

Modelling Price Dynamics Through Fundamental Relationships in Electricity and Other Energy Markets



Michael Coulon
Mansfield College
University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2009

Acknowledgements

I am extremely grateful to my supervisor Prof Sam Howison for his friendly and helpful guidance throughout my DPhil studies, from the initial motivation for the topic to the final write-up. By combining useful advice and direction with flexibility and willingness to explore different ideas, he has played a large part in making it a very enjoyable, stimulating and rewarding time in Oxford. I am also grateful to Ben Hambly, Christoph Reisinger and Alvaro Cartea for helping to arrange teaching opportunities during my studies, which have been both enjoyable experiences and excellent preparation for the future. I am very thankful for the financial support from the EPSRC, as well as the more general support and help from Mansfield College.

I am also appreciative of the various mathematical finance groups at Oxford which have provided many opportunities for interesting seminars and conversations about a wide range of relevant topics. I have very much enjoyed the environment of Dartington House including the opportunity for frequent interaction with faculty members and my fellow DPhil students, who have provided useful help on various occasions. I thank my officemates in particular for a very friendly and cheerful working environment with many biscuits and palm trees, and well-controlled productivity disturbances.

Through the chance to attend numerous conferences and events, I am very pleased to have been able to make frequent contact with other researchers, both from the energy industry and academia. I thank EnergyForum for inviting me to speak at a very interesting conference early in my studies which provided excellent insight. Industry perspective has often been particularly enlightening and useful for my topic, and I'm grateful for recent conversations with Daniel Crispin, Christian Jacobsen and Jeremy Lin. More general commodity related interactions with SFA Oxford have also been a valuable experience. In terms of data, PJM, NEPOOL and NYMEX have all been very helpful throughout my studies, providing crucial information for my modelling methodology.

Finally, a very special thank you goes to my parents, Denis and Carol, and brother Anthony, for all their love and support over the years. You have always given me all the opportunities in the world and encouraged me to chase any dream. Thank you so much for everything. Last but by no means least, a massive thank you to my lovely new wife and best friend Neha, for marrying me (twice in the space of 100 days) and for making me happy all the time. Your love and support in the final stretch of my thesis has been invaluable.

Abstract

Energy markets feature a wide range of unusual price behaviour along with a complicated dependence structure between electricity, natural gas, coal and carbon, as well as other variables. We approach this broad modelling challenge by firstly developing a structural framework to modelling spot electricity prices, through an analysis of the underlying supply and demand factors which drive power prices, and the relationship between them. We propose a stochastic model for 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 spot electricity price. Based on the intuition of cost-related bids from generators, the model describes mathematically how different fuel prices drive different portions of the bid stack (i.e., the merit order) and hence influence power prices at varying levels of demand. Using actual bid data, we find high correlations between the movements of bids and the corresponding fuel prices (coal and gas). We fit the model to the PJM and New England markets in the US, and assess the performance of the model, in terms of capturing key properties of simulated price trajectories, as well as comparing the model's forward prices with observed data. We then discuss various mathematical techniques (explicit solutions, approximations, simulations and other numerical techniques) for calibrating to observed fuel and electricity forward curves, as well as for pricing of various single and multi-commodity options. The model reveals that natural gas prices are historically the primary driver of power prices over long horizons in both markets, with shorter term dynamics driven also by fluctuations in demand and reserve margin. However, the framework developed in this thesis is very flexible and able to adapt to different markets or changing conditions, as well as capturing automatically the possibility of changes in the merit order of fuels. In particular, it allows us to begin to understand price movements in the recently-formed carbon emissions markets, which add a new level of complexity to energy price modelling. Thus, the bid stack model can be viewed as more than just an original and elegant new approach to spot electricity prices, but also a convenient and intuitive tool for understanding risks and pricing contracts in the global energy markets, an important, rapidly-growing and fascinating area of research.

Contents

1	Introduction to Electricity Markets	7
1.1	Supply and Demand Factors	9
1.2	Power Price Features	10
1.2.1	Spot Price Behaviour	10
1.2.2	Forward Price Behaviour	13
1.2.3	Other Price Observations	15
1.3	Review of Modelling Approaches	16
1.3.1	Pure Spot/Forward Price Models	16
1.3.2	Optimisation or Equilibrium Methods using Generation Costs	18
1.3.3	Fundamental or Hybrid Models	18
1.3.4	Comparison of Approaches	25
2	Modelling Bid Stack Dynamics	27
2.1	PJM and NEPOOL dataset	27
2.2	Price Setting Mechanism - the Bid Stack	29
2.3	Bid Stack Modelling - First Steps	31
2.4	Separate Bid Stacks for Different Fuels	34
2.4.1	Linear Bid Stacks	36
2.4.2	Exponential Bid Stacks	37
2.5	Distribution-based Model	38
2.5.1	The General Case of N Fuel Types	38
2.5.2	The One-Fuel Case: Results for NEPOOL	40
2.5.3	The Two-Fuel Case: Results for PJM	41
2.5.4	Choice of Distributions	44
2.5.5	Choice of Weights and Bid Truncation Points	46
2.5.6	Error Analysis	50
2.5.7	Implied Generation Volume	52
2.6	Discussion	55
3	From Underlying Drivers to Electricity Prices	57
3.1	Natural Gas Prices	57
3.1.1	Parameter Estimation	58
3.1.2	Calibration to Forward Curve	60
3.2	Coal Prices	62

3.2.1	Parameter Estimation	63
3.2.2	Calibration to Forward Curve	64
3.3	Demand	65
3.3.1	Parameter Estimation	65
3.4	Capacity and Margin	66
3.4.1	Parameter Estimation	69
3.4.2	Alternative Processes for Margin	72
3.5	Demand over Capacity - Invariant approach	73
3.6	Empirical Results for Power Prices - NEPOOL and PJM	76
3.6.1	Power Price Simulations	76
3.6.2	Forward Curve Analysis	78
3.7	Discussion	80
4	Pricing Power Derivatives - One Fuel Case	82
4.1	Calibration to the Power Forward Curve	83
4.2	Approximation for Longer Maturities	85
4.3	Option Pricing Techniques	86
4.4	Empirical Results	89
4.4.1	NEPOOL - Results for Forwards	90
4.4.2	NEPOOL - Results for Options	90
5	Pricing Power Derivatives - Multi-Fuel Case	97
5.1	Bid Stack Approximation - non-interchanging fuel types	98
5.2	Pricing Power Forwards	103
5.3	Error Analysis for the Bid Stack Approximation	105
5.4	Calibration to the Power Forward Curve	107
5.5	Forward Pricing Techniques and Results	114
5.5.1	PJM - Computational Techniques for Forwards	114
5.5.2	PJM - Results for Forwards	117
5.6	Pricing Power Options	118
5.7	Option Pricing Techniques and Results	122
5.7.1	PJM - Computational Techniques for Options	122
5.7.2	PJM - Results for Options	125
5.8	Adding Zero Bids (Must-run fuel types)	130
6	Carbon Emissions Allowances	133
6.1	Introduction to Cap-and-trade markets and models	134
6.2	A Simple Equilibrium Price Model	137
6.2.1	Model Assumptions	137
6.2.2	Solving for Carbon Prices	139
6.2.3	Numerical Methods - Trinomial Trees	141
6.2.4	The Impact of Fuel Switching and Exogenous Demand	144
6.3	Generalising the Simple Framework	146
6.3.1	Model Assumptions	146
6.3.2	Solving for Carbon Prices	147

6.3.3	The Impact of Fuel Switching	150
6.4	Carbon in the full Bid Stack Model	151
6.4.1	Model Assumptions	151
6.4.2	Solving for the Carbon Price	152
6.4.3	Simulation Analysis	155
6.5	Discussion	158
7	Conclusion	161
A	Taylor Series Approximation for the Two-Fuel Bid Stack	166
A.1	Various Cases of the Bid Stack Approximation	166
A.2	Error Analysis for the Bid Stack Approximation	170
B	Trinomial Tree Techniques for Coal and Gas Prices	175

Chapter 1

Introduction to Electricity Markets

In recent years, the deregulation of many electricity markets has created a sudden need for power generators and suppliers to understand and model the often unusual behaviour of electricity prices. Furthermore, while the world's attention is being drawn increasingly towards the uncertainty of climate change and future energy sources, dramatic movements have been observed throughout energy markets. From a fairly stable multi-year level of \$20-\$30 per barrel, crude oil prices steadily increased from 2004 onwards to set numerous new record highs and surpass \$140 last year, before dropping suddenly to below \$40 and then recovering to around \$70 currently. Meanwhile, natural gas prices reached remarkable record highs both in late 2005 and again in mid 2008, before falling to six year lows recently, along with many other commodity prices. Being highly dependent on these fuel prices, global electricity markets responded with huge jumps in prices and very high volatility. In addition, new markets have continued to spring up, such as the recently created carbon emissions trading market in Europe, now an integral component in forecasting power prices. The growing number of financial institutions involved in trading energy and power derivatives has further increased interest in this area and sparked a drive in both industry and academia to find suitable mathematical models. Investment banks have benefitted enormously from the new opportunities, with 2006 figures showing that up to 50% of total trading profits on Wall Street were due to energy trading! Goldman Sachs even spent \$4 billion to buy 30 power plants in 2004, and financial institutions in 2006 owned 5% of total power-generating capacity in North America (Carmona and Ludkovski (2006)). While some banks have scaled back these operations since the 2008 financial crisis, others continue to expand, both through conventional trading activities and through projects traditionally reserved for energy companies. Some invest in pipelines, storage facilities and power plants, and some even trade physical contracts for crude oil, which all suggests that the importance of energy markets for banks and hedge funds is likely to continue to grow.

To control such heightened activity in energy and electricity markets, new pricing and hedging techniques are clearly required. The US hedge fund Amaranth should have reluctantly agreed, following its announcement in September 2006 that in just a few weeks it lost \$6 billion of its \$9 billion in assets on natural gas positions. This recent collapse serves as a timely reminder that the task of choosing a reliable model for energy prices should not be underestimated! While techniques developed in the more traditional financial mathematics seem like a reasonable starting point for

modelling, we soon realise the nature of the challenge, particularly with regard to electricity. During the famous California power crisis of 2000-01, monthly average prices rose as high as \$580/MWh, only to return to stable levels below \$50 for the following years. The cost of the crisis ultimately came to \$20 billion (Carmona and Ludkovski (2006)). Although this was exceptional, milder and shorter versions of these ‘price spikes’ are apparent throughout all electricity markets, with jumps of several hundred percent a fairly regular occurrence, typically followed by returns to normal price levels only hours or days later. In contrast, intra-day patterns can sometimes be very regular for several days, while off-peak hours can often lead to remarkably low or even negative prices. As we shall see, these are just a few of the issues to face when constructing a power price model. Of course, the need to understand such volatile behaviour should be viewed as an excellent opportunity for new research in this fascinating area of finance.

Using standard approaches from financial mathematics, various stochastic processes have been suggested to describe both spot and forward prices. As we shall discuss later in this chapter, many of these are fairly simple low dimensional models which consider the spot (or forward) price directly, can often be made to match current market prices, while also providing useful analytic solutions for pricing and hedging complex derivative products. However, the unusual features of electricity prices, together with a lack of liquid derivative contracts and limited historical data, present many obstacles in the construction, calibration and testing of these models. A natural question to ask is whether a meaningful solution to these problems might be obtained by looking one level below the price series, and modelling the principal microeconomic drivers of the price. In fact, utility companies have for many years been modelling and forecasting power prices using a much more fundamental method based on local supply and demand characteristics. These models consider inputs such as generation capabilities and costs of specific power stations, transmission constraints and outage possibilities, local demand and weather patterns, and many other factors. Due to the very detailed local knowledge required, they can be computationally intensive, too inflexible to apply to more than one market, and they can also often fail to capture important overall features of price dynamics. Our goal is instead to create a modelling framework which is complex and realistic enough to intuitively incorporate prominent features of electricity market structure, while retaining the flexibility to adapt to different markets and the tractability to price derivative products easily.

Unlike a typical financial asset or even a typical commodity, electricity cannot be bought, held and then sold at a later date. As a result, the total supply and demand for electricity must be in balance at all times. Furthermore, the price of electricity on one day or one hour does not automatically affect the price in the next time period as the two cannot be substituted. Nonetheless, prices in consecutive time periods are clearly linked by the autocorrelations of underlying factors, such as weather conditions (driving demand) and fuel prices (driving supply). This suggests the use of continuous stochastic processes (plus perhaps jumps), as is standard for other financial markets such as equities. However, due to the unstorable nature of electricity, we cannot think of it as a traded asset in the usual sense, and many of the fundamental ‘no arbitrage relationships’ break down. For example, deterministic seasonal spot price behaviour cannot be risklessly exploited. (Though one could imagine stockpiling fuel or water in dams as a means of storing potential electricity, the ability in practice to quickly convert this to power is limited as the process is inefficient and costly.) Transmission constraints also lead to different prices for different locations within the same market.

Hence, apparent arbitrage opportunities through both space and time are in practice unfeasible. Furthermore, the cost-of-carry relationship between futures and spot prices no longer holds. When hedging derivatives, positions in spot electricity cannot be used. The relaxation of all of these classical conditions allows a great deal of flexibility in creating a model for electricity prices, so it is useful firstly to understand both which underlying factors drive prices and how power prices are determined in practice.

1.1 Supply and Demand Factors

The most significant underlying economic and physical factors which determine power price movements can vary greatly from market to market. Nonetheless, supply-side factors can generally be categorised as relating to either the cost or availability of electricity generation, while demand-side factors typically relate to either weather or business activity. In markets with significant renewable energy (either hydro, wind or solar), weather conditions can be an important driver of both supply and demand, creating an additional modelling challenge.

In many power markets (including the US markets we will primarily analyse in this work), fuel prices are the key factors determining the cost of power production. These can include coal, natural gas and oil prices, as well as those of more unusual fuels such as lignite. Heat rates are important supply-side information, as they reveal the amount of fuel needed per MWh of power for a certain generator. Other supply costs such as fuel transport, operational, maintenance and transmission costs should also be included. For a hydropower plant, there may also exist an opportunity cost to generating electricity at certain times, due to limited water stocks in the reservoir. Finally, carbon emissions prices are now rapidly becoming a crucial fundamental cost of power production.

As well as cost factors, the supply of power is driven by availability issues which are often less straightforward to analyse. Both scheduled outages (for maintenance) and unexpected outages continually impact the number of generators available, while specific constraints can also force individual units offline. For example, some units require minimum down-times, ramp-up rates, maximum run-times, and minimum or maximum generation levels. In addition, overall and regional constraints on transmission apply to the electricity grid to prevent congestion. Imports and exports also increase or reduce overall availability. Finally, weather-related factors such as rainfall levels, wind speed and sunshine determine the availability of renewable generators.

On the demand side, we have fewer factors to consider, and power demand consequently often reveals a stable periodic behaviour. Heating and lighting needs in the winter and air-conditioning in summer drive annual periodicities, while business activity leads to regular daily and weekly patterns of demand. Since much of the short to medium term variation around cyclical patterns stems from temperature changes, temperature data are sometimes used directly to model and forecast power demand. In the much longer term, overall economic conditions can clearly also drive power demand, as energy use in periods of recession can drop noticeably.

1.2 Power Price Features

The delivery of electricity is split between the spot and forward markets, with many utilities and other companies choosing to lock in their future power purchases in the forward market. ‘Spot’ markets typically refer to short-term power delivery in hourly or half-hourly blocks and can vary from day-ahead delivery to same day delivery to immediate delivery in rebalancing markets. Often all of these different markets exist, and each of the corresponding price series may behave somewhat differently. Nonetheless, all spot electricity price series share common features, and are easily distinguishable from forward prices, which typically correspond to an obligation to deliver power over a weekly, monthly or seasonal period in the future.¹

1.2.1 Spot Price Behaviour




Figure 1.1: In order from top left to bottom right, sample spot price series from Nordpool (Scandinavia), APX NL (Netherlands), EEX (Germany) and NEPOOL (New England, USA).

¹Thus retail power prices (for delivery to individual consumers) which move only occasionally are much more similar to longer-term forward prices in the wholesale market, and in fact the hedging needs of utilities are a natural consequence of the fact that retail prices are not adjusted very often.

Electricity spot prices have many unique features when compared with other financial time series and even other commodity markets, including periodicity (at annual, weekly and daily horizons), mean-reversion, very high volatility and sudden price spikes. Figure 1.1 illustrates these spot price features (using average daily prices) in various markets: Nordpool (Scandinavia), APX NL (Netherlands), EEX (Germany) and NEPOOL (New England, USA). Emphasising the roles of supply and demand, we can identify the primary causes of this behaviour and also how best to incorporate it into a modelling framework. For further details on the impact of supply and demand related factors, see for example textbooks by Burger *et al* (2007) and Eydeland and Wolyniec (2003).

- *Seasonality.* It is clear that strong seasonal patterns exist in most markets, though the timing and magnitude of peaks obviously depend on local conditions such as air-conditioning and heating requirements. Intra-day and intra-week patterns also exist, and clearly originate from higher levels of demand during business hours. Supply side factors can also exhibit seasonal characteristics, through for example the total available capacity of the market (as a percentage of max capacity). This may be seasonal, either if hydroelectricity is dominant and capacity depends on water levels and hence rainfall, or alternatively if maintenance schedules for power plants follow seasonal patterns. (This can often dampen seasonality in prices if maintenance is targeted at low-demand months.) Finally, seasonality may be present in fuel prices such as natural gas, and can filter through to power prices. However, the most obvious seasonality is due to load (or demand) which is an easily observable variable.

Lucia and Schwartz (2002) were perhaps the first to incorporate a deterministic or seasonal component $f(t)$ into a spot price model, and this approach is now present in virtually all models. A common choice is to use a combination of trigonometric functions with periods to match weekly, semi-annual or annual patterns. An alternative is the use of dummy or indicator variables for day of the week or month of the year, leading to a piecewise constant representation for $f(t)$. Cartea and Figueroa (2005) use a Fourier series of order 5, while wavelet decomposition can also be applied.

- *Volatility and Price Spikes.* Electricity spot markets can be very volatile and typically experience large and frequent price spikes, particularly during times of high demand. As a result, price distributions tend to have high levels of kurtosis, as well as skew, since spikes are upward jumps followed by rapid returns to previous levels. These spikes can be linked a posteriori both to demand shocks due to extreme weather, and to supply shocks caused by unexpected outages, transmission failures or spikes in fuel prices. Outages can become more significant when looking at real-time prices, as they often only last a few hours. As in the case of seasonality, it is important to identify the primary source of this behaviour in a particular market, though often it is a combination of factors. A key modelling challenge is to determine whether large price changes should be attributed to discontinuities in the process or large diffusive moves.

The most common approach to incorporating spikes is probably the use of jump-diffusion processes, in which jumps arrive following a Poisson process. One difficulty is the fact that these

‘up jumps’ are then permanent and we cannot easily force a ‘down jump’ to follow shortly, without making the model non-Markovian. We can use fast mean-reversion instead, but this does not allow a jump to last more than one time period, which would certainly be a weakness in the case of hourly prices. It also damps down the ordinary noise, for which additional factors would then be required. A regime switching model with a waiting time between switches is one solution which would allow spikes to last longer. Realistically, we would also need the intensity of the Poisson process or the probability of regime switches to be seasonal, or dependent on demand, but this adds extra complexity. Finally, volatility is often argued to be seasonal and hence a deterministic function $\sigma(t)$ is used by various authors (e.g., Cartea and Figueroa (2005)). Alternatively, as we shall advocate, using a supply and demand approach, we can induce some spikes in the price process simply as a result of the shape of the supply curve, particularly for high demand. Furthermore, this automatically means that both spikes and higher volatility occur during periods of high demand, as is observed in the data.

- *Mean Reversion.* There is widespread consensus that electricity spot prices are mean-reverting. Unlike for equities, Geometric Brownian Motion is clearly an inappropriate model, since prices return to similar (possibly seasonal) levels after shocks. It can be argued that this characteristic stems from several underlying mean-reverting factors. Firstly, temperature, and hence demand is strongly mean-reverting to its seasonal level over a timescale of a few days. Indeed, Pirrong and Jermakyan (2008) find load (demand) to be rapidly mean-reverting with a half-life of only 11.3 hours for the PJM market. Secondly, capacity availability is obviously mean-reverting as outages are followed by repairs, and these should correspond to an even faster speed of mean reversion. Finally, other energy prices are often considered mean-reverting, as in the well-known work of Schwartz (1997). This is mean reversion over a period of months rather than days, and therefore should not be (implicitly or explicitly) captured by the same parameter as the short-term mean reversion of demand or capacity.

Interestingly, Weron (2008) has used various techniques such as Hurst R/S analysis to test for mean reversion in electricity prices (after removing spikes and seasonality). The results indicate the presence of mean reversion for time intervals ranging from one day to four years, but not for intra-day movements. This would indicate that intra-day shocks persist through the day, though it should be noted that outages causing spikes were already removed from this analysis, creating a test of the diffusive component of price movements only.

- *Non-stationarity.* Though perhaps less obvious, there is evidence of non-stationarity in historical spot electricity prices in some markets. This behaviour could be seen as capturing stochastic changes in the long-term equilibrium level of prices, perhaps caused by changes in the long-term equilibrium level of fuel prices or demand. Other influences could include global economic changes, technological changes in power generation, psychological factors and the impact of speculators entering the market. Indeed there is an extra degree of uncertainty about the long term behaviour of power prices, especially since most electricity markets are very young and still developing. This factor can be particularly hard to estimate given the limited data. As will be discussed in the next section, the behaviour of long maturity forward

prices gives stronger evidence of this non-stationarity.

Incorporating non-stationarity into a model is most often achieved by adding an additional factor which follows an arithmetic Brownian Motion (possibly with drift). This approach is used in Schwartz’s two-factor model for commodities (Schwartz (1997)), as well as electricity specific models suggested by Lucia and Schwartz (2002), among others. Crucially, unlike a mean-reverting Ornstein-Uhlenbeck (OU) process, the conditional distribution of Brownian Motion does not converge to a stationary distribution, since the variance increases forever.

1.2.2 Forward Price Behaviour

In the power markets, derivatives are very important predominantly for both generators and utilities to hedge their risk, but increasingly also for speculation by other players. In particular, forward contracts are certainly the most liquid and actively traded derivatives, and the accurate pricing of forwards, together with appropriate dynamics of forward prices through time, are obvious requirements for any model. Forward curve models will automatically satisfy this requirement as they model the forward curve dynamics directly. Starting instead with a spot price model and using the standard methods of derivative pricing, a forward price is simply a conditional expectation of the spot price in the future, taken under an appropriate risk-neutral probability measure $\mathbb{Q}$:

$$F(t, T) = \mathbb{E}_t^{\mathbb{Q}}[S_T] \quad (1.1)$$

where we use the notation $\mathbb{E}_t[\cdot]$ to mean expectation conditional on the information set (filtration generated by all factors in the model) at time t . While the existence of a unique risk-neutral probability measure $\mathbb{Q}$ is sometimes simply assumed or stated upfront, it is important to understand the weaknesses of such an assumption particularly for electricity markets. Firstly, the replication argument for traditional stock price models requires the underlying asset to be tradeable, storable and liquid such that a dynamic hedging portfolio (using a stock and a bank account) can be created. However, spot electricity is non-storable, implying that S_t cannot be used in a replicating portfolio, similarly to the short-rate r_t in stochastic models for interest rates. The common alternative is to look for a liquid derivative product which can be used for hedging purposes and can therefore determine for us the risk preferences of the market. The natural candidate in power markets is the forward contract, which is actively traded for many different maturities. However, the existence of delivery periods and the lack of a continuous forward curve for all maturities suggest that an infinite number of equivalent martingale measures (risk-neutral measures) exist since the market is incomplete. It is an open question how best to price the residual risk following hedging using available forwards, and is not discussed in this thesis. Instead we focus on the more practical question of choosing a risk-neutral measure which correctly reproduces the observed forward curve in order to then price other derivatives. Thus, the measure $\mathbb{Q}$ should be interpreted as the appropriate risk-neutral measure chosen based on our method of calibrating the model to the observed forward curve via typical market price of risk assumptions (to be discussed more in Chapters 3-5). Such assumptions are quite common and necessary in a framework which combines a description of observed price dynamics under $\mathbb{P}$ with derivative pricing techniques.

Given the close link described mathematically by (1.1), the following key features of forward prices match closely with the spot price behaviour discussed above, although visually the forward price history looks quite different. Figure 1.2 shows for US markets PJM and NEPOOL that even short maturity forwards (next month delivery) are much less volatile than spot prices and contain no spikes. Note that jumps in these forward prices series correspond to changes in the delivery month of the contract plotted and are just due to seasonality effects.




Figure 1.2: Sample forward price series from PJM (North East, USA, left graph) and NEPOOL (New England, USA, right graph) for various times to maturity (1 month, 1 year and 2.5 years).

- *Seasonality.* We would expect seasonality in spot prices to carry through to seasonality in the forward curve, with higher prices for forwards maturing in high demand months. This pattern is well supported by data such as the Nordpool data in Lucia and Schwartz (2002) or the US data we use in Section 4.4 and later chapters. There is also the possibility of additional seasonality in the market price of risk, which determines the drift of our process under the risk-neutral pricing measure. This feature could be attributed to the hedging needs and risk attitudes of both generators and utility companies during different times of the year, or even day. It would lead to an additional seasonality term in forward prices, as compared to spot prices.
- *Term structure of volatility.* When considering the evolution of a forward price with a fixed maturity date, our definition above implies that the forward price must be a martingale under the measure $\mathbb{Q}$. Geometric Brownian Motion is a common model, with the attractive feature of lognormal forward prices, and is a standard assumption in forward curve models. However, it is important to note that forward prices typically become more volatile as their maturity approaches, as illustrated by Figure 1.2, where 2.5 year forwards are clearly the least volatile. This decreasing term structure of forward volatility is a consequence of Samuelson's hypothesis² and is observed throughout energy and commodity markets. It is clearly linked to the

²Samuelson (1965) was the first to formally propose that futures prices typically become more volatile as they approach maturity, an effect now known as Samuelson's hypothesis which has been widely studied and examined in energy markets.

mean reversion of spot prices or demand, since new information today has little effect on the expected spot price in the distant future.

- *Need for multiple factors.* Evidence from the Nordpool market indicates that long term forwards appear to be driven by different factors from short term forwards. Koekebakker and Ollmar (2005) use Principal Component Analysis to show that only 75% of the forward price variation can be explained by two factors, while this number is closer to 95% in other markets such as interest rates. Audet *et al* (2004) propose a forward price structure with decreasing correlation as difference between maturity increases. When considering spot price models, a single factor model will imply perfect correlation between forward prices of different maturities, so it is clear that multiple stochastic factors are necessary. Furthermore, as emphasised by Clewlow and Strickland (2000), it can be observed that the volatility of very long maturity forwards does not seem to go to zero, as would be implied by a purely mean-reverting and stationary model. This behaviour can be attributed to the same causes as the non-stationarity of spot prices mentioned above, reflecting uncertainty in long-term levels of supply and demand.

1.2.3 Other Price Observations

In addition to our observations of electricity prices themselves, these price series reveal clear correlation with other energy prices, such as oil and natural gas, though the amount of dependence of course corresponds to the local power sources for a given market. As we shall see in Chapter 2, for the PJM and New England markets, the relationship between power and natural gas prices over long horizons is particularly striking. It is also visible in Figure 1.1, where all power markets except hydro-dominated Nordpool reveal significant upwards movements in late 2005, matching precisely with the several months of higher natural gas prices. Burger *et al* (2007) and Emery and Liu (2002) discuss the apparent cointegration of power and fuel prices, as long-term levels seem to move together. Cointegration implies that a linear combination of the two processes is stationary, whereas correlated Brownian Motions cannot produce this effect. Even for unrealistically high levels of correlation, long-term simulations will show that price levels eventually drift apart. Considering the role of gas in power production, this may not be appropriate. 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.

Most modelling approaches from the traditional ‘financial mathematics’ literature choose to directly model electricity prices, as opposed to the fundamental drivers of supply and demand. It is then difficult to realistically incorporate a fuel price into the model, as would be necessary to accurately value derivatives such as spark or dark spread options (options on the power-to-fuel spread), which are actively traded in practice. Treating fuel price instead as one of the factors in our power price model allows these correlations to be naturally implied by the structure of the model, instead of artificially added afterwards. A similar argument can be made for incorporating demand as a factor, since we might be interested in valuing derivatives whose payoff is a function of demand. For example, Pirrong and Jermakyan (2008) point out that the revenue of a utility or the value of a power plant can be viewed as a function of price multiplied by load. Hence, these companies are often interested in hedging ‘volumetric risks’, and even use weather derivatives to achieve this (since

temperature and load are so closely related). A supply and demand based approach can create a modelling framework to handle both standard electricity price dependent options and claims which depend on fuel prices, load (volume), or possibly even temperature.

1.3 Review of Modelling Approaches

There exists a vast literature of approaches to electricity price modelling, with models ranging from pure spot price models similar to typical stock price models, to high-dimensional market-replicating tools which incorporate the costs, constraints and possibly strategies of all individual power generators. Between these two extremes, we find a rich spectrum of models tailored to a variety of electricity markets around the world, and serving a variety of different purposes. Of course, our choice of model should always be at least partly determined by what we are trying to analyse, simulate, price or hedge, and the need for realism should be balanced against the need for simplicity and rapid calculations.

1.3.1 Pure Spot/Forward Price Models

The desire for analytic formulas and efficient pricing techniques has led to a large literature on pure spot or forward price modelling, in which underlying supply and demand factors are ignored in favour of directly fitting historical price series (and current derivative prices). 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. The following simple model clearly leads to lognormal spot and forward prices and Black-Scholes like solutions for options:

$$\begin{aligned}
dX_t &= -\kappa X_t dt + \sigma_X dW_t \\
dY_t &= \mu dt + \sigma_Y d\tilde{W}_t \\
dW_t d\tilde{W}_t &= \rho dt \\
S_t &= \exp(f(t) + X_t + Y_t)
\end{aligned} \tag{1.2}$$

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, as illustrated for example by the work of Deng (1999), and Culot *et al* (2006). In particular, analytic formulas for forward and option prices can be derived using methods from Duffie *et al* (2000). Let X_t be a n -vector of state variables in $\mathbb{R}^n$, W_t a standard n -dimensional Brownian Motion, and Z_t an m -vector of pure jump processes, with intensity λ and jump size distribution J on $\mathbb{R}^n$. Then let

$$\begin{aligned}
dX_t &= \mu(X_t)dt + \sigma(X_t)dW_t + dZ_t \\
\mu(X_t) &= A_1 + A_2 X_t \\
\sigma(X_t)\sigma(X_t)^\top &= A_3 + A_4 X_t \\
\lambda(X_t) &= A_5 + A_6 X_t,
\end{aligned} \tag{1.3}$$

where $A_1 \in \mathbb{R}^n$, $A_2, A_3 \in \mathbb{R}^{n \times n}$, $A_4 \in \mathbb{R}^{n \times n \times n}$, $A_5 \in \mathbb{R}^m$, and $A_6 \in \mathbb{R}^{m \times n}$. Under these assumptions of affine dependence on the state variables X_t , the conditional characteristic function has the solution form, for any $u \in \mathbb{C}$,

$$\psi(u) = \mathbb{E}_t[e^{u^\top X_T}] = e^{\alpha_t + \beta_t^\top X_t} \quad \text{for } t \leq T \quad (1.4)$$

where $\alpha_t \in \mathbb{R}$ and $\beta_t \in \mathbb{R}^n$ satisfy the following Riccati ODEs:

$$\frac{d}{dt}\alpha_t = -A_1^\top \beta_t - \frac{1}{2}\beta_t^\top A_3 \beta_t - A_5^\top [\zeta(\beta_t) - 1] \quad (1.5)$$

$$\frac{d}{dt}\beta_t = -A_2^\top \beta_t - \frac{1}{2}\beta_t^\top A_4 \beta_t - A_6^\top [\zeta(\beta_t) - 1] \quad (1.6)$$

with $\alpha_T = 0$ and $\beta_T = u$, where $\zeta(c) = \int_{\mathbb{R}^n} e^{c^\top z} dv(z)$.

Deng (1999) considers three cases of these AJDs (affine jump-diffusions), all mean-reverting two-factor models but with deterministic volatility, then stochastic volatility, then additional regime-switching jumps. Exploiting the results above, derivative prices are calculated throughout, including cross-commodity spread options and locational spread options. Although Deng therefore discusses incorporating fuel prices such as natural gas prices, the correlation with fuel prices is achieved only through the matrix $\sigma(X_t)$, as opposed to the power price actually being a function of fuel prices.

Another model embedded in the AJD framework is the model of Culot *et al* (2006), applied to the Amsterdam Power Exchange. The authors propose a three-factor mean-reverting component X_t , (different reversion speeds), combined with an independent three-factor jump component $\tilde{X}_t$. With spot price $S_t = \exp(\gamma^\top X_t + \tilde{\gamma}^\top \tilde{X}_t)$, this approach allows log forward prices to be affine functions of the state variables, and hence the Kalman Filter can easily be implemented for calibration. Derivative prices are calculated using a Fourier transform technique based on the work of Culot (2003) and Carr and Madan (1999). The jump (or spike) component involves regime-switching ideas, as $\tilde{\gamma}^\top \tilde{X}_t$ can only equal zero or one of three possible spike levels, so jump sizes are fixed and a Markov chain transition matrix governs the intensities of all the possible jumps.

At least partial separation of jumps (or spikes) from diffusion factors is necessary due to the large difference in time horizons over which spike recoveries occur relative to other mean-reverting behaviour. Possible approaches include the use of multiple factors with many speeds of mean reversion, regime-switching jumps (which lead to downwards jumps to recover from spikes) or pure regime-switching models. The last of these has been studied for example by De Jong and Huisman (2003) and Weron *et al* (2004), where independent dynamics are given for the ‘spike’ and ‘non-spike’ regimes. Kholodnyi (2001) retains a closer connection between the two regimes in his model, instead suggesting that the price jumps from X_t to λX_t for some constant λ when there is a regime switch.³ Regime switching models benefit from the fact that high prices can last for several time periods, reflecting for example periods of generator outages. The recovery from an outage can be as sudden as the outage itself, a characteristic difficult to mimic with mean-reverting jump-diffusions. A variation proposed by Geman and Roncoroni (2006) is a jump-diffusion model which forces jumps to be downwards when prices are above a certain threshold.

³He refers to his model as non-Markovian but this is only under the assumption of not observing which regime we are in. If an indicator of the regime is considered a state variable (as it often is) then the model would be Markovian.

While many of the models discussed above 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 (or regimes). In addition, and perhaps more importantly from an industry perspective, they typically fail to capture the important correlations between power prices and other energy prices.

1.3.2 Optimisation or Equilibrium Methods using Generation Costs

At the other end of the modelling spectrum, there exist optimisation or equilibrium models based on detailed knowledge of every power generator’s production costs and constraints. While these are less common in academic literature, they appear to be fairly widespread in the energy industry due to the large quantities of market data used as inputs. Bessembinder and Lemmon (2002) provide a somewhat simplified version of a production cost minimisation model. More detailed versions may include generator-specific data such as minimum and maximum daily run-times, start-up costs, shut-down costs, ramp-up rates and outage probabilities as well as grid data on reserve requirements, transmission constraints and spatial variations in demand. If operational constraints exist over long horizons (e.g., maximum emissions per year), then complex optionality problems are introduced. This may also be the case for markets with pumped hydro storage facilities, as there is no fuel cost per se, but instead a value of water corresponding to the opportunity to pump and then release water at the optimal time.

Ultimately, these approaches typically require large-scale optimisation methods to forecast power prices which correspond to profit maximisation or cost minimisation while fulfilling the system load. Although they provide insight into market mechanisms and are useful for generator scheduling, they clearly require large quantities of data to implement, and hence it is likely that various assumptions will be needed to fill gaps. Furthermore, their benefit for questions of risk management or derivative pricing is quite limited, as they typically provide poor price distributions and fail to fit higher moments, as discussed by Eydeland and Wolyniec (2003). One reason for this is the mismatch between generation costs and bids or offers made in the power market auction. Attempts to explain this mismatch include agent-based models of market power and bidding strategy (e.g. Ruibal and Mazumdar (2008), Supatgiat *et al* (2001)). While strategic bidding may be limited in large competitive markets with many different bidders, Eydeland and Wolyniec also stress the importance of using generator outage probabilities to construct a bid curve from an initial cost curve and suggest a parametric approach to this transformation. They thereby propose a fundamental hybrid model which is more manageable than most large-scale optimisation models, but still relies on fairly detailed local market knowledge and large simulations.

1.3.3 Fundamental or Hybrid Models

In addition to the extremes of direct price modelling or complex bottom-up production cost approaches, there exists a spectrum of alternatives in the category of fundamental, structural, hybrid, or supply and demand based models. While terminology is used differently by different authors and

various categorisations are possible, we refer here to models which analyse and capture the relationship between electricity prices and underlying drivers, but stop short of a full description of market intricacies in order to improve tractability and emphasise dominant relationships. These approaches range from high-frequency analyses of the most significant spot price drivers (e.g. Karakatsani and Bunn (2008)), to studies of relationships with physical variables such as temperature, rainfall patterns and other demand-side factors (e.g. Huisman (2008), Vehvilainen and Pyykkonen (2004)), as well as work focused on specifying the correct shape of the electricity supply function (e.g. Kanamura and Ohashi (2007)). These models all attempt to provide a compromise or bridge the gap between pure spot power price models and the complex equilibrium models often used in industry.




