z_t^{SP}). Then, from (4), we have

$$\begin{aligned}
F^P(t, T) &= E_t^{\mathbb{Q}}[S_T] \\
&= E_t^{\mathbb{Q}} \left[\alpha_0 + \alpha_1 G_T + (\beta_0 + \beta_1 G_T) \left(\log \tilde{D}_T - \log \tilde{M}_T \right) \right] \\
&= \alpha_0 + \alpha_1 F^G(t, T) + (\beta_0 + \beta_1 F^G(t, T)) \left\{ f(T) + Y_t e^{-\kappa_Y(T-t)} + \mu_Y \left(1 - e^{-2\kappa_Y(T-t)} \right) \dots \right. \\
&\quad \left. - (1 - p_i) \left(Z_t^{OU} e^{-\kappa_Z(T-t)} + \mu_Z \left(1 - e^{-2\kappa_Z(T-t)} \right) \right) - p_i \left(\alpha - \frac{1}{\lambda_i} \right) \right\},
\end{aligned}$$

for seasons $i = 1, 2, 3, 4$, where $F^G(t, T)$ is the forward gas price for the same maturity, and given by (7). We can choose either to calculate $F^G(t, T)$ based on our model or instead use observed Henry Hub gas forwards.²⁵

The top row of Figure 9 compares the model's two-year forward curve with the observed two-year peak forward curve for NEPOOL for the chosen historical dates.²⁶ As expected, the use of observed gas forward prices improves the fit of the model, lifting the model's power forward curves

²⁵Note that even when the gas forward curve is not matched exactly, we do calculate the gas model's seasonality function $h(t)$ from the observed forward curve for that date.

²⁶The end of the observed forward curve sometimes flattens out, because the forwards maturing in that particular calendar year still have a one year delivery period, instead of monthly delivery periods. For the chosen dates, this occurs for forwards maturing in Jan 2008 and beyond. For PJM, this effect occurs only at longer maturities.

slightly. This suggests either higher future risk premiums or higher future expected gas prices, versus the historical period used to calibrate the model. Overall, the results are encouraging, as both the seasonality and upwards or downwards trend in the power forward curve are fairly well captured. Note that the seasonality is a combination of gas price seasonality $h(t)$ (annual with peaks in winter) and demand seasonality $f(t)$ (semi-annual with higher summer peaks), as well as small differences in the margin spike regime. The model performs worst for forwards maturing in January or February, perhaps because the market has priced in the risk of sudden spikes in bids, as was observed in Jan 2004 and 2005, and to a lesser extent in Feb 2007. Upwards trends (eg right graph) and downwards trends (eg left graph) in the power forward curve match with trends in the natural gas forward curve, as they are explained purely by whether G_t (and in particular X_t^1) is higher or lower than the mean reversion level. Finally, the model suggests that a small demand or margin risk premium is present in the New England market, with forward power prices generally slightly above expected spot prices, on average by about \$6.20 for the case of matched gas forwards. This is consistent with the fact that the observed mean in Table 6 is only about \$68 for 2005-07, while the NEPOOL forward curves in Figure 9 generally lie above this level.