Figure 1.3: Traditional supply and demand curves from economic theory (left graph); Inverting, capping supply and setting demand to be inelastic produces the most appropriate framework for modelling electricity prices (right graph).

The two key ingredients in a fundamental model are the choice of underlying factors (and how to model them) and some sort of transformation, which can be referred to as the inverse supply curve. Before introducing a progression of fundamental approaches recently proposed for the electricity market, we should discuss briefly how the use of a supply curve here relates to the traditional economic modeling of supply and demand for any good. Recall that in a basic economic model, the price associated with a certain level of supply is assumed to be an increasing (and often linear) function of quantity, while the price associated with a certain level of demand is a decreasing (and often linear) function of quantity. The equilibrium price and quantity pair is of course set at the unique point where supply intersects demand. Outside factors which drive supply and demand are typically assumed to cause parallel shifts in the each of the lines, thus producing a new equilibrium price and quantity, as illustrated in Figure 1.3. In electricity markets, the first significant feature is the inelasticity of demand to price, since most consumers are very inflexible and willing to pay high prices to ensure having the necessary power. We shall justify this further later, but for now it is reasonable to treat the demand curve as a horizontal line (as a function of price). On the other hand, the supply curve (quantity against price) stems from the ordering of costs and availabilities of different generators, which itself is often referred to as the merit order, with fuel types typically

falling into a familiar order.⁴ In addition, the upper bound on the supply curve in power markets is clearly defined by the total installed capacity of power generators and is important for modelling purposes, while for other markets this may be assumed unbounded or left unspecified. Finally, we note that with the goal of modelling electricity prices (as opposed to modelling supply and demand as outputs), it is more suitable here to work with the inverse supply curve, which treats price as a function of quantity. As shown in Figure 1.3, power demand is then a vertical line moving via horizontal shifts, and the inverse supply curve is precisely the transformation from demand to price, with other factors potentially incorporated along the way.

Barlow's Model

The natural starting point is the very simple model of Barlow (2002), where the demand is the only factor and the transformation is described by a simple one-parameter function. Barlow assumes an OU process for demand D_t , (without seasonality as the Alberta data he uses show little seasonality):

$$dD_t = \kappa(\mu - D_t)dt + \sigma dW_t$$

He then assumes that the inverse supply curve is constant over time, but allows flexibility in its shape. Instead of simple linear and exponential shapes, he generalises to the following family of functions:

$$f_\alpha(x) = (1 + \alpha x)^{1/\alpha}, \quad \alpha \neq 0, \quad f_0(x) = e^x$$

Thus

$$S_t = B_t(D_t) = f_\alpha(D_t) = (1 + \alpha D_t)^{1/\alpha}$$

By varying the choice of α , we can clearly vary the steepness of the curve. In particular, any $\alpha < 0$ corresponds to a function steeper than the exponential, while the special cases of $\alpha = 0$ and $\alpha = 1$ are linear and exponential respectively. In order to choose a value of α for the Alberta and California data, Barlow uses a maximum likelihood approach based on spot price data for a fixed α , then maximises numerically over all α . As a result of the very spiky data used, the optimal α is negative for almost all time periods, implying the need for a steep inverse supply curve. Another key result is that when using these optimal α 's, the mean-reversion rate κ is less than in the $\alpha = 0$ or $\alpha = 1$ case. Therefore, Barlow's model has the advantage of being able to capture extreme spikes with a pure diffusion process, and without excessively large parameters κ or σ . This is demonstrated by his simulated price process when $\alpha = -1.08$. Considering the price spikes visible in Figure 1.1, it is not surprising to observe a significantly negative α in Barlow's model for most electricity markets, though the graphs also suggest significant variation in 'spikiness' from market to market.

It should be noted that Barlow's argument for this inverse supply curve shape is based on a finite total supply and a supply curve $g(x) = a_0 - b_0 x^\alpha$, which means a_0 is the limit on supply (in the $\alpha < 0$ case). Therefore, Barlow also sets a maximum price of A_0 for any demand above $a_0 - \epsilon_0 b_0$ (for some small ϵ_0). This condition limits the availability of closed-form solutions for forward prices

⁴In traditional economics, the cheapest supply is assumed to be used first simply because of economic arguments that this equilibrium will eventually be reached, while in electricity markets this results from the actions of the market administrator, collecting bids and ordering them from every day from lowest to highest. For this reason the supply curve (or more precisely, the inverse supply curve) will come to be replaced by the 'bid stack' in our work, which can be understood as the highest bid needed to satisfy a given demand. We shall discuss in detail in Chapter 2.

and other derivative prices. Furthermore, while allowing flexibility in the shape of $B(x)$, Barlow allows no movement in the curve over time, which is certainly a big approximation. While it is clear from most power data that demand is a key driver of spot prices, it is also clear that they are not perfectly correlated. In any case, it seems clear that a strength of Barlow's method is the basic idea of a flexible relationship between S_t and an underlying OU process, through which local information about the supply curve can be incorporated.

Extensions to Barlow

In contrast, Skantze *et al* (2000b) assume an exponential inverse supply curve, but allow parallel horizontal shifts determined by a supply variable often referred to as capacity available. Very similar approaches are put forward by Eydeland and Geman (1999), and Cartea and Villaplana (2008), although Cartea and Villaplana focus more on the model's implications of risk premiums in the forward curve. In each of these cases, S_t has the form

$$S_t = e^{aD_t + bC_t} \quad (1.7)$$

where demand (or load) is D_t and 'supply' (or capacity) is C_t . In the work of Skantze, both load and 'supply' follow the two-factor OU plus ABM model in (1.2) (but with zero correlation), while outages are also incorporated into C_t using an independent Bernoulli process. The supply factor C_t is not observed but instead calibrated as a residual and the model is in discrete time and focuses on simulation techniques as opposed to analytic formulas for derivatives.⁵ Villaplana proposes a similar model in continuous time, and chooses to interpret C_t as 'generation capacity', which he then estimates from hydro reservoir levels in the Nordpool market in Villaplana (2004). He models C_t as a jump-diffusion process like in (1.4) and demand as a simple OU process with seasonality and seasonal volatility. Cartea and Villaplana adopt a similar approach, using observed capacity data and correlated OU processes for D_t and C_t . These approaches all raise the question of how changes in available capacity affect the inverse supply curve, as a result of new power plants, changed reservoir levels or outages. A horizontal parallel shift as above implicitly assumes that all new capacity (or removed capacity) will be positioned to the far left of the inverse supply curve (or at least below the price setting point), and the justification is unclear. Intuitively, it seems more appropriate for the effect to be multiplicative, implying that the change in capacity is spread evenly through the curve. Equivalently, we can think of the inverse supply curve function $f(x)$ as acting on D_t/C_t instead of D_t , where D_t/C_t is the ratio of demand to total capacity. This is the approach which we shall favour in our model, and shall discuss further for PJM and NEPOOL in Chapter 2.

In Burger *et al* (2004), the authors propose a three-factor model which incorporates the ideas above. Capacity is a deterministic function $c(t)$, reflecting seasonal patterns in power plant maintenance in Germany, which is typically carried out in the summer. Demand (or load) D_t is stochastic but includes a deterministic component $d(t)$ as usual. Finally, a 'price-load curve' $f(t, D_t/c(t))$ is introduced to represent the inverse supply curve, but no parametric form is assumed. The complete

⁵In Skantze and Ilic (2000a), the relationship in (1.7) is maintained, but instead relate the shifts C_t to the existence of multiple markets in California for generators to bid into.

model is written in discrete time and is given by

$$S_t = \exp \left\{ f \left(t, \frac{D_t}{c(t)} \right) + X_t + Y_t \right\}, \quad t = 0, 1, \dots,$$

where both X_t and Y_t are attributed primarily to “psychological aspects of the behaviour of speculators and other influences”. However, X_t is a short-term process affecting spot prices and short futures, while Y_t is a long term process (Arithmetic Brownian Motion plus drift but in discrete time) affecting long-term futures. The processes D_t and X_t are calibrated from historic load and price data using SARIMA time series models, while f is fitted using a cubic spline technique. Note also that the shape of f is allowed some time dependence, since a different curve is used for peak and off-peak hours. Finally, it is argued that the dynamics of long maturity futures are essentially only driven by Y_t , since the distributions of the other factors (X_t and D_t) far into the future can be approximated by their stationary distributions. This approximation has several advantages. Firstly it allows a simple method for exactly matching the forward curve using the drift of Y_t . Secondly, it allows easy calibration of the volatility of Y_t to futures volatility either implied from option prices or from historical data. Finally, it simplifies option pricing by reducing the forward curve dynamics (for long maturities) to simple Geometric Brownian Motion, so that Black’s method can be applied.

Instead of using D_t or C_t directly as state variables, a slightly different perspective is gained by considering reserve margin (extra capacity available beyond demand) to be a key driver, which is similar to using D_t/C_t . In particular, margin tells us more about spike occurrence, 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), where either mean price levels or transition probabilities between regimes are allowed to depend on the margin. Boogert and Dupont (2008) analyse the relationship between margin and spot price as well as margin and spike probability, and suggest a non-parametric approach. Finally, Cartea *et al* (2008) use forward-looking margin information (or similarly, demand over capacity) as an indicator of when a spike is likely to occur. Forward gas prices are also incorporated, but only as the deterministic (or seasonal) component of future power prices, so gas is not included as a risk factor. Note that by using a regime-switching framework, these last few models do not necessarily require the notion of a supply curve, since underlying factors can instead be used to determine which of two or more spot price processes is most likely to apply at a given time.

Pirrong-Jermakyan Model

The authors above focus mainly on load (demand) and capacity fluctuations, known to be the main short-term drivers of power prices due to the required hourly matching of supply and demand. However, they choose not to directly model fuel prices as underlying factors, despite the clear importance of production costs in determining electricity prices, particularly for medium or long-term dynamics. The difficulty arises in how to incorporate fuel price movements into supply curve movements, particularly in markets with multiple power sources and complicated merit orders. Hydropower, renewables and nuclear all require slightly different considerations as well, since the quantity of power generated from these sources is driven not by fuel price movements but instead by resource

availability (in the case of hydro and renewables) and the need to avoid shut-down costs (for nuclear).

A useful model for heavily gas-based markets is introduced by Pirrong and Jermakyan (2008) (and Pirrong (2006)), with application to the PJM market in the US, a market which we will discuss in detail soon. They assume that power prices are driven by two factors, both observable: fuel prices (natural gas specifically) and demand. Demand D_t is assumed to be driven by an exponential OU process with seasonality, while gas prices G_t follow Geometric Brownian Motion. Under the risk-neutral measure, we have

$$\left. \begin{aligned} dD_t &= \kappa(\mu(t) - \ln D_t - \lambda(D_t, t)\sigma_D) D_t dt + \sigma_D D_t dW_t \\ dG_t &= rG_t dt + \sigma_G G_t d\tilde{W}_t \\ dW_t d\tilde{W}_t &= \rho dt \end{aligned} \right\} \quad (1.8)$$

where $\lambda(D_t, t)$ is the market price of log-demand risk, and is assumed to be dependent on demand and time. Note that Pirrong and Jermakyan in fact use gas forwards as the second factor, with $f_t^{[T]} = \sigma_G f_t^{[T]} d\tilde{W}_t$. These two approaches are equivalent for derivative pricing, because the maturity T of the gas forward is chosen to match the maturity of the power derivative that we are pricing, so that at maturity $f_T^{[T]} = G_T$ anyway. Hence our payoff can be written in terms of either variable, but the evolution of the inverse supply curve should intuitively be a function of spot prices G_t , instead of forward prices with a particular maturity.

Instead of explicitly modelling spot electricity prices, they focus their attention on pricing forwards and options via a PDE approach. The usual application of Ito's Lemma and no arbitrage arguments yields the following pricing PDE for any traded power derivative $V(t, G_t, D_t)$:

$$\frac{\partial V}{\partial t} + \kappa(\mu(t) - \ln D - \lambda(D, t)\sigma_D) D \frac{\partial V}{\partial D} + rG \frac{\partial V}{\partial G} + \frac{1}{2}\sigma_D^2 D^2 \frac{\partial^2 V}{\partial D^2} + \frac{1}{2}\sigma_G^2 G^2 \frac{\partial^2 V}{\partial G^2} + \sigma_D \sigma_G \rho DG \frac{\partial^2 V}{\partial D \partial G} = rV$$

Forward prices F (using simplified notation for $F(t, T)$), not traded assets themselves, satisfy the same equation but with right-hand side equal to zero:

$$\frac{\partial F}{\partial t} + \kappa(\mu(t) - \ln D - \lambda(D, t)\sigma_D) D \frac{\partial F}{\partial D} + rG \frac{\partial F}{\partial G} + \frac{1}{2}\sigma_D^2 D^2 \frac{\partial^2 F}{\partial D^2} + \frac{1}{2}\sigma_G^2 G^2 \frac{\partial^2 F}{\partial G^2} + \sigma_D \sigma_G \rho DG \frac{\partial^2 F}{\partial D \partial G} = 0 \quad (1.9)$$

As the terminal condition of an electricity forward is $F(T, T) = S_T$, the spot price, we must make an assumption about the relationship between S_T and the two underlying factors in order to solve the PDE.⁶ This relationship is simply the shape of the inverse supply curve and its behaviour as gas prices move. Pirrong and Jermakyan make the simplifying assumption that the inverse supply curve is multiplicative in fuel price, meaning that $S_t = G_t \phi(\ln D_t)$, or more generally $S_t = G_t^\gamma \phi(\ln D_t)$, for $\gamma \geq 0$. They then suggest several methods for determining the function $\phi(x)$ (and possibly γ too). One option is to use specific data on marginal costs of power production to construct the 'generation stack', but this assumes perfect competition in the sense that generators simply bid at cost. As discussed earlier, both profit taking and strategic bidding (the exercise of market power) could make the inverse supply curve different from the generation stack. A second option is an econometric approach based on regression techniques and requires assuming a parametric form of some sort for ϕ . Instead,

⁶In practice, we will have a delivery period $[T_1, T_2]$ instead of simply T , but we can price this as a sum or integral of simple forwards $F[T]$.

they favour a third option: the use of historical bid data, which are available at www.pjm.com, but with a six month lag. Assuming that the function ϕ is stable over time, a historical bid stack can be chosen as the inverse supply curve (from an appropriate season of the year to allow for possible seasonality) and then divided by the historical gas price for the same date. Note that in the case of $\gamma = 1$, ϕ can be thought of as a ‘heat rate’ curve, since it describes the ratio between electricity and fuel prices as a function of demand. The simplifying bid stack assumptions allow (1.9) to be reduced to a two-dimensional PDE, which is then solved by finite differences. This can only be done when the payoff function is linear in G_T (or G_T^γ), as for forwards and spark spread options. Spark spread options with strike heat rate H have payoff $\max(S_T - HG_T, 0) = G_T \max(\phi(\ln D_T) - H, 0)$. Slightly more complicated numerical methods in Pirrong (2006) are used to price other options when $\rho = 0$, and Monte Carlo simulation is suggested for $\rho \neq 0$. It should also be noted that the market price of risk $\lambda(t, D_t)$ is calibrated using market data on forwards, and inverse problem techniques. The authors minimise the sum of squared differences between market forward prices and model-implied prices, with a regulator function used to penalise over-fitting, and to ensure λ varies smoothly with demand.

As well as introducing several interesting ideas in their approach, their two papers Pirrong (2006) and Pirrong and Jermakyan (2008) include noteworthy results as well. Firstly, the market price of risk $\lambda(D_t)$ is found to be very significant in PJM, particularly in the earlier years of the study. It is also found to be particularly large (and negative) during periods of high D_t . Another key result of their study is the fast mean-reversion of demand, which implies that the conditional distribution of demand just a few days ahead is essentially equal to the stationary distribution. Consequently, the model implies that demand shocks should not affect forward prices beyond a few days, or options on long maturity forwards. ‘Daily strike options’ are typically options on day-ahead prices, so these options should not react to demand until very near maturity. Before this time, they behave as options on the fuel price alone. Further implications include the result that ‘monthly strike options’ (with delivery period of one month starting soon after option maturity) should never be affected by demand shocks, even just before maturity, because they are options on a series of forward contracts with an average maturity of about half a month. Finally, Pirrong mentions the possibility of short-run factors such as outages, asserting that these should only affect the valuation of very short maturity forwards or options on the real-time spot price.

The Pirrong-Jermakyan model is important as it introduces the idea of using fuel price as a state variable, while retaining tractability and not resorting to a full simulation based model as in the fundamental hybrid models of Eydeland and Wolyniec (2003). However, there are several assumptions which could be modified or extended in various ways. Firstly, the assumption of GBM for natural gas is definitely questionable, as illustrated for example in the Pyndyck (1999) study mean reversion in various commodities (including 75 years of natural gas prices), which tends to support Schwartz’s two-factor model. More importantly, the relationship between electricity spot prices and gas prices is obviously more complicated than assumed, particularly in a market with a variety of power sources like PJM. In particular, as we shall investigate in Chapter 2, gas prices typically drive only the upper half of the PJM bid stack. Nonetheless, despite the overly simplistic approach of reusing last years’ bids, the use of observed historical bid stack data is another important contribution of Pirrong and Jermakyan. In addition to automatically capturing bidding strategies

of generators which could otherwise be problematic to model, it has the great benefit of providing us with an entire curve of data for each point in time. This is invaluable information for understanding the transformation between underlying factors and power price, since we effectively observe not only the power price S_t set by the marginal fuel at time t given demand D_t , but we also observe what the price would have been for any other value of demand.

1.3.4 Comparison of Approaches

The wide range of underlying factors together with the numerous features of price dynamics have naturally led to a remarkably broad class of electricity price models. While some are more suitable for short-term forecasting, others cater more to long-term risk management or derivative pricing, and some specialise in understanding multi-commodity risks. Furthermore, unlike most financial markets, certain regions or particular markets are more amenable to certain modelling approaches, owing both to market structure and data availability. In fact, the approaches discussed in this section can easily be categorised in terms of the amount of data they require, ranging from a single spot price history to detailed knowledge of all individual generators. Clearly our choice of model should be tailored to our market features and data availability as well as modelling goals. Though no model can suit all purposes, the electricity markets nevertheless provide a number of reasons to favour a fundamental or hybrid approach.

As discussed in Sections 1.1 and 1.2, it is well acknowledged in the literature that electricity prices primarily reflect supply and demand for power on a particular hour or day. This relationship is stronger than in other markets due to fact that electricity is non-storable, meaning that inventories cannot be used to buffer against supply and demand shocks. Thus, the unusual features of price dynamics can be more readily traced back to changes in underlying factors which in turn are often easily observable. In light of the need for multiple factors, particularly to capture clear differences between short, medium and long-term dynamics, the use of known observable price drivers (instead of unintuitive unobservable factors) is especially appealing. Moreover, given the short price histories and changing structure of some markets (since deregulation), a fundamental framework can greatly expand the pool of relevant historical data available.

Incorporating demand, fuel prices and possibly even temperature also allows us to value a large variety of financial products or other assets in a single framework. These include spark and dark spread options (both clean and dirty, following the introduction of carbon emissions markets), utility company revenues (both demand and price dependent payoffs), new power plants and possibly weather derivatives. Considering direct impact on profitability, energy companies often have more interest in understanding the risk of power-fuel spread movements than of power or fuel price movements themselves. Another advantage of a supply and demand framework is the ability to capture specific local information about the future in a simple manner. For example, the knowledge of new generation capacity planned for a certain date, or known maintenance schedules can be easily accounted for in a model. As electricity prices are very sensitive to regional and local conditions, this information can be crucial.

A hybrid or fundamental approach to power price modelling allows an alternative to convenient

but overly-simplistic models or more realistic but complicated models. Our aim is to reach a suitable compromise which produces a flexible and intuitive framework while retaining sufficient tractability. As we have discussed, a key challenge in this task is the construction of an appropriate bid stack function (or inverse supply curve) which ideally should capture most of the unusual features of electricity price distributions while keeping fairly simple distributions for the underlying factors. To this end, we shall propose a realistic, parametric approach to the bid stack function, allowing for the overlap of bids from generators of different fuel types. Using observed bid data as a starting point, we investigate the rather complex dependence structure of power and fuel prices, as is crucial for the pricing and hedging of risks throughout all energy markets.

Chapter 2

Modelling Bid Stack Dynamics

2.1 PJM and NEPOOL dataset

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 by PJM and NEPOOL: see www.pjm.com and www.iso-ne.com for data. Unless otherwise stated, all data used in this chapter including figures and tables are available on these sites. In particular, all electricity bid data are available at <http://www.pjm.com/markets-and-operations/energy/real-time/historical-bid-data.aspx> and at http://www.iso-ne.com/markets/hstdata/mkt_offer_bid/index.html. While free public access to bid data in a convenient form is still uncommon in electricity markets, the overall availability of power auction data is reasonably good and industry participants often make use of such data for bidding strategy or internal modelling purposes.

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 bid-related parameter estimation, and we use PJM West price data throughout.¹ The PJM market has an interesting mix of fuel types (power sources), with a significant proportion of capacity coming from each of coal, gas or oil, and nuclear. Table 2.1 illustrates how this fuel type breakdown has changed slightly from year to year over the period considered.

The New England region 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

¹Power 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 contracts traded on NYMEX.

Table 2.1: Fuel breakdown (percent) in the PJM market (2002-06), and NEPOOL (2006). PJM data have been taken from annual generating capacity reports while NEPOOL data from Eydeland and Wolyniec (2003). Note that gas generators that also have oil-based generation capability are listed separately, so the proportion of gas generators is really the sum of these two rows.

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

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 2.5, 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 2.1: 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 since 2002 only the first 3 days of August 2006, with a maximum of \$327.) Data from www.pjm.com and www.iso-ne.com.

Figure 2.1 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 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 2.1. 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 2.2: 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.) Data from www.pjm.com and www.iso-ne.com.

2.2 Price Setting Mechanism - the Bid Stack

In Section 2.5, 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.² 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 \$80.³ 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 2.2. The market administrator then determines the hourly spot price by matching with the total demand for power (as well as including

²Note that only daily observations of the bid stack are available for PJM, whereas hourly stacks are available for NEPOOL, though intra-day variation is extremely 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.

³In PJM, generators have the alternative of connecting these bid points linearly instead of using a step function.

other constraints into an optimisation routine). 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.⁴ 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 2.2 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 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.⁵ In order to understand this relationship, it is useful to consider the bid stack as a histogram of bids, as shown in Figure 2.3. We simply add up the total amount of capacity in MW that has been bid within each price bin.⁶ This provides an alternative perspective on the same data shown in Figure 2.2, 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 2.1). 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. Our approach to modelling the entire bid stack focuses on the movement of these clusters of bids as fuel prices change. 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.

⁴In 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.

⁵In PJM, there is an independent monitoring organisation which specifically investigates the possibility of excessive strategic bidding and reports units whose bidding appears irregular or too far out of line with costs or with bids from previous days. (I am grateful to a private conversation with Jeremy Lin of PJM for this information.) Therefore, market power can be assumed to be kept to a minimum.

⁶In 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 NEPOOL, 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.




Figure 2.3: 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. Data from www.pjm.com and www.iso-ne.com.

2.3 Bid Stack Modelling - First Steps

Let S_t represent the spot price at time t . 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. Thus let $D_t \in [0, 1]$ be the demand at time t , assumed to be inelastic with respect to price, as is often the case for electricity. For generator i of n , let x_i be quantity supplied (again, normalised by total capacity) and let $b_t^{(i)}(x_i)$ be the bid curve at time t , where $b_t^{(i)} : [0, c_i^{\max}/c^{\max}] \rightarrow [0, p^{\max}]$. 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

$$S_t = \max_{1 \leq i \leq n} \left\{ b_t^{(i)}(x_i^*) \right\}, \quad \text{where} \quad \{x_1^* \dots x_n^*\} = \underset{x_1, \dots, x_n}{\operatorname{argmin}} \left\{ \max_{1 \leq i \leq n} \left(b_t^{(i)}(x_i) \right) : \sum_{i=1}^n x_i = D_t \right\}. \quad (2.1)$$

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 $B_t^{\text{obs}}(\cdot)$ denote the exact bid stack observed in the market at time t , we can write

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

As explained in Chapter 2, 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. Despite these complications and the fact that many of these factors are not easily observable, we can broadly think of C_t simply as the percentage of maximum capacity available.

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 the 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 by $c^{\max}$), 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), as mentioned in Section 1.3.3. We will discuss further the behaviour of the capacity process C_t in Chapter 3, but for now it suffices to interpret it intuitively as the proportion of power generating capacity online at time t and to imagine that it is primarily driven by outages. A fundamental requirement of this framework is that demand and capacity are both positive and 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 extra capacity from abroad). The spot price is now given by⁷

$$S_t = B_t \left(\frac{D_t}{C_t} \right), \quad \text{for } 0 < \frac{D_t}{C_t} < 1, \quad (2.2)$$

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. We then search for an continuous non-decreasing function $B_t(\cdot)$ that best approximates the step function $B_t^{\text{obs}}(\cdot)$. 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. A more complicated approach would be to remove sections of the bid stack piece by piece as individual generator outages occur, but this would severely limit the tractability of the model.




Figure 2.4: Historical progression of daily installed capacity, total capacity from bids, daily peak average available capacity and daily peak average demand for PJM (left) and NEPOOL (right). Data from www.pjm.com and www.iso-ne.com.

It is important to remember that the methodology presented here is appropriate for the PJM and NEPOOL markets and the bid data as they are presented for these markets. In particular, historical bid data are observed prior to the incorporation of scheduled or unscheduled outages (i.e., the removal of these bids), and all units are forced to submit bids every day.⁸ Therefore, we expect

⁷We 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.

⁸Again, I am grateful to Jeremy Lin of PJM (private communication) for confirming this at first somewhat sur-

that the total quantity bid should closely follow the total installed capacity of the market, remaining quite stable and jumping only significantly at times of market expansion. Figure 2.4 illustrates that this is accurate for both PJM and NEPOOL, with total capacity from bids remaining well above available capacity C_t and then demand D_t , both of which reveal significant intra-year (and even intra-week) fluctuation.⁹ Other markets do not necessarily follow this same pattern of bids adding up to approximately match installed capacity and therefore require some modifications when modelling. In particular, quantities of bids in European markets more closely follow available capacity, as we might expect intuitively.¹⁰ Furthermore, as bidding is not mandatory, generators with high costs may for example opt out of bidding if they expect demand to be very low, suggesting that capacity changes will not be evenly spread through the stack. In addition, the German market reveals a significant amount of price elasticity on the demand side, further complicating the modelling of bids. Nonetheless, for the US markets studied, we assume that unavailable capacity is equally spread through the stack (so that $B_t(\cdot)$ is a function of D_t/C_t) and that the demand bid curve is vertical. We investigate the validity of these assumptions in detail in Section 2.5.7.




Figure 2.5: Series of histograms investigating the frequency and magnitude of changes to individual generators' bids, both for the highest quantity bid (left graph) and the highest price bid (right graph) among the daily price-quantity pairs.

Additional evidence of the stability of total capacity from bids can be obtained by analysing changes in the bids of individual generators over time. Although generator names are not observed, code names typically remain unchanged at least through a calendar year. We choose the year 2004 in PJM as a sample and investigate the 829 generators present for at least one month of the year. The left-hand graph of Figure 2.5 shows three histograms from the data, all assessing the amount of

prising information as well as providing further insight into the PJM market.

⁹Note that for NEPOOL, available capacity data can be found online from 2002 onwards (and is hence labeled 'obs' for observed), while for PJM they are only available from mid-2006. Hence, for PJM the available capacity line corresponds to implied (not observed) capacity available, as we shall describe in Section 3.4. Daily data for total installed capacity are also available from earlier for NEPOOL. However, in annual PJM reports (those used to create Table 2.1), we do observe annual values of total installed capacity throughout the time period which match well with total capacity from bids.

¹⁰I am grateful to Daniel Crispin and Christian Jacobsen (private conversations), who trade energy in the German market, for providing information and insight into the European power markets and how they differ from the American markets.

movement in the highest quantity bid (total capacity offered) by individual generators. We can see that over 80% of generators changed their total capacity bid in less than 5% of days (and 45% never changed). The mean of this distribution is in fact 7.2% of days, though the weighted average (using total capacity bid as weights) is 3.1% of days. Similarly, the other histograms show that generators tend to repeat many times a particular highest quantity bid (the mode), and that when they do change, this tends to be only a small drop relative to their maximum for the year. As a means of comparison, the right-hand graph of Figure 2.5 shows the same histograms but for highest price bid instead of highest quantity. As expected, there is much more variation in prices than quantities, with only about a third of generators changing their highest bid less than 5% of days, a substantial number changing roughly every other day and only a handful (5%) changing almost every day. Also, the sizes of the changes is now bigger, as the histogram for minimum over maximum shows a broad range of bidding behaviour. Interestingly, the mean of the distribution of frequency of changes is now 27.8% of days, just over one in four. While this is not as often as we might expect, it is sufficient to lead to substantial overall movement in the total bid stack on a daily basis.¹¹

2.4 Separate Bid Stacks for Different Fuels

The bid stack function $B_t(\cdot)$ can be constructed or estimated in a number of different ways and using different datasets. The models discussed in Chapter 1 often use these methods typically without referring to the ‘bid stack’, but simply to the transformation from underlying factors (especially demand) to spot electricity prices. Firstly we can choose between a parametric function (e.g., Cartea and Villaplana (2007)) and a non-parametric approach (e.g., Burger *et al* (2004)). Secondly, we can choose to use either cost and heat rate data, price and demand data only, or actual bid data itself. The first of these corresponds to creating a ‘generation stack’ of generation costs, which should then be converted to a bid stack using some assumptions, as discussed in Eydeland and Wolyniec (2003). The effect of outages as well as possible strategic bidding should then be carefully considered. The alternative approach of using only historical price and demand data (and perhaps capacity and/or fuel prices) to imply the transformation has the significant disadvantage that we only observe one set of values $(S_t, D_t, \dots)$ for each time period t , telling us where in the stack the price was set, but giving no information about the rest of the curve at time t . In contrast, using observed bid data (when available) provides us with a large quantity of data for each t , as well as avoiding the tricky issue of creating a bid stack from a generation stack. Thus, when we introduce our model for this function $B_t(\cdot)$ in Section 2.5, we estimate directly from available bid data for PJM and NEPOOL.

Unlike any of the articles discussed above, our primary requirement for an appropriate choice of bid stack function is to capture the key observation that different fuel prices drive different sections of the stack, due to the merit order. While this was first illustrated in Figure 2.2, further evidence is presented in Figure 2.6, which shows the evolution of the different points in the PJM bid stack over time. (So $\frac{D_t}{C_t} = 30\%, 40\%, 50\%$ in the first graph and $\frac{D_t}{C_t} = 60\%, 70\%, 80\%$ in the second.) Here we

¹¹One can speculate that an individual generator’s decision on how frequently to submit changes in bids is related to its cost changes, but also to its expectations of prices movements and finally to its hedged positions. For companies which have already entered into forward contracts to buy (or sell) power, there may be less need to frequently adjust bids to ensure profitability.

observe a clear relationship between the lower region of the stack and the coal price, and similarly between the upper region of the stack and the natural gas price.




Figure 2.6: Dynamics of different points on the PJM bid stack over time (top row); Coal and gas price history (bottom row)

Although it seems clear that on the whole natural gas bids remained higher than coal bids during the historical period considered, our model must also reflect the possibility of a merit order change for very high coal price P_t and low gas price G_t . Typically, with coal below gas, S_t should be most sensitive to G_t during peak demand periods, and most sensitive to P_t during off-peak demand periods. The challenge is therefore to find an appropriate function $B(\cdot)$ to capture these features, while ensuring that it remains strictly increasing and defined over the range $(0, 1)$. We perhaps face a common choice between oversimplifying the bid stack function and retaining closed form solutions for forwards (and possibly options), or retaining more realism and local market information, but resorting to numerical methods. (The models of both Burger *et al* (2004) and Pirrong and Jermakyan (2008) require numerical methods in conjunction with a non-parametric bid stack function.)

A natural starting point is to imagine the existence of two different bid stacks, one which includes only coal bids, and the other only natural gas bids. Interestingly, while an overall linear bid stack seems entirely inappropriate, two linear ‘sub bid stacks’ can lead to more realistic dynamics, as we shall now investigate.

2.4.1 Linear Bid Stacks

Let coal and gas generators have weights w_1 and $w_2 = 1 - w_1$ respectively. Let $x = D_t/C_t$. Suppose the two bid stacks $B_1(D_t/C_t) = f_1(x)$ and $B_2(D_t/C_t) = f_2(x)$ for coal and natural gas respectively are given by

$$\begin{aligned} f_1(x) &= b_1 x, & 0 \leq x \leq w_1 \\ f_2(x) &= a_2 + b_2 x, & 0 \leq x \leq w_2. \end{aligned}$$

We require $a_2, b_1, b_2 > 0$ and typically would expect $b_2 > b_1$, as natural gas has a steeper bid stack. Now, in practice, two bid stacks are combined by ordering all bids by price, alternating between natural gas and coal when necessary. Mathematically, this corresponds to inverting each bid stack, adding them, and then inverting the sum, as is illustrated in the progression from the first to the second graph in Figure 2.7. (The inverse bid stacks $h_1(y)$ and $h_2(y)$ can be called supply curves, as it is in common in economics to draw supply and demand as functions of price.) Here we have

$$\begin{aligned} h_1(y) = f_1^{-1}(y) &= \begin{cases} y/b_1 & \text{if } 0 \leq y \leq b_1 w_1 \\ w_1 & \text{if } y \geq b_1 w_1 \end{cases} \\ h_2(y) = f_2^{-1}(y) &= \begin{cases} 0 & \text{if } 0 \leq y \leq a_2 \\ \frac{y-a_2}{b_2} & \text{if } a_2 \leq y \leq a_2 + b_2 w_2 \\ w_2 & \text{if } y \geq a_2 + b_2 w_2 \end{cases} \end{aligned}$$

The overall bid stack is then a piecewise linear function given below. Note that we assume here that $a_2 \leq b_1 w_1$, implying that the lowest natural gas bid is no higher than the highest coal bid. This ensures continuity of the bid stack.

$$B(x) = (h_1 + h_2)^{-1}(x) = \begin{cases} b_1 x = B_1(x) & \text{if } 0 \leq x \leq \frac{a_2}{b_1} \\ \frac{b_1 a_2}{b_1 + b_2} + \frac{b_1 b_2 x}{b_1 + b_2} & \text{if } \frac{a_2}{b_1} \leq x \leq \frac{w_1(b_1 + b_2) - a_2}{b_2} \\ a_2 - b_2 w_1 + b_2 x = B_2(x - w_1) & \text{if } \frac{w_1(b_1 + b_2) - a_2}{b_2} \leq x \leq 1 \end{cases} \quad (2.3)$$

This function consists of three linear sections of different slopes. The first, with slope b_1 , (blue in Figure 2.7) consists entirely of coal bids. The second, with slope $\frac{b_1 b_2}{b_1 + b_2}$ (less than b_1), (purple in Figure 2.7) consists of both coal and gas bids. The third, with slope b_2 , (red in Figure 2.7) consists only of natural gas bids and is the steepest section of the curve. The next step is to decide how changes in coal price P_t and gas price G_t affect $f_2(x)$, and consequently $B(x)$. The simplest choice is to let the slope of each bid stack be dependent on fuel price, such that $b_1 = \alpha P_t$ and $b_2 = \beta G_t$. Note that for low coal prices ($P_t < a_2/\alpha w_1$), we may have a discontinuous stack.¹²

Equation (2.3) and Figure 2.7 illustrate what happens when G_t (and hence b_2) moves. An upward move increases the slope of both the second and third line segments, as well as shifting their point of intersection to the left. In other words, prices are higher for any level of demand beyond a_2/b_1 and the influence of coal in the bid stack now runs out at a lower level of capacity. Figure 2.2 partially supports this approach as there is a clear steepening of slope following the gas price increase in February 2003, and only beyond 40,000MW in the stack. However, crucially, the small

¹²Also choosing $a_2 = \alpha G_t$ would be intuitive but would increase the chance of potential discontinuities in $B(x)$, since high gas prices could also lead to $a_2 > b_1 w_1$.




Figure 2.7: Construction of a market bid stack from two linear ‘sub bid stacks’

section in Figure 2.2 which should correspond to the overlap of fuels (immediately after 40,000MW) is not the flattest section of the stack as Figure 2.7 would imply, but instead one of the steepest.

2.4.2 Exponential Bid Stacks

A more realistic alternative is for the two ‘sub bid stacks’ to be exponential instead of linear. Hence suppose the aggregation of all coal bids produces the stack

$$B_1(x) = P_t e^{a_1 + b_1 x}, \text{ for } 0 \leq x \leq w_1.$$

Similarly,

$$B_2(x) = G_t e^{a_2 + b_2 x}, \text{ for } 0 \leq x \leq w_2.$$

Then the total market bid stack is again given by

$$\begin{aligned} B(x) &= (h_1 + h_2)^{-1}(x), \quad \text{for } 0 \leq x \leq 1 \\ \text{where } h_1(x) &= (B_1)^{-1}(x) \\ \text{and } h_2(x) &= (B_2)^{-1}(x) \end{aligned}$$

Typically $P_t e^{a_1} < G_t e^{a_2} < P_t e^{a_1 + b_1 w_1} < G_t e^{a_2 + b_2 w_2}$, implying again that coal bids are generally lower than gas bids but there is some overlap. In this case we have three regions in the overall market bid stack, again corresponding to only coal bids, overlapping bids, and only gas bids:

$$B(x) = \begin{cases} B_1(x) & \text{for } 0 \leq x \leq x_1 \text{ (= first impact from gas)} \\ P_t^\alpha G_t^\beta e^{\gamma + \delta x} & \text{for } x_1 \leq x \leq x_2 \text{ (= last impact from coal)} \\ B_2(x - w_1) & \text{for } x_2 \leq x \leq 1 \end{cases}$$

where

$$x_1 = \frac{\log\left(\frac{G_t}{P_t}\right) + a_2 - a_1}{b_1}, \quad x_2 = w_1 + \frac{\log\left(\frac{P_t}{G_t}\right) + a_1 + b_1 w_1 - a_2}{b_2}$$

The constants α , β , γ , and δ can be found to be as follows:

$$\alpha = \frac{b_2}{b_1 + b_2}, \quad \beta = \frac{b_1}{b_1 + b_2}, \quad \gamma = \frac{b_1 b_2}{b_1 + b_2}, \quad \delta = \frac{a_1 b_2 + a_2 b_1}{b_1 + b_2}$$

While this model is still overly simplistic to accurately fit observed bid data, it provides a useful tool to understand the movement of the overall bid stack when coal and gas prices change. As in the linear case, the lowest section of the stack is given simply by the coal stack function $B_1(x)$, and the highest section by the shifted gas stack function $B_2(x - w_1)$, with a more complicated formula for the region of overlap. The overall stack is piecewise exponential, just as we had piecewise linear previously. At first glance the expression above appears promising for calculating forward prices as an expectation of $S_T = B(P_T, G_T, D_T/C_T)$, but we must note firstly that indicator functions of the form $1_{x < x_1}$ would be functions of all factors. Secondly, in order to take an expectation over an infinite range of coal and gas prices (as necessary for typical fuel price processes), we would need to relax our assumption that $P_t e^{a_1} < G_t e^{a_2} < P_t e^{a_1 + b_1 w_1} < G_t e^{a_2 + b_2 w_2}$ and instead incorporate min and max functions to cover all cases, including discontinuous bid stacks.

Therefore, these simple examples illustrate the significant challenge in obtaining convenient formulas using a bid stack approach for two fuels, particularly when aiming to price derivatives in closed form. Our next suggestion will ultimately provide an encouraging alternative and will also form the core of our power price model.

Parts of this chapter have been published in "Stochastic Behaviour of the Electricity Bid Stack: From Fundamental Drivers to Power Prices", Journal of Energy Markets, Vol. 2, no. 1 (spring 2009). Available at <http://www.journalofenergymarkets.com/>

2.5 Distribution-based Model

As illustrated in Figure 2.3, a histogram of bids provides a useful alternative to simply observing the bid stack directly in Figure 2.2, 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.

2.5.1 The General Case of N Fuel Types

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 $B(\cdot)$ is simply the inverse of the cdf, $F(x)$, of the mixture distribution:

$$S_t = B\left(\frac{D_t}{C_t}\right) = F^{-1}\left(\frac{D_t}{C_t}\right).$$

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 $P[b_L < D_t/C_t < b_U] = 1$ and $P[D_t > b_U] = 1$. These requirements are typically met for historical peak data with choices of $b_L = 0.2$ or 0.3 and $b_U = 0.9$ or 0.95 to eliminate low nuclear bids and very high bids.¹³ 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 $D_t/C_t \in (b_L, b_U)$ adds complications 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 non-truncated 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

$$F(S_t) = \sum_{i=1}^N w_i F_i(S_t) = \frac{\tilde{D}_t}{\tilde{C}_t}, \quad \text{where} \quad \tilde{D}_t = \frac{1}{b_U - b_L} (D_t - b_L), \quad \frac{\tilde{D}_t}{\tilde{C}_t} = \frac{1}{b_U - b_L} \left(\frac{D_t}{C_t} - b_L \right), \quad (2.4)$$

and similarly to above (now with tildes added),

$$S_t = B \left(\frac{\tilde{D}_t}{\tilde{C}_t} \right) = F^{-1} \left(\frac{\tilde{D}_t}{\tilde{C}_t} \right). \quad (2.5)$$

We fit distributions to the bid data independently each day 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 the 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 $\sum_{i=1}^N w_i = 1$), 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

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

For example, for a mix of N logistic distributions, the log-likelihood function is given by

$$L(w_1, \dots, w_N, m_1, \dots, m_N, s_1, \dots, s_N) = \sum_{j=1}^M q_j \log \left\{ \sum_{i=1}^N \frac{w_i}{4s_i} \operatorname{sech}^2 \left(\frac{p_j - m_i}{2s_i} \right) \right\}. \quad (2.6)$$

In Section 2.5.4, we compare results using the following distributions: Gaussian, logistic, Cauchy, and Weibull. These distributions 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

¹³In Section 2.5.5 we discuss how best to choose values for these truncation points, while in Section 3.4 we discuss a remedy to the only one of our conditions (i.e., $D_t/C_t < b_U$) which is very occasionally broken by observed historical data for PJM and NEPOOL.

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 outperform 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.

2.5.2 The One-Fuel Case: Results for 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 one-fuel model is reasonable. We therefore estimate the bid stack parameters for each historical date as follows. Firstly, we ignore bids of below \$10 (which are roughly 30% throughout and consist almost entirely of the zero bids¹⁴) 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 $[0.3, 0.85]$.¹⁵ Then, for a chosen probability distribution, we maximise with respect to m_1 and s_1 the appropriate likelihood function for $N = 1$. For the logistic, we use (2.6). 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 2.8 illustrates the maximum likelihood estimation (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 January 2005, the correlation with gas prices is very high, and as high as 95% in the more recent data.¹⁶ The two spikes could perhaps have been caused by strategic bidding, though no particular information has been found to explain these events.¹⁷ Nonetheless, the results strongly suggest that we assume a linear dependence structure between m_1 and s_1 and the natural gas price G_t :

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

With all four distributions considered, when $N = 1$, using (2.5) the spot price $S_t = B_t(\tilde{D}_t/\tilde{C}_t)$ can be written explicitly as a function of the factors $G_t, \tilde{D}_t$ and $\tilde{C}_t$:

- For one Gaussian, $S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t)\sqrt{2} \operatorname{erf}^{-1}\left(2\frac{\tilde{D}_t}{\tilde{C}_t} - 1\right)$.
- For one logistic, $S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t)\left(\log\left(\tilde{D}_t/\tilde{C}_t\right) - \log\left(1 - \tilde{D}_t/\tilde{C}_t\right)\right)$.
- For one Cauchy, $S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t)\tan\left(\pi\left(\frac{\tilde{D}_t}{\tilde{C}_t} - \frac{1}{2}\right)\right)$.
- For one Weibull, $S_t = -(\alpha_0^G + \alpha_1^G G_t)\left(\log(1 - \tilde{D}_t/\tilde{C}_t)\right)^{(\beta_0^G + \beta_1^G G_t)^{-1}}$.

¹⁴Ignoring only zero bids but not those just above zero causes a problem primarily for the Weibull distribution which has extremely low probabilities in this range, since it has support $[0, \infty)$, and density 0 at 0.

¹⁵Prior to this, we first remove all bids between \$994 and \$1000, as there are sometimes clusters of irrelevant bids at these levels which complicate matters if they are considered to be part of total capacity.

¹⁶Note 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.

¹⁷The 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.




Figure 2.8: NEPOOL estimation results (left graph) for the logistic distribution's 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 $\{\alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G\}$ 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 2.2 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. (We discuss this point further in Section 2.5.5). 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, the one-fuel logistic model for NEPOOL leads to a convenient equation for the bid stack, and hence the spot electricity price, which can be written as follows:

$$S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t) \left(\log(\tilde{D}_t) - \log(\tilde{C}_t - \tilde{D}_t) \right). \quad (2.7)$$

Thus we obtain a spot electricity price S_t which is linear in the natural gas price G_t (as for the Gaussian and Cauchy), but also in log of demand and log of reserve margin, the extra capacity not needed to satisfy demand. The fairly simple form of (2.7) is very appealing, particularly for the pricing of forwards presented in Section 3.6.2 and further developed in Chapters 4-5. Note that the linearity in G_t mirrors the linearity in the model for PJM prices developed by Pirrong and Jermakyan (2008). However, this only occurs in the one-fuel case in our model, so for NEPOOL, but not PJM.

2.5.3 The Two-Fuel Case: Results for PJM

As discussed briefly in Section 2.1 and illustrated in Figure 2.6, 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




Figure 2.9: 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).

move in the same manner as the neighbouring coal bids.¹⁸ Next, calculate fixed weights w_1 and $w_2 = 1 - w_1$ based on the split between coal and gas or oil in Table 2.1. These weights change only at a few discrete points in time to reflect market changes.¹⁹ Finally (for the logistic distribution), maximise the likelihood function (2.6) for $N = 2$ with respect to $\{m_1, s_1, m_2, s_2\}$.




Figure 2.10: 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).

Figure 2.9 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 2.9 over the corresponding

¹⁸Tests 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.

¹⁹Tests using fully flexible weights in the optimisation procedure suffer from greater instability particular at times of significant overlap between the two clusters of bids. An alternative approach to the maximum likelihood estimation procedure is the use of the expectation-maximisation (EM) algorithm common for mixture models. However, even with this method, we find significant instability in the weights and favour instead the use of installed capacity data. We discuss the validity of these fixed weights further in Section 2.5.5.

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 2.10, 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.

Table 2.2: 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

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

$$m_1 = \alpha_0^C + \alpha_1^C P_t, \quad s_1 = \beta_0^C + \beta_1^C P_t, \quad m_2 = \alpha_0^G + \alpha_1^G G_t, \quad s_2 = \beta_0^G + \beta_1^G G_t,$$

where P_t and G_t are stochastic processes for coal and gas prices respectively. However, when $N \geq 2$, none of the distributions in general lead to explicitly invertible bid stack functions. Instead, for $N = 2$, we find $S_t = x$ such that $B_t^{-1}(x) = \tilde{D}_t / \tilde{C}_t$, where

- for two Gaussians, $B_t^{-1}(x) = 1 + w_1 \operatorname{erf} \left(\frac{x - (\alpha_0^C + \alpha_1^C P_t)}{(\beta_0^C + \beta_1^C P_t)\sqrt{2}} \right) + (1 - w_1) \operatorname{erf} \left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{(\beta_0^G + \beta_1^G G_t)\sqrt{2}} \right)$,
- for two logistics, $B_t^{-1}(x) = \frac{1}{2} + \frac{w_1}{2} \tanh \left(\frac{x - (\alpha_0^C + \alpha_1^C P_t)}{2(\beta_0^C + \beta_1^C P_t)} \right) + \frac{1 - w_1}{2} \tanh \left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{2(\beta_0^G + \beta_1^G G_t)} \right)$,
- for two Cauchys, $B_t^{-1}(x) = \frac{1}{2} + \frac{w_1}{\pi} \arctan \left(\frac{x - (\alpha_0^C + \alpha_1^C P_t)}{\beta_0^C + \beta_1^C P_t} \right) + \frac{1 - w_1}{\pi} \arctan \left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{\beta_0^G + \beta_1^G G_t} \right)$,
- for two Weibulls, $B_t^{-1}(x) = 1 - w_1 \exp \left\{ - \left(\frac{x}{\alpha_0^C + \alpha_1^C P_t} \right)^{\beta_0^C + \beta_1^C P_t} \right\} - (1 - w_1) \exp \left\{ - \left(\frac{x}{\alpha_0^G + \alpha_1^G G_t} \right)^{\beta_0^G + \beta_1^G G_t} \right\}$.

Regression can again be used to estimate the parameters $\{\alpha_0^C, \alpha_1^C, \beta_0^C, \beta_1^C, \alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G\}$. Results are shown in the middle and lower sections of Table 2.2, and are very encouraging, 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 2.1. 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 (2.5) can be written for PJM in the logistic case as follows:

$$S_t = x \quad \text{such that} \quad B_t^{-1}(x) = \frac{\tilde{D}_t}{\tilde{C}_t}, \quad (2.8)$$

where

$$B_t^{-1}(x) = \frac{1}{2} + \frac{w_1}{2} \tanh\left(\frac{x - (\alpha_0^C + \alpha_1^C P_t)}{2(\beta_0^C + \beta_1^C P_t)}\right) + \frac{1 - w_1}{2} \tanh\left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{2(\beta_0^G + \beta_1^G G_t)}\right)$$

or

$$B_t^{-1}(x) = \frac{w_1}{1 + \exp\left\{-\left(\frac{x - (\alpha_0^C + \alpha_1^C P_t)}{\beta_0^C + \beta_1^C P_t}\right)\right\}} + \frac{1 - w_1}{1 + \exp\left\{-\left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{\beta_0^G + \beta_1^G G_t}\right)\right\}}.$$

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

2.5.4 Choice of Distributions




Figure 2.11: Comparison of fitting bids with four different distributions for sample dates 1st Feb 2006 (for PJM, left graph) and 1st Oct 2004 (for NEPOOL, right graph)

Before moving on to the application of the bid stack model for electricity price modelling and derivative pricing in the following chapters, we make a slight diversion to further explore the details of bid stack construction. Although less important from a general mathematical modelling perspective, these issues are certainly relevant in practice when seeking the best fit to the available bid data. In particular, in this and the next section we focus on the choice of bid distribution, weights and truncation points, while in the following two sections we evaluate the bid stack model in terms of direct error analysis as well as its ability to match observed power generation patterns.

In this thesis, we primarily advocate the use of the logistic distribution to best capture the clusters of bids which form the observed bid stack. Here we compare the results of all four distributions considered, namely the Gaussian, logistic, Cauchy and Weibull. Table 2.3 lists the parameters of each distribution, as well as their log-likelihood functions in the general case of choosing a mix of N distributions. In addition to each having two parameters and appropriate humped shapes, our four chosen distributions also share the useful property of having a fairly simple inverse cumulative distribution function, and hence closed-form solutions for S_t in the one-fuel case, as given in Section 2.5.2. The simplest and most convenient relationship from a mathematical perspective remains the logistic one, as given in (2.7). However, it is interesting to observe that in the one-fuel case, for all but the Weibull distribution, we obtain a spot electricity price which is linear in the natural gas

Table 2.3: Log-likelihood functions for the four distributions considered.

Distribution	Parameters	Log-likelihood function $L(w_1, \dots, w_N, m_1, \dots, m_N, s_1, \dots, s_N)$
N Gaussians	means m_i standard deviations s_i	$L = \log \left[\prod_{j=1}^M \left\{ \sum_{i=1}^N \frac{w_i}{s_i} \phi \left(\frac{p_j - m_i}{s_i} \right) \right\}^{q_i} \right]$
N Logistics	means m_i scale parameters s_i	$L = \sum_{j=1}^M q_j \log \left\{ \sum_{i=1}^N \frac{w_i}{4s_i} \operatorname{sech}^2 \left(\frac{p_j - m_i}{2s_i} \right) \right\}$
N Cauchys	location parameters m_i scale parameters s_i	$L = \sum_{j=1}^M q_j \log \left\{ \sum_{i=1}^N \frac{w_i s_i}{\pi [s_i^2 + (p_j - m_i)^2]} \right\}$
N Weibulls	scale parameters m_i shape parameters s_i	$L = \sum_{j=1}^M q_j \log \left\{ \sum_{i=1}^N \frac{w_i s_i}{m_i} \left(\frac{p_j}{m_i} \right)^{s_i - 1} e^{-\left(\frac{p_j}{m_i} \right)^{s_i}} \right\}$

price. On the other hand, in the two-fuel model for PJM, power price is not linear in either fuel price, and none of the distributions in general lead to explicitly invertible bid stack functions.

As all four distributions considered are fairly similar in overall shape, their performance in capturing the shape and behaviour of the bid stack is also similar. Figure 2.11 confirms the similarity for each market on some sample date, while also revealing how closely the bid histogram can be represented by a continuous density function. For different distributions, the plots of $\hat{m}_1$ and $\hat{s}_1$ (and $\hat{m}_2, \hat{s}_2$ for PJM) share very similar overall patterns to those in Figures 2.8-2.10, although the plots for $\hat{s}_i$ can vary more than those for $\hat{m}_i$, particularly as b_U increases.²⁰ For example, a heavier tailed distribution such as the Cauchy distribution outperforms others when the higher bids are included, while quite erratic results begin to appear for the thin-tailed Gaussian when $b_U = 0.95$, but it performs well if less of the tail is included. The estimates $\hat{s}_1$ for NEPOOL and $\hat{s}_2$ for PJM are of course most affected by including too much of the tail, although other parameters are somewhat affected too. Further discussion of the cutoff point for bids b_U is given in Section 2.5.5. Overall, we can compare distributions using several methods. Firstly the log-likelihood can be compared directly for each date in order to rank the four distributions. Alternatively, higher values of R^2 in the regressions against fuel prices can be taken as an indication of a better fit. Table 2.4 compares the distributions for each market using both of these methods for different values of b_U by showing: (i) the percentage of dates for which each distribution ranks first by likelihood; (ii) the percentage of dates for which each distribution ranks second by likelihood; and (iii) the value of R^2 observed for $\hat{m}_1$ for the date range March 05 - August 07 for New England, or $\hat{m}_2$ for June 03 - July 07 for PJM.

Interestingly, Table 2.4 reveals that for both markets the Weibull distribution usually ranks highest by MLE but lowest or second lowest by R^2 . This is probably because it has the advantage of a fairly thick right tail as well as being the only non-symmetric distribution, making it more flexible in shape. On the other hand, regression results are weaker since m_1 corresponds less closely to the mean or location of bids than for the other distributions. Furthermore, the shape parameter s_1 in the Weibull distribution has very little correlation with gas prices, instead remaining fairly stable over time (e.g. $\hat{s}_1 \approx 3$ for $b_U = 0.9$ for NEPOOL, and $\hat{s}_1 \approx 4, \hat{s}_2 \approx 2$ for PJM). However,

²⁰We note again that the rescaling of demand and capacity to $\tilde{D}_t$ and $\tilde{C}_t$ must correspond to the cutoff points of the bid data, b_L and b_U .

Table 2.4: Comparison of the performance of the four distributions, for different values of b_U (NEPOOL and PJM).

	Distribution	$b_U = 0.85$			$b_U = 0.9$			$b_U = 0.95$		
		1 st	2 nd	R^2	1 st	2 nd	R^2	1 st	2 nd	R^2
NE	Gaussian	5.7%	80.3%	0.897	2.8%	51.1%	0.892	0.1%	3.7%	0.819
	Logistic	18.2%	7.3%	0.912	19.9%	25.0%	0.908	5.3%	42.7%	0.889
	Cauchy	0.2%	1.5%	0.914	9.9%	3.0%	0.912	42.4%	12.0%	0.912
	Weibull	75.9%	10.9%	0.887	67.4%	21.0%	0.884	52.2%	41.6%	0.821
PJM	Gaussian	43.2%	53.8%	0.915	31.2%	64.4%	0.909	0.7%	36.0%	0.883
	Logistic	3.0%	15.1%	0.916	3.4%	13.1%	0.915	5.0%	51.9%	0.899
	Cauchy	0.0%	0.0%	0.861	0.0%	0.0%	0.881	0.1%	7.4%	0.899
	Weibull	53.9%	31.1%	0.912	65.4%	22.4%	0.905	94.2%	4.7%	0.889

this raises the appealing possibility of letting $\beta_1^G = 0$ and assuming all gas price dependence is incorporated through the location parameter m_1 or m_2 . In contrast, the other three distributions are all symmetric, and hence distinguishable primarily based on the thickness of their tails. The very thick-tailed Cauchy distribution generally performs well for NEPOOL in terms of R^2 but not by likelihood, suggesting that the thicker tails prevent unusual bids from disturbing the path of $\hat{m}_1$ over time, at the expense of the overall shape of the distribution. In contrast for PJM it is outperformed throughout. The Gaussian distribution suffers from thin tails, and hence only compares favourably with the others for low values of b_U . Finally, the logistic distribution falls in the middle in terms of tail thickness, and it therefore performs consistently well throughout. For NEPOOL it ranks second best overall under both criteria, while for PJM it consistently gives the highest values of R^2 . In light of these observations and its mathematical simplicity, we favour the logistic distribution in practice.

2.5.5 Choice of Weights and Bid Truncation Points

We now address several more details related to the modelling of the bid stack, in particular the choice of range $[b_L, b_U]$, and weight w_1 for PJM. Firstly, the choice of b_U , the upper truncation point for bids, can clearly have a significant impact on the parameter estimates for the gas distribution as witnessed in the previous section. Thus, even ignoring the issue of price spikes, the model's spot price distribution for mid to high levels of D_t/C_t can be distorted by including too many of these very high bids. Although this distortion may be evident in lower values of R^2 , it is not necessarily clear what value of b_U to choose to avoid this. Figure 2.12 shows how values of R^2 for the distribution nearest the tail vary as b_U changes, plotted for all four distributions (L=Logistic, G=Gaussian, C=Cauchy, W=Weibull). While it is generally safer to eliminate more of the tail to avoid large inaccuracies, we would ideally prefer to fit as much of the stack as possible. There appear to be points on the graphs beyond which the values of R^2 begin to drop rapidly for the Gaussian and logistic distributions, particularly for $\hat{s}_1$ (NEPOOL) and $\hat{s}_2$ (PJM). (For the thicker-tailed Cauchy and Weibull distributions, we sometimes observe higher values of R^2 when more of the tail is included.) For the logistic distribution, the choice of $b_U = 0.9$ for NEPOOL and $b_U = 0.95$ for PJM appears reasonable, as the R^2 plots drop off earlier for NEPOOL than for PJM.




Figure 2.12: Effect of varying cutoff b_U on the values of R^2 for regressions of bid distribution parameters versus gas prices. Results for PJM (for $\hat{m}_2$ and $\hat{s}_2$) and NEPOOL (for $\hat{m}_1$ and $\hat{s}_1$) are plotted on the left and right respectively.

An extra complication is that the appropriate cutoff b_U clearly varies with time, as is illustrated by the first graph in Figure 2.13. Here we use the logistic distribution with $b_U = 0.9$ for NEPOOL and $b_U = 0.95$ for PJM, and plot against time the number of standard deviations K between the mean of the gas distribution and the highest bid used for estimation. So for the logistic distribution,

$$K = \frac{\sqrt{3}}{\pi} \left(\frac{b_U - \hat{m}_1}{\hat{s}_1} \right) \quad \text{or} \quad K = \frac{\sqrt{3}}{\pi} \left(\frac{b_U - \hat{m}_2}{\hat{s}_2} \right)$$

We find that at various points in time, particularly in the earliest years and some recent periods, the number of standard deviations reaches as high as 4 or 5. This indicates that our estimates $\hat{m}_1$ (or $\hat{m}_2$ for PJM) and especially $\hat{s}_1$ (or $\hat{s}_2$ for PJM) are probably unrealistically high for these dates. This also suggests that a time-dependent b_U would improve results if chosen correctly.

The lower cutoff for the bids, b_L , is less of an issue for several reasons. Firstly, peak hour prices are much more likely to be set near the right tail than the left tail of the bid stack. Secondly, small changes in b_L have small effects because the lowest bids (bounded below by \$0) are far less spread out than the highest bids (bounded above by \$1000). Finally, we have a better understanding of the source of these bids - nuclear and must-run hydro generators - and thus a good idea of an appropriate weight in each market. Bid data for NEPOOL reveal that the percentage of zero bids is fairly stable, fluctuating between 25% and 35% over the entire dataset. In contrast, PJM has fewer zero bids and the histogram in Figure 2.3 shows that it is difficult to identify a clear transition point from nuclear to coal bids. We therefore assume $b_L = 0.2$ based on the fuel type data in Table 2.1.

Table 2.1 also reveals that in reality New England is not simply a one-fuel market, while PJM is clearly more than a two-fuel market. For example, we could extend the NEPOOL model to incorporate a small percentage of coal with bids lower than the gas bids. Alternatively, we could split oil and gas dependence, introducing a third driving fuel price. However, ultimately we choose to focus on retaining a tractable model which captures not all factors, but the most dominant factors driving power prices.




Figure 2.13: Number of standard deviations K between the mean of the gas distribution and the highest bid used for estimation (left graph); Coal distribution weight w_1 for PJM, as estimated by SSE using (2.10) or from market data in Table 1 using (2.9) (right graph).

For the case of PJM, a key parameter is the weight w_1 corresponding to the proportion of coal generators in the region $[b_L, b_U]$ of the bid stack. As demand over capacity, D_t/C_t , often falls near w_1 (i.e. in the region containing a mix of coal and gas bids), the distribution for S_t (and hence the model's forward prices) is quite sensitive to this parameter. Unfortunately, the maximum likelihood estimates for m_1, s_1, m_2, s_2 are quite insensitive to the choice of w_1 , so high values of R^2 in our regressions do not confirm that w_1 is accurate. Fortunately, we do not need to estimate w_1 , because we take it to be given by the observed breakdown of generator types in the market, as in Table 2.1. However, it is not so obvious how to categorise the bids from 'Hydro' and 'Other', and also what assumptions to make regarding the bids in the far right. For example, run-of-the-river hydro power facilities (must-run plants) essentially make bids of zero as they cannot choose when to dispatch their power, while storage hydro power plants bid higher up in the stack based on the opportunity cost of using reservoir water at a particular point in time (see Burger *et al* (2007) for more details). Our assumptions for PJM are that nuclear capacity is grouped with coal (before removing $b_L = 20\%$), hydro capacity is grouped with natural gas, while 'Other' is split equally between coal and gas, and the tail of bids beyond b_U is also equally likely to be coal or gas. Thus

$$w_1 = \frac{1}{0.8} (w_{\text{nuclear}} + w_{\text{coal}} + 0.5 w_{\text{other}} - 0.2). \quad (2.9)$$

The weights calculated using Table 2.1 are assumed to change in sudden jumps mostly at points in time when the capacity of PJM increased significantly, which translates to following seven changes in w_1 : June 2000 ($w_1 = 0.469$), June 2002 ($w_1 = 0.451$), June 2003 ($w_1 = 0.424$), May 2004 ($w_1 = 0.521$), June 2005 ($w_1 = 0.509$), June 2006 ($w_1 = 0.503$), December 2006 ($w_1 = 0.496$). We test the accuracy of these weights over time by comparing them to the optimal weights found by minimising the sum of squared errors (SSE) of the observed stack and model implied stack. Unlike MLE, this approach is particularly sensitive to w_1 , as it determines the location of the sudden kink in the middle of the stack, caused by the region of relatively few bids between coal and gas. We only consider the middle of the stack, $B_t(x) : x \in [0.25, 0.75]$, to concentrate on w_1 and avoid the

impact of the tail. Thus for each date t we find

$$\hat{w}_1 = \underset{w_1 \in (0,1)}{\operatorname{argmin}} \sum_{j=0}^{20} B_t^{\text{mod}}(0.25 + 0.025j; w_1, \hat{m}_1, \hat{s}_1, \hat{m}_2, \hat{s}_2) - B_t^{\text{obs}}(0.25 + 0.025j), \quad (2.10)$$

where B_t^{mod} is the estimated model bid stack ($[b_L, b_U] = [0.2, 0.95]$) with estimates $\hat{m}_1, \hat{s}_1, \hat{m}_2, \hat{s}_2$, and B_t^{obs} is the observed bid stack, rescaled to $[b_L, b_U]$. The results for $\hat{w}_1$ are plotted in the second graph in Figure 2.13, along with the fixed weights based on our calculations using Table 2.1. The graph is quite encouraging, as $\hat{w}_1$ remains fairly stable over time and roughly follows the path of the assumed fixed weight, in particular with a gradual decrease in coal between 2000-03, followed by a rapid increase in 2004.

The final point to discuss is the role of the right-hand tail of the bid stack. The bid distributions suggested above ignore this portion of the bid stack as they assume that D_t/C_t never reaches b_U . Even if this is true, the prices produced by the model may become inaccurate for the region just below b_U , as the chosen bid distributions generally cause the bid stack to become too steep too suddenly. The importance of this region depends both on the market in question and on the goal of our model. For example, for PJM, the daily average (or peak average) real-time value of D_t/C_t never reaches 0.9 in any of the historical data, whereas the hourly real-time value does surpass 0.9 in 1.5% of the peak hours observed and surpasses 0.85 in 4.4%. In contrast, NEPOOL day-ahead hourly peak prices only reach 0.85 in 0.5% and 0.8 in 2.4% of observations.²¹ The modelling of the right tail of the bid stack is crucial in capturing the magnitude of these spikes in hourly spot prices. Since forward electricity contracts are often settled by averaging all hourly prices in a given delivery period, working with daily average demand directly is likely to lead to an underestimation of forward prices, as we may miss out on spikes caused by the tail of the bid stack.

Modelling hourly prices therefore requires a modification to the model to better capture the upper tail of the price distribution. Closer examination of the far right of historical bid stacks in both markets suggests a fairly uniform distribution of bids between the tail of the gas distribution and the maximum of \$1000. A simple extension to the model is therefore the addition of a uniform distribution over the entire range $[\$0, \$1000]$, or alternatively over just the tail. This complicates the model by adding an extra weight parameter which can be tricky to estimate. However, tests show that adding this uniform distribution can allow us to set $b_U = 1$ while still retaining the strong correlations between fuel prices and bid stack parameters. This comes unfortunately at the expense of model simplicity and closed-form solutions for derivative prices. Furthermore, it does not guarantee an accurate description of price spikes, as these are very sensitive to both the tail of the bid stack and the tail of the D_t/C_t distribution. Therefore, instead of tackling both of these tails, we propose retaining the simpler framework for the bid stack, and capturing the tail of the price distribution solely through the tail of the D_t/C_t distribution, as we shall describe Section 3.4.

²¹The far left of the bid stack is less significant for both markets, with observations of $D_t/C_t < 0.4$ for 1.3% of PJM data and 0.3% of NEPOOL data.

2.5.6 Error Analysis

As mentioned in the previous sections, $L(w_1, \dots, w_N, m_1, \dots, m_N, s_1, \dots, s_N)$ (the daily likelihood function of the bid stack fitting procedure) is only one possible method of model assessment and comparison over time or over parameter choice. In Section 2.5.4, we emphasised the need to use both likelihood estimates and R^2 values from regressions to compare different distributions, while in 2.5.5 we mentioned the use of an SSE (sum of squared errors) approach to check the parameter value w_1 . We now extend these analyses by assessing the error in the bid stack function in terms of average or maximum absolute deviation in dollar terms of the stack (similarly to using SSE, but more meaningful intuitively), and compare the results for PJM and NEPOOL.

We note firstly that the use of SSE or absolute deviation as a means of estimating parameters $m_1, \dots, m_N, s_1, \dots, s_N$ is highly unstable, due to the right tail of the bid stack. As this tail of the observed stack consists of a sprinkling of bids producing small steep steps all the way to the maximum level of a \$1000, it is inevitable that errors in this section will overwhelm the rest of the error in the stack, even after truncating at b_U . In addition, the fact that $B(\cdot)$ is unbounded both above and below for our chosen distributions (except the Weibull case, which is bounded below) will cause us problems depending on how close to the tails we calculate the deviation. Nonetheless, SSE and average or maximum absolute deviation are very useful tools for analysing the middle of the bid stack, particularly near the typical kink in the stack caused by the region of fewer bids. As in Section 2.5.5, we thus consider only the range $\tilde{D}_t/\tilde{C}_t \in [1/4, 3/4]$. Letting $B^{\text{obs}}(\tilde{D}_t/\tilde{C}_t)$ represent the observed bid stack, and $B^{\text{mod}}(\tilde{D}_t/\tilde{C}_t)$ the model bid stack, define

$$\text{Err}^{\text{AD}(\text{mid})} = 2 \int_{1/4}^{3/4} |B_t^{\text{obs}}(x) - B_t^{\text{mod}}(x)| dx$$

and

$$\text{Err}^{\text{MD}(\text{mid})} = \max_{x \in (1/4, 3/4)} |B^{\text{obs}}(x) - B^{\text{mod}}(x)|$$

to be the average and maximum absolute deviation (over the middle of the stack) respectively.




Figure 2.14: Historical analysis of error in bid stack fitting for PJM (left) and NEPOOL (right) showing $\text{Err}^{\text{AD}(\text{mid})}$ and $\text{Err}^{\text{MD}(\text{mid})}$ with and without the use of observed fuel prices.

We discretise the range $[1/4, 3/4]$ into 100 bins of width 0.005 to estimate these error measures,

as well as to calculate SSE as before. Figure 2.14 illustrates the results over time for PJM and NEPOOL. The results for average absolute deviation (red lines) are quite encouraging in both cases, with a historical average of \$3.65 for PJM and \$2.76 for NEPOOL. Maximum absolute deviation (dark blue line) has a historical average of \$12.16 for PJM and \$6.17 for NEPOOL. While these numbers may seem high, we should remember that we are approximating a step function by a strictly increasing continuous function, so we will of course underestimate at some points and overestimate at others. Furthermore, if we are interested in pricing or hedging derivative contracts, we are always taking an expectation over part or all of the bid stack, so some of the error is likely to average out. On the other hand, using the model as a tool for short-term forecasting of S_t given a precise demand forecast is potentially more dangerous. Figure 2.14 also illustrates the difference between using observed fuel prices to find m_1, s_1, m_2, s_2 from regression results (calculated over the entire dataset) and taking MLE results $\hat{m}_1, \hat{s}_1, \hat{m}_2, \hat{s}_2$. For PJM there is little difference, while for NEPOOL the results with observed fuel prices are generally worse with some periods particularly weak.




Figure 2.15: Location of the maximum absolute deviation point (within $[0.25, 0.75]$) for PJM and NEPOOL (left graph); Maximum deviation for PJM using fixed (from (2.9)) and flexible (from (2.10)) weights w_1 (right graph).

The observation that $\text{Err}^{\text{MD}(\text{mid})}$ is twice as high for PJM as for NEPOOL reflects the extra difficulty in capturing the overlap region of a two-fuel market. There is often a narrow strip of the stack for which $\text{Err}^{\text{MD}(\text{mid})}$ is briefly quite large because of the steeper slope required as well as the restriction of having a fixed weight w_1 . The left graph of Figure 2.15 illustrates this by showing that for much of PJM's history, the maximum deviation in the range $[1/4, 3/4]$ occurs near 0.5, and thus near $\tilde{D}_t/\tilde{C}_t = w_1$. Otherwise, in about a quarter of dates, it occurs above 0.7. In contrast for NEPOOL, the location of maximum deviation is much more varied over time. In both markets, only about 8% of observations occur at the boundaries 0.25 and 0.75. (If we increase the upper boundary 0.75, this proportion will significantly increase.) The right graph of Figure 2.16 shows how much the maximum deviation can be reduced by choosing the optimal weight w_1 using (2.10) for each date instead of fixed weights. In some cases this is substantial but for most dates the difference is small, and overall the average is reduced to \$7.91. Finally, Table 3.5 compares different distributions similarly to Table 2.4 (with $b_U = 0.95$ throughout), but using our error measure for the middle of the stack as a new criteria for ranking performance. Interestingly the Cauchy distribution

now performs well for PJM, ranking first for around half of the dates. The Weibull also performs well as before, with the logistic reasonably strong and well ahead of the Gaussian. For NEPOOL, the logistic distribution dominates, ranking either first or second in around 95% of cases by these measures of performance.

Table 2.5: Comparison of the four distributions by maximum and average absolute deviation.

Distribution	PJM				NEPOOL			
	Err ^{AD} (mid)		Err ^{MD} (mid)		Err ^{AD} (mid)		Err ^{MD} (mid)	
	1 st	2 nd	1 st	2 nd	1 st	2 nd	1 st	2 nd
Gaussian	2.7%	16.0%	1.0%	13.0%	37.4%	3.7%	37.8%	7.1%
Logistic	16.6%	48.0%	12.3%	44.2%	35.2%	60.9%	38.3%	57.4%
Cauchy	44.0%	10.6%	54.1%	10.7%	16.7%	22.5%	13.2%	13.7%
Weibull	36.3%	25.1%	32.2%	31.8%	10.3%	12.6%	10.2%	21.4%

2.5.7 Implied Generation Volume

The error analysis presented above provides a measure of the bid stack model's performance purely in terms of how well the probability densities match the histograms of observed bids, as illustrated in Figure 2.11. We now provide an alternative measure of performance which assesses model performance as a whole, including the fit of the densities as well as the underlying model assumptions from Section 2.5.1. This can be achieved by comparing the observed historical power generation volumes by fuel type with those predicted by the model. In light of data availability, we focus only on the PJM market in this section. As a two-fuel market, it is also arguably a more interesting test of the model.