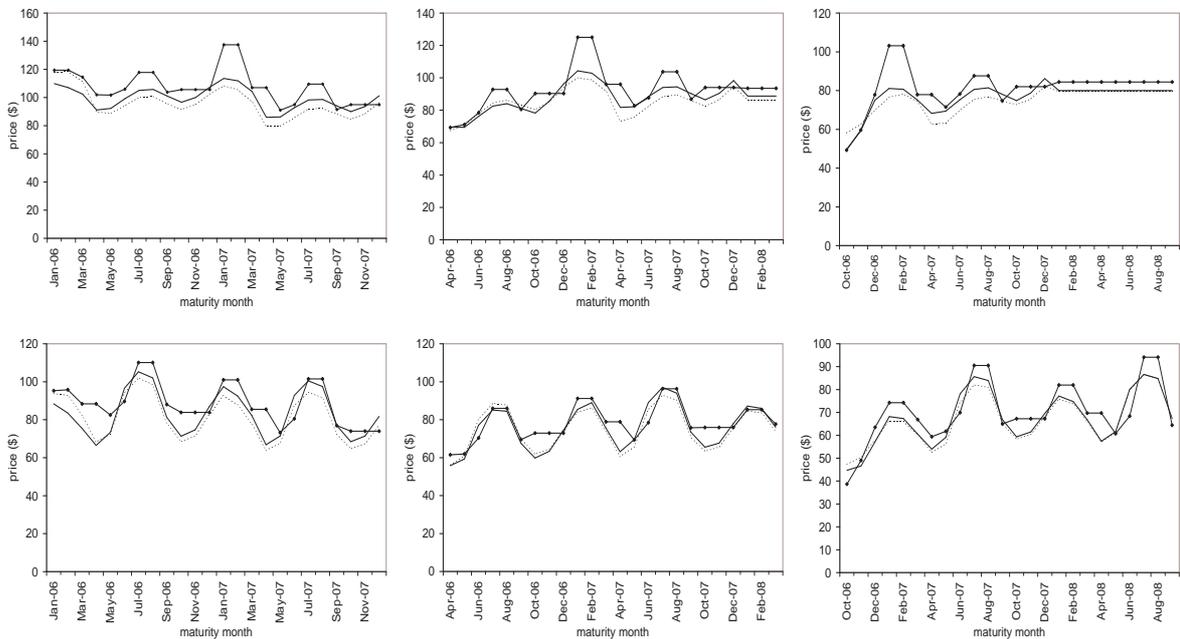


Figure 9: Sample forward curves for NEPOOL (top row) and PJM (bottom row), corresponding to the dates 30 December 2005 (left), 31 March 2006 (middle) and 29 Sept 2006 (right). The solid lines with data points indicated by circles are observed prices. The other lines are model results, with the smooth solid lines corresponding to matching the observed gas forward curve, and the dotted lines corresponding to using the model's gas forward curve.

The bottom row of Figure 9 shows the forward curves for PJM for the same dates. As no closed-form solution for $F^P(t, T)$ is available in the two-fuel case, these results have all been generated by Monte Carlo simulation with 2000 runs. The mean reversion level μ_X is allowed to be time dependent to match current gas forwards. The results are perhaps slightly stronger than for NEPOOL, with both the seasonality and trends quite well captured. While the gas forward curve remains the main determinant of the trend in the power forward curve, there is also some contribution from the direction of the coal forward curve. Furthermore, unlike for NEPOOL, the differences in margin spike frequency (particularly in the summer) play quite a significant role in determining the seasonality of power forwards. There are certain months that are less well captured than others, as for example the observed forward price in June is typically closer to May prices than July prices, while October and

November tend to equal December. Observed data reveals other patterns which the model cannot capture, such as the fact that neighbouring March and April forwards also tend to have identical prices. However, we can definitely conclude that overall the model performs remarkably well in terms of forward pricing. Finally, though less convincing than for NEPOOL, there is also some evidence of a small demand or margin risk premium in PJM, since observed forward prices are on average \$3.60 higher than simulated prices (using matched gas forwards).

6 Conclusion

The analysis of both the behaviour of simulated trajectories and the model's prices for forward contracts lends strong support to the methodology introduced in Sections 3 and 4, and its ability to capture the dominant features of electricity prices. The bid stack model can now be implemented to price a variety of different options and other derivative contracts. While we leave the details of option pricing to future work, we outline here the benefits of the approach and the variety of extensions possible using the basic underlying framework.

One of the biggest challenges in modelling electricity prices is to capture the variety of different behaviour observed at different timescales, particularly if we are interested in hourly as well as daily prices. It is well acknowledged in the literature (see eg Culot *et al* (2006) or Kluge (2006)) that multiple factors are necessary with differing speeds of mean reversion. In our model, this conclusion can be drawn directly from the parameter estimates in Section 4, with X_t^2 , X_t^1 , \tilde{D}_t , Z_t^{OU} and Z_t^{SP} all varying greatly in terms of mean reversion. Long-term behaviour (over monthly and annual horizons) is driven by gas prices (and coal prices to a lesser extent), while medium-term behaviour (over daily and weekly horizons) is largely driven by trends in demand and margin, and finally very short-term behaviour (over hourly horizons) is primarily affected by the spikes in the margin process, which we can think of as sudden outages or transmission constraints. As well as providing welcome intuition for the price dynamics, these results can also simplify the pricing of certain derivative contracts. For example, the value of long-term forwards (or options on long-term forwards) becomes independent of current values of demand and margin. We only require current fuel prices and the invariant distribution of \tilde{D} and \tilde{M} , (or more specifically of \tilde{D}/\tilde{C}) to price these claims.²⁷

In addition to describing the variety of different timescales in the market, the bid stack model captures well the striking relationships presented earlier in Figure 1, as well as Tables 8 and 9. In particular, by viewing power prices as a function of gas prices (via the bid stack), we capture a more complex dependence structure than would be possible for example using correlated Brownian Motions. While the longer-term trends of correlated processes will tend to diverge eventually, the levels of cointegrated processes will continue to follow each other. The bid stack model explains these characteristics while providing us with a more sophisticated approach than assuming a simple linear relationship between power and gas. For example, in the two-fuel PJM model, the dependence between power and gas becomes stronger during higher demand periods when the coal bids are less relevant, whereas higher demand volatility can weaken it. Even in the one-fuel NEPOOL model, where (4) leads to a linear dependence, the slope of this relationship is constantly changing due to the seasonality of D_t/C_t . Thus we end up with a relationship which mimics cointegration, but is somewhat more involved.