The PJM website²² publishes power generation volumes by fuel type at a monthly frequency from 2005 onwards, with approximately 15-20 categories of fuel type, including multiple coal and oil categories. Given hourly observed demand D_t and price S_t (with D_t normalised by $c^{\max}$, the total capacity from bids), we can calculate as follows the bid stack model's prediction of total monthly power produced by coal or gas generators. Letting $\chi_G(T_1, T_2)$ represent total electricity (in MWh) produced from gas in the period $[T_1, T_2]$, and using the logistic distribution for bids,²³

$$\chi_G(T_1, T_2) = \sum_{\text{peak } t \in [T_1, T_2]} (D_t - b_L) c^{\max} \frac{\frac{1-w_1}{2} + \frac{1-w_1}{2} \tanh\left(\frac{x-m_2}{2s_2}\right)}{\frac{1}{2} + \frac{w_1}{2} \tanh\left(\frac{x-m_1}{2s_1}\right) + \frac{1-w_1}{2} \tanh\left(\frac{x-m_2}{2s_2}\right)}$$

Note that the denominator of the fraction above would equal D_t/C_t if the model exactly reproduced the observed bid stack. The factor $(D_t - b_L)$ assumes implicitly here that none of the lowest portion of truncated bids come from gas generators. In addition, we sum only over peak hours (weekday

²²More specifically, an independent subsidiary of PJM, known as PJM EIS (PJM Environmental Information Services, Inc.) owns and administers the Generation Attribute Tracking System (GATS) which collects the relevant data. See <https://gats.pjm-eis.com/%5Cmodule%5Cmypage.asp>

²³We choose parameters w_1, b_U, b_L as before, and take $\{m_1, s_1, m_2, s_2\}$ from daily MLE estimates. Using regression results together with actual gas and coal prices also leads to very similar results.

hours 8-23) in the period $[T_1, T_2]$, assuming that the quantity of gas production for off-peak hours is negligible.²⁴ Typically off-peak demand (nights and weekends) constitutes 25% to 30% of total monthly power demand. While a small proportion of this off-peak demand is likely to come from gas generators, the error should be small and in fact offset by a small overestimation of $\chi_G(T_1, T_2)$ due to the grouping of hydro and gas together in weight $1 - w_1$ (see (2.9)). Dividing $\chi_G(T_1, T_2)$ by total monthly power demand we obtain the fraction

$$\frac{\chi_G(T_1, T_2)}{\sum_{\text{all } t \in [T_1, T_2]} D_t c^{\max}},$$

which we compare in Figure 2.16 (left graph) with the actual proportion of gas-based generation in PJM over the period 2005-2007. A similar exercise can of course be performed for coal, though we should then take into account the residual capacity from nuclear in the left of the stack but above b_L . The results of the initial analysis are fairly encouraging as the shape of the blue line roughly matches that of the green line below, including small jumps from month to month. These movements reflect changes in both monthly levels in demand and of gas prices. However, the model clearly overestimates the proportion of gas generation for all months except the summer, when gas generation peaks.




Figure 2.16: Analysis of the model's performance in capturing the historical proportion of electricity generation volume originating from natural gas in PJM (left graph); Proportion of demand considered inelastic in PJM (right graph)

In order to explain this overestimation, we analyse several of the model's key assumptions. Firstly, we investigate the possibility of demand elasticity, which could in some markets distort the simple relationship provided by the function $B(\tilde{D}_t/\tilde{C}_t)$. In particular, the traditional economic demand curve as a function of price may not be completely flat (see earlier comment in Chapter 1), in which case a transformation $B(\cdot)$ should be constructed by shifting the demand curve and recalculating its intersection with supply for all D_t . However, the right graph of Figure 2.16 suggests that our

²⁴This assumption is needed in part because the bid stack model performs less well for off-peak data than for peak data. For example, prices around zero dollars may unrealistically imply a significant proportion of gas generation relative to coal generation due to the much thicker left tail of the second distribution relative to the first. (i.e., $s_2 \gg s_1$). Ultimately, this weakness can only be remedied by choosing a non-symmetric distribution (such as the Weibull, as discussed in Section 2.5.4) since the right tail of the bid stack is likely to always be much heavier than the left tail.

assumption for PJM is reasonable, as about 98% of demand bids in 2006 (and 97% in 2007) can be considered inelastic with respect to price, meaning they are either ‘any price bids’ or bids at \$999, just below the maximum possible price. Like available and installed capacity numbers, this dataset is also available from the PJM website at an hourly frequency from mid-2006 onwards.

Rejecting demand elasticity as a cause, we now discuss two more likely sources of the discrepancy in generation volumes in Figure 2.16. Firstly, a typical weakness of the bid stack model is the overestimation the left tail of the gas bid distribution, due to our use of symmetric probability distributions (except the Weibull distribution). For example in the PJM fit of Figure 2.11, we might expect the lowest gas bids to feature around \$50, but instead the tail of gas distribution continues to be noticeable even at \$0. Secondly, we can question the assumption that capacity becomes unavailable with equal probability through the stack (or more precisely, throughout the rescaled section of the stack (b_L, b_U)). Low demand seasons in particular may lead to a greater proportion of gas generators undergoing maintenance, or alternatively the rate of unexpected outages may differ between fuels. If capacity is initially removed from the gas section of the stack at a greater rate than from the coal section, the result should be a decrease in generation volume from gas especially for low demand months.

An alternative approach to the bid stack model is therefore to allow capacity changes to affect the coal and gas distributions differently, although we assume that within each fuel type, capacity changes are equally spread. Therefore the fixed weights w_1 and $1 - w_1$ in our model can be replaced by functions $w_c(\tilde{C}_t)$ and $w_g(\tilde{C}_t)$ for gas and coal respectively. Retaining the logistic distribution for simplicity, the power price now solves

$$w_c(\tilde{C}_t) \left[\frac{1}{2} + \frac{1}{2} \tanh \left(\frac{S_t - m_1}{2s_1} \right) \right] + w_g(\tilde{C}_t) \left[\frac{1}{2} + \frac{1}{2} \tanh \left(\frac{S_t - m_2}{2s_2} \right) \right] = \tilde{D}_t, \quad (2.11)$$

where we require $w_c, w_g : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ to be non-decreasing non-negative functions satisfying²⁵

$$w_c(x) + w_g(x) = x, \quad \text{and} \quad w_c(1) = w_1,$$

so that at full capacity our initial weights are restored.

For example, let $\xi \in \mathbb{R}^+$ represent the rate at which the gas bids ‘disappear’ relative to coal bids (as capacity decreases from $\tilde{C}_t = 1$). Thus if $\xi = 2$, then once all gas bids have been removed, half of the initial coal bids remain. Therefore

$$\begin{aligned} w_g(x) &= \begin{cases} 0 & \text{if } 0 \leq x \leq \frac{\xi-1}{\xi} - \frac{\xi-1}{\xi}(1-w_1) \\ \frac{1-w_1}{\frac{1}{\xi} + \frac{\xi-1}{\xi}(1-w_1)} \left(x - \frac{\xi-1}{\xi} + \frac{\xi-1}{\xi}(1-w_1) \right) & \text{if } x \geq \frac{\xi-1}{\xi} - \frac{\xi-1}{\xi}(1-w_1) \end{cases}, \\ w_c(x) &= x - w_g(x). \end{aligned}$$

Note that the special case of $\xi = 1$ in this more general framework returns us to the original bid stack model of Section 2.5.1, since $w_c(\tilde{C}_t) = w_1 \tilde{C}_t$ and $w_g(\tilde{C}_t) = (1 - w_1) \tilde{C}_t$, so capacity can return

²⁵The logical interpretation of $w_c(\tilde{C}_t)$ or $w_g(\tilde{C}_t)$ for $\tilde{C}_t > 1$ is the quantity of additional capacity from imports for each fuel type. However, we are primarily interested in region $\tilde{C}_t \in [0, 1]$, as it is rare for available capacity to exceed total installed capacity.

to the right-hand side of (2.11). The red line in Figure 2.16 illustrates the impact of setting $\xi = 2$ (instead of $\xi = 1$) when calculating generation volumes from gas.²⁶ Clearly, there is a significant improvement in the fit, particularly as the summer months show only small changes relative to the standard model, while other seasons show a greater decrease.

Overall, the analysis of historical generation volumes lends further support to the modelling methodology, as the model is clearly able to identify times of higher or lower gas-based generation, purely through the bid structure and demand levels. These times in turn correspond to times of higher or lower levels of electricity to gas price dependence, since the frequency of gas being the marginal fuel type is clearly linked to its generation volumes. At least in broad terms, the model is able to correctly identify regions of the stack dependent on different fuel prices and consequently approximate well the power-fuel dependence structure. However, the above analysis also points out possible improvements to the model, including either non-symmetric bid distributions or changes to capacity assumptions. While both of these are realistic options and easy to implement, they remove the possibility of explicit expressions for forward and option prices, as we shall present in later chapters. Hence, while acknowledging its potential limitations, we stick with the standard model and argue that for typical modelling purposes, it provides a sufficiently strong approximation to reality.

2.6 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. The fit to historical generation volumes is also reasonably convincing, lending further support not just to the fitting of the bid stack but also to the resulting relationship between electricity price and demand. 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 type, 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 $\alpha_0^G = A, \alpha_1^G = \mu_H, \beta_0^G = 0$ and $\beta_1^G = \sigma_H$. With fixed costs $A > 0$ we expect $\alpha_0^G > 0$, as we observe in our regressions for both coal and gas. The observed mix of positive and negative values of β_0^G could be reproduced by letting A instead have a Gaussian distribution correlated with H .²⁷ 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^G (7.67 and 8.51 for NEPOOL and PJM respectively) also correspond closely to average PJM

²⁶Note that this involves calculating a model-implied capacity number from S_t , D_t and the model bid stack, as we shall introduce formally in Section 3.4.

²⁷Of 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.

heat rates for gas generators, listed by PJM²⁸ as approximately 7.3 MBtu/MWh in 2004. Similarly, $\alpha_1^C = 0.408$ 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.

In conclusion, the primary advantage of the framework introduced in this Section 2.5 is the flexible and intuitive manner in which electricity prices are linked to underlying factors, together with the wealth of information provided by using observed bid data. Taking advantage of an entire curve of data points (i.e., the bid stack) for each date immediately eases the problem of discovering relationships between a single power price S_t and its numerous fundamental drivers. Moreover, it provides greater justification and accuracy to generation cost based models which are otherwise typically require detailed knowledge and/or numerous assumptions (including the role of or lack of strategic bidding). Nevertheless, the bid stack model itself requires us to make a number of approximations and choices in order to best describe the movement of bids. For some markets such as PJM and NEPOOL, we may be able to safely make enough assumptions (such as truncation of bid data, ignoring of minor fuel-types or those which rarely set the price) in order to obtain convenient mathematical expressions, while in other markets we need to be satisfied with a more complicated model based entirely on numerical methods. However, in either case, the approach provides an innovative new way of analysing the electricity markets and in particular their relationship with other energy markets.

²⁸Data from "Impacts of the PJM RTO Market Expansion", Energy Security Analysis, Inc., November 2005.

Chapter 3

From Underlying Drivers to Electricity Prices

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 sometimes coal prices P_t , demand (or load) D_t , and capacity available C_t . We now introduce models and calibration techniques for each of these factors. 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 choice of bid stack function. In terms of pricing or hedging contracts dependent on S_t , this relegates much of the complexity of the model to the payoff function.

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 primarily impact a flatter and lower part of the bid stack and they are less significant in setting prices particularly during peak hours. Thus, we propose both a simple deterministic model to incorporate some coal price information without increasing the complexity of the model, and a more realistic stochastic coal price model. In a more general setting, other fuel prices and even carbon emissions prices can be included, as is clearly necessary for European power markets today. We shall discuss carbon price modelling in detail in Chapter 6.

3.1 Natural Gas Prices

We choose a standard approach to modelling gas prices, by fitting log gas prices with the Schwartz two-factor model described by Schwartz (1997),¹ capturing the mean-reversion which is widely believed to exist in most commodity prices, as well as changes in the long term equilibrium level of prices. The term structure of volatility in the natural gas market is shown in Figure 3.1, and justifies the use of the Schwartz two-factor model in order to capture the decreasing volatility with time to maturity (Samuelson effect) as well as the fact that it levels off significantly above zero. Let X_t^1 and

¹This 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.

X_t^2 be the two correlated stochastic factors driving the spot gas price G_t , and let $g(t)$ be a seasonal component. Dynamics under an appropriate risk-neutral measure $\mathbb{Q}$ are given by

$$\begin{aligned} dX_t^1 &= \kappa(\mu_1 - X_t^1)dt + \sigma_1 dW_t^1, \\ dX_t^2 &= \mu_2 dt + \sigma_2 dW_t^2, \\ dW_t^1 dW_t^2 &= \rho_{12} dt \\ G_t &= \exp(g(t) + X_t^1 + X_t^2). \end{aligned} \quad (3.1)$$

3.1.1 Parameter Estimation

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

$$\begin{aligned} \log(F^G(t, T)) &= \log\left(\mathbb{E}_t^{\mathbb{Q}}\left[e^{g(T) + X_T^1 + X_T^2}\right]\right) \\ &= g(T) + X_t^1 e^{-\kappa(T-t)} + \mu_1 \left(1 - e^{-\kappa(T-t)}\right) + X_t^2 + \mu_2(T-t) \\ &\quad + \frac{\sigma_1^2}{4\kappa} \left(1 - e^{-2\kappa(T-t)}\right) + \frac{1}{2}\sigma_2^2(T-t) + \frac{\rho_{12}\sigma_1\sigma_2}{\kappa} \left(1 - e^{-\kappa(T-t)}\right) \end{aligned} \quad (3.2)$$

Note that, in practice, gas forwards are contracts for delivery over a period $[T_U, T_L]$ (typically a month), split into n units (typically about 21 business days) of length $\Delta t = (T_U - T_L)/n$. However, using the midpoint of the month as an approximate delivery time is a reasonable assumption in most cases, particularly for maturities more than one month away. More precisely, for a delivery period $[T_U, T_L]$, let $T = \frac{T_U + T_L}{2}$ and $t_i = T_L - \frac{\Delta t}{2} + i\Delta t$, for $i = 1, \dots, n$. Then,

$$F^P(t, [T_L, T_U]) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_t^{\mathbb{Q}}[S_{t_i}] \approx \mathbb{E}_t^{\mathbb{Q}}[S_T].$$

As X_t^1 and X_t^2 are unobservable factors, we use the Kalman Filter to calibrate our model to available Henry Hub² gas data (see e.g. 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 stochastic differential equations in (??). As liquid futures or forward prices are widely available for the US gas market, (3.2) 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, while more recent forward curves have been obtained from daily factsheets available at <http://www.cmegroup.com/trading/energy-metals/nymex-daily-reports.html>. 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 midpoint of each month, thus ignoring month-long delivery periods.

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 (and is given in Figure 3.1), the amplitude varies significantly over time for both the forward curve and log-forward curve. Thus, we

²While we use Henry Hub as the most well-known and easily available US gas price data, others exist and in particular there may be some benefit from using North Eastern prices if available.




Figure 3.1: Volatility term structure of forward gas prices (left graph); Average seasonality of the natural gas log-forward curve (2006-08) (right graph)

cannot easily remove the seasonality using the same function $g(T)$ throughout. Instead, we deseasonalise each daily 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 (e.g., January) equal the base month, and calculate for every other month the average difference between forwards with maturity in that 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 standard deviation of observation noise σ_{noise} . Hence we perform the optimisation in two stages, using a method similar to that of Culot *et al* (2006). For fixed σ_{noise} , 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 January 2000, and setting 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_{noise} 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 (3.2), call option prices (and similarly for puts) with strike K , option maturity T_1 , forward maturity T_2 , constant interest rate r are given in closed form by

$$\begin{aligned} V^G(t, T_1, T_2) &= e^{-r(T_1-t)} E_t \left[(F^G(T_1, T_2) - K)^+ \right] \\ &= e^{-r(T_1-t)} \left[e^{\tilde{\mu}_G + \frac{1}{2}\tilde{\sigma}_G^2} \Phi \left(\frac{-\log(K) + \tilde{\mu}_G + \tilde{\sigma}_G^2}{\tilde{\sigma}_G} \right) - K \Phi \left(\frac{-\log(K) + \tilde{\mu}_G}{\tilde{\sigma}_G} \right) \right] \end{aligned} \quad (3.3)$$

where Φ is the standard Gaussian cdf,

$$\begin{aligned} \tilde{\mu}_G &= g(T_2) + X_t^1 e^{-\kappa(T_2-t)} + \mu_1 \left(1 - e^{-\kappa(T_2-t)} \right) + X_t^2 + \mu_2(T_2 - t) \\ &\quad + \frac{1}{2}\sigma_2^2(T_2 - T_1) + \frac{\sigma_1^2}{4\kappa} \left(1 - e^{-2\kappa(T_2-T_1)} \right) + \frac{\rho_{12}\sigma_1\sigma_2}{\kappa} \left(1 - e^{-\kappa(T_2-T_1)} \right), \\ \tilde{\sigma}_G^2 &= \frac{\sigma_1^2}{2\kappa} \left(1 - e^{-2\kappa(T_1-t)} \right) e^{-2\kappa(T_2-T_1)} + \sigma_2^2(T_1 - t) + \frac{2\rho_{12}\sigma_1\sigma_2}{\kappa} \left(1 - e^{-\kappa(T_1-t)} \right) e^{-\kappa(T_2-T_1)}. \end{aligned}$$

The results are shown in Table 3.1 for different date ranges. The parameters seem fairly stable over

Table 3.1: Kalman Filter results for gas and coal parameters

	Natural Gas parameters							Coal parameters			
Date range	ρ_{12}	κ	μ_1	σ_1	μ_2	σ_2	σ_{noise}	μ_3	σ_3	ρ_{13}	ρ_{23}
Jan00 - Dec05	0.018	0.781	1.009	0.566	-0.030	0.137	0.0592	0.043	0.231	0.158	0.065
Jan00 - Dec06	0.098	1.137	1.049	0.646	-0.044	0.129	0.0694	-0.016	0.228	0.208	0.101
Jan00 - Dec07	0.044	1.221	1.088	0.567	-0.046	0.122	0.0786	0.024	0.221	0.221	0.111
Jan00 - Nov08	0.024	1.157	1.028	0.695	-0.041	0.133	0.0637	0.051	0.285	0.258	0.141

time, though the more recent data are characterised by a slightly higher volatility and faster speed of mean reversion for X^1 and larger negative drift for X^2 . Correlation ρ_{12} is very low throughout. Interestingly, the overall volatility $\tilde{\sigma}_G$ of forward gas prices is greater in parameter estimates from December 2005 and November 2008, than from December 2006 or 2007. This makes sense given the history of G_t in Figures 2.5 and 3.2, since late 2005 and mid 2008 correspond to the two most prominent periods of sudden gas price rises. Note also that $\mu_2 < 0$ is necessary to account for the fact that the gas forward curve was in backwardation (downward sloping) in 81.5% of the observations in the dataset.

3.1.2 Calibration to Forward Curve

Historical observations of natural gas forward curves are used directly in the estimation of parameters $\{\kappa, \mu_1, \sigma_1, \mu_2, \sigma_2, \rho_{12}\}$. However, in order to price other derivatives such as options on forwards, we require that our model reproduce (or match) the current observed natural gas forward curve $F^G(t, T), \forall T$. In fact, this step is also necessary in the estimation technique above before options on forwards are priced by (3.3). We can easily achieve this by allowing the mean-reversion level μ_1 to be a time dependent function $\mu_1(T)$, a method similar to the Hull and White approach for yield curve matching. In doing so, we implicitly assume a time varying market price of risk for gas, and also that this time varying risk premium is attached only to X_t^1 , not X_t^2 . For simplicity, we assume that $\mu_1(T)$ is a step function with jumps at the beginning of each new month. We use the notation $\mu_1(T_i), i = 1, \dots, N$ where T_1 corresponds to the midpoint of the delivery month for the shortest maturity contract, then T_2 the midpoint of the next month, and so on. For any time between the current time t and the beginning of the first delivery period, μ_1 also equals $\mu_1(T_1)$, and hence we have the same number of values of μ_1 to find as we have observable forward contracts.

Note that this calibration method provides us with risk-neutral dynamics for G_t which lead to correct forward prices for all maturities. These dynamics (i.e., the market prices of risk) are unique under the assumptions and approximations of the model, such as the ignoring of delivery periods and the piecewise constant nature of $\mu_1(T)$. However they do not define a unique risk-neutral measure in the strict sense of perfect derivative replication. In order to do this, as with bond prices needed in the Hull and White model, we would strictly require a separate derivative contract (typically forward contract) of every possible maturity in order to hedge all risk. As a monthly forward contract depends on an average of spot prices over the month, it can only provide an approximate

hedge for that period's spot price movement, and hence the parameter $\mu_1(T)$ can be thought of as a sort of average over possible risk-neutral measures $\mathbb{Q}$ given our market data. Similarly, additional derivatives would be necessary in order to determine whether a time-varying market price should also apply to X_t^2 , as well as X_t^1 . However, we focus here on the common practical issue of calibrating a model to forward prices in order to then consistently price other derivatives within the framework of the model. The above assumptions are necessary for an empirical analysis of this sort.

Solving the SDE (3.1) with a time-dependent μ_1 gives the conditional distribution

$$\log G_{T,t} \sim N \left(g(T) + X_t^1 e^{-\kappa(T-t)} + \int_t^T \kappa \mu_1(s) e^{-\kappa(T-s)} ds + X_t^2 + \mu_2(T-t), \sigma_G^2(T-t) \right),$$

where

$$\sigma_G^2(T-t) = \frac{\sigma_1^2}{2\kappa} \left(1 - e^{-2\kappa(T-t)} \right) + \sigma_2^2(T-t) + \frac{2\sigma_1\sigma_2\rho_{12}}{\kappa} \left(1 - e^{-\kappa(T-t)} \right).$$

Thus for the shortest maturity contract

$$\mu_1(T_1) = \frac{1}{1 - e^{-\kappa(T_1-t)}} \left\{ \log F^G(t, T_1) - g(T_1) - \frac{1}{2} \sigma_G^2(T_1-t) - X_t^1 e^{-\kappa(T_1-t)} - X_t^2 - \mu_2(T_1-t) \right\}$$

Letting $\Delta T = T_{i+1} - T_i \quad \forall i$, assuming that all months are the same length (not a necessary assumption but a useful approximation), then for $i \geq 2$ we can write

$$\int_t^{T_i} \kappa \mu_1(s) e^{-\kappa(T_i-s)} ds = \mu_1(T_1) e^{-(i-1)\kappa\Delta T} (1 - e^{-\kappa(T_1-t)}) + \sum_{k=2}^i \mu_1(T_k) e^{-(k-1)\kappa\Delta T} (1 - e^{-\kappa\Delta T}).$$

Hence to reproduce the observed forward curve (for $i = 2$ the sum term does not appear)

$$\begin{aligned} \mu_1(T_i) = \frac{1}{1 - e^{-\kappa\Delta T}} & \left\{ \log F^G(t, T_i) - g(T_i) - \frac{1}{2} \sigma_G^2(T_i-t) - X_t^1 e^{-\kappa(T_i-t)} - X_t^2 - \mu_2(T_i-t) \right. \\ & \left. - \mu_2(T_i-t) - \mu_1(T_1) e^{-(i-1)\kappa\Delta T} (1 - e^{-\kappa(T_1-t)}) - \sum_{j=2}^{i-1} \mu_1(T_j) e^{-(j-1)\kappa\Delta T} (1 - e^{-\kappa\Delta T}) \right\}. \end{aligned}$$

Under the lognormal model presented here, pricing forwards with delivery periods involves finding the expectation of a sum of (hourly) lognormal random variables. While this cannot be written explicitly without the summation, authors such as Culot *et al* (2006) suggest the use of a Taylor Series expansion to further simplify the expression. In this case the simple calibration method can still be adopted. Similarly, in the case that gas prices (instead of log-gas prices) are linear in the state variables, an explicit expression can be written (i.e., using the sum of normal random variables), however the performance of the model is quite poor in this case, particularly as gas price is no longer bounded below by zero and options are very poorly priced. Instead, we maintain the approximation of using the midpoint of the month as a maturity, and stress that, for most forward contracts, the length of the delivery period Δt is small compared to the length of time to maturity $T_1 - t$. Moreover, even for short maturities the approximation is reasonable since the expectation and variance of $X_{T_1}^1$ and $X_{T_1}^2$ are typically quite near the average of the expectations and variances of these same factors at the delivery endpoints $T_1 - \Delta T/2$ and $T_1 + \Delta T/2$, while $g(T)$ is modelled as constant through the delivery period.

3.2 Coal Prices

As PJM is roughly 40% fueled by coal, changes in coal prices can have a significant impact on the level of power prices. However, as we have seen (e.g. Figure 2.6), coal prices moved only gradually over the time period 2001-07, 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 initially 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.




Figure 3.2: Gas versus coal over the period (2006-08) illustrated as both a line graph and scatter plot. Data from NYMEX daily fact sheets.

While this simple model is suitable for most years in the dataset, Figure 3.2 shows that in 2008 coal price volatility increased dramatically, as well as prices reaching record highs near \$140 before falling equally rapidly. Figure 3.3 also confirms this change, with roughly a doubling of volatility observed throughout the forward curve. This implies that for the most recent dates and for future periods, it is necessary to consider the more general case of stochastic coal prices. Coal forward data reveal that (unlike gas) the term structure of volatility is very flat (see Figure 3.3), and also no seasonality is present. Hence we propose a much simpler model for coal, using a single factor Geometric Brownian Motion. 2008 data in Figure 3.2 also suggest that coal and gas price levels (and commodity prices in general) can often show significant co-movement, though it can be difficult to distinguish between cointegration and correlation over short time horizons. In the right graph of Figure 3.3, we plot recent correlation estimates calculated for a one-year rolling window using either the shortest forward contract available or an average of the first two years of the curve. Of course correlation estimates using prices are unstable and often meaningless (as the series are non stationary), but do at least demonstrate the recent co-movement. On the other hand, correlations using returns are much more stable and useful. Interestingly, while the correlation of the short end of the curves has remained near 0.25 throughout 2006-08, the correlation of the entire curve's average

has steadily increased towards 0.5, suggesting that there is evidence of correlation of coal both with factors X_t^1 and X_t^2 .

The discussion above suggests that our process for P_t should at least be correlated to G_t . Thus, we choose the following one-factor model for coal:

$$\begin{aligned} dX_t^3 &= \mu_3 dt + \sigma_3 dW_t^3, \\ dW_t^1 dW_t^3 &= \rho_{13} dt \\ dW_t^2 dW_t^3 &= \rho_{23} dt \\ P_t &= \exp(X_t^3), \end{aligned} \tag{3.4}$$

where the drift μ_3 can be interpreted as the interest rate adjusted by storage costs and convenience yield, as is common for storable commodities like coal. As all of these can be time dependent, it is logical for μ_3 to be time dependent, as we shall suggest below for calibration purposes.




Figure 3.3: Volatility term structure of forward coal prices (left graph); Rolling one-year correlation estimates for coal and gas prices and returns, using both the shortest forward contract and an average over the first two years of the forward curve.

3.2.1 Parameter Estimation

The coal price parameters $\{\mu_3, \sigma_3, \rho_{13}, \rho_{23}\}$ can be estimated by MLE using the histories of X_t^1 and X_t^2 obtained from the Kalman Filter approach, together with the observed history of P_t . As coal is not a function of unobservable factors, changes to the filtering algorithm are not necessary. However, maximum likelihood estimation (for step size δt) requires calculating the conditional distribution of $X_{t+\delta t}^3$ given the set $\{X_{t+\delta t}^1, X_{t+\delta t}^2, X_t^3\}$. First note that for all t and T , (X_T^1, X_T^2, X_T^3) is multivariate Gaussian with mean Ω and covariance matrix Σ , where

$$\Omega = \begin{pmatrix} X_t^1 e^{-\kappa(T-t)} + \mu_1 (1 - e^{-\kappa(T-t)}) \\ X_t^2 + \mu_2(T-t) \\ X_t^3 + \mu_3(T-t) \end{pmatrix}, \tag{3.5}$$

and

$$\Sigma = \begin{pmatrix} \frac{\sigma_1^2}{2\kappa} (1 - e^{-2\kappa(T-t)}) & \frac{\rho_{12}\sigma_1\sigma_2}{\kappa} (1 - e^{-\kappa(T-t)}) & \frac{\rho_{13}\sigma_1\sigma_3}{\kappa} (1 - e^{-\kappa(T-t)}) \\ \frac{\rho_{12}\sigma_1\sigma_2}{\kappa} (1 - e^{-\kappa(T-t)}) & \sigma_2^2(T-t) & \rho_{23}\sigma_2\sigma_3(T-t) \\ \frac{\rho_{13}\sigma_1\sigma_3}{\kappa} (1 - e^{-\kappa(T-t)}) & \rho_{23}\sigma_2\sigma_3(T-t) & \sigma_3^2(T-t) \end{pmatrix}. \quad (3.6)$$

Thus, the conditional distribution of $X_{t+\delta t}^3$ given $\{X_{t+\delta t}^1, X_{t+\delta t}^2, X_t^3\}$ is $N(\bar{\mu}, \bar{\sigma})$ where (using Ω_i and Σ_{ij} to represent components of the above vector and matrix with $T = t + \delta t$)

$$\bar{\mu} = \Omega_3 + \begin{pmatrix} \Sigma_{31} & \Sigma_{32} \end{pmatrix} \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}^{-1} \begin{pmatrix} X_{t+\delta t}^1 - \Omega_1 \\ X_{t+\delta t}^2 - \Omega_2 \end{pmatrix} \quad (3.7)$$

and

$$\bar{\sigma} = \Sigma_{33} - \begin{pmatrix} \Sigma_{31} & \Sigma_{32} \end{pmatrix} \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix}^{-1} \begin{pmatrix} \Sigma_{13} \\ \Sigma_{23} \end{pmatrix}. \quad (3.8)$$

The parameter estimates for coal prices over various time periods were given in Table 3.1 along with those for gas. As discussed, 2008 clearly witnessed a significant increase in coal volatility σ_3 , so recent estimates would be higher if estimated over a shorter time horizon. In addition, correlation estimates between gas and coal have steadily increased in recent years, and would again be higher over a shorter time horizon. The correlation estimates in Table 3.1 are in line with the correlation of one-month (nearest maturity) forward contracts, as shown in Figure 3.3, but lower than recent correlation witnessed in the entire curve. Correlation can be difficult to estimate reliably particularly since gas and coal markets may perhaps be better modelled as co-integrated processes than as processes with correlated returns. In energy or commodity markets, there are typically some common macroeconomic factors which keep long-term levels of prices moving in similar directions, even when short term returns are less closely linked. These points suggest that our values ρ_{13} and ρ_{23} should perhaps be somewhat higher than those in Table 3.1 in order to capture realistic joint distributions of (C_T, G_T) .

3.2.2 Calibration to Forward Curve

Calibration of the coal forward curve $F^C(t, T_i), \forall i$ can be achieved similarly to gas. Since

$$\log F^C(t, T_1) = X_t^3 + \mu_3(T_1)(T_1 - t) - \frac{1}{2}\sigma_3^2(T_1 - t), \quad (3.9)$$

we set

$$\begin{aligned} \mu_3(T_1) &= \frac{\log F^C(t, T_1) - \frac{1}{2}\sigma_3^2(T_1 - t) - X_t^3}{T_1 - t}, \\ \mu_3(T_i) &= \frac{\log F^C(t, T_1) - \frac{1}{2}\sigma_3^2(T_1 - t) - X_t^3 - \mu_1(T_1)(T_1 - t) - \sum_{j=2}^{i-1} \mu_1(T_j)\Delta T}{\Delta T}, \quad \text{for } i \geq 2. \end{aligned}$$

Thus, just as we ultimately ignore our parameter estimate for μ_1 , we also ignore our initial estimate of μ_3 , and instead deduce this parameter from observed forwards. Again, we assume $\mu_3(T)$ to be piecewise constant over each monthly delivery period, and to implicitly reflect a time-dependent market price of risk in the coal market.

3.3 Demand

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 3.4 show average daily peak (real-time) demand in PJM and New England. As we use peak data only, we firstly remove weekdays and public holidays, but account for the longer step sizes (e.g., from Friday to Monday, as temperature should vary more than over one day) when performing the likelihood estimation. Using daily data, we have averaged out intra-day effects, though these can easily be captured by a smooth and regular intra-day profile. The primary seasonality remaining is annual, and 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 of one year and six months. However, data show 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 the trend (parameter a_2 below) to be zero for both markets. Finally, the deseasonalised process is fitted by maximum likelihood estimation using an exponential Ornstein-Uhlenbeck (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

$$\left. \begin{aligned} \log(\tilde{D}_t) &= f(t) + Y_t \\ f(t) &= a_1 + a_2 t + a_3 \cos(2\pi t + a_4) + a_5 \cos(4\pi t + a_6) \\ dY_t &= \kappa_Y(\mu_Y - Y_t)dt + \sigma_Y dB_t \end{aligned} \right\} \quad (3.10)$$

where B_t is a Brownian Motion independent of W_t^1 , W_t^2 and W_t^3 in the gas and coal processes. 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 (2008) also suggest this to be a reasonable assumption over most markets, though of course a prolonged cold spell over a large region is likely to impact both power demand and gas prices. However, power demand is primarily driven by short-term local weather patterns while gas and coal prices are driven largely by longer term and more global factors. Historical correlation between log-returns of gas prices and those of PJM demand is found to be 0.10 over the three year estimation period, while for NEPOOL demand we find 0.07. Using coal prices instead we obtain 0.09 and 0.08 respectively. These results lend support to the simplifying assumption of independence.

3.3.1 Parameter Estimation

The maximum likelihood estimation for the OU process is a simple procedure. Table 3.2 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

Table 3.2: 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)	52.31	-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

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.

3.4 Capacity and Margin

As explained in Chapter 2, 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. Despite these complications and the fact that many of these factors are not easily observable, we can broadly think of C_t simply as the percentage of maximum capacity available. Since we observe hourly historical values for price, demand and all generator bids, we can calculate the implied capacity available C_t^{imp} which allows (2.2) to hold as closely as possible.³ As $B^{\text{obs}}(\cdot)$ is non-decreasing, C_t^{imp} is uniquely defined by⁴

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

The upper lines in Figure 3.4 show the historical evolution of the daily peak average implied capacity availability C_t^{imp} for both PJM and NEPOOL.⁵ It has clear seasonality matching roughly

³Recall that B^{obs} is typically a step function, though some generators in PJM may submit continuous curves of bids, making B^{obs} a combination of step functions and piecewise linear functions.

⁴Occasionally, we observe $D_t > D_t/C_t^{\text{imp}}$, implying excess capacity available in the market relative to normal maximum capacity ($C_t^{\text{imp}} > 1$ in Figure 3.4). 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. Therefore if $C_t^{\text{imp}} > 1$, 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). This change corresponds to assuming that capacity changes cause horizontal shifting instead of tilting of the stack, an alternative mentioned earlier in Section 1.3.3. For these rare occasions, we replace C_t^{imp} by an adjusted value C^{adj} where $C^{\text{adj}} = D_t + (\text{extra capacity to the right of the marginal bid}) = D_t + 1 - D_t/C_t^{\text{imp}}$.

⁵Note that this same PJM data is plotted earlier in Figure 2.4 before rescaling by total capacity from bids, c^{max} , also given in Figure 2.4. The slight discrepancy between Figures 2.4 and 3.4 stems from the fact that Figure 2.4 plots c^{max} before removing a cluster of bids between \$994 and \$1000, as mentioned in a footnote in Section 2.5.2.




Figure 3.4: 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 are plotted for PJM, while DA are used for NEPOOL.)

with demand seasonality, thus dampening the seasonality of prices. This is due to generators' maintenance schedules which are designed to avoid high demand periods.⁶ 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 $\tilde{C}_t^{\text{imp}}$ using the rescaling in (2.4) without breaking the condition $0 < \tilde{D}_t/\tilde{C}_t^{\text{imp}} < 1$. This holds as long as $B^{\text{obs}}(b_L) < S_t < B^{\text{obs}}(b_U)$ for all historical data, or equivalently, $b_L < D_t/C_t^{\text{imp}} < b_U$. However, for PJM, we find that $D_t/C_t^{\text{imp}} > 0.95$ for 0.41% of recent hourly data, while for NEPOOL $D_t/C_t^{\text{imp}} > 0.9$ for 0.6%. For these observations, we cannot define $\tilde{C}_t^{\text{imp}}$ as above, without breaking the restriction $\tilde{D}_t < \tilde{C}_t^{\text{imp}}$. However, we can always define the (rescaled) model-implied capacity available $\tilde{C}_t^{\text{mod}}$ as the unique solution to

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

where $B(\cdot)$ is our model bid stack, as defined in (2.5). In general, we expect $\tilde{C}_t^{\text{mod}}$ to remain close to $\tilde{C}_t^{\text{imp}}$ for the majority of data but to behave differently in the tails.

Given our model for $\tilde{D}_t$ in (3.10) 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.

⁶As discussed in Chapter 2, while it might be expected that bid data should not include generators known to be undergoing maintenance, historical bid data are in fact observed prior to the incorporation of scheduled outages.

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

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

Hourly historical data for $\tilde{M}_t^{\text{imp}}$ 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. This is confirmed by Figure 3.5 which shows that for both PJM and NEPOOL, hourly changes in capacity are typically much larger than hourly changes in demand, with a very fat-tailed distribution. We are interested less in describing the precise timing or auto-correlation 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’:




Figure 3.5: Histogram of hourly changes in real-time peak demand D_t (blue line) and implied capacity available C_t^{imp} (red line), for PJM (left) and NEPOOL (right).

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

where the normal regime is given by

$$dZ_t^{\text{OU}} = \kappa_Z (\mu_Z - Z_t^{\text{OU}}) dt + \sigma_Z d\tilde{B}_t, \quad dB_t d\tilde{B}_t = \rho dt \quad (3.11)$$

and the spike regime is given by

$$Z_t^{\text{SP}} = \alpha - J, \quad J \sim \text{Exp}(\lambda_i), \quad \text{for seasons } i = 1, 2, 3, 4. \quad (3.12)$$

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. Data suggest 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

Table 3.3: 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

$\{p_i, \lambda_i : i = 1, 2, 3, 4\}$ 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).

3.4.1 Parameter Estimation

In order to estimate the parameters above, we firstly use daily average implied margin $\tilde{M}_t^{\text{imp}}$, which averages over intra-day spikes, to help better capture 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 $\tilde{D}_t$.⁷ 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 the fact that Z_t^{OU} has a Gaussian distribution and hence a skew of zero. In contrast, the historical distribution of $\log \tilde{M}_t^{\text{imp}}$ 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 $\log \tilde{M}_t^{\text{imp}} < \alpha$ 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 μ_Z to the mean of the truncated distribution, and σ_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,

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

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

⁷We cap $\tilde{D}_t/\tilde{C}_t^{\text{imp}}$ at 0.99 to avoid the rare cases of negative $\tilde{M}_t^{\text{imp}}$ when $\tilde{D}_t/\tilde{C}_t^{\text{imp}} > b_U$ and reduce the largest downwards spikes in margin. This ultimately has little impact as the left tail of the margin is captured in the second stage of estimation.

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 of $\tilde{M}^{\text{imp}}$ to that of $\tilde{M}^{\text{mod}}$. 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 3.6 shows log-histograms of $\log \tilde{M}^{\text{mod}}$ 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 from 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.⁸ Table 3.3 lists all the estimated parameters for the margin process (with RT data for PJM, and DA data for NEPOOL), 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).




Figure 3.6: 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 Chapter 2. Firstly, for the one-fuel case of NEPOOL, our equation for electricity spot prices, (2.7), can be rewritten as follows:

$$S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t) \left(\log \tilde{D}_t - \log \tilde{M}_t \right) \quad (3.13)$$

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

⁸Firstly, 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, note that 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.

delivery period $[T_1, T_2]$, which has little impact for monthly forwards. The same heuristic justification can be made as in Section 3.2.1, and is further strengthened by the rapid mean reversion of demand and capacity, implying that very little difference exists between their distributions at different points in the delivery period. Secondly, we assume that the market prices of demand risk and margin risk are both zero. In other words, (3.10) and (3.4)-(3.12) hold under both the real world measure $\mathbb{P}$ and the appropriate risk-neutral measure $\mathbb{Q}$. (Note again that by using the Kalman Filter, the dynamics of gas prices G_t in (3.1) 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 (3.13), we have

$$\begin{aligned}
F^{\mathbb{P}}(t, T) &= \mathbb{E}_t^{\mathbb{Q}}[S_T] \\
&= \mathbb{E}_t^{\mathbb{Q}} \left[\alpha_0 + \alpha_1 G_T + (\beta_0 + \beta_1 G_T) \left(\log \tilde{D}_T - \log \tilde{M}_T \right) \right] \\
&= \alpha_0 + \alpha_1 F^{\text{G}}(t, T) + (\beta_0 + \beta_1 F^{\text{G}}(t, T)) \left\{ f(T) + Y_t e^{-\kappa_Y(T-t)} + \mu_Y \left(1 - e^{-\kappa_Y(T-t)} \right) \right. \\
&\quad \left. - (1 - p_i) (Z_t^{\text{OU}} e^{-\kappa_Z(T-t)} + \mu_Z \left(1 - e^{-\kappa_Z(T-t)} \right)) - p_i \left(\alpha - \frac{1}{\lambda_i} \right) \right\}, \tag{3.14}
\end{aligned}$$

for seasons $i = 1, 2, 3, 4$, where $F^{\text{G}}(t, T)$ is the forward gas price for the same maturity, and is given by (3.2).

In addition, the logistic distribution provides an appealing description of the right tail of the price distribution. 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.⁹ Unlike typical regime-switching models for power prices (e.g., De Jong and Huisman (2003), 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 2.1 suggests that while an hourly price of \$150 could reasonably have been considered a spike in early 2004, it could not have been 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

$$\log \tilde{\tilde{M}}_t \approx \log \tilde{\tilde{D}}_t - \frac{S_t - m_2}{s_2}, \tag{3.15}$$

where $\tilde{\tilde{M}}_t$ and $\tilde{\tilde{D}}_t$ (and $\tilde{\tilde{C}}_t$ below) represent a second rescaling of demand and margin from (b_L, b_U)

⁹Ignoring 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.

to the gas-dominated portion of the bid stack $(b_L + w_1(b_U - b_L), b_U)$. The rescaling technique introduced in (2.4) 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 (3.15),

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

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.

3.4.2 Alternative Processes for Margin

As is the case for any of the state variables in the model, some markets may require different processes for these same variables. In particular, margin $\tilde{M}_t$ is a likely candidate for alternative modeling approaches, as it is the most erratic of the variables, due in part to its construction as an implied variable in our case, capturing in a sense the mismatch between $B^{\text{obs}}(D_t)$ and price S_t . Here we present a brief aside to investigate some alternatives.

Firstly, it is worth mentioning the more formal case of a continuous time Markov Chain with a full transition matrix between the two regimes introduced earlier (i.e., between Z_t^{OU} in (3.11) and Z_t^{SP} in (3.12)), instead of independently choosing the state for each hour or each time t . This is clearly a more realistic description of outages and outage repairs (or spikes and spike recoveries), but has virtually no impact on pricing results due to the frequency of switches. In particular, assuming a Poisson process for outages with intensity λ_1 and outage duration exponentially distributed with parameter λ_2 , then the infinitesimal generator matrix of the regime switching process is simply

$$\begin{pmatrix} -\lambda_1 & \lambda_2 \\ \lambda_2 & -\lambda_1 \end{pmatrix}.$$

Tests of historical data using this approach lead to an expected duration of stay in the spike regime of only approximately 2 hours for PJM and 3 hours for NEPOOL, implying a very rapid recovery from outages or other grid problems. Thus the need for correctly capturing the timing of these short stays in the spike regime is low, and far inferior to a correct description of spike size as given by Z_t^{SP} .

Alternatively, the volatile and erratic behaviour of the margin process can also be represented by a jump-diffusion process or perhaps more realistically by a combination of a jump process and a

diffusion process (since the speed of mean reversion for the diffusive component seems to be similar to that of the correlated demand movements, but quite different from the speed of spike recoveries). Either of these cases falls into the affine jump-diffusion (AJD) framework introduced in (1.4) in Chapter 1, and thus we can still obtain closed-form solutions for forward prices in the one-fuel logistic case. Moreover we can make use of equations (1.4)-(1.6) as well as other other results for option pricing from Duffie *et al* (2000). The slight variation of regime switching jumps (which jump back instead of mean-reverting) can also be incorporated into the AJD framework by first creating a discrete jump size distribution and then using an indicator state variable for each different jump size, as is discussed by Culot *et al* (2006). Although potentially useful for some markets, we do not provide further details, since, having tested various MLE and moment matching techniques, we conclude that the methods do not fit the PJM and NEPOOL margin data very well.

3.5 Demand over Capacity - Invariant approach

Our bid stack model assumes throughout that power prices depend on demand and capacity (or margin) only through the ratio D_t/C_t . This suggests firstly the possibility of reducing our dimensions by one with a model for D_t/C_t directly. However, finding a process to capture both the dynamics of demand and capacity with a single factor, while also ensuring that $0 < D_t/C_t < 1$, is probably impossible. On the other hand, fitting the invariant distribution of D/C while ignoring its dynamics is much more feasible. Furthermore, the mean-reversion rates of Y_t and Z_t^{OU} suggest that both demand and margin fluctuate about seasonal levels only over short horizons of hours, days or at most weeks. In other words, the conditional distribution of D_T/C_T rapidly approaches its invariant distribution, as T increases. Consequently, for most power derivatives (long-term forwards or options on forwards), we can safely price contracts without knowledge of current values of D_t and C_t , or even their dynamics.




Figure 3.7: Histograms of D_t/C_t^{imp} for PJM (RT) over the period June 03 - May 06: an aggregation of all months (left graph) followed by a comparison of August and November (right graph).

Figure 3.7 illustrates the historical invariant distribution of D_t/C_t for PJM. When aggregating all hourly data points together (as in the left graph), we obtain a distribution which is remarkably

close to a triangle distribution centered around 0.6. The right graph indicates that this triangular shape shifts from left to right depending on the season, and in particular the shift between August and November is approximately 0.1. However these particular months represent the highest and lowest points in seasonal cycle, so it is much less pronounced than the cycle for D_t . For PJM there is a clear summer peak, but no clear winter peak, as we would expect from the seasonal parameter estimates for Z_t^{OU} given in Table 3.3 earlier. As discussed, the seasonality of capacity somewhat offsets the seasonality of demand.

While a non-parametric approach is possible (e.g., simulating by drawing independent samples from historical data), we propose instead fitting a symmetric triangle distribution with endpoints $a(T)$ and $b(T)$ ($0 < a(T) < b(T) < 1$, and peak at $c(T) = \frac{a(T)+b(T)}{2}$) to the seasonal invariant distribution of D_T/C_T , meaning that it depends on T but not on current values of D_t or C_t . So for long enough time to maturity $T - t$,

$$P_t \left\{ \frac{D_T}{C_T} < x \right\} \approx P \left\{ \frac{D_T}{C_T} < x \right\} = \begin{cases} 0 & \text{if } 0 < x \leq a(T) \\ \frac{2(x-a(T))^2}{(b(T)-a(T))^2} & \text{if } a(T) \leq x \leq \frac{a(T)+b(T)}{2} \\ 1 - \frac{2(b(T)-x)^2}{(b(T)-a(T))^2} & \text{if } \frac{a(T)+b(T)}{2} \leq x \leq b(T) \\ 1 & \text{if } b(T) \leq x < 1 \end{cases}$$

where $P_t\{\dots\}$ indicates a conditional probability given time t information (the filtration generated by all factors), and $P\{\dots\}$ is an unconditional probability (as is used throughout this thesis, and also for the expectation operator $\mathbb{E}$).




Figure 3.8: Unconditional and conditional distributions of $\tilde{D}_t/\tilde{C}_t$ for PJM (RT) (left graph) and for NEPOOL (DA) (right graph) for August or November maturity using model parameters. Here conditional distributions (1 week and 2 week) are calculated assuming initial conditions $Y_t = \mu_Y - 0.2$ and $Z_t^{\text{OU}} = \mu_Z + 0.2$, such that $\tilde{D}_t/\tilde{C}_t$ is somewhat below its invariant mean.

The distribution above provides an alternative to using the invariant distribution of D_T/C_T produced by the processes for demand and margin in Sections 3.3 and 3.4. Although this distribution has no simple form, it can be generated numerically either by simulation or by numerical integration

techniques using the following calculation:

$$\begin{aligned}
P_t \left\{ \frac{D_T}{C_T} < x \right\} &\approx P \left\{ \frac{D_T}{C_T} < x \right\} \\
&= P \{ D_T \leq D_T x + M_T x \} \\
&= P \{ D_T(1-x) \leq M_T x \} \\
&= P \{ f(T) + Y_T + \log(1-x) \leq \log(M_T) + \log(x) \} \\
&= (1-p_i)P \left\{ Y_T - Z_T^{\text{OU}} \leq \log \left[\frac{x}{1-x} \right] - f(T) \right\} + p_i P \left\{ Y_T \leq Z_T^{\text{SP}} + \log \left[\frac{x}{1-x} \right] - f(T) \right\} \\
&= (1-p_i)\Phi_1 \left(\log \left[\frac{x}{1-x} \right] - f(T) \right) + p_i \int_0^\infty \Phi_2 \left(\alpha - z + \log \left[\frac{x}{1-x} \right] - f(T) \right) \lambda_i e^{-\lambda_i z} dz,
\end{aligned}$$

where i is the season corresponding to time T and where $\Phi_1(\cdot)$ is the Gaussian cumulative distribution function (cdf) for $N \left(\mu_Y - \mu_Z, \frac{\sigma_Y^2}{2\kappa_Y} + \frac{\sigma_Z^2}{2\kappa_Z} - \frac{2\rho\sigma_Y\sigma_Z}{\kappa_Y + \kappa_Z} \right)$, the unconditional distribution of $Y_T - Z_T^{\text{OU}}$, and $\Phi_2(\cdot)$ is the Gaussian cdf for $N \left(\mu_Y, \frac{\sigma_Y^2}{2\kappa_Y} \right)$, the unconditional distribution of Y_T . Note that the fifth line above follows from conditioning on M_T lying in each of the possible regimes.

Figure 3.8 illustrates the conditional distributions (for $Y_t = \mu_Y - 0.2$, $Z_t^{\text{OU}} = \mu_Z + 0.2$ and $T - t$ equal to one week and two weeks) and unconditional distribution for PJM (RT) and NEPOOL (DA) with the parameters from Tables 3.2 and 3.3. As in Figure 3.7, the red lines correspond to November (autumn season, $i = 4$) while the blue lines correspond to August (summer season, $i = 3$). We can see clearly that the unconditional distribution rapidly approaches the invariant distribution due to the high speeds of mean reversion $\kappa_Y = 64.2$ and $\kappa_Z = 133.6$ for PJM, and similarly for NEPOOL. The two-week distribution is barely distinguishable from the invariant distribution. However, we should also note that for PJM the invariant distribution does not match particularly well with the triangular distribution observed in Figure 3.7, especially in the right tail, and especially for August. This is due to our calibration approach for M_t which involves the calculation of M_t^{mod} for this region, and implicitly corrects for errors in fitting the tails of the generator bids. Noting that Figure 3.7 plots D_t/C_t (not $\tilde{D}_t/\tilde{C}_t$), the far right tail of the original August distribution (Figure 3.7) has essentially been redistributed into the final bins of the model's distribution (Figure 3.8) such that large and frequent summer spikes are accurately captured. This explains the increase in the histogram for $\tilde{D}_t/\tilde{C}_t$ above about 0.9. Therefore we note that the use of triangle distributions may not be sufficiently accurate unless the tail of the bid stack is better captured.¹⁰ For PJM, it will underestimate the magnitude of spikes in the peak summer months. On other hand, for NEPOOL we don't face much of a problem as historical implied margin never becomes quite so low, due in part to the use of day-ahead (instead of real-time) data, and the lower frequency of spikes relative to PJM.

Despite this disadvantage, an obvious advantage of using this approach is that we do not have to worry about finding appropriate stochastic processes to describe D_t and M_t . In many markets this is in itself a challenging task, and even here, the method of correlating D_t and M_t is perhaps a relatively weak point in the model. Although using the invariant distribution is inappropriate when

¹⁰For example, as mentioned in Section 2.5.5, addition of a uniform distribution to the two logistic distributions shows promising results for improving our fit to the highest bids. However, we do not elaborate on this alternative here as we focus mainly on the more tractable case of two logistics, combined with the processes for D_t and M_t introduced in Sections 3.3 and 3.4.

simulating a path of S_t , the value of medium to long-term forwards (or options on these forwards) can reasonably be assumed independent of current values of demand and margin. We then only require current fuel prices and the invariant distribution of $\tilde{D}/\tilde{C}$ (or alternatively $\tilde{D}$ and $\tilde{M}$) to price these claims.¹¹

3.6 Empirical Results for Power Prices - NEPOOL and PJM

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, as it is particularly important in power markets for a model to capture the forward curve accurately. We make this comparison 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 2.2. For the PJM gas distribution, we use the regression results for $\{\alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G\}$ 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 $\{\alpha_0^C, \alpha_1^C, \beta_0^C, \beta_1^C\}$ over the most recent period June 2005 to July 2007. For NEPOOL, we use the parameters $\{\alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G\}$ from regression over the entire period March 2003 to August 2007. For the gas process G_t , we use parameters in Table 3.1 for January 2000 - December 2006. Finally, for demand and margin, we use the parameters in Tables 3.2 and 3.3, noting that forward prices (taken from NYMEX fact sheets) are settled using real-time prices for PJM, but day-ahead prices for NEPOOL.

3.6.1 Power Price Simulations

Tables 3.4 and 3.5 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

¹¹Pirrong (2006) reaches 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. However, neither coal prices nor margin are incorporated as factors in his approach.

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, as are other statistics to a lesser extent. As discussed, part of the problem lies in the non-stationarity of gas prices, while other issues such as growth and development of the market may also play a part. 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 .¹² 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 appropriate 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 NEPOOL, 04-07 for PJM).

Table 3.4: 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

¹²These 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 for the mean-reverting component X_0^1 .

Table 3.5: 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

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 the historical path. 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 3.4 and 3.5 underscore the difficulty in comparing statistics such as moments of price distributions, as these tend to vary significantly for different observed time periods (and for real-time vs day-ahead prices), as well as for different starting values of simulations.

3.6.2 Forward Curve Analysis

We now assess the performance of the model in terms of capturing the power forward curve. In the one-fuel case we can use (3.14) for $F^P(t, T)$ under the assumptions discussed above the expression. We can choose either to calculate $F^G(t, T)$ based on our model from (3.2) or instead use observed Henry Hub gas forwards.¹³ 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_1 can be allowed to be time dependent to reproduce observed gas forwards, as discussed in Section 3.1.2.

The left column of Figure 3.9 compares the model's two-year forward curve with the observed two-year peak forward curve for NEPOOL for the chosen historical dates.¹⁴ As expected, the use of observed gas forward prices improves the fit of the model, lifting the model's power forward curves

¹³Note that even when the observed gas forward curve is not matched exactly, we do calculate the gas model's seasonality function $g(t)$ from the observed forward curve for that date.

¹⁴The 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.




Figure 3.9: Sample forward curves for NEPOOL (left column) and PJM (right column), corresponding to the dates 30 December 2005 (top), 31 March 2006 (middle) and 29 Sept 2006 (bottom). The solid blue lines are observed prices. The other lines are model results, with the solid red lines corresponding to reproducing the observed gas forward curve, and the dotted red lines corresponding to using the model's gas forward curve.

slightly. This suggests either higher future risk premiums or higher future expected gas prices, compared with the historical period used to calibrate the model. Overall, the results are encouraging, as both the seasonality and the upwards or downwards trend in the power forward curve are fairly well captured. Note that the seasonality is a combination of gas price seasonality $g(t)$ (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 Jan 2005, and to a lesser extent in Feb 2007. Upwards trends (e.g. bottom graph) and downwards trends (e.g. top 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 3.5 is only about \$68 for 2005-07, while the NEPOOL forward curves in Figure 3.9 generally lie above this level.

The right column of Figure 3.9 shows the forward curves for PJM for the same dates. 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's and November's tend to equal December's. Observed data reveal 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).

3.7 Discussion

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 Chapters 2 and 3, 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, as we shall now investigate. While we tend to choose components of the model with either the most convenient mathematical expressions, or the best fit for PJM and NEPOOL, it is important to remember that the bid stack model is in fact a broad and flexible framework. Many different modelling choices can be made for the underlying factors or for the distributions of bids, while retaining the overall approach. Therefore the methodology can certainly be tailored to markets other than the US ones studied. In addition, the model can and should also be tailored to different modelling goals, such as risk management

versus derivative pricing, short-term versus long-term analysis or peak versus baseload power. Here we have chosen for simplicity to primarily focus on peak power prices, and in the coming chapters we shall concentrate on derivative pricing, often for medium to long maturities. As suggested by the striking relationships presented earlier in Figure 2.1 (as well as Tables 3.4 and 3.5), it is at these timescales that electricity price dynamics reveal their great dependence on fuel prices.

Thus, 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 e.g. Culot *et al* (2006) or Kluge (2006)) that multiple factors with differing speeds of mean reversion are necessary. In our model, this conclusion can be drawn directly from the parameter estimates in Tables 3.1, 3.2 and 3.3, 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 speed. 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, as we shall now discuss.

Chapter 4

Pricing Power Derivatives - One Fuel Case

In Chapters 2 and 3, we proposed a fundamental and structural approach to modelling peak hourly spot electricity prices by analysing both the behaviour of underlying price drivers and the corresponding movement of generator bids. We introduced a model for the bid stack in which distributions are fit by maximum likelihood to bids from generators of each fuel type. In this and the following chapters, we shall focus on applications of this model to the pricing of derivatives common in energy markets. We shall always be working with rescaled demand and margin variables, but suppress the tilde in the notation D_t and M_t for simplicity. Furthermore, as we no longer need to refer to a process for capacity available, we instead use the notation C_t to represent the coal price process (avoiding the confusion of using $F^P(t, T)$ for power forwards, and P_t for coal spot prices). We either refer to D_t and M_t separately when needed, or more often to a new variable $R_t = D_t/(D_t + M_t)$, representing the ratio of demand to capacity.

We choose to use the logistic bid stack model throughout, and in this chapter we consider only the one-fuel case (the model for New England), where a simple and explicit formula exists for S_t :

$$S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t) (\log(D_t) - \log(M_t)), \quad (4.1)$$

or

$$S_t = \alpha_0^G + \alpha_1^G G_t + (\beta_0^G + \beta_1^G G_t) (\log(R_t) - \log(1 - R_t)).$$

As a reminder, natural gas prices are given by the lognormal model

$$\begin{aligned} dX_t^1 &= \kappa(\mu_1 - X_t^1)dt + \sigma_1 dW_t^1, \\ dX_t^2 &= \mu_2 dt + \sigma_2 dW_t^2, \\ dW_t^1 dW_t^2 &= \rho_{12} dt \\ G_t &= \exp(g(t) + X_t^1 + X_t^2), \end{aligned}$$

where $g(t)$ is a seasonality function. X_t^1 and X_t^2 primarily drive the short-term and long-term behaviour of gas prices respectively.

Coal prices are also lognormal and described by the one-factor model

$$\begin{aligned} dX_t^3 &= \mu_3 dt + \sigma_3 dW_t^3, \\ dW_t^1 dW_t^3 &= \rho_{13} dt \\ dW_t^2 dW_t^3 &= \rho_{23} dt \\ C_t &= \exp(X_t^3). \end{aligned}$$

Next, we model (rescaled) power demand (or load) using a seasonality function combined with an OU process:

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

Finally, to capture short-term outages and recoveries, we choose for (rescaled) log-margin a regime-switching process with ‘normal’ and ‘spike’ regimes:

$$\begin{aligned} \log(M_t) &= \begin{cases} Z_t^{\text{OU}} & \text{with probability } 1 - p_i \\ Z_t^{\text{SP}} & \text{with probability } p_i \end{cases} \\ \text{where } dZ_t^{\text{OU}} &= \kappa_Z (\mu_Z - Z_t^{\text{OU}}) dt + \sigma_Z d\tilde{B}_t, \quad dB_t d\tilde{B}_t = \rho dt, \\ \text{and } Z_t^{\text{SP}} &= \alpha - J, \quad J \sim \text{Exp}(\lambda_i), \quad \text{for seasons } i = 1, 2, 3, 4. \end{aligned}$$

While the equations above describe appropriate choices for the PJM and NEPOOL markets we calibrate the model to, many of the derivative pricing results presented below can easily be adapted to other processes, particularly to different models for D_t and M_t .

In the one-fuel case, we benefit from a simple expression for forward power prices in terms of forward gas prices of the same maturity (ignoring delivery periods by choosing the midpoint of each month). From (4.1), we have

$$\begin{aligned} F^P(t, T) &= E_t^{\mathbb{Q}}[S_T] \\ &= E_t^{\mathbb{Q}}[\alpha_0^G + \alpha_1^G G_T + (\beta_0^G + \beta_1^G G_T) (\log D_T - \log M_T)] \\ &= \alpha_0^G + \alpha_1^G F^G(t, T) + (\beta_0^G + \beta_1^G F^G(t, T)) \left\{ f(T) + Y_t e^{-\kappa_Y(T-t)} + \mu_Y \left(1 - e^{-\kappa_Y(T-t)} \right) \right. \\ &\quad \left. - (1 - p_i) \left(Z_t^{\text{OU}} e^{-\kappa_Z(T-t)} + \mu_Z \left(1 - e^{-\kappa_Z(T-t)} \right) \right) - p_i \left(\alpha - \frac{1}{\lambda_i} \right) \right\}, \end{aligned} \quad (4.2)$$

for seasons $i = 1, 2, 3, 4$, where $F^G(t, T)$ is the forward gas price for the same maturity, and given by (3.2). We can choose either to calculate $F^G(t, T)$ based on our model or to instead use observed Henry Hub gas forwards. Using the observed gas forward curve allows us to better evaluate the bid stack model without the influence of model error in the fuel price model. Calibrating the model to observed forward curves of course means that both methods will be identical.

4.1 Calibration to the Power Forward Curve

As a result of (4.2), calibration to the power forward curve can be achieved using a method analogous to that of gas forward curve calibration in Section 3.1.2. Again, we allow a time-dependent

mean-reversion level, but now for Y_t . Note that here the parameters estimated earlier by maximum likelihood correspond to the physical probability measure $\mathbb{P}$, whereas our pricing equation (4.2) is under the appropriate risk-neutral measure $\mathbb{Q}$ corresponding to our chosen calibration method. Consequently, the process of choosing $\mu_Y(T)$ for $T > t$ can be understood as a means of identifying a risk-neutral measure which correctly prices power forwards. As discussed before, technical conditions for the existence of a unique measure $\mathbb{Q}$ (such as the hedging of intra-delivery period risk) are ignored in favour of a robust and applicable calibration method. In addition to the previous comments for gas curve calibration, we should also now note that we are fitting a time-dependent market price of demand risk, but essentially setting the market price of margin risk to be zero. Therefore, this is clearly not a unique choice of $\mathbb{Q}$ but is a reasonable assumption for our purposes. Power prices depend only on demand and margin through the ratio $R_t = D_t/(D_t + M_t)$, and the distribution of this ratio will be shifted upwards or downwards regardless of which of the factors $Y_t, Z_t^{\text{OU}}, Z_t^{\text{SP}}$ we attach the risk premium to. A more complete assessment of the dynamics of each process under $\mathbb{Q}$ would require additional information such as prices of derivatives on demand or margin alone.¹

The forward power curve is exactly reproduced by the model for all maturities $T_1, T_2, \dots$ if and only if

$$\mu_Y(T_1) = \frac{\log F^P(t, T_1) - \alpha_0^G + \alpha_1^G F^G(t, T_1) + (\beta_0^G + \beta_1^G F^G(t, T_1)) \left(f(T_1) + Y_t e^{-\kappa_Y(T_1-t)} - \mathbb{E}_t^{\mathbb{Q}}[\log(M_{T_1})] \right)}{(\beta_0^G + \beta_1^G F^G(t, T_1)) (1 - e^{-\kappa_Y(T_1-t)})},$$

where for all T

$$\mathbb{E}_t^{\mathbb{Q}}[\log(M_T)] = (1 - p_i) \left(Z_t^{\text{OU}} e^{-\kappa_Z(T-t)} + \mu_Z \left(1 - e^{-\kappa_Z(T-t)} \right) \right) + p_i \left(\alpha - \frac{1}{\lambda_i} \right),$$

and for $k \geq 2$ (no sum term for $k = 2$),

$$\begin{aligned} \mu_Y(T_k) = & \frac{1}{(\beta_0^G + \beta_1^G F^G(t, T_k)) (1 - e^{-\kappa_Y \Delta T})} \left\{ \log F^P(t, T_k) - \alpha_0^G + \alpha_1^G F^G(t, T_k) \right. \\ & + (\beta_0^G + \beta_1^G F^G(t, T_k)) \left(f(T_k) + Y_t e^{-\kappa_Y(T_k-t)} - \mathbb{E}_t^{\mathbb{Q}}[\log(M_{T_k})] \right. \\ & \left. \left. - \mu_Y(T_1) e^{-(k-1)\kappa \Delta T} (1 - e^{-\kappa(T_1-t)}) - \sum_{j=2}^{k-1} \mu_Y(T_j) e^{-(j-1)\kappa \Delta T} (1 - e^{-\kappa \Delta T}) \right) \right\}. \end{aligned}$$

The calibration approach described above can be applied to any forward curve $F^P(t, T_k)$, $k = 1, 2, \dots$, and there exists a corresponding unique set of mean-reversion levels $\mu_Y(T_k)$, $k = 1, 2, \dots$. This is a simple consequence of the fact that $F^P(t, T_k)$ is strictly increasing and unbounded (both above and below) in $\mu_Y(T_k)$ (since $B(x)$ is strictly increasing and unbounded in $x = \frac{D}{D+M}$, itself strictly increasing and unbounded in μ_Y). Note that we implicitly assume here that $\mathbb{E}[s_1] = \beta_0^G + \beta_1^G F^G(t, T_k) > 0, \forall T_k$, as the model produces nonsensical downward sloping bid stack functions $B(\cdot)$ for $s_1 < 0$. For NEPOOL, this is never an issue since all of our parameter estimates satisfy $\alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G > 0$, ensuring both m_1 and s_1 are positive for any $G_t \geq 0$, as required.

¹In PJM, products exist such as FTRs (Financial Transmission Rights), which hedge against price uncertainty caused by transmission losses and constraints, and thus depend on only M_t . Alternatively, we could consider weather derivatives that depend only on temperature, which is typically assumed to be the main factor driving D_t .

4.2 Approximation for Longer Maturities

Equation (4.1) reveals that S_t has a probability distribution created from products of normally and lognormally distributed random variables, as well as an exponential distribution from Z_t^{SP} . Similarly forward prices $F^P(t, T)$ in (4.2) depend on products of normal and lognormal random variables. Therefore, the pricing of both options on forwards and options on spot prices does not lead to simple closed-form expressions. Numerical methods are typically required for pricing purposes. However, the observation that driving factors operate on very different timescales can lead to useful results when we are interested in medium to long term maturities. In particular, as seen in Chapter 2, the state variables $\{Y_t, Z_t^{\text{OU}}, Z_t^{\text{SP}}\}$ relating to demand and capacity changes revert to their stationary distributions over periods of a few days or weeks, while fuel-related variables $\{X_t^1, X_t^2, X_t^3\}$ are either non-stationary or have slow speeds of mean-reversion. As in (4.2) above, forward power prices $F^P(t, T)$ can be written as linear functions of forward gas prices $F^G(t, T)$:

$$F^P(t, T) = \mathbb{E}_t^{\mathbb{Q}} [\alpha_0^G + \alpha_1^G G_T + (\beta_0^G + \beta_1^G G_T) (\ln(R_T) - \ln(1 - R_T))] \quad (4.3)$$

$$= A_0(t, T, Y_t, Z_t^{\text{OU}}) + A_1(t, T, Y_t, Z_t^{\text{OU}}) F^G(t, T) \quad (4.4)$$

Now, for fast mean-reverting $\frac{D_t}{D_t + M_t}$ and for large enough time to maturity $T - t$, the coefficients A_0 and A_1 are no longer stochastic processes but depend only on T :

$$A_0(D_t, M_t, t, T) \approx \bar{A}_0(T) \quad \text{and} \quad A_1(D_t, M_t, t, T) \approx \bar{A}_1(T), \quad \text{where} \quad (4.5)$$

$$\bar{A}_i(T) = \alpha_i^G + \beta_i^G \mathbb{E}^{\mathbb{Q}} \left[\log \left(\frac{D_T}{D_T + M_T} \right) - \log \left(1 - \frac{D_T}{D_T + M_T} \right) \right], \quad \text{for } i = 0, 1$$

$$\text{or } \bar{A}_i(T) = \alpha_i^G + \beta_i^G \mathbb{E}^{\mathbb{Q}} [\log(D_T) - \log(M_T)], \quad \text{for } i = 0, 1.$$

Hence the expectations conditional on time- t information become unconditional expectations. In other words, forward prices do not depend on the current values of demand D_t or margin M_t . Instead the distribution of $R_T = \frac{D_T}{D_T + M_T}$ depends (seasonally) on maturity T only.

In order to calculate these unconditional expectations $\bar{A}_0$ and $\bar{A}_1$, various approaches are possible. While the results are easily obtained from (3.10) and (3.4), we may also choose to use the historical invariant distribution of R_T directly, as discussed in Section 3.5. A symmetric triangle distribution with different parameters for each month of the year appears to be an appropriate choice and can also allow for fairly convenient calibration to power forwards. Assume (for calibration purposes) that the invariant distribution of R_T under the appropriate risk-neutral measure $\mathbb{Q}$ also has a symmetric triangular shape, but different endpoints $a(T)$ and $b(T)$. In this case, greater risk premiums in forward prices (corresponding to greater risk-aversion from buyers of forward contracts) translates to a shift of the triangle distribution towards the right. However, as $b(T) < 1$, we eventually require a narrowing of the distribution to accompany this shift to the right. Similarly, to reproduce very low forward prices (corresponding intuitively to high risk-aversion from sellers of forward contracts), we require the triangle to shift to the left, while still maintaining $0 < a(T) < b(T) < 1$.

Hence let $\varphi_1, \varphi_2 : \mathbb{R} \rightarrow (0, 1)$ be any two strictly increasing functions such that $\varphi_1(y), \varphi_2(y) \rightarrow 0$ as $y \rightarrow -\infty$ and $\varphi_1(y), \varphi_2(y) \rightarrow 1$ as $y \rightarrow \infty$, and also $\varphi_1(y) < \varphi_2(y), \forall y \in \mathbb{R}$. Then for any forward power price $F^P(t, T) \in \mathbb{R}$, there exists a unique $y \in \mathbb{R}$ such that equations (4.4) and (4.5)

exactly reproduce the observed forward price when the invariant distribution of R_T under $\mathbb{Q}$ follows a symmetric triangle distribution with endpoints $a(T) = \varphi_1(y)$ and $b(T) = \varphi_2(y)$.

This follows from the fact that the bid stack function is strictly increasing with $B(x) \rightarrow -\infty$ as $x \rightarrow 0$ and $B(x) \rightarrow \infty$ as $x \rightarrow 1$, while $a(T)$ and $b(T)$ defined as above squeeze our distribution closer and closer to the tails of the stack. More specifically, as $y \rightarrow -\infty$ (so $a(T), b(T) \rightarrow 0$, with $a(T) < b(T)$ and $c(T) = \frac{a(T)+b(T)}{2}$ as before),

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}[\log(R_T)] &= \int_a^c \log(x) \frac{4(x-a)}{(b-a)^2} dx + \int_c^b \log(x) \frac{4(b-x)}{(b-a)^2} dx \\ &= \frac{4}{(b-a)^2} \left\{ c^2 \log(c) - \frac{a^2}{2} \log(a) - \frac{b^2}{2} \log(b) + \frac{1}{4}(a^2 + b^2 - 2c^2) - \frac{a}{c} - \frac{b}{c} + 2 \right\} \\ &\rightarrow -\infty \quad \text{due to } \log \text{ terms.}\end{aligned}$$

while $\mathbb{E}^{\mathbb{Q}}[\log(1-R_T)]$ is finite. Similarly, as $y \rightarrow \infty$, $\mathbb{E}^{\mathbb{Q}}[\log(1-R_T)] \rightarrow -\infty$ while $\mathbb{E}^{\mathbb{Q}}[\log(R_T)]$ is finite. (With a symmetric triangle distribution X , $\mathbb{E}[X] = \mathbb{E}[1-\hat{X}]$, where $\hat{X}$ is also symmetric triangular but with endpoints $\hat{a} = 1-b$ and $\hat{b} = 1-a$.) Finally, uniqueness of y is guaranteed by the fact that $F^P(t, T)$ is strictly increasing² in $\mathbb{E}^{\mathbb{Q}}[\log(R_T)]$ which is strictly increasing in $a(T)$ and $b(T)$, while strictly decreasing in $\mathbb{E}^{\mathbb{Q}}[\log(1-R_T)]$ which is strictly decreasing in $a(T)$ and $b(T)$.

An obvious choice of functions for $\varphi_1(y)$ and $\varphi_2(y)$ is the cdf of a probability distribution with support $\mathbb{R}$. We propose to again use the logistic distribution with scale parameter 1 and mean chosen such that $a(T, y) = a_0$ and $b(T, y) = b_0$, where a_0 and b_0 are the endpoints of the seasonal invariant distribution under the physical measure $\mathbb{P}$ (fit to historical data). Therefore

$$a(T, y) = \frac{1}{1 + \exp \left\{ - \left(y - \log \left[\frac{1}{a_0} - 1 \right] \right) \right\}}, \quad b(T, y) = \frac{1}{1 + \exp \left\{ - \left(y - \log \left[\frac{1}{b_0} - 1 \right] \right) \right\}}. \quad (4.6)$$

The calibration method discussed above clearly relies on the strong and rather arbitrary assumption that the symmetric triangular distribution under $\mathbb{P}$ leads to a shifted symmetric triangular distribution under $\mathbb{Q}$. This type of assumption is necessary to extend the model's uses from pure simulation and risk management (under $\mathbb{P}$) to forward and option pricing. Note that while the assumption seems sensible, it is only one of many ways of capturing the market's risk preferences given the observed data. As mentioned before, the triangle-based approach does not allow us to compensate for errors stemming from the truncation of the bid data, so the right tail of our distribution for S_t under $\mathbb{P}$ (and thus also under $\mathbb{Q}$) is likely to be less accurate than in the standard SDE-based approach.