Another advantage of the bid stack model is that it exploits well-known relationships with underlying observable driving factors in the power markets. This allows us to take advantage of more available data for calibration, as well as to price power and gas derivatives, spark and dark spread options, demand and weather dependent contracts, and even to value power stations in a single unified framework. Finally, this framework is also flexible enough to incorporate local market information which may or may not be available. For example, knowledge about anticipated structural changes

²⁷Pirrong and Jermakyan (2005,2005) reach the similar conclusion that PJM options or forwards depend only on gas prices until near maturity, as a result of the fast mean reversion rate of demand.

(such as an anticipated increase in coal generators versus gas generators) can easily be added to the model in a manner which would be very difficult for a pure spot price model to match. Alternatively, if available, forward looking capacity or margin data could be incorporated through a time-dependent mean reversion level μ_Z of the process Z_t^{OU} . The use of such data has been shown to be particularly important by Cartea *et al* (2008).

The weaknesses of the model include the large number of parameters to be estimated, and the reliance on a large quantity of data. However, the complexity of electricity markets justifies a complicated model and the use of more data avoids a large reliance on historical spot price data. The estimation of risk-neutral dynamics for \tilde{D}_t and \tilde{M}_t has not been fully discussed here and is an extra complication of the model, though calibration of the market prices of risk to available derivative data is possible. In addition, the model simplifies the often complex structure of electricity markets (including rebalancing effects, agent behaviour, learning and market power) by treating all factors as exogenous, though possibly correlated, whereas in practice a much more involved set of interdependencies may exist. Finally, the simplifying assumptions for the shape of the bid stack may be unrealistic in some markets with many different overlapping fuel types or large amounts of imports from abroad. In this sense the model shares some characteristics with reduced-form models, yet retains strong intuition and clear links to underlying factors. The above issues in power markets imply that there is of course an inevitable compromise between realism and model tractability, and we aim to strike a balance between the two.

We conclude that the bid stack framework serves as an intuitive and realistic approach to modelling electricity prices, tractable enough to obtain some closed-form expressions for forward prices and simple numerical methods for pricing other derivatives. It is also flexible enough to adapt to different markets, and to make use of a variety of data for calibration. We feel that its ability to describe well the many features of the PJM and NEPOOL markets highlights the advantages of a more complicated fundamental model over a pure reduced-form spot price model. In particular, the recent record price levels and high volatility in global oil, natural gas and coal markets further emphasise the need to treat power and other energy price dynamics through a single unified framework. We believe that the bid stack model can be a useful tool in understanding and capturing the important interdependencies that drive energy prices.

References

- [1] C.L. Anderson, M. Davison, *A Hybrid System-Econometric Model for Electricity Spot Prices: Considering Spike Sensitivity to Forced Outage Distributions*, IEEE Transactions on Power Systems, **23(3)**, 927-937 (2008)
- [2] M.T. Barlow, *A diffusion model for electricity prices*, Mathematical Finance **12**, 287-98 (2002)
- [3] H. Bessembinder, M. Lemmon, *Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets*, Journal of Finance, **57(3)** (2002)
- [4] A. Boogert, D. Dupont, *When supply meets demand: the case of hourly spot electricity prices*, IEEE Transactions on Power Systems **23(2)**, 389-398 (2008)
- [5] M. Burger, B. Klar, A. Muller, G. Schindlmayr, *A spot market model for pricing derivatives in electricity markets*, Quantitative Finance **4**, 109-122 (2004)
- [6] M. Burger, B. Graeber, G. Schindlmayr, *Managing Energy Risk: An Integrated View on Power and Other Energy Markets* Chichester: John Wiley & Sons Ltd. (2007)
- [7] A. Cartea, M.G. Figueroa, *Pricing in Electricity Markets: A mean reverting jump diffusion model with seasonality*, Applied Mathematical Finance, **12(4)**, 313-335 (2005)