4.3 Option Pricing Techniques

While both options on spot prices and options on forward contracts exist in electricity markets, the latter are more widely traded and for example quoted on NYMEX. The relationship derived above in (4.4) can be used to provide simple relationships between options on power forwards and options on gas forwards for long enough maturity forward contracts. Letting T_o equal the maturity of the

²with the same implicit assumption that $s_1 > 0$

option, T_f equal the maturity of the underlying forward contract and K equal the option's strike, call options on power forwards $V_t^P(T_o, T_f, K)$ are related to call options on gas forwards $V_t^G(T_o, T_f, K)$ by the following:

$$\begin{aligned}
V_t^P(T_o, T_f, K) &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(F^P(T_o, T_f) - K, 0)] \\
&= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(\bar{A}_0(T_f) + \bar{A}_1(T_f) F^G(T_o, T_f) - K, 0)] \\
&= \bar{A}_1(T_f) e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} \left[\max \left(F^G(T_o, T_f) - \frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)}, 0 \right) \right] \\
&= \bar{A}_1(T_f) V_t^G \left(T_o, T_f, \frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)} \right). \tag{4.7}
\end{aligned}$$

The corresponding relationship for put options is also given by

$$V_t^P(T_o, T_f, K) = \bar{A}_1(T_f) V_t^G \left(T_o, T_f, \frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)} \right).$$

Similarly, spark spread options $V_t^{SS}(T_o, T_f, H, K)$ (which depend on both power forwards and gas forwards) with heat rate H and strike K are given by

$$\begin{aligned}
V_t^{SS}(T_o, T_f, H, K) &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(F^P(T_o, T_f) - H F^G(T_o, T_f) - K, 0)] \\
&= (\bar{A}_1(T_f) - H) V_t^G \left(T_o, T_f, \frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f) - H} \right).
\end{aligned}$$

Finally, calendar spread options on power $V_t^{CS,P}$ can be written in terms of calendar spread options on gas $V_t^{CS,G}$ as follows in the case of a general payoff $\max(K_0 F^P(T_o, T_3) - K_1 F^P(T_o, T_f) - K_2, 0)$ with parameters K_0 , K_1 and K_2 :

$$\begin{aligned}
V_t^{CS,P}(T_o, T_f, T_3, K_0, K_1, K_2) &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(K_0 F^P(T_o, T_f) - K_1 F^G(T_o, T_3) - K_2, 0)] \\
&= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(K_0 \bar{A}_1(T_f) F^G(T_o, T_f) - K_1 \bar{A}_1(T_3) F^G(T_o, T_3) + K_0 \bar{A}_0(T_f) - K_1 \bar{A}_1(T_3) - K_2, 0)] \\
&= V_t^{CS,G}(T_o, T_f, T_3, K_0 \bar{A}_1(T_f), K_1 \bar{A}_1(T_3), K_2 - K_0 \bar{A}_0(T_f) + K_1 \bar{A}_1(T_3)).
\end{aligned}$$

Note that all of the expressions above assume $\bar{A}_1(T_f) > 0$, meaning that power forwards are increasing in gas forwards as we would expect. While this is not guaranteed by the model, typical values of α_1^G and β_1^G make this very likely unless expected values of R_{T_f} are very low, again highly unlikely especially for peak prices.³ Similar expressions exist for $\bar{A}_1(T_f) < 0$.

Note that the above results also hold for American options, which is useful as most options in the power markets are American, not European. This can be deduced for (4.7) by observing the fact that holding $\bar{A}_1(T_f)$ units of American gas options with strike $\frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)}$ exactly replicates both the final payoff and the early exercise payoffs of the American power option in all possible states and at all times. Hence the optimal exercise times of these options will be identical, and their prices given by the relationship above. The same argument holds for all the linear relationships exploited above.

³The anomalous case of power forwards decreasing in gas forwards (despite $\alpha_1^G, \beta_1^G > 0$) results from the widening of the left tail of the distribution as G_t increases. If demand over capacity is concentrated in the far left tail, then this widening may offset the upwards shift in the mean, leading to lower prices. While this is an unrealistic effect of the model, it is also very unlikely to impact results, and is therefore not a significant weakness.

For the results above, it is important to note that our assumptions rely on a medium to long maturity for the forward contract at the option's maturity. So $T_f - T_o$ should not be small, while the time to option maturity $T_o - t$ is not important. For contracts typically traded on NYMEX, options mature only about a week before the beginning of the delivery period of the underlying monthly forward. Nonetheless, this approximation is still quite effective, as we shall investigate in the empirical results of Section 4.4. For options on very short forward contracts (at maturity) or options on spot prices, the simple linear power to gas relationship is inappropriate. However, by independence G_{T_o} from other factors, we may write power option prices as the integral over a range of gas option prices, using the conditional density of D_{T_o} and M_{T_o} (or R_{T_o}) to weight these. Moreover, if $T_o - t$ is of medium to long duration, option prices should still not depend on the current values of the fast mean-reverting processes D_t and M_t , and so the conditional density can be replaced by the invariant (but seasonal) density.

Firstly, for the case of very short maturity forwards, we can use (4.2) to obtain

$$\begin{aligned} V_t^P(T_o, T_f, K) &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(F^P(T_o, T_f) - K, 0)] \\ &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(A_0(T_o, T_f, Y_{T_o}, Z_{T_o}^{\text{OU}}) + A_1(T_o, T_f, Y_{T_o}, Z_{T_o}^{\text{OU}}) F^G(T_o, T_f) - K, 0)] \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} A_1(T_o, T_f, y, z) V_t^G\left(T_o, T_f, \frac{K - A_0(T_o, T_f, y, z)}{A_1(T_o, T_f, y, z)}\right) f_{Y,Z}(y, z) dy dz, \end{aligned}$$

where

$$\begin{aligned} A_j(T_o, T_f, Y_{T_o}, Z_{T_o}^{\text{OU}}) &= \alpha_j^G + \beta_j^G \left\{ f(T_f) + Y_{T_o} e^{-\kappa_Y(T_f-T_o)} + \mu_Y \left(1 - e^{-\kappa_Y(T_f-T_o)}\right) \right. \\ &\quad \left. - (1 - p_i) \left(Z_{T_o}^{\text{OU}} e^{-\kappa_Z(T_f-T_o)} + \mu_Z \left(1 - e^{-\kappa_Z(T_f-T_o)}\right) \right) - p_i \left(\alpha - \frac{1}{\lambda_i} \right) \right\}, \end{aligned}$$

for $j = 0, 1$, where $i \in \{1, \dots, 4\}$ corresponds to the season in which the forward contract matures, and $f_{Y,Z}(y, z)$ is the bivariate Gaussian joint probability distribution (conditional on time t values) of $\{Y_{T_o}, Z_{T_o}^{\text{OU}}\}$.

Secondly, for options on spot prices a similar formula exists (here we reuse notation V^P , A_0 , etc, slightly differently in each case):

$$\begin{aligned} V_t^P(T, K) &= e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(S_T - K, 0)] \\ &= e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(A_0(D_T, M_T) + A_1(D_T, M_T) G_T - K, 0)] \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} A_1(e^{f(T)+y}, e^z) V_t^G\left(T, \frac{K - A_0(e^{f(T)+y}, e^z)}{A_1(e^{f(T)+y}, e^z)}\right) \tilde{f}_{Y,Z}(y, z) dy dz, \end{aligned}$$

where

$$A_j(d, m) = \alpha_j^G + \beta_j^G (\log(d) - \log(m)), \quad \text{for } j = 0, 1,$$

$V_t^G(T, K)$ is the price of an option on gas spot prices, and $\tilde{f}_{Y,Z}(y, z)$ is the joint probability distribution (conditional on time t values) of $\{Y_T, \log(M_T)\}$. Note that $\tilde{f}_{Y,Z}(y, z)$ is therefore not the same as $f_{Y,Z}(y, z)$ in the equation above, and in particular is no longer bivariate Gaussian since we have yet to separate out the densities for the two regimes of margin M_t . The advantage of having serially and spatially independent shocks to the margin process is that we can separate this problem

into two integrals, effectively creating a separate pricing problem for each regime:

$$\begin{aligned} V_t^P(T, K) = & (1 - p_i) \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} A_1(e^{f(T)+y}, e^z) V_t^G \left(T, \frac{K - A_0(e^{f(T)+y}, e^z)}{A_1(e^{f(T)+y}, e^z)} \right) f_{Y,Z}(y, z) dy dz, \\ & + p_i \int_0^{\infty} \left(\int_{-\infty}^{\infty} A_1(e^{f(T)+y}, e^z) V_t^G \left(T, \frac{K - A_0(e^{f(T)+y}, e^z)}{A_1(e^{f(T)+y}, e^z)} \right) f_Y(y) dy \right) f_{\alpha_i, \lambda_i}(z) dz, \end{aligned}$$

where $A_j(d, m)$, $j = 0, 1$ and $f_{Y,Z}(y, z)$ (with $T_o = T$) are as before, $f_Y(y)$ is the marginal probability distribution (conditional on the time t value) of Y_t , and $f_{\alpha_i, \lambda_i}(z)$ is the (unconditional) probability distribution function for Z_T^{SP} . In the more general case of a transition matrix governing regime switches, the extra calculation needed is simply the transition probabilities between times t and T .

Alternatively, for large enough $T - t$, we can again make use of the invariant triangle distribution for $D_t/(D_T + M_T)$ and reduce the problem to a one-dimensional integral:

$$V_t^P(T, K) = \int_{a(T)}^{b(T)} A_1(x) V_t^G \left(T, \frac{K - A_0(x)}{A_1(x)} \right) f_X^{\text{tri}}(x) dx,$$

where

$$A_j(x) = \alpha_j^G + \beta_j^G (\log(x) - \log(1 - x)), \quad \text{for } j = 0, 1,$$

and $f_X^{\text{tri}}(x)$ is the invariant density of $D_T/(D_T + M_T)$, fitted by a symmetric triangle distribution with endpoints $a(T)$ and $b(T)$.

More generally, for any derivative $U_t^P(T, K)$ on power prices with payoff of the form either $h(S_T)$ or $h(F^P(T_o, T_f))$, we can represent its value as the weighted integral over a suitable range of gas option prices (on either spot or forward gas prices), where weights are proportional to the second derivative of $h(\cdot)$ as follows:

$$U_t^P(T, K) = \int_{a(T)}^{b(T)} \left(\int_0^{\infty} A_1^2(x) h''(A_0(x) + A_1(x)K) V^G(T, K) dK \right) f_X^{\text{tri}}(x) dx$$

with $A_j(x)$, $j = 0, 1$ as above.

A key advantage of all of the expressions above (assuming a liquid market of gas options as in the US) is that we can automatically capture the volatility smiles and term structure of volatility in the natural gas market. Without using the observed gas option prices directly, it may be difficult to incorporate this additional information into the pricing of power derivatives.

4.4 Empirical Results

For the New England market, we have collected from daily NYMEX fact sheets a large quantity of forward and option data over the period 2006-2008. In particular, we have 723 daily observations from January 2006 to Nov 2008 (labeled $n = 1, \dots, 723$) of monthly NEPOOL peak forward curves (typically up to about 3 years) and 399 daily observations from January 2007 to July 2008 of options on forwards (for dates $251 \leq 649$). All of the options have maturity about one week before

the beginning of the monthly delivery period of the forward contract. Thus, $T_f - T_o \approx 22$ days in all cases. The estimated mean-reversion rates $\kappa_Y = 132$ and $\kappa_Z = 76$ (from Tables 3.2 and 3.3), imply that over a 22 day horizon, the conditional expectations of both demand (through Y_t) and margin (through Z_t^{OU}) have reverted to their (seasonal) mean level by 98% or more. Figure 3.8 illustrated this same feature for a shorter two-week horizon. Thus the impact of D_{T_o} and M_{T_o} on the distribution of $D_{T_f}/(D_{T_f} + M_{T_f})$ is minimal, and the use the invariant distribution for pricing forwards as in (4.5) or options as in (4.7) seems reasonable, as we shall confirm with numerical examples. However, we should remember that 22 days corresponds to the midpoint of the delivery period, so in fact a forward contract will have some small dependence on $M_{T_o+7/365}$, the conditional expectation of which will only have reverted to its mean by 70%. For all of the results here, we use gas parameters from Table 4.1 estimated over the data range ending on January 1st 2006. As parameter estimates are fairly stable, the general patterns in the results remain true using any of the estimation periods.

4.4.1 NEPOOL - Results for Forwards

Firstly we assess forward risk premiums defined by the difference between observed power forwards and the model's forward prices (following calibration to the gas forward curve). This definition of forward risk premiums (excluding the step of gas curve calibration) is common in the electricity literature (see e.g. Bunn and Karakatsani (2005) or Cartea and Figueroa (2008)), and can be interpreted either as a measure of error in the model (if we expect low market prices of risk for demand over capacity), or simply as an indication of risk aversion levels in the market. Figure 4.1 illustrates that this risk premium is almost always positive, with market prices lying above the model's forward prices (where dynamics under $\mathbb{P}$ are used for demand and margin, as discussed in Section 3.5.1). The surface plot shows that there is clear seasonality in these forward risk premiums, with the peak seasons of winter and summer leading to the highest values. In order to look past the seasonality and assess whether short or long maturity contracts typically carry higher risk premiums, we average over annual periods corresponding to short ($< 1\text{yr}$), medium (1-2yr) and long (2-3yr) maturity forwards. The results indicate no clear dependence of forward risk premium on time to maturity, as the ordering varies significantly through 2006-08. However there is some suggestion that periods of spikes in the gas market (such as early-mid 2008) tend to raise the risk premium attached to short contracts more than for long contracts. This is intuitive considering the hedging needs of utilities and their concerns that energy price spikes will persist in the near term.

4.4.2 NEPOOL - Results for Options

Secondly, we move on to the analysis of options on these monthly forwards. Figure 4.2 shows that NEPOOL options only began trading on NYMEX gradually through 2007 before stabilising at around 400 contracts, whereas the larger PJM market has grown to as many as 1400 daily prices, and data are available throughout 2006-08. Note that both of these numbers are small in comparison to the 8000 European gas option prices available daily by end of 2008. For each of the NEPOOL contracts (and each observation date n), we can evaluate its price in several ways:

1. Using simulation of forward prices $F^P(T_o, T_f)$ (via simulation of $F^G(T_o, T_f)$, Y_{T_o} and $Z_{T_o}^{\text{OU}}$).




Figure 4.1: Surface of forward risk premiums for different observation dates n and maturities m (left); Progression of average forward risk premiums for short, middle and long end of the curve (right).

We have a choice between using equation (4.2) or ignoring Y_{T_o} and $Z_{T_o}^{OU}$ and using equation (4.5) for the invariant distribution case. In addition we can choose between averaging over the delivery period, or ignoring it by using the midpoint of the month.

2. Using equation (4.7) and the Black-Scholes like solution (3.3) for $V^G(\cdot, \cdot)$.
3. Using equation (4.7) and observed gas option prices for $V^G(\cdot, \cdot)$. Here we take advantage of the large number of gas options available to find one with strike near $\frac{K - \bar{A}_0(T_f)}{A_1(T_f)}$. We interpolate linearly between the prices of available options with strikes just above and just below this value. In only a very few cases, $\frac{K - \bar{A}_0(T_f)}{A_1(T_f)}$ falls outside of range the available strikes, in which case we skip this particular option.




Figure 4.2: Left graph shows the growth of trade in option contracts in NEPOOL and PJM, while right graph gives a sample of NEPOOL option price trajectories, (observed data vs the two possible model prices) for the following contracts: a call with strike \$86 on a Dec 09 forward (upper lines); a put with strike \$80 on a Dec 08 forward (lower lines)

First we observe in Figure 4.2 the price dynamics over the year 2007 of a pair of options, one put and one call. The lines labeled ‘market’ are observed prices, while ‘model’ gives model prices using the standard approach (method 2 above) and ‘gas opt’ gives model prices when using observed gas option prices to find $V^P(t, T_o, T_f)$ (method 3 above). The upper lines representing the call option all stay reasonably close together throughout the one year period, with the ‘gas opt’ case appearing to perform better overall and particularly at times of high call prices. For the put prices in the lines below, there is a more significant gap noticeable between market prices and model prices, particularly at times further from maturity. It suggests that the model tends to underprice puts far from maturity but that this gap narrows near maturity. One consideration here is that these options on power forwards are American options, not European, so there should be some premium related to early exercise rights which should mostly impact puts. However, surprisingly, the observed prices of American and European puts on gas forwards show very little difference, and the use of American gas options with method 3 above leads to little improvement.⁴ If however, there is a longer term component of demand, capacity or any other factor which has not been fully captured, this could increase the early exercise premium of power options relative to gas options.

We use these same two sample options to investigate the convergence of the simulation approaches mentioned above (method 1). In particular we aim to further justify the use of the invariant distribution discussed already in Section 3.5 and at the beginning of Section 4.4. Table 4.1 illustrates the results of pricing these two options on Aug 20th, 2007 (date chosen arbitrarily) by four different simulation methods, as well as by the explicit solutions of methods 2 and 3. Moving along the columns from left to right, the number of simulations N^{sim} increases. The first four rows of the table represent the four choices of simulation approach, either using the conditional or invariant distribution of demand over capacity, and either correctly modelling delivery periods or not. In all cases, the same random numbers are used to obtain the four different simulated prices. The second row corresponding to using the invariant distribution and no delivery period must converge to the fifth row, the closed-form price using equation (4.7) and (3.3). Indeed, in both cases this is confirmed by the numbers with an error of about \$0.01 for 10,000 simulations. However, the more important conclusions to make from this table are that the price differences between the four different simulations are very low throughout, and particularly low for high N^{sim} . In particular, considering the third row of the table as the most accurate simulations, we can argue that for 10,000 or more simulations, the error caused by using the invariant distribution is typically only around \$0.001. Similarly we can claim that the error caused by ignoring delivery periods is typically less than \$0.02. Of course these errors do not converge to zero as $N^{\text{sim}} \rightarrow \infty$, because they correspond to changes to the assumptions of either the demand and margin processes or to the option contract. Furthermore, we have only looked at two particular options. However, Table 4.1 serves to give us confidence in the two key approximations made in order to use the approach of method 2 above. Thus we no longer resort to simulation for the remainder of the results presented here.

⁴This is confirmed by calculations of the average percentage difference between European and American options over the year 2007. Using all daily observations of all options priced above 10 cents for which both American and European versions exist, we find that American calls are on average 0.55% higher than their European counterpart, while for puts this increases to 1.04%. (Looking at only the second half of 2007, these numbers actually drop to 0.24% and 0.61% respectively.)

Table 4.1: Option pricing results from Aug 20th, 2007 for sample put and call

	Sample Put			Sample Call		
Number of simulations $N^{\text{sim}} =$	100	1,000	10,000	100	1,000	10,000
Sims - cond dist; no delivery	5.9195	5.7924	6.3574	12.5787	13.2298	13.2452
Sims - inv dist; no delivery	5.9126	5.7908	6.3571	12.5737	13.2332	13.245
Sims - cond dist; with delivery	5.9535	5.8106	6.3727	12.5759	13.2036	13.2286
Sims - inv dist; with delivery	5.9256	5.8043	6.3714	12.5564	13.2163	13.2278
Method 2 - standard model	6.3474	6.3474	6.3474	13.1328	13.1328	13.1328
Method 3 - using gas options	6.7771	6.7771	6.7771	12.2262	12.2262	12.2262
Observed market prices	7.42	7.42	7.42	12.72	12.72	12.72

Of course these two sample options are not sufficient evidence on which to evaluate the model. We now analyse the error (both in absolute and relative or percentage terms) of all options observed for different maturities, strikes and observation dates.⁵ We group our results for each observation date n by time to maturity, month of maturity and by the ratio of strike to current forward price (a measure of ‘moneyness’). Specifically, we use the following categorisation by moneyness:

- Deep In-the-Money (DITM): For calls, $\frac{K}{F^G(t, T_f)} \leq 0.84$. For puts, $\frac{K}{F^G(t, T_f)} > 1.17$.
- In-the-Money (ITM): For calls, $0.84 < \frac{K}{F^G(t, T_f)} \leq 0.96$. For puts, $1.05 < \frac{K}{F^G(t, T_f)} \leq 1.17$.
- At-the-Money (ATM): For both calls and puts, $0.96 < \frac{K}{F^G(t, T_f)} \leq 1.05$.
- Out-of-the-Money (OTM): For calls, $1.05 < \frac{K}{F^G(t, T_f)} \leq 1.17$. For puts, $0.84 < \frac{K}{F^G(t, T_f)} \leq 0.96$.
- Deep Out-of-the-Money (DOTM): For calls, $\frac{K}{F^G(t, T_f)} > 1.17$. For puts, $\frac{K}{F^G(t, T_f)} \leq 0.84$.

Figure 4.3 illustrates our results, shown as time series over the entire dataset of option prices. The first row of graphs includes only ATM options, with errors measured firstly in percentage terms, and secondly in dollar terms. In both cases, we average over the absolute value of all options (of any maturity) which fall into each category. While it is difficult to establish clear patterns, we can conclude that the model tends to price calls better than puts, and errors seem to be lowest in late 2007 before increasing again, perhaps due to the sudden volatility and price surges witnessed in energy markets. As well as the early exercise issue discussed above, we could also attribute the weaker pricing of puts to the fact that the bid stack model (as we have calibrated it) generally places more emphasis on correctly modelling the upper portion of the bid stack. For the lower portion, we have failed to capture the zero bids correctly as well as the 10% of coal bids which may be increasingly relevant with rising coal prices. Another observation we can make is that for most dates the use of gas options to price power options through (4.7) improves results, as we might expect since it should give us a better indication of the true (non-lognormal) distribution of gas forwards. Finally, the top left figure also includes two lines for the average relative error of all observed gas options, which can only be compared to power when using relative terms. Interestingly, relative errors in pricing power options tend to be lower than for gas options, particular in the case of using observed gas

⁵Note that we exclude options with market prices less than or equal to 10 cents, as these are typically illiquid and difficult to price accurately. They also have a disproportionately large impact on error estimates when calculated in relative (percent) terms.

options and (4.7). This suggests that a significant portion of the error in the power model stems from the gas model, and that by using observed gas option prices, we effectively eliminate this error source.




Figure 4.3: Absolute and relative errors of option prices using the standard model (method 2 above) and when using gas options (method 3 above): Top row gives relative (left) and absolute (right) errors for at-the-money options. Left graph also compares with relative errors of pricing gas options. Bottom row gives absolute errors for all levels of moneyness, namely deep in-the-money (DITM), in-the-money (ITM), at-the-money (ATM), out-of-the-money (OTM), and deep out-of-the-money (DOTM). The bottom left graph uses only call options.

The second row of graphs in Figure 4.3 illustrates that for ITM and OTM options (and particularly DITM and DOTM options), errors are more volatile. However, much of this volatility can be attributed to the fact that our sample of options is constantly shifting, with options jumping between categories from date to date. Thus, one of the difficulties of the above analysis is of course the varying number of options in each category for each date. For example, on some dates, most short maturity contracts may be in the money with longer maturities mostly out of the money, while on other dates the situation may be reversed. In order to separate these effects further, as well as to identify the direction of the pricing errors (i.e., overpricing or underpricing) we can instead look at implied volatility curves or even surfaces. However, unlike for the typical geometric Brownian Motion model for stocks, we now have multiple volatility parameters which we could choose between, namely $\sigma_1, \sigma_2, \sigma_Y, \sigma_Z$. Since $\rho_{12} > 0$, both the mean $\tilde{\mu}_G$ and the standard deviation $\tilde{\sigma}_G$ of

$\log F^G(T_o, T_f)$ in (3.3) are strictly increasing in both σ_1 and σ_2 , and so call prices must also be. For put prices, this is not the case, and for large enough values of σ_1 or σ_2 , $V^G(t, T_o, T_f)$ is decreasing in volatility. However, this is remedied by allowing $\mu_1(T)$ to vary as σ_1 or σ_2 vary, such that we always maintain $\mathbb{E}^\mathbb{Q}[F^G(T_o, T_f)] = F^G(t, T_f)$, the observed forward gas price. In other words, define implied volatility $\sigma_1^{\text{imp,G}}(t, T_o, T_f, K)$ by

$$\sigma_1^{\text{imp,G}}(t, T_o, T_f, K) = \sigma \quad \text{such that} \quad V_{\text{obs}}^G(t, T_o, T_f; K) = V_{\text{mod}}^G(t, T_o, T_f; K, \sigma),$$

where

$$V_{\text{mod}}^G(t, T_o, T_f; K, \sigma) = e^{-r(T_o-t)} \left[e^{\tilde{\mu}_G + \frac{1}{2}\tilde{\sigma}_G^2} \Phi\left(\frac{-\log(K) + \tilde{\mu}_G + \tilde{\sigma}_G^2}{\tilde{\sigma}_G}\right) - K \Phi\left(\frac{-\log(K) + \tilde{\mu}_G}{\tilde{\sigma}_G}\right) \right],$$

with

$$\begin{aligned} \tilde{\sigma}_G^2 &= \frac{\sigma^2}{2\kappa} \left(1 - e^{-2\kappa(T_o-t)}\right) e^{-2\kappa(T_f-T_o)} + \sigma_2^2(T_o-t) + \frac{2\rho_{12}\sigma\sigma_2}{\kappa} \left(1 - e^{-\kappa(T_o-t)}\right) e^{-\kappa(T_f-T_o)} \\ \tilde{\mu}_G &= \log(F^G(t, T_f)) - \frac{1}{2}\tilde{\sigma}_G^2, \end{aligned}$$

and where $V_{\text{obs}}^G(t, T_o, T_f; K)$ are the market observed gas option prices. (The expressions above are identical to (3.3) in the case that $\sigma = \sigma_1$, but otherwise require an additional shift in μ_1 .) While a similar implied volatility expression could be written for σ_2 , we find that using σ_1 performs better, particularly for shorter maturities.

For the purpose of comparing with power options, we can also define implied volatility using observed power option prices, but still in terms of σ_1 , the volatility of the mean-reverting component of gas, X_t^1 . We denote this new version of implied volatility by $\sigma_1^{\text{imp,P}}$ and set

$$\sigma_1^{\text{imp,P}}(t, T_o, T_f, K) = \sigma \quad \text{such that} \quad V_{\text{obs}}^P(t, T_o, T_f; K) = V_{\text{mod}}^P(t, T_o, T_f; K, \sigma),$$

where, as in (4.7),

$$V_{\text{mod}}^P(t, T_o, T_f; K, \sigma) = \bar{A}_1(T_f) V_{\text{mod}}^G\left(T_o, T_f, \frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)}\right),$$

where $V_{\text{obs}}^P(t, T_o, T_f; K)$ are the market observed electricity option prices and V_{mod}^G is given by the expression above. Clearly V_{mod}^P is strictly increasing in σ_1 and implied volatility is uniquely defined by the above formula.

Interestingly, since $\bar{A}_i(T) = \alpha_i^G + \beta_i^G \mathbb{E}^\mathbb{Q}[\log(D_T) - \log(M_T)]$, (for $i = 0, 1$) from (4.5), and $\log(\cdot)$ is a concave function, $\bar{A}_0(T)$ and $\bar{A}_1(T)$ must both be strictly decreasing in σ_Z and increasing in σ_Y . The same is true of power forwards. Hence $\frac{K - \bar{A}_0(T_f)}{\bar{A}_1(T_f)}$ is strictly increasing in σ_Z and decreasing in σ_Y , and so call options on power forwards are strictly decreasing in σ_Z and increasing in σ_Y . These results may be somewhat surprising, particularly for σ_Y , but reflect the fact that the lower part of the bid stack is in fact concave while the upper part is convex. Furthermore, an increase in demand volatility widens the right tail of the distribution of R_{T_f} slightly more than the left, while an increase in margin volatility does the opposite.




Figure 4.4: Progression of average implied volatilities over time for different levels of moneyness: in-the-money, at-the-money and out-of-the-money for power options (left), and gas options (right).

The implied volatility results are shown in Figure 4.4, using power options (left graph) to determine $\sigma_1^{\text{imp,P}}$ as well as gas options (right graph) to find $\sigma_1^{\text{imp,P}}$. Recall that $\sigma_1 = 0.566$ is our estimated constant volatility. These graphs reveal again that puts are generally underpriced by the model, particularly relative to calls. They also demonstrate that market prices of options have decreased over the period 2007-2008 relative to our estimated parameters. Though at first surprising, we should recall the observation near Table 3.1 that parameters estimated for December 2005 lead to relatively high values of gas volatility due to the spike in gas prices at the time. This would suggest that the decreases in the two graphs of Figure 4.4 are linked to each other. There is some evidence for this hypothesis, although the decrease in σ_1^{imp} is greater for power and the precise shape of the two graphs does not match as well as we might expect. Another plausible explanation could be a greater illiquidity in power options in the early years, as 2007 marked the beginning of exchange traded NEPOOL options. Finally, another interesting point to discuss is the existence of implied volatility smiles or skews, which can be gauged at least roughly from the distance between the ITM, ATM and OTM lines above. These are much closer together for gas than power options, suggesting that a lognormal fit is better for gas forwards than for power forwards. However, the shape of the power option smile, skew (or frown) is unstable through time, with different categories leading to the highest implied volatility in different cases. This can again be due in part to the varying maturities of options in each category at each date. In conclusion, the overall results for NEPOOL are fairly reasonable, and suggest that the bid stack model can become a useful tool both for forward and option pricing.

Chapter 5

Pricing Power Derivatives - Multi-Fuel Case

We now consider the (logistic) bid stack model with two or more fuel types and thus have, in general,

$$S_t = x \quad \text{such that} \quad B_t^{-1}(x) = \frac{D_t}{D_t + M_t} = R_t,$$

where

$$B_t^{-1}(x) = \sum_{i=1}^N \left\{ \frac{w_i}{2} + \frac{w_i}{2} \tanh \left(\frac{x - m_i(\text{fuel prices})}{2s_i(\text{fuel prices})} \right) \right\}, \quad (5.1)$$

or in the two-fuel case ($N = 2$) for the PJM market (coal and gas driven),

$$B_t^{-1}(x) = \frac{1}{2} + \frac{w_1}{2} \tanh \left(\frac{x - (\alpha_0^C + \alpha_1^C C_t)}{2(\beta_0^C + \beta_1^C C_t)} \right) + \frac{1 - w_1}{2} \tanh \left(\frac{x - (\alpha_0^G + \alpha_1^G G_t)}{2(\beta_0^G + \beta_1^G G_t)} \right). \quad (5.2)$$

No matter which of the bid distributions we choose, the convex combination of cumulative bids $w_1 F_1(x) + (1 - w_1) F_2(x)$ from different fuel types is not explicitly invertible, and hence $B(\cdot)$ cannot be written explicitly for the case of two or more fuel types. The pricing of any electricity derivative therefore requires the use of numerical methods and may be computationally intensive as we have at least five state variables involved. For PJM with stochastic coal prices we have a total of six: $\{X_t^1, X_t^2, X_t^3, Y_t, Z_t^{\text{OU}}, Z_t^{\text{SP}}\}$. However, noting the independence of fuel prices from demand and margin changes, as well as the independence of the spike regime for margin, we can simplify the computation, similarly to the one-fuel case. For any derivative $V_t(T)$ with maturity T and payoff $h(S_T)$ (writing $B(\cdot)$ now as a function of all underlying factors),

$$\begin{aligned} V_t(T) &= e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}}[h(S_T)] \\ &= e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}}[h(B(C_T, G_T, D_T, M_T))] \\ &= e^{-r(T-t)} \int_0^\infty \int_0^\infty \mathbb{E}_t^{\mathbb{Q}}[h(B(c, g, D_T, M_T))] f_{C,G}(c, g) dc dg \\ &= e^{-r(T-t)} \iiint_{-\infty}^\infty \mathbb{E}_t^{\mathbb{Q}}[h(B(e^{x_3}, e^{g(T)+x_1+x_2}, D_T, M_T))] f_X(\mathbf{x}) d\mathbf{x} \\ &= e^{-r(T-t)} \iiint_{-\infty}^\infty \left\{ (1 - p_i) \iint_{-\infty}^\infty h(B(e^{x_3}, e^{g(T)+x_1+x_2}, e^{f(T)+y}, e^z)) f_{Y,Z}(y, z) dy dz \right. \\ &\quad \left. + p_i \int_0^\infty \left(\int_{-\infty}^\infty h(B(e^{x_3}, e^{g(T)+x_1+x_2}, e^{f(T)+y}, e^z)) f_{Y,Z}(y, z) dy dz \right) f_{\alpha_i, \lambda_i}(z) dz \right\} f_X(\mathbf{x}) d\mathbf{x} \end{aligned} \quad (5.3)$$

where $f_{C,G}(c, g)$ is the joint probability distribution (conditional on time t values) of coal and gas prices, $f_X(\mathbf{x})$ is the joint probability distribution (conditional on time t values) of fuel price variables $\{X_t^1, X_t^2, X_t^3\}$, $f_{Y,Z}(y, z)$ is the joint probability distribution (conditional on time t values) of $\{Y_t, Z_t^{\text{OU}}\}$, $f_Y(y)$ is the marginal probability distribution (conditional on the time t value) of Y_t , and $f_{\alpha_i, \lambda_i}(z)$ is the (unconditional) probability distribution function for Z_t^{SP} . Note that $f_X(\mathbf{x})$ is a trivariate Gaussian density, $f_{Y,Z}(y, z)$ is a bivariate Gaussian density, $f_Y(y)$ is a univariate Gaussian density and $f_{\alpha_i, \lambda_i}(z)$ is a shifted exponential density with parameters corresponding to the season of our maturity T .

Alternatively, we can of course price any derivative using a PDE approach, again splitting the problem by conditioning on knowledge of the regime of margin at time T . Since the problem is easily split between the short term factors (demand and margin), and the long term Gaussian process for fuel prices, a combination of finite differences and Gaussian quadrature, similar to the suggestion of Pirrong and Jermakyan (2005), can be used to decrease computation time.

In addition, for large enough time to maturity $T - t$ (more than a week or so), the argument presented earlier allows us to replace $f_{Y,Z}(y, z)$ and $f_Y(y)$ with the appropriate (maturity T) unconditional distribution of the state variables. For each maturity T , we can therefore calculate $\mathbb{E}^Q[h(B(c, g, D_T, M_T)) | C_T = c, G_T = g]$ for a range of c and g , and reuse these when re-pricing a derivative on different observation dates, or for compound derivatives such as options on forwards. However, the independence of $R_T = D_T/(D_T + M_T)$ (from D_t , M_t as well as fuel prices) is only a very small benefit compared with the large computational expense of the pricing of most power derivatives, including for example the numerical inversion of the bid stack for every simulation or for every point (C_T, G_T, R_T) .¹ Moreover, as for NEPOOL, we typically would like to calibrate to the power forward curve before pricing other derivatives, especially for options on forwards. Calibration itself is a complicated and expensive inverse problem procedure, which we aim to avoid by the approximation technique of the following section.

5.1 Bid Stack Approximation - non-interchanging fuel types

Clearly the lack of explicit expression for the bid stack is a significant disadvantage of the bid stack model for two or more fuels. Therefore we now propose an approximation to the shape of the bid stack for the case of two logistic distributions (denoted $B(x)$ and given by (5.2)), based on piecing together two one-fuel bid stacks. While this approach may appear rather complicated or inelegant, it ultimately leads to an expression for the two-fuel case similar in form (if not simplicity) to that of (4.4) and is hence very useful for calibration and pricing purposes. In particular, we firstly obtain a

¹A three or four dimensional look-up table (combined with interpolation between grid points) is one possible alternative to repeated inversion of the stack within each simulation. However, if we use a three dimensional table (with R_T as a dimension, instead of D_T and M_T) we must be careful to employ fine enough grid spacing in the upper region of the bid stack, as very small changes in R_T can lead to large changes in electricity prices. The decision to create the three dimensional look-up table is of course dependent on the particular pricing problem, as well as our tolerance to errors created by linear interpolation.

relationship between S_t and its underlying factors of the form

$$S_t = \sum_i \psi_i(R_t) \Psi_i(C_t, G_t).$$

For all i , ψ_i is an indicator function multiplied by either a polynomial or a logarithmic function. More importantly, for all i , Ψ_i is either of the form $C_t^\eta G_t^\xi$ (for some $\eta, \xi \in \mathbb{R}$), or simply the product of two of the following functions: C_t , G_t , $\log C_t$, $\log G_t$. As a result, we can then obtain in closed form a very similar relationship between forward prices as between spot prices:

$$F^P(t, T) = \sum_j \bar{A}_j(T) \Psi_j(F^C(t, T), F^G(t, T)),$$

where (averaging over the unconditional distribution of R_T) we have ‘coefficients’

$$\bar{A}_j(T) = c_0 + \sum_i c_i \mathbb{E}^\mathbb{Q}[\psi_i(R_T)] \quad \text{for some constants } c_i,$$

multiplying the same functions Ψ_j . (The use of indices i and j , as well as constants c_i , simply reflects the fact that by taking expectations of S_t there is a change in the weighting of functions Ψ_j , although overall the expressions are very similar in form.)

The expression above is convenient firstly because it separates the dependence of $F^P(t, T)$ on its short and long-term drivers, and secondly because the functions of the long-term drivers are simple enough to obtain closed-form solutions in terms of fuel forward prices. Furthermore, though we focus on two fuels for simplicity and for calibration to the PJM market, this approximation technique can be extended to as many fuel types as we would like. However, the important restriction is that the approximation is only valid for certain ranges of fuel prices, (or equivalently bid parameters $\{m_i, s_i : i = 1, 2, \dots\}$) and in particular only while the current merit order of fuel types is preserved. For example, if there is a significant probability that over our time period of interest the majority of coal bids will move above the natural gas bids in PJM, then the approach is no longer reasonable. Thus, it is most suitable for markets where there is a clear and stable ordering by price of bids from different types of generators.²

We begin the approximation with a function $B_0(x)$ which combines two one-fuel bid stacks for $x = R_t$ over the ranges $x \in (0, w_1)$ and $x \in (w_1, 1)$:

$$B_0(x) = \begin{cases} m_1 + s_1 (\log(x) - \log(w_1 - x)) & \text{if } 0 < x < w_1 \\ m_2 + s_2 (\log(x - w_1) - \log(1 - x)) & \text{if } w_1 < x < 1. \end{cases}$$

This function is clearly unreasonable as a bid stack, since it goes to ∞ and $-\infty$ as we approach w_1 from the left or right respectively. Nonetheless, it captures quite accurately the shape of the bid stack in the upper and lower regions of the range $x \in (0, 1)$. In order to fix the middle, we now

²The difficulty in constructing a function which approximates $B(x)$ more generally for all fuel prices is that for one case (low C_t , high G_t) we have a double-humped function with coal driving only the lower part and gas only the upper part, while for another case (abnormally high C_t and low G_t), we could have just the opposite. In between we have the case of complete mixing of coal and gas bids which leads to a single-humped function driven by both fuels!

replace terms going to $\pm\infty$ at $x = w_1$ with a 3rd order Taylor expansion of $\log x$ about x_0 :

$$B_1(x) = \begin{cases} m_1 + s_1 \left(\log(x) - \log x_0 - \frac{1}{x_0}(w_1 - x - x_0) \right. \\ \quad \left. + \frac{1}{2x_0^2}(w_1 - x - x_0)^2 - \frac{1}{3x_0^3}(w_1 - x - x_0)^3 \right) & \text{if } 0 < x \leq w_1 \\ m_2 + s_2 \left(-\log(1-x) + \log x_0 + \frac{1}{x_0}(x - w_1 - x_0) \right. \\ \quad \left. - \frac{1}{2x_0^2}(x - w_1 - x_0)^2 + \frac{1}{3x_0^3}(x - w_1 - x_0)^3 \right) & \text{if } w_1 < x < 1. \end{cases}$$

As seen in Figure 5.1, by an appropriate choice of x_0 , the function $B_1(x)$ (red dashed line, coinciding




Figure 5.1: Analysis of two-fuel bid stack approximation approach for PJM in the case of unusually low C_t and high G_t (the right-hand graph is a zoomed-in version of the left-hand graph). Here $w_1 = \tilde{w}_1$ for all but the yellow line, $b = 0$ for all but the pink line, and $N = 3$ for all but the light blue line. Note that the red, light blue and pink lines almost exactly coincide throughout the zoomed-in graph, while for $R_t < w_1 = 0.5$ they also coincide with the case $\tilde{w}_1 = 0.52$ (yellow line).

with pink and light blue lines in the middle region of the stack) diverges just enough from the earlier $B_0(x)$ (green line) on either side of the point $x = w_1$ such that $B_1(x)$ is continuous throughout $x \in (0, 1)$. In order to ensure exact continuity at $x = w_1$, we require:

$$\begin{aligned} B_1(w_1^-) &= B_1(w_1^+) \\ m_1 + s_1 \left[\log w_1 - \log x_0 + 1 + \frac{1}{2} + \frac{1}{3} \right] &= m_2 + s_2 \left[\log x_0 - 1 - \frac{1}{2} - \frac{1}{3} - \log(1 - w_1) \right] \\ \implies x_0 &= w_1^{\frac{s_1}{s_1+s_2}} w_2^{\frac{s_2}{s_1+s_2}} \exp \left\{ \frac{11}{6} + \frac{m_1 - m_2}{s_1 + s_2} \right\}. \end{aligned} \quad (5.4)$$

As $\{m_1, s_1, m_2, s_2\}$ depend on fuel prices, x_0 is a stochastic process. For the typical range of coal prices C_t and gas prices G_t (and the bid stack parameters $\{\alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G, \alpha_0^C, \alpha_1^C, \beta_0^C, \beta_1^C\}$), the function $x_0(C_t, G_t)$ can be well approximated by

$$x_0 = \lambda C_t^\delta G_t^{-\gamma} \quad \text{for suitable constants } \lambda > 0, \delta > 0, \gamma > 0.$$

The optimal choice will vary with weight w_1 , bid stack parameters listed above, and also the range over which we search for a best fit.³ For all but the longest maturity forwards (i.e., those with the

³In the special case that coal prices are treated as deterministic instead of stochastic, we can simplify our approach by using $x_0 = \lambda G_t^{-\gamma}$, but must then recalculate the optimal λ and γ for each maturity with a different forward coal price.

widest fuel price distribution), the fit is generally very strong.⁴

The next obstacle that we face when using this approximation is the risk that the upper and lower sections of $B_1(x)$ become very steep and diverge from the exact bid stack $B(x)$. This happens in particular when there is a large amount of separation between gas and coal bids, as illustrated by as the red line in Figure 5.1 where $C_t = 60$ and $G_t = 15$. In these cases the steep section of the stack near $x = w_1$ forces a smaller value for x_0 and consequently a divergence in the expansion before we reach $x = 0$ and $x = 1$. A natural way to remedy this problem is to increase the number of terms of the Taylor Series expansion (as shown by the light blue line which uses five terms instead of three), which introduces additional terms in all the formulas.⁵ To reduce the number of terms required, we therefore propose an alternative modification where the terms $\log(w_1 - x)$ and $\log(x - w_1)$ are now replaced by linear functions in the lowest and highest sections of the bid stack (for $x \leq w_A$ and $x > w_B$), instead of being replaced by the Taylor Series. These linear functions are chosen to equal the original logarithmic function at $x = 0$ and $x = 1$, but to equal the Taylor Series expansion at $x = w_A$ and $x = w_B$, thus ensuring continuity of the bid stack. Though the proportion of the stack for which we can safely make this approximation clearly varies with x_0 , we can typically apply it to at least half of the stack. Let b equal the percentage of the stack for which we apply the approach. Then with $w_A = bw_1$ and $w_B = 1 - bw_2$, our new approximation $B_2(x)$ is

$$B_2(x) = \begin{cases} m_1 + s_1 \left(\log(x) - \log(w_1) - \frac{x}{w_A} \left[\log(x_0) + \frac{1}{x_0}(w_1 - w_A - x_0) - \frac{1}{2x_0^2}(w_1 - w_A - x_0)^2 + \frac{1}{3x_0^3}(w_1 - w_A - x_0)^3 - \log(w_1) \right] \right) & 0 < x \leq w_A \\ B_1(x) & w_A < x \leq w_1 \\ B_1(x) & w_1 < x \leq w_B \\ m_2 + s_2 \left(\log(1 - w_1) - \log(1 - x) + \frac{1-x}{1-w_B} \left[\log(x_0) + \frac{1}{x_0}(w_B - w_1 - x_0) - \frac{1}{2x_0^2}(w_B - w_1 - x_0)^2 + \frac{1}{3x_0^3}(w_B - w_1 - x_0)^3 - \log(1 - w_1) \right] \right) & w_B < x < 1. \end{cases}$$

The advantage of the approximation $B_2(x)$ is that we now need only ensure that the significant divergence in the Taylor Series expansion begins outside the range (w_A, w_B) instead of $(0, 1)$. Again, we may if necessary increase the number of terms of the expansion to achieve this criterion, but fewer terms will be necessary than for $B_1(x)$. The purple line in Figure 5.1 illustrates the improvement (relative to the red line) achieved in the approximation with $b = 0.6$ while keeping a third order Taylor expansion. Therefore, we should view increasing b as an alternative to increasing N in our approximation. When adding extra terms the choice of x_0 must of course change slightly, but conveniently the only difference in (5.4) is that the fraction $11/6$ is replaced by $\sum_{i=1}^N \frac{1}{i}$, where N is the number of terms. Hence the best fit to $x_0(C_t, G_t)$ will retain the same values for δ and γ as we increase N , while the coefficient λ will gradually increase but we need not rerun the optimisation.⁶

⁴Note that if the approximation for x_0 overestimates the exact value from (5.4), we still retain a strictly increasing bid stack, which is therefore preferable to the case of underestimating the exact x_0 .

⁵Note that by choosing an odd number of terms we can ensure that the bid stack remains strictly increasing throughout. However, for large enough N , whether N is odd or even, the expansion is increasing for the relevant range $(0, 1)$, as the divergence in the expansion occurs outside of the range $(0, 1)$. This is the more important consideration in practice.

⁶Here we assume for simplicity that the number of terms used in the expansion for the region below w_1 equals the number of terms used above w_1 . This is not a necessary condition, and indeed in markets where the price is never set in the lowest region of the stack, we may choose to increase the number of the terms only in the expansion above w_1 . However, in this case the optimisation must be repeated as N varies.

Finally, the zoomed-in image of Figure 5.1 reveals that there is typically a narrow region in the middle of the stack (for example for $x \in (w_1, w_1 + 0.05)$) for which the approximation $B_1(x)$ (and $B_2(x)$) can significantly overestimate the true price (by as much as \$10 in the fairly extreme case of Figure 5.1). As this is an important region of the stack in terms of how often it sets the spot power price, we suggest one final modification to our approximation. Again, we require a simple adjustment to retain the advantages of the approximation as we shall discuss below. Thus, we propose shifting the meeting point of the two one-fuel approximations from w_1 to $\tilde{w}_1$.⁷ (In the PJM market, tests show that $\tilde{w}_1 \approx w_1 + 0.015$ appears to be a good choice⁸ for typical fuel prices (based on historical ranges), and is shown by the yellow line in Figure 5.1. However, as we can see, the side-effect of this modification is earlier divergence at the top of the bid stack and hence more terms needed to compensate.) Our new approximation is denoted $B_3(x)$ and we now allow for two different values x_0 and $\tilde{x}_0$ for the lower and upper expansions respectively. Thus,

$$B_3(x) = \begin{cases} m_1 + s_1 \left(\log(x) - \log(w_1) - \frac{x}{w_A} \left[\log(x_0) + \frac{1}{x_0}(w_1 - w_A - x_0) - \frac{1}{2x_0^2}(w_1 - w_A - x_0)^2 + \frac{1}{3x_0^3}(w_1 - w_A - x_0)^3 - \log(w_1) \right] \right) & 0 < x \leq w_A \\ m_1 + s_1 \left(\log(x) - \log x_0 - \frac{1}{x_0}(w_1 - x - x_0) + \frac{1}{2x_0^2}(w_1 - x - x_0)^2 - \frac{1}{3x_0^3}(w_1 - x - x_0)^3 \right) & w_A < x \leq \tilde{w}_1 \\ m_2 + s_2 \left(-\log(1 - x) + \log \tilde{x}_0 + \frac{1}{\tilde{x}_0}(x - w_1 - \tilde{x}_0) - \frac{1}{2\tilde{x}_0^2}(x - w_1 - \tilde{x}_0)^2 + \frac{1}{3\tilde{x}_0^3}(x - w_1 - \tilde{x}_0)^3 \right) & \tilde{w}_1 < x \leq w_B \\ m_2 + s_2 \left(\log(1 - w_1) - \log(1 - x) + \frac{1-x}{1-w_B} \left[\log(\tilde{x}_0) + \frac{1}{\tilde{x}_0}(w_B - w_1 - \tilde{x}_0) - \frac{1}{2\tilde{x}_0^2}(w_B - w_1 - \tilde{x}_0)^2 + \frac{1}{3\tilde{x}_0^3}(w_B - w_1 - \tilde{x}_0)^3 - \log(1 - w_1) \right] \right) & w_B < x < 1 \end{cases}$$

To ensure continuity of the bid stack throughout, we now require x_0 and $\tilde{x}_0$ chosen such that $B_3(\tilde{w}_1^-) = B_3(\tilde{w}_1^+)$. Hence,

$$\begin{aligned} & m_1 + s_1 \left[\log w_1 - \log x_0 - \frac{1}{x_0}(w_1 - x - x_0) + \frac{1}{2x_0^2}(w_1 - x - x_0)^2 - \frac{1}{3x_0^3}(w_1 - x - x_0)^3 \right] \\ &= m_2 + s_2 \left[\log x_0 - \frac{1}{\tilde{x}_0}(x - w_1 - \tilde{x}_0) - \frac{1}{2\tilde{x}_0^2}(x - w_1 - \tilde{x}_0)^2 + \frac{1}{3\tilde{x}_0^3}(x - w_1 - \tilde{x}_0)^3 - \log(1 - w_1) \right]. \end{aligned} \quad (5.5)$$

As we have not included the restriction $x_0 = \tilde{x}_0$ used in $B_1(x)$, a range of solutions exists for x_0 and $\tilde{x}_0$. For the lower region of the stack, we let x_0 be given by (5.4) as before. We can then find $\tilde{x}_0$ numerically from (5.5) above, and choose appropriate parameters $\tilde{\lambda}$, $\tilde{\delta}$ and $\tilde{\gamma}$, such that $\tilde{x}_0 = \tilde{\lambda} C_t^{\tilde{\delta}} G_t^{-\tilde{\gamma}}$ gives the best possible fit for an appropriate range of fuel prices. Again, results suggest that both x_0 and $\tilde{x}_0$ can be very accurately approximated by such curves. However, as this introduces an additional numerical procedure, it is an optional improvement to the approximation which should be ignored if the error using $B_2(x)$ is tolerable.

⁷Note that w_1 retains its position in the functions $\log(x - w_1)$ and $\log(w_1 - x)$ as this is necessary to ensure the correct shape of the approximation.

⁸A good choice for $\tilde{w}_1$ can be made by minimising the maximum absolute deviation in the middle region of the stack. This is easy to find for a particular pair (C_t, G_t) , but when pricing derivatives require firstly determining a relevant region of the fuel price space as we shall discuss in detail in Section 5.4.

5.2 Pricing Power Forwards

Using the approximation $B_3(x)$ above, we have $S_t = B_3(R_t)$ and can now calculate forward electricity prices $F^P(t, T)$ in the two-fuel case as follows:

$$F^P(t, T) = \mathbb{E}_t^{\mathbb{Q}}[S_T] = E_1 + E_2 + E_3 + E_4, \quad (5.6)$$

where

$$\begin{aligned} E_1 &= \mathbb{E}_t^{\mathbb{Q}} \left[1_{R_T \leq w_A} \left\{ \alpha_0^C + \alpha_1^C C_T + (\beta_0^C + \beta_1^C C_T) \left(\log(R_T) - \log(w_1) - \frac{R_T}{w_A} \left[-\log(w_1) + \log(\lambda C_T^{\delta} G_T^{-\gamma}) \right. \right. \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{1}{\lambda} C_T^{-\delta} G_T^{\gamma} (w_1 - w_A - \lambda C_T^{\delta} G_T^{-\gamma}) - \frac{1}{2\lambda^2} C_T^{-2\delta} G_T^{2\gamma} (w_1 - w_A - \lambda C_T^{\delta} G_T^{-\gamma})^2 + \frac{1}{3\lambda^3} C_T^{-3\delta} G_T^{3\gamma} (w_1 - w_A - \lambda C_T^{\delta} G_T^{-\gamma})^3 \right] \right) \right\} \right] \\ E_2 &= \mathbb{E}_t^{\mathbb{Q}} \left[1_{w_A < R_T \leq \bar{w}_1} \left\{ \alpha_0^C + \alpha_1^C C_T + (\beta_0^C + \beta_1^C C_T) \left(\log(R_T) - \log(\lambda C_T^{\delta} G_T^{-\gamma}) - \frac{1}{\lambda} C_T^{-\delta} G_T^{\gamma} (w_1 - R_T - \lambda C_T^{\delta} G_T^{-\gamma}) \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{1}{2\lambda^2} C_T^{-2\delta} G_T^{2\gamma} (w_1 - R_T - \lambda C_T^{\delta} G_T^{-\gamma})^2 - \frac{1}{3\lambda^3} C_T^{-3\delta} G_T^{3\gamma} (w_1 - R_T - \lambda C_T^{\delta} G_T^{-\gamma})^3 \right) \right\} \right] \\ E_3 &= \mathbb{E}_t^{\mathbb{Q}} \left[1_{\bar{w}_1 < R_T \leq w_B} \left\{ \alpha_0^G + \alpha_1^G G_T + (\beta_0^G + \beta_1^G G_T) \left(\log(\tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}}) + \frac{1}{\tilde{\lambda}} C_T^{-\tilde{\delta}} G_T^{\tilde{\gamma}} (R_T - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}}) \right. \right. \right. \\ &\quad \left. \left. \left. - \frac{1}{2\tilde{\lambda}^2} C_T^{-2\tilde{\delta}} G_T^{2\tilde{\gamma}} (R_T - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}})^2 + \frac{1}{3\tilde{\lambda}^3} C_T^{-3\tilde{\delta}} G_T^{3\tilde{\gamma}} (R_T - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}})^3 - \log(1 - R_T) \right) \right\} \right] \\ E_4 &= \mathbb{E}_t^{\mathbb{Q}} \left[1_{R_T > w_B} \left\{ \alpha_0^G + \alpha_1^G G_T + (\beta_0^G + \beta_1^G G_T) \left(\log(1 - w_1) - \log(1 - R_T) + \frac{R_T - 1}{w_B - 1} \left[-\log(1 - w_1) + \log(\tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}}) \right. \right. \right. \right. \\ &\quad \left. \left. \left. + \frac{1}{\tilde{\lambda}} C_T^{-\tilde{\delta}} G_T^{\tilde{\gamma}} (w_B - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}}) - \frac{1}{2\tilde{\lambda}^2} C_T^{-2\tilde{\delta}} G_T^{2\tilde{\gamma}} (w_B - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}})^2 + \frac{1}{3\tilde{\lambda}^3} C_T^{-3\tilde{\delta}} G_T^{3\tilde{\gamma}} (w_B - w_1 - \tilde{\lambda} C_T^{\tilde{\delta}} G_T^{-\tilde{\gamma}})^3 \right] \right) \right\} \right]. \end{aligned}$$

Expanding the squared and cubed terms above, we can regroup all terms by their dependence on C_t and G_t . We then have terms consisting of expectations of $\log C_T$, $\log G_T$, $C_T \log C_T$, $G_T \log G_T$, $C_T \log G_T$, $G_T \log C_T$, and $C_T^{\eta} G_T^{\xi}$ for various pairs of $\eta, \xi \in \mathbb{R}$. Thanks to the independence of R_T with respect to fuel prices, the coefficients of all of the functions of C_T and G_T can be written as $\mathbb{E}_t^{\mathbb{Q}}[\psi_i(R_T)]$ for various functions $\psi_i(\cdot)$. Now we observe that conditional expectations of each of these expressions are available in simple closed form since $(\log C_t, \log G_t)$ is a Gaussian process (multivariate Gaussian $\forall t$). Specifically, from (3.1) and (3.4), the joint conditional distribution of log-fuel prices (after calibration to the fuel forward curves) is

$$\begin{pmatrix} \log C_T \\ \log G_T \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_C \\ \mu_G \end{pmatrix}, \begin{pmatrix} \sigma_C^2 & \rho_{CG} \sigma_C \sigma_G \\ \rho_{CG} \sigma_C \sigma_G & \sigma_G^2 \end{pmatrix} \right), \quad (5.7)$$

where

$$\begin{aligned} \mu_C &= \log F^C(t, T) - \frac{1}{2} \sigma_C^2, \\ \mu_G &= \log F^G(t, T) - \frac{1}{2} \sigma_G^2, \\ \sigma_C^2 &= \sigma_3^2(T - t), \\ \sigma_G^2 &= \frac{\sigma_1^2}{2\kappa} \left(1 - e^{-2\kappa(T-t)} \right) + \sigma_2^2(T - t) + \frac{2\rho_{12}\sigma_1\sigma_2}{\kappa} \left(1 - e^{-\kappa(T-t)} \right), \\ \rho_{CG} &= \frac{1}{\sigma_C \sigma_G} \left(\frac{\rho_{13}\sigma_1\sigma_3}{\kappa} \left(1 - e^{-\kappa(T-t)} \right) + \rho_{23}\sigma_2\sigma_3(T - t) \right). \end{aligned}$$

Therefore, we have the following solutions:

$$\begin{aligned}
\mathbb{E}_t^Q[C_T] &= \exp(\mu_C + \frac{1}{2}\sigma_C^2) &= F^C(t, T) \\
\mathbb{E}_t^Q[G_T] &= \exp(\mu_G + \frac{1}{2}\sigma_G^2) &= F^G(t, T) \\
\mathbb{E}_t^Q[\log C_T] &= \mu_C &= \log F^C(t, T) - \frac{1}{2}\sigma_C^2 \\
\mathbb{E}_t^Q[\log G_T] &= \mu_G &= \log F^G(t, T) - \frac{1}{2}\sigma_G^2 \\
\mathbb{E}_t^Q[C_T \log C_T] &= (\mu_C + \sigma_C^2)e^{\mu_C + \frac{1}{2}\sigma_C^2} &= (\log F^C(t, T) + \frac{1}{2}\sigma_C^2)F^C(t, T) \\
\mathbb{E}_t^Q[G_T \log G_T] &= (\mu_G + \sigma_G^2)e^{\mu_G + \frac{1}{2}\sigma_G^2} &= (\log F^G(t, T) + \frac{1}{2}\sigma_G^2)F^G(t, T) \\
\mathbb{E}_t^Q[C_T \log G_T] &= (\mu_G + \rho_{CG}\sigma_C\sigma_G)e^{\mu_C + \frac{1}{2}\sigma_C^2} &= (\log F^G(t, T) - \frac{1}{2}\sigma_G^2 + \rho_{CG}\sigma_C\sigma_G)F^C(t, T) \\
\mathbb{E}_t^Q[G_T \log C_T] &= (\mu_C + \rho_{CG}\sigma_C\sigma_G)e^{\mu_G + \frac{1}{2}\sigma_G^2} &= (\log F^C(t, T) - \frac{1}{2}\sigma_C^2 + \rho_{CG}\sigma_C\sigma_G)F^G(t, T),
\end{aligned}$$

and, for any $\eta, \xi \in \mathbb{R}$,

$$\begin{aligned}
\mathbb{E}_t^Q[C_T^\eta G_T^\xi] &= \exp \left\{ \eta\mu_C + \frac{\eta^2\sigma_C^2}{2} + \xi\mu_G + \frac{\xi^2\sigma_G^2}{2} + \eta\xi\rho_{CG}\sigma_C\sigma_G \right\} \\
&= (F^C(t, T))^\eta (F^G(t, T))^\xi \exp \left\{ \frac{1}{2}(\eta^2 - \eta)\sigma_C^2 + \frac{1}{2}(\xi^2 - \xi)\sigma_G^2 + \eta\xi\rho_{CG}\sigma_C\sigma_G \right\}.
\end{aligned}$$

Using the results above and the independence of R_T from both C_T and G_T , all of the terms in the long expression (5.6) for $F^P(t, T)$ above can be written in terms of forward coal and gas prices, as well as the parameters $\{\sigma_C, \sigma_G, \rho_{CG}\}$. The form of the conditional expectations listed above implies that terms in $\mathbb{E}_t^Q[\Psi(C_T, G_T)]$ always lead to terms in $\Psi(F^C(t, T), F^G(t, T))$, while sometimes producing an additional term as well (in cases involving log-fuel prices). We shall shorten the notation of forward prices to be F_t^P , F_t^C and F_t^G , for power, coal and gas respectively, with maturity T throughout. After much algebra, the expression for F_t^P reduces to

$$\begin{aligned}
F_t^P &= A_1 + A_2^C \log(F_t^C) + A_2^G \log(F_t^G) + A_3^C F_t^C + A_3^G F_t^G + A_4(F_t^C)^{-\delta}(F_t^G)^\gamma + \tilde{A}_4(F_t^C)^{-\tilde{\delta}}(F_t^G)^{\tilde{\gamma}} \\
&\quad + A_5(F_t^C)^{-2\delta}(F_t^G)^{2\gamma} + \tilde{A}_5(F_t^C)^{-2\tilde{\delta}}(F_t^G)^{2\tilde{\gamma}} + A_6(F_t^C)^{-3\delta}(F_t^G)^{3\gamma} + \tilde{A}_6(F_t^C)^{-3\tilde{\delta}}(F_t^G)^{3\tilde{\gamma}} \\
&\quad + A_7^C(F_t^C)^{-\delta+1}(F_t^G)^\gamma + \tilde{A}_7^G(F_t^C)^{-\tilde{\delta}}(F_t^G)^{\tilde{\gamma}+1} + A_8^C(F_t^C)^{-2\delta+1}(F_t^G)^{2\gamma} + \tilde{A}_8^G(F_t^C)^{-2\tilde{\delta}}(F_t^G)^{2\tilde{\gamma}+1} \\
&\quad + A_9^C(F_t^C)^{-3\delta+1}(F_t^G)^{3\gamma} + \tilde{A}_9^G(F_t^C)^{-3\tilde{\delta}}(F_t^G)^{3\tilde{\gamma}+1} \\
&\quad + A_{10}^C F_t^C \log F_t^C + A_{10}^G F_t^G \log F_t^G + A_{11} F_t^C \log F_t^G + A_{12} F_t^G \log F_t^C.
\end{aligned} \tag{5.8}$$

where the coefficients are all in the form of conditional expectations of functions of R_T , and are given in Appendix A. In particular, they require the computation of terms of the form $P^Q\{R_T < w\}$, $\mathbb{E}^Q[R_T \mathbf{1}_{R_T \leq w}]$, $\mathbb{E}^Q[\log(R_T) \mathbf{1}_{R_T \leq w}]$, $\mathbb{E}^Q[\log(1 - R_T) \mathbf{1}_{R_T > w}]$, and $\mathbb{E}^Q[(R_T - w_1)^{k-3} \mathbf{1}_{R_T \leq w}]$, for either $w = w_A$, $w = \tilde{w}_1$ or $w = w_B$.

If we require more than the third order Taylor Series expansion (to avoid the divergence near the top and bottom of the bid stack, as discussed earlier), then only fairly simple modifications are required. The new coefficients are also provided in Appendix A, along with the more general version of (5.8).

For the simpler case of deterministic coal prices, we set $\delta = 0$ and $m_1 = \alpha_0^C + \alpha_1^C F_t^C$ and $s_1 = \beta_0^C + \beta_1^C F_t^C$ to match the current coal forward price. Then for the third order Taylor expansion as in (5.8), we have instead (incorporating m_1 and s_1 into the coefficients):

$$\begin{aligned} F_t^P = & A_1 + A_2 \log(F_t^G) + A_3 F_t^G + A_4 (F_t^G)^\gamma + \tilde{A}_4 (F_t^G)^{\tilde{\gamma}} + A_5 (F_t^G)^{2\gamma} + \tilde{A}_5 (F_t^G)^{2\tilde{\gamma}} \\ & + A_6 (F_t^G)^{3\gamma} + \tilde{A}_6 (F_t^G)^{3\tilde{\gamma}} + \tilde{A}_7 (F_t^G)^{\tilde{\gamma}+1} + \tilde{A}_8 (F_t^G)^{2\tilde{\gamma}+1} + \tilde{A}_9 (F_t^G)^{3\tilde{\gamma}+1} + A_{10} F_t^G \log F_t^G. \end{aligned} \quad (5.9)$$

where $A_1, A_2, \dots$ are again given in Appendix A. Note that these expressions further simplify in the case that $\tilde{w}_1 = w_1$ (i.e., we skip the final modification to the approximation), since most terms with coefficients $\tilde{A}_i$ combine with the corresponding ones in A_i .

Although both (5.8) and (5.9) (and (A.1)) appear long and complicated, they provide a very useful relationship between power forwards and coal and gas forwards of the same maturity. Crucially, they explicitly separate the dependence on gas and coal from the dependence on demand and margin as was achieved automatically in the one-fuel case in (4.3). Therefore, for medium to long maturity forwards, we can now use the same approach as before, replacing all the conditional expectations $\mathbb{E}_t^Q[\psi_i(R_T)]$ in the coefficients listed above by unconditional expectations dependent only on T (due to seasonality) but not t . This greatly simplifies both the calculation and calibration of the power forward curve, as we shall discuss below.

5.3 Error Analysis for the Bid Stack Approximation

As mentioned above, the bid stack approximation based on Taylor Series expansions is not valid for all values of C_t and G_t . Moreover, the adjustments made between the early approximation $B_1(x)$ and the final version $B_3(x)$ are more effective for certain ranges of coal and gas price than for others. Thus the error in the best possible approximation is a function of C_t and G_t . While there are several ways of measuring this error, we consider the average absolute deviation (over the range⁹ $x \in (0, 1)$) between the exact stack $B(x)$ and the approximation $B_3(x)$ to be the most suitable measure. Thus we define the error as

$$\text{Err}^{\text{AD}}(C_t, G_t) = \int_0^1 |B_3(x; C_t, G_t) - B(x; C_t, G_t)| dx.$$

In addition, it is useful to consider the total deviation and the maximum absolute deviation of the two bid stacks $B(x)$ and $B_3(x)$, as given by

$$\text{Err}^{\text{TD}}(C_t, G_t) = \int_0^1 (B_3(x; C_t, G_t) - B(x; C_t, G_t)) dx$$

⁹Since the bid stack is defined in $(0, 1)$, the average absolute deviation is also equal to the total absolute deviation in the stack. However, for PJM, in practice we calculate the average absolute deviation of the stack only over the range $x \in (0.2, 1)$ for several reasons. Firstly, the lowest fifth of the bid stack (even after previous truncation of bids) very rarely ever sets the peak power price. Secondly, the very lowest section of the bid stack is badly represented by the model's bid stack $B(x)$ in the first place, particularly as $B(x)$ becomes negative for very small x , while the observed stack does not. Finally, since typically $s_2 \gg s_1$, the initial approximation $B_0(x)$ diverges from $B(x)$ near $x = 0$ since the contribution from the left tail of gas bids (eg. below zero) in $B(x)$ actually exceeds the contribution from the left tail of coal bids for small enough x , while $B_0(x)$ considers only coal bids. In fact, $B_0(x)$ tends to better represent $B^{\text{obs}}(x)$ since it does not become negative until even closer to $x = 0$. Fortunately, we do not have these issues in the right tail of the bid stack, again because typically $s_2 \gg s_1$.

and

$$\text{Err}^{\text{MD}}(C_t, G_t) = \max_{x \in (0,1)} \{|B_3(x; C_t, G_t) - B(x; C_t, G_t)|\}.$$

While the third of these is the strictest overall error measure (requiring a very precise fit throughout the entire bid stack), the second is the most lenient, deliberately allowing slight downwards errors in some sections of the stack to offset slight upwards errors in others (e.g., on either side of $x = \tilde{w}_1$). Bearing in mind that for pricing forwards (or other derivatives), we integrate over the bid stack, this may be a reasonable measure if the probability of D_t/C_t lying in these regions is similar. Nonetheless, overall we tend to favour $\text{Err}^{\text{AD}}(C_t, G_t)$ as the most balanced of the three error measures.




Figure 5.2: Average absolute deviation $\text{Err}^{\text{AD}}(C_t, G_t)$ (top left), maximum absolute deviation $\text{Err}^{\text{MD}}(C_t, G_t)$ (top right), and number of terms required (bottom left) for the fuel price space; Historical analysis based on fuel price movements in 2006-08 (bottom right).

Figure 5.2 illustrates how these error measures vary over the fuel price space for the ranges $C_t \in [\$13.5, \$270]$ and $G_t \in [\$2.2, \$33]$, with $b = 0$, $\tilde{w}_1 = w_1$ and all of the bid stack parameters $\{\alpha_0^C, \alpha_1^C, \beta_0^C, \beta_1^C, \alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G, w_1\}$ fixed to be the values estimated for PJM. Note that we have chosen a much larger ranger of coal and gas prices than have been observed historically, with upper endpoints around double recent record levels. The top row of graphs shows that the average absolute deviation remains below about \$3 for typical values of coal and gas prices, and reaches \$15 only for extreme cases of coal near \$250 and gas near \$3. Similarly, $\text{Err}^{\text{MD}}(C_t, G_t)$ remains below about

\$10 for the most relevant region. The lower left graph of Figure 5.2 indicates the number of terms required such that the deviation at the top of the stack (at $x = 0.995$) remains below \$5. (Note that the results are not very sensitive to the tolerance parameter \$5 here, because when the deviation occurs, it is very sudden and rapid.) Finally, the lower right graph sheds more light on the three other graphs in Figure 5.2 by finding both average and maximum absolute deviation as well as number of terms required over the historical path of fuel prices in 2006-08 (linearly interpolating between points in each surface). Recall from Figure 4.2 that over this period gas and coal prices showed significant co-movement, suggesting that the more problematic regions of the fuel price space were less relevant. As a result, the error measure $\text{Err}^{\text{AD}}(C_t, G_t)$ averages only \$1.34, while $\text{Err}^{\text{MD}}(C_t, G_t)$ averages \$9.39, and the number of terms required remains low (below 6) throughout.¹⁰ We should remember throughout that the ‘exact’ stack $B(x)$ from (5.2) is itself an approximation to the observed step function $B^{\text{obs}}(x)$, so the magnitude of deviation between these functions throughout history (in Figure 2.14) provides a useful comparison to $\text{Err}^{\text{AD}}(C_t, G_t)$ plotted here.¹¹ Clearly the approximation error typically remains well below the error in the bid stack fit itself (which averages \$3.65 as discussed in Section 2.5.6), suggesting that the Taylor Series Approximation is a reasonable approach to adopt based on likely fuel price ranges. However, for derivative pricing purposes, we must ensure that the approximation holds sufficiently well even for unlikely fuel price movements. In Appendix A, we investigate further the sources of error in $B_3(x)$, how the surface plots of Figure 5.2 vary with b and $\tilde{w}_1$, and thus how the error can be reduced. We now move on to the step of forward curve calibration and define tolerance parameters for determining when the approximation technique can be safely implemented.

5.4 Calibration to the Power Forward Curve

The error in the bid stack approximation (as a function of C_t and G_t) as defined above corresponds to the immediate application of the approach to find spot prices S_t . However, if we are interested in using the approach to price forward contracts (or to calibrate to observed forwards), we must ensure that the approximation holds (with $\text{Err}^{\text{AD}}(C_t, G_t)$ below a chosen error tolerance level ϵ) over the joint distribution of future coal and gas prices, instead of at just one point (C_t, G_t) . Of course, as these are lognormally distributed, the range of possible values is the entire positive quadrant $\mathbb{R}^+ \times \mathbb{R}^+$, so instead we must choose a quantile of the joint distribution for which to ensure validity.

Calibration Region in Fuel-Price Space

Let q be our quantile tolerance parameter, meaning that we only implement the approximation for pricing $F^P(t, T)$ if the probability of the pair (C_T, G_T) falling outside of the approximation’s range of validity is less than q . As T increases, this criterion becomes harder to satisfy as the distributions

¹⁰Interestingly, correlation of maximum deviation with coal price is very high (93.9%) while correlation of average deviation with gas price is very high (97.6%). Higher fuel prices lead to higher errors partly because of the use of lognormal processes, since increasing initial values widens distributions in the fuel price space even as they remain constant in the log fuel price space.