- [8] A. Cartea, M.G. Figueroa, H. Geman, *Modelling Electricity Prices with Forward Looking Capacity Constraints*, working paper (2008)
- [9] A. Cartea, P. Villaplana, *Spot Price Modelling and the Valuation of Electricity Forward Contracts: The Role of Demand and Capacity*, working paper (2007)
- [10] L. Clewlow, C. Strickland *Energy Derivatives: Pricing and Risk Management*, London: Lacima Productions (2000)
- [11] M. Culot, V. Goffin, S. Lawford, S. de Menten, Y. Smeers *An Affine Jump Diffusion Model for Electricity* working paper, (2006)
- [12] C. de Jong, R. Huisman, *Option Formulas for Mean-Reverting Power Prices with Spikes*, ERIM Report Series (2002)
- [13] S. Deng, *Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes*, working paper (1999)
- [14] D. Duffie, J. Pan, K. Singleton *Transform Analysis and Asset Pricing for affine jump-diffusions*, *Econometrica* **68(6)**, 1343-76 (2000)
- [15] G. Emery, Q. Liu, *An Analysis of the Relationship between Electricity and Natural-Gas Futures Prices*, *The Journal of Futures Markets* **22(2)**, (2002)
- [16] A. Eydeland, H. Geman, *Fundamentals of Electricity Derivatives*, Energy Modelling & the Management of Uncertainty, Risk Books (1999)
- [17] A. Eydeland, K. Wolyniec, *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging* Wiley Finance (2003)
- [18] M. Fehr, J. Hinz, *A quantitative approach to carbon price risk modeling*, working paper (2006)
- [19] H. Geman, A. Roncoroni, *Understanding the Fine Structure of Electricity Prices*, *Journal of Business* **79(6)**, (2006)
- [20] Hortacsu, Puller, *Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Market*, *The RAND Journal of Economics* **39(1)**, 86-114 (2007)
- [21] R. Huisman, *The influence of temperature on spike probability in day-ahead power prices*, *Energy Economics*, **30**, 2697-2704 (2008)
- [22] T. Kanamura, K. Ohashi, *A structural model for electricity prices with spikes: measurement of jump risk and optimal policies for hydropower plant operation*, working paper (2004)
- [23] N. Karakatsani, D. Bunn, *Intra-day and regime-switching dynamics in electricity price formation*, *Energy Economics*, **30**, 1776-1797 (2008)
- [24] T. Kluge, *Pricing Swing Options and Other Electricity Derivatives*, DPhil Thesis, University of Oxford (2006)
- [25] S. Koekebakker, F. Ollmar, *Forward Curve Dynamics in the Nordic Electricity Market*, *Managerial Finance*, **31** (2005)
- [26] J.J. Lucia, E.S. Schwartz, *Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange*, *Review of Derivatives Research* **5**, 5-50 (2002)
- [27] T.D. Mount, Y. Ning, X. Cai, *Predicting Price Spikes in Electricity Markets using a Regime-Switching Model with Time-Varying Parameters*, *Energy Economics* **28**, 62-80 (2006)
- [28] C. Pirrong, M. Jermakyan, *The Price of Power: The Valuation of Power and Weather Derivatives*, working paper (2005)

- [29] C. Pirrong, *The Valuation of Power Options in a Pirrong-Jermakyan Model*, working paper (2005)
- [30] C. Ruibal, M. Mazumdar, *Forecasting the Mean and the Variance of Electricity Prices in Deregulated Markets*, IEEE Transactions on Power Systems, **23(1)**, 25-32 (2008)
- [31] E.S. Schwartz, *The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging*, The Journal of Finance **3**, 923-973 (1997)
- [32] E. Schwartz, J. Smith, *Short-Term Variations and Long-Term Dynamics in Commodity Prices*, Management Science **46** 893-911 (2000)
- [33] P. Skantze, M. Ilic, A. Gubina, *Bid-based Stochastic Model for Electricity Prices: Impact of Fundamental Drivers on Market Dynamics*, MIT Energy Laboratory Publication (2000)
- [34] C. Supatgiat, R. Zheng, J. Birge, *Equilibrium Values in a Competitive Power Exchange Market* Computational Economics **17** 93-121 (2001)
- [35] I. Vehvilainen, T. Pyykkonen, *Stochastic factor model for electricity spot price - the case of the Nordic Market*, Energy Economics, **27**, 351-367 (2004)
- [36] P. Villaplana, *Valuation of Electricity Forward Contracts: The Role of Demand and Capacity*, working paper (2004)
- [37] R. Weron, M. Bierbrauer, S. Truck, *Modeling Electricity Prices: jump diffusion and regime switching* Physica A 336 39-48 (2004)
- [38] *Impacts of the PJM RTO Market Expansion* Energy Security Analysis, Inc., November 2005