¹¹However, as discussed in Section 2.5.6, the comparison is most appropriate for the middle region of the bid stack, due the inevitable large deviations in the right tail of the stack. The error function to compare $B(x)$ and $B^{\text{obs}}(x)$ should be considered only for $x \in [1/4, 3/4]$.

of C_T and G_T continue to widen without bound, implying the existence of an upper bound for the maturity of forwards we can price. In addition to time to maturity $T - t$ and current fuel prices, the parameters $\{\kappa, \sigma_1, \sigma_2, \sigma_3, \rho_{12}, \rho_{13}, \rho_{23}\}$ of the coal and gas processes play an important role in determining which forwards we can price. In particular, a high positive correlation ρ_{CG} (due to high values of ρ_{13} and/or ρ_{23}) can greatly improve the applicability of the methodology, since it implies a low probability of reaching the problematic regions for the approximation. Finally, observed forward gas and coal prices $F^G(t, T)$ and $F^C(t, T)$ also play an important role, since they provide the expected values of the marginal conditional distributions of gas and coal respectively (following calibration to the fuel forward curves). So for chosen tolerance parameters ϵ and q , and for a chosen number of terms N in the Taylor expansion, we can find the upper bound M_U and lower bound M_L on the length of maturity forward contract that can be priced by such methods:¹²

$$m \in [M_L, M_U] \implies P_t^Q \{ \text{Err}^{\text{AD}}(C_{t+m}, G_{t+m}) > \epsilon \} \leq q. \quad (5.10)$$

In order to ease computation of these values, we can define two-dimensional quantiles of the joint (Gaussian) distribution of log coal and log gas prices using contours of their joint probability density function. Letting $f_{C,G}(c, g; m)$ be the joint probability density function of $c = \log C_{t+m}$ and $g = \log G_{t+m}$ (conditional on time t), define for any length of maturity m and quantile tolerance q the contour level $\theta(q, m)$ by

$$\theta(q, m) = x > 0 \quad \text{such that} \quad P_t^Q \{ f_{C,G}(c, g; m) < x \} = q.$$

For the bivariate Gaussian distribution, contours are elliptical with

$$f_{C,G}(c, g; m) > x \iff \frac{(c - \mu_C)^2}{\sigma_C^2} + \frac{(g - \mu_G)^2}{\sigma_G^2} - \frac{2\rho_{CG}(c - \mu_C)(g - \mu_G)}{\sigma_C\sigma_G} < 2(1 - \rho_{CG}^2) \log \left(2x\pi\sigma_C\sigma_G\sqrt{1 - \rho_{CG}^2} \right), \quad (5.11)$$

where the parameters $\{\mu_C, \mu_G, \sigma_C, \sigma_G, \rho_{CG}\}$ are as given below (5.7) but with $T = t + m$. Then find M_L and M_U such that

$$m \in [M_L, M_U] \implies \max_{c, g: f_{C,G}(c, g; m) > \theta(q, m)} (\text{Err}^{\text{AD}}(e^c, e^g)) < \epsilon. \quad (5.12)$$

This condition is clearly stricter on the range of acceptable maturities, so the new value of M_U from (5.12) will be less than or equal to the value defined in (5.10). One advantage of this new definition is that the contour level $\theta(q, m)$ does not depend on $F^G(t, t + m)$ or $F^C(t, t + m)$, as varying the mean of the joint density does not vary its shape. Hence, $\theta(q, m)$ can be reused for different calibration dates. Using this approach, we calculate for PJM the upper bound M_U to the length of maturity forward contract that can be priced during the historical period 2006-08 (dates chosen at weekly intervals to simplify graphs), using gas and coal parameters from Table 3.1 through end of 2007. Results are shown in Figure 5.3 for varying values of ϵ , q , ρ_{13} and σ_3 . We find that $M_L = 0$ throughout. The left graph of Figure 5.3 fixes $\epsilon = \$5$, while the right graph sets $\epsilon = \$10$. As expected,

¹²Note that while an upper bound always exists, there will only exist a lower bound M_L in cases of steepness in the short end of the gas forward curve, and particularly low or high current gas prices. Otherwise set $M_L = 0$, as seems to be the case throughout our dataset. It could also be possible that the approximation no longer holds at some maturity $m > M_U$ but then holds again at some later maturity $m' > m$ due to the shape of the coal or gas forward curves and perhaps seasonality effects. However, again, we did not encounter any such case for PJM.

higher tolerance parameters ϵ and q can significantly increase the approximation's region of validity, as described by M_U . In particular, with $\epsilon = 5$, during much of 2008 (with unusually high energy prices), the approximation could only be used for the first few months of the forward curve, though it is worth noting that 2009 data should show a return to values of M_U closer to earlier years as commodity prices have dropped. An increase to $\epsilon = 10$ is helpful but is not recommended for typical levels of error tolerance.¹³ Finally, it is also useful to note on Figure 5.3 that, by adjusting our parameter estimates for coal and gas, we can sometimes significantly extend the range of validity of the approximation. An increase in correlation from $\rho_{13} \approx \rho_{23} \approx 0.2$ to $\rho_{13} = \rho_{23} = 0.5$ can significantly increase M_U , as can a halving of coal volatility to $\sigma_3 = 0.1$. Both of these adjustments could be realistic, since coal volatility was much lower in early years, while correlation may be underestimated by our model due to the difficulty of capturing long-term co-integration effects stemming from the general dependencies in the wider energy and commodity markets (see discussion in Section 3.2.1).




Figure 5.3: Historical analysis of longest maturity M_U of forward contracts which may be priced accurately (within tolerance levels set by ϵ and q) for varying parameters q , σ_3 and ρ_{13} . In left graph, $\epsilon = 5$, while in right graph $\epsilon = 10$.

The above analysis assumes that as many terms as necessary may be included into the expansion to avoid the divergence for high G_t and low C_t . Therefore it measures the maturity at which problems in the low G_t and high C_t region breach the tolerance levels. If we are also interested in limiting the number of terms to a fixed N (e.g. $N = 3$), then our values of M_U may be significantly less. To investigate the minimum number of terms required to price the entire forward curve (with a maximum maturity M_U as defined above), we calculate the number of terms needed to satisfy the above requirement for each historical date. Thus

¹³While even $\epsilon = 5$ sounds rather large, one should remember (from (5.12)) that this represents the largest tolerated value of the average deviation over the appropriate elliptical region. Hence, for most of the likely fuel price movements, the deviation will be very small relative to this parameter. There is also a dependence on the distribution of R_T since the deviation varies through the bid stack. These complications prevent us from describing the typical error for example in a forward price $F^P(t, T)$ as a function of ϵ and q , but tests show that with $\epsilon = 5$ and $q = 0.02$, the maximum difference (typically at $m = M_U$) between $F^P(t, T)$ using N terms and using $N - 1$ terms is almost always below \$0.2 (and often more like \$0.05), which suggests that the divergence has been avoided. The results with $\epsilon = 10$ are much more erratic. Of course, slight errors are also less significant if we are first calibrating to observed forward prices and then pricing other derivatives, since the error will be compensated for by the exact calibration.

$$N^{\min} = \min \left\{ N \geq 3 \in \mathbb{N} : \max_{c, g: f_{C, G}(c, g; m) > \theta(q, m)} (\text{Err}^{\text{AD}}(e^c, e^g; N)) < \epsilon \quad \forall m \in [0, M_U] \right\}.$$




Figure 5.4: Historical analysis of minimum number of terms $N^{\min}$ needed to price forwards contracts with maturity $m \leq M_U$, for varying parameters b , q , σ_3 and ρ_{13} . In left graph, M_U is equal to values in Figure 5.3, while in right graph, we fix $M_U = 1/2$ (6 months). For dates where 6 months is beyond the valid region of the approximation, no data point is plotted.

The left-hand graph of Figure 5.4 plots $N^{\min}$ as defined above (again at weekly intervals through 2006-08), while the right-hand graph plots the number of terms required specifically for the 6 month forward contract, thus providing a more stable comparison between different lines on the graph. Both graphs clearly indicate that the number of terms $N^{\min}$ required to match the entire forward curve has decreased in recent years, but (as shown in Figure 5.3) at the expense of a smaller set of forwards which we can price. This can be attributed to the rapid growth of coal prices illustrated earlier in Figure 3.2, narrowing the gap between coal and gas bids. However, the recent increase in coal price volatility requires us to use more terms, as can be seen by comparing with $\sigma_C = 0.1$, particularly in the right graph. The plots also indicate clearly the benefit provided by using $b = 0.6$ as opposed to $b = 0$, greatly reducing the number of terms required to six or fewer throughout 2006-08. Finally, note that while varying b has relatively little impact on M_U , varying ϵ has relatively little impact on $N_{\min}$. We can now better understand all of these effects by plotting the fuel-price space itself.

The left-hand graph of Figure 5.5 illustrates the regions of the coal-gas plane which cause us difficulties, namely the upper left and lower right. In the upper left, the lines indicate the limits to the region of validity with various choices of ϵ (equal to 5 or 10). In the lower right, the lines represent the limits to the region of validity for a fixed number of terms N (also equal to 5 or 10). Finally, the curves (ellipses in log space) in the middle are contours $\theta(q, m)$ for $q = 0.02$, m equal to 3 months, 1 year and 2 years, and with forward gas and coal prices fixed at \$10 and \$100 respectively. The case of stronger positive correlation with $\rho_{13} = \rho_{23} = 0.5$ is also shown in red, and we can clearly see the advantage it provides in avoiding the problematic corners. Thus we can visualise the expanding elliptical contours $f_{C, G}(c, g; m) = \theta(q, m)$ as m increases, with the values M^U and $N^{\min}$ determined by their intersection with the lines in the upper left and lower right respectively. Note that these lines correspond to exact fitting of the function $x_0(C_t, G_t)$. To be more precise,




Figure 5.5: Historical analysis of minimum number of terms $N^{\min}$ needed to price forwards contracts with maturity $m \leq M_U$, for varying parameters b, q, σ_3 and ρ_{13} . In left graph, M_U is equal to values in Figure 5.3, while in right graph, we fix $M_U = 1/2$ (6 months). For dates where 6 months is beyond the valid region of the approximation, no data point is plotted.

as discussed in the error analysis of Appendix A, for longer maturities there may sometimes be a significant underestimation or overestimation of $x_0(C_t, G_t)$ in some regions of the fuel price space, which will further limit the overall region of validity. Setting $\nu = \hat{x}_0/x_0$, the right-hand graph of Figure 5.5 illustrates by how much the contour limits shift in the case of an overestimation by 40% or an underestimation by 20%. Interestingly, the 20% underestimation is the more troublesome case for the methodology, as the contours for number of terms as well as absolute average deviation both move significantly inwards in the graph.

Calibration Algorithm

We now move from the issue of when we can calibrate the model to reproduce observed power forwards to how this is achieved. We describe the procedure in the most general case. Firstly, fix tolerance parameters ϵ and q throughout, as well as a linear cutoff parameter b to avoid large expansions.¹⁴ For a simpler calibration method, also fix $\tilde{w}_1$ throughout, and in particular setting $\tilde{w}_1 = w_1$ (i.e., using $B_2(x)$, not $B_3(x)$) further simplifies. Now, for maturities which fall within the range of acceptability $[M_L, M_U]$ as discussed above, the power forward curve can be calibrated as follows for each maturity m :

1. After calibration to $F^C(t, t+m)$ and $F^G(t, t+m)$, identify the appropriate elliptical region $\mathcal{A} = \{(e^c, e^g) : f_{C,G}(c, g; m) > \theta(q, m)\}$ for which we need the bid stack approximation to hold.

¹⁴Technically, the optimal value of b is zero in terms of error minimisation, so the choice reflects a choice to improve computation time by reducing the number of terms in the expansion. In particular, some feasible values of (C_t, G_t) can lead to an extremely large N if $b = 0$, at which point we may need to also be careful about issues such as rounding error and the propagation of errors from our calculation of terms such as $\mathbb{E}^Q[(R_T - w_1)^n 1_{R_T \leq w}]$.

2. If we are using a non-fixed weight $\tilde{w}_1$, then find a suitable value of $\tilde{w}_1$ which minimises $\max_{(c,g) \in \mathcal{A}} \{\text{Err}^{\text{AD}}(c,g)\}$ (for the chosen b and as many terms N as necessary to avoid errors in the top and bottom of the stack).
3. Calculate the number of terms N needed to fit the top and bottom of the stack within the required tolerance level (for the chosen b). Check that $\text{Err}^{\text{AD}}(c,g) < \epsilon, \forall (c,g) \in \mathcal{A}$. (If this condition fails, then we are beyond the region of validity of the approximation, so $m \notin [M_L, M_U]$.)
4. Find values $\{\lambda, \delta, \gamma\}$ such that $x_0(C_t, G_t)$ is best approximated over the region $\mathcal{A}$. This can be done by minimising the sum of squared errors between $\hat{x}_0$ and x_0 (discretising the integral over the elliptical region). In the case that $\tilde{w}_1 = w_1$,

$$\{\lambda, \delta, \gamma\} = \underset{\{\lambda, \delta, \gamma\}}{\text{argmin}} \left\{ \int_{(c,g) \in \mathcal{A}} \left(\lambda c^\delta g^{-\gamma} - w_1^{\frac{s_1(c)}{s_1(c)+s_2(g)}} w_2^{\frac{s_2(g)}{s_1(c)+s_2(g)}} e^{\sum_{i=1}^N \frac{1}{i} + \frac{m_1(c)-m_2(g)}{s_1(c)+s_2(g)}} \right)^2 dcdg \right\},$$

while for $\tilde{w}_1 \neq w_1$, we require additional numerical techniques to first solve (5.5). Note that if $\tilde{w}_1 = w_1$, we can calculate $\{\lambda, \delta, \gamma\}$ for the region $\mathcal{A}$ before even identifying N , as the choice of N simply rescales δ and leaves λ and γ unchanged.

5. Find $\nu^{\max} = \max_{(c,g) \in \mathcal{A}} \frac{\hat{x}_0(c,g)}{x_0(c,g)}$ and $\nu^{\min} = \min_{(c,g) \in \mathcal{A}} \frac{\hat{x}_0(c,g)}{x_0(c,g)}$.
 - (a) If $\tilde{w}_1$ is fixed, then repeat Step 3 with the assumption that $\hat{x}_0$ under or overestimates x_0 by these proportions. Ensure $\text{Err}^{\text{AD}}(c,g;\nu) < \epsilon, \forall (c,g) \in \mathcal{A}$ still holds in both cases (i.e., $\nu = \nu^{\max}, \nu^{\min}$) and choose N sufficiently large for both cases.
 - (b) If $\tilde{w}_1$ is not fixed, first repeat Step 2, but now minimise the worst case over the appropriate range of ν ; i.e. minimise $\max_{(c,g) \in \mathcal{A}} \{\max(\text{Err}^{\text{AD}}(c,g;\nu^{\max}), \text{Err}^{\text{AD}}(c,g;\nu^{\min}))\}$. Then repeat Steps 3 (again, checking both cases in Step 3) through 5 and continue the iterative process until $\tilde{w}_1$ converges sufficiently.
6. As for the one-fuel model, allow the mean-reversion level of demand μ_Y to be time-dependent in order to reproduce the observed power forward curve for maturity m . For all but very short maturity forwards, we can use the same argument as in the one-fuel case that the coefficients $A_1, A_2^C, A_2^G, \dots$ of (5.8) depend only on the seasonal unconditional distribution of $R_T = \frac{D_T}{D_T + M_T}$, not on current values of demand and margin. This greatly simplifies the computation of $\mu_Y(T_i)$ needed to match the forward curve, as we shall now discuss.

Existence and Uniqueness of Calibration Parameters

As in the one-fuel case for NEPOOL, the existence and uniqueness of the value $\mu_Y(T_i)$ to match $F^P(t, T_i)$ (for $i = 1, 2, \dots$) depends on the bid stack being a strictly increasing function of R_T throughout, and reaching sufficiently low or high values in the range $(0, 1)$. Using the Taylor series approximation of this section, we retain the feature of the model that $B(x) \rightarrow -\infty$ as $x \rightarrow 0$ and $B(x) \rightarrow \infty$ as $x \rightarrow 1$, implying that $F^P(t, T_i) \rightarrow \pm\infty$ as $\mu_Y(T_i) \rightarrow \pm\infty$, as for the one-fuel case. Furthermore, even though $B(x)$ may be discontinuous at its joint (at $x = \tilde{w}_1$), $F^P(t, T_i)$ is still continuous in $\mu_Y(T_i)$ since it's an integral over the continuous distribution of R_T (as well as that of X_T^1, X_T^2 and X_T^3). Therefore, the existence of a $\mu_Y(T_i)$ to reproduce any possible $F^P(t, T_i)$ is

guaranteed. Similarly, in the case of using triangle distributions for the invariant distribution of R_T , the existence of $y \in \mathbb{R}$ (to determine $a(T_i)$ and $b(T_i)$ as in (4.6)) is also ensured.

While in practice we always also find a unique solution, in theory uniqueness requires several conditions in order to be strictly proven. Essentially, we must ensure that $B(x)$ is strictly increasing throughout $(0, 1)$. Therefore assume that the following hold:

1. The number of terms N is an odd number greater than or equal to 3.
2. $\{\lambda, \delta, \gamma\}$ are chosen such that $\hat{x}_0 \geq x_0, \forall (c, g) \in \mathcal{A}$, and if $\tilde{w}_1 \neq w_1$, $\{\tilde{\lambda}, \tilde{\delta}, \tilde{\gamma}\}$ are chosen such that $\hat{\tilde{x}}_0 \geq \tilde{x}_0, \forall (c, g) \in \mathcal{A}$ (i.e., we always overestimate x_0)
3. $s_1 = \beta_0^C + \beta_1^C c > 0$ and $s_2 = \beta_0^G + \beta_1^G g > 0, \quad \forall (c, g) \in \mathcal{A}$
4. $P_t^{\mathbb{Q}} \{(C_{T_i}, G_{T_i}) \notin \mathcal{A}\} = 0$

Under such conditions $B(x)$ is strictly increasing for all possible fuel prices (C_{T_i}, G_{T_i}) at maturity (though not continuous except when $\hat{x}_0 = x_0$ exactly), implying that it is also strictly increasing in $\mu_Y(T_i)$ (or parameter y for triangles). Thus the calibration method is unique.

While the role of conditions three and four is quite straightforward, some derivation is needed to show the relevance of conditions one and two. Firstly, we prove the claim that an odd number $N \geq 3$ is sufficient to guarantee a strictly increasing Taylor Series approximation to $\log(x - w_1)$. The expansion is given by

$$f(x) = \log \tilde{x}_0 + \sum_{k=1}^N (-1)^{k+1} \frac{1}{k \tilde{x}_0^k} (x - w_1 - \tilde{x}_0)^k,$$

and its derivative with respect to x is

$$f'(x) = \sum_{k=1}^N (-1)^{k+1} \frac{1}{\tilde{x}_0^k} (x - w_1 - \tilde{x}_0)^{k-1}.$$

Let $N \geq 3$ be an odd number and note that $f'(x)$ has N terms. Writing $y = x - w_1$ and noting that $\tilde{x}_0 > 0$, it suffices now to show that

$$\frac{f'(x)}{\tilde{x}_0} = \frac{f'(y + w_1)}{\tilde{x}_0} = \sum_{k=1}^N (-1)^{k-1} \left(\frac{y - \tilde{x}_0}{\tilde{x}_0} \right)^{k-1} > 0 \quad \forall y \in \mathbb{R}.$$

We consider different ranges of y in turn. For $y < \tilde{x}_0$, all terms in the expression above are positive, while for $y = \tilde{x}_0$, all are zeros except the first, equal to one. For $\tilde{x}_0 < y < 2\tilde{x}_0$, the k -th term is always greater in absolute value than the $(k+1)$ -th term. Since N is odd and the first and last terms are both positive, the sum is also positive. Similarly, for $y > 2\tilde{x}_0$, the k -th term is always less in absolute value than the $(k+1)$ -th term. As before, since N is odd and the first and last terms are both positive, the sum is also positive. Finally, for $y = 2\tilde{x}_0$ all terms are equal to one. Hence for all $x = (y - w_1) \in \mathbb{R}$, $f(x)$ is strictly increasing. Analogously, the N term Taylor expansion of $\log(w_1 - x)$ is strictly decreasing. Hence, $B_1(x)$ is strictly increasing in each of the regions $x \in (0, w_1)$ and $x \in (w_1, 1)$. This is also always true of $B_2(x)$ from its construction.

Condition two is necessary at $x = \tilde{w}_1$ as it implies that $B_3(\tilde{w}_1^-) < B_3(\tilde{w}_1^+)$. If $w_1 = \tilde{w}_1$, then the expansion $f(x)$ above equals $\log \tilde{x}_0 + \sum_{k=1}^N 1/k$ at $x = \tilde{w}_1$, and is clearly increasing in x_0 , so $B_3(\tilde{w}_1^-) < B_3(\tilde{w}_1^+)$. If $\tilde{w}_1 > w_1$, then the derivation above implies that $B_3(x)$ is increasing for $x \in [w_1, \tilde{w}_1)$, while the condition $\hat{\tilde{x}}_0 \geq \tilde{x}_0$ ensures that $B_3(\tilde{w}_1^-) < B_3(\tilde{w}_1^+)$. A similar argument holds for $\tilde{w}_1 < w_1$ (though this would be an unusual choice of $\tilde{w}_1$). Along with conditions three and four to restrict coal and gas price movements to an appropriate region of the fuel price space, this suffices to ensure uniqueness of $\mu_Y(T_i)$.

Note, however, that the above criteria are sufficient but not necessary for uniqueness, suggesting that for a sufficiently small quantile tolerance parameter q , uniqueness is virtually always secured. In particular, the fourth condition is never satisfied in our model, since coal and gas prices are log-normally distributed, implying that they can always take values throughout the positive quadrant. However, very low probability fuel price movements will have only a very small impact, and do not cause us calibration difficulties in practice. Similarly, the first condition is important to avoid a downwards divergence of $B_3(x)$ in the right tail of the stack, and an upwards one in the left, which in some extreme cases (of high G_t , low C_t) could cause $B(x)$ to decrease over certain ranges. However, in practice N can be even, as long as we choose a large enough value of N to avoid this divergence altogether. The second condition ensures that $B(\tilde{w}_1^-) < B(\tilde{w}_1^+)$ whenever there is a discontinuity at $x = \tilde{w}_1$, hence preserving the stack's strictly increasing property. In practice we choose a sufficiently small tolerance parameter ϵ to ensure that $B(x)$ can only ever be decreasing over very short regions of the stack. This argument also applies to the third condition, though it is worth noting that the parameters for $\hat{s}_1$ from Table 2.2 limit the range of coal prices for which the model makes sense.¹⁵ Nonetheless we conclude that in practice it is not important to think about the strict conditions above, as long as the calibration is performed only for regions where it is valid based on a strict enough choice of tolerance parameters q and ϵ .

5.5 Forward Pricing Techniques and Results

5.5.1 PJM - Computational Techniques for Forwards

The calibration method discussed above is very general. For PJM, we generally favour setting $b = 0.6$ and $\tilde{w}_1 = w_1$ from the outset. The use of a shifted weight $\tilde{w}_1$ is less important when using the average absolute deviation error measure $\text{Err}^{\text{AD}}(c, g)$, as opposed to the maximum absolute deviation $\text{Err}^{\text{MD}}(c, g)$. This choice greatly simplifies (and reduces computation time for) the calibration of the forward curve as well as the pricing of other derivatives.

Both in the case of the normal SDE-based model for D_t and M_t (as in (3.10) and (3.4)), and in the case of direct modelling of R_T , we can simplify the final calibration step above by creating a

¹⁵As for NEPOOL, in PJM there is no problem for gas prices, since $\beta_0^G, \beta_1^G > 0$. For coal prices, the parameters fit to the entire dataset allow coal prices as low as \$13, while for the most recent period the lower limit is \$27. If the elliptical contours $f_{C,G}(c, g; m) = \theta(q, m)$ reach this bound, the bid stack model in general should not be used, and for that reason we tend to favour the parameter set estimated over the entire dataset. With coal parameter σ_3 given by the first row of Table 2.1, a low forward price \$40, and long time to maturity of 3 years, $P_t\{F^C(T_o, T_f) \leq \$13\} = 0.0045$, so the risk is quite small.

set of re-usable matrices which contain the values $\mathbb{E}^{\mathbb{Q}}[\psi_i(R_T)]$ for the necessary functions ψ_i and for different values of $\mu_Y(T) + f(T)$, the seasonal unconditional mean of log demand.¹⁶ In particular, for each of $w = w_A, \bar{w}_1, w_B$ and each season of the year (due to margin seasonality), we require a matrix $\mathbf{A}^{(w)}$ whose j -th row corresponds to some log demand mean level x_j (increasing gradually in j) and whose columns are given as follows:

$$\begin{aligned}\mathbf{A}^{(w)}(j, 1) &= P^{\mathbb{Q}}\{R_T < w | x_j = \mu_Y(T) + f(T)\} \\ \mathbf{A}^{(w)}(j, 2) &= \mathbb{E}^{\mathbb{Q}}[R_T 1_{R_T \leq w} | x_j = \mu_Y(T) + f(T)] \\ \mathbf{A}^{(w)}(j, 3) &= \mathbb{E}^{\mathbb{Q}}[\log(R_T) 1_{R_T \leq w} | x_j = \mu_Y(T) + f(T)] \\ \mathbf{A}^{(w)}(j, 4) &= \mathbb{E}^{\mathbb{Q}}[\log(1 - R_T) 1_{R_T > w} | x_j = \mu_Y(T) + f(T)],\end{aligned}$$

and for columns $k \geq 5$:

$$\mathbf{A}^{(w)}(j, k) = \mathbb{E}^{\mathbb{Q}}[(R_T - w_1)^{k-3} 1_{R_T \leq w} | x_j = \mu_Y(T) + f(T)].$$

Note that the number of columns required here depends on the number of terms N used in the Taylor expansion, as terms of the form $\mathbb{E}^{\mathbb{Q}}[(R_T - w_1)^n 1_{R_T \leq w}]$ appear for $n = 1 \dots N$.

When D_t and M_t follow (3.10) and (3.4), all of these expected values can be calculated easily either by simulation or other numerical integration techniques. We simply need to integrate over a joint Gaussian density, a univariate Gaussian density and an exponential density. For any of the functions $\psi_i(R_T)$ above:

$$\begin{aligned}\mathbb{E}^{\mathbb{Q}}[\psi_i(R_T) | x_j = \mu_Y(T) + f(T)] &= (1 - p_i) \int_{-\infty}^{\infty} \psi_i\left(\frac{e^{x_j+y}}{e^{x_j+y} + e^z}\right) f_{Y,Z}(y, z) dy dz \\ &\quad + p_i \int_0^{\infty} \left(\int_{-\infty}^{\infty} \psi_i\left(\frac{e^{x_j+y}}{e^{x_j+y} + e^z}\right) f_Y(y) dy \right) f_{\alpha_i, \lambda_i}(z) dz\end{aligned}$$

where $f_{Y,Z}(y, z)$ is a bivariate Gaussian density, $f_Y(y)$ is a univariate Gaussian density and $f_{\alpha_i, \lambda_i}(z)$ is a shifted exponential density with parameters correspond to the season of our maturity T . (The densities are the invariant versions of those in (5.3) except with mean of log demand shifted to zero, since it is captured by x_j instead.)

Letting x_j vary over a large enough range with small enough step sizes, we can now use these matrices to find the values $\mu_Y(T_i), i = 1, 2, \dots$ to reproduce the observed forward curve. (For example, let $x_j = \log(0.3 + 0.001(j - 1))$ for $j = 1 \dots 1001$.) The key observations here are firstly that all coefficients $A_1, A_2^C, A_2^G \dots$ in (5.8) can be written (as vectors where we still let x_j vary) as linear functions of $\mathbf{A}^{(w)}(:, k)$ for various choices of w and k , and secondly that the resulting vector of power forward prices from (5.8) is strictly increasing in x_j .¹⁷ Hence, calibration to the power forward curve is achieved by finding the values of x_j and x_{j+1} (or rows j and $j + 1$) such that the observed power forward price falls in between the model's prices conditional on an expected log demand of x_j and x_{j+1} . Then for each T_i , linear interpolation can be used to find an approximate value for $x \in [x_j, x_{j+1}]$ such that (5.8) reproduces the observed power forward $F^P(t, T_i)$ conditional

¹⁶Recall that $\mu_Y(T)$ implicitly incorporates a time-dependent market price of demand risk, and thus describes the dynamics of D_t under $\mathbb{Q}$.

¹⁷As discussed earlier, this is very close to accurate for the region of validity for calibration.

on $\mathbb{E}^{\mathbb{Q}}[\log D_{T_i}] = x$. Then, returning to the conditional expectation of the demand process, we can solve for $\mu_Y(T_i)$ as follows. For $i = 1$,

$$\mu_Y(T_1) = \frac{x - f(T_1) - Y_t e^{-\kappa(T_1-t)}}{1 - e^{-\kappa(T_1-t)}},$$

while for $i \geq 2$ (for $i = 2$ the sum term does not appear),

$$\mu_Y(T_i) = \frac{x - f(T_i) - Y_t e^{-\kappa(T_i-t)} - \mu_Y(T_o) e^{-(i-1)\kappa\Delta T} (1 - e^{-\kappa(T_o-t)}) - \sum_{j=2}^{i-1} \mu_Y(T_j) e^{-(j-1)\kappa\Delta T} (1 - e^{-\kappa\Delta T})}{1 - e^{-\kappa\Delta T}}.$$

In addition to avoiding the numerical inversion of the bid stack, the computational advantage of the above approach is that all calculations (either simulation or numerical integration) involving the factors demand and margin can be performed at the beginning of the code. These can then be re-used for any historical calibration date, for any maturity m , any observed forward coal and gas prices $F^C(t, t+m)$ and $F^G(t, t+m)$, and for any bid stack parameters (other than w_1 , $\tilde{w}_1$ and b which should be fixed throughout). The alternative of simultaneous simulation of all underlying factors (or other numerical techniques) combined with bid stack inversion to evaluate an expression of the form (5.3), is a high-dimensional computationally intensive process. To then solve the inverse problem of varying $\mu_Y(T_i)$ to match $F^P(t, t+m)$ is even more demanding. Thus the fundamental advantage of the approximation technique (and the calibration approach above) is the separation of the impact of the short-term driving factors from long-term driving factors. This also conveniently allows us to use observed forward fuel prices directly, and to treat any remaining differences between model and observed forward prices as risk premiums due to demand and/or margin. Although at first glance messy and restricted in its range of validity, the Taylor Series approximation to the bid stack therefore provides a very useful pricing and calibration tool, transforming the bid stack model from an elegant mathematical approach to a practical market model.

As for NEPOOL earlier, for PJM we have collected from daily NYMEX reports a large quantity of forward and option data over the period 2006-2008. In particular, we have 723 daily observations (labeled $n = 1, \dots, 723$) of monthly PJM peak forward curves (typically three years) for which we calculate M_U and then calibrate the bid stack model for all maturities $m < M_U$. This process involves the repetition of similar calculations for each observation date and for each maturity. Therefore, we reduce computation time by taking advantage of the following initial calculations which can be stored and re-used throughout:

- Create the matrices (or look-up tables) $\mathbf{A}^{(w)}$ for $w = w_A, \tilde{w}_1, w_B$ as discussed above.¹⁸ These can be reused for all dates n , maturities m , and forward prices, though particularly low or high values of prices may require enlarging the initial matrices with more rows.
- For the chosen q and forward contracts with maturity $m_1 = T_o - t, m_2 = T_f - t, \dots$, calculate the contour value $\theta(q, m)$ such that $P_t^{\mathbb{Q}}\{f_{C,G}(c, g; m) < \theta(q, m)\} = q$. For each observation date n , the contract with shortest maturity m_1 (assuming middle of month delivery) has a

¹⁸This method of look-up tables can naturally be compared to the approach of using a three or four dimensional look-up table for $S_t = B(C_t, G_t, R_t)$ itself, as mentioned in a footnote at the beginning of this chapter. Even though $\mathbf{A}^{(w)}$ has many columns and is repeated for three different values of w , it is still a one-dimensional look-up table as it only discretises one state variable, R_t . Therefore, reducing interpolation error by reducing step size is much more expensive for the three-dimensional look-up table than for the one-dimensional one.

different maturity, varying from 8 to 40 days. Erring on the side of caution, we therefore, set $m_1 = 40/365$, with $m_i = m_1 + (i - 1)/12$, for $i = 2, \dots, 36$. The vector of contour levels $\theta(q, m_i), i = 1 \dots, 36$ can then be re-used for any observation date to identify the elliptical regions which we denote $\mathcal{A}_{n,m}$. Crucially, the symmetry of the Gaussian distribution implies that these ellipses retain their shape and size for different fuel forward prices, simply shifting to be centred on the observed values $(F^C(t, T_i), F^G(t, T_i))$.

- For a large range of coal and gas prices (e.g., $\log C_t \in [2.6, 5.6]$, $\log G_t \in [0.8, 3.8]$, with step size 0.05), calculate (for the chosen ϵ) the number of terms required to avoid divergence in the upper and lower stack, as well as the average absolute deviation $\text{Err}^{\text{AD}}(C_t, G_t)$, as were plotted in Figure 5.2. These matrices can be reused for any date n and maturity m in conjunction the region $\mathcal{A}_{n,m}$ and in Steps 2 and 3 of the calibration. By finding the maximum value of either number of terms or error over $\mathcal{A}_{n,m}$, we determine when and how the approximation may be employed (assuming no error from fitting x_0).
- Repeat the calculation of the above matrices for a series of scenarios of under or overestimating $x_0(C_t, G_t)$. For example, choose $\nu = 0.7, 0.8, 0.9, 1, 1.2, 1.4, 1.6, 1.8, 2$, which is sufficient to cover all values of $\nu^{\min}$ and $\nu^{\max}$ observed for PJM for any date n and maturity m . Then for each n and m (in Step 5 of the calibration), erring on the side of caution again, we choose values ν_1, ν_2 from the range above such that $\nu_1 \leq \nu^{\min} \leq \nu^{\max} \leq \nu_2$ and then reassess the use of the approximation technique (both validity and number of terms required).

In summary, much of the complexity of the two-fuel methodology for forward pricing can be incorporated into a series of initial calculations re-usable for each observation date n and maturity m . For our purposes this is particularly useful as we test the performance of the model throughout the period 2006-08, while in industry the benefit lies in much faster computation time when re-calibrating to the new observed forward curve each day. Clearly for the pricing of one or two options on isolated occasions, this rather involved methodology should not be implemented, but for the repeated application of the model to option pricing, it can provide a great benefit.

5.5.2 PJM - Results for Forwards

As for NEPOOL, an initial test of the model on empirical data is through the pricing of monthly power forwards, introduced earlier in Section 4.5.1. We have 723 daily observations from January 2006 to Nov 2008 (labeled $n = 1, \dots, 723$) of monthly PJM peak forward curves (typically up to about 3 years) and 649 daily observations from January 2007 to July 2008 of options on forwards (for dates $1 \leq 649$). However, as we choose to focus on the bid stack approximation technique of this section, we are limited in the range of forwards which we can price and calibrate to, as shown earlier in Figure 5.3. Using coal and gas parameters from Table 2.1 estimated through 2007, and tolerance parameters $\epsilon = 5$, $q = 00.2$ we can only price about a year or so of the forward curve for the early years, and only a few months for the more recent dates in 2008. Adopting the scenario of higher coal to gas correlation ($\rho_{13} = \rho_{23} = 0.5$) significantly increases M_U (see Figure 5.3) and allows us to test the model on more data, while having only a small impact on forward prices. Thus for illustrative purposes, and also considering the evidence from Section 3.2.1 that correlation may

be underestimated, we shall use this case throughout this section and some of Section 5.7.2.




Figure 5.6: Surface of forward risk premiums for different observation dates and maturities m (left); Progression of average forward risk premiums for short-end, middle and long-end of the curve (right) Note that we average over one-year portions of the curve to avoid seasonal effects, but use overlapping periods (0-1yr, 0.5-1.5yrs, 1-2yrs) due to the shorter ranger of forward prices we can calculate, relative to the unrestricted NEPOOL case earlier.

As before, we assess the forward risk premiums implied by the model (following calibration to the fuel forward curves as usual), and plot these as a surface in Figure 5.6. As for NEPOOL, there is typically a positive risk premium which averages at \$3.35 for the 2,500 forward prices plotted. Note also, that there is a clear seasonal effect as there was for NEPOOL, with higher risk premiums for peak months. The right graph of Figure 5.6 illustrates more clearly the increase in risk premiums over the period 2006-2008. This could either be a true reflection of increasing risk aversion in very volatile times, or otherwise an indication that parameters of factors estimated over the entire dataset are no longer so suitable for recent dates. The large increase in coal price volatility is one such example. Finally, separating the short, middle and long end of the forward curve shows no strong pattern, though perhaps some small evidence that longer term forwards tend to contain smaller risk premiums.

5.6 Pricing Power Options

As discussed for the one-fuel case, the calibration of the power forward curve can be viewed as a first step to pricing options and other more complicated derivative contracts. While the simple relationships derived in Section 4.3 are not available for the two-fuel case, both simulation and approximation techniques can be used, and can be combined with the Taylor Series approximation for the bid stack. Recall that forward power prices $F_t^P = F^P(t, T_f)$ (with maturity fixed at T_f throughout, but dropped from notation to save space) can be approximated in terms of forward fuel

prices F_t^G and F_t^C using (with $\tilde{w}_1 = w_1$ for simplicity)

$$\begin{aligned}
F_t^P &= A_1 + A_2^C \log(F_t^C) + A_2^G \log(F_t^G) + A_3^C F_t^C + A_3^G F_t^G + A_4(F_t^C)^{-\delta}(F_t^G)^\gamma \\
&+ A_5(F_t^C)^{-2\delta}(F_t^G)^{2\gamma} + A_6(F_t^C)^{-3\delta}(F_t^G)^{3\gamma} + A_7^C(F_t^C)^{-\delta+1}(F_t^G)^\gamma + A_7^G(F_t^C)^{-\delta}(F_t^G)^{\gamma+1} \\
&+ A_8^C(F_t^C)^{-2\delta+1}(F_t^G)^{2\gamma} + A_8^G(F_t^C)^{-2\delta}(F_t^G)^{2\gamma+1} + A_9^C(F_t^C)^{-3\delta+1}(F_t^G)^{3\gamma} + A_9^G(F_t^C)^{-3\delta}(F_t^G)^{3\gamma+1} \\
&+ A_{10}^C F_t^C \log F_t^C + A_{10}^G F_t^G \log F_t^G + A_{11} F_t^C \log F_t^G + A_{12} F_t^G \log F_t^C,
\end{aligned} \tag{5.13}$$

where as before coefficients $A_i = \mathbb{E}_t^\mathbb{Q}[\psi_i(R_{T_f})]$ for some ψ_i . We assume these to be constants (for the given T_f) due to fast mean reversion as discussed.

The approximation above is only valid for certain ranges of F_t^C and F_t^G , as well as time to maturity $T_f - t$. Therefore, we cannot simply insert this expression into the payoff function $\max(F^P(T_o, T_f) - K, 0)$ of an option on a forward unless it is valid for a sufficient range of forward prices $F_{T_o}^C$ and $F_{T_o}^G$ (and for forward maturity $T_f - T_o$). More specifically, we require a similar analysis to that used in calibrating the forward curve, but looking at the joint distribution of $(F_{T_o}^C, F_{T_o}^G)$ conditional on time t information, instead of (C_{T_o}, G_{T_o}) . From (3.2) and (3.9), we observe that $(\log F_{T_o}^C, \log F_{T_o}^G)$ is multivariate Gaussian $\forall T_o$ with

$$\begin{pmatrix} \log F_{T_o}^C \\ \log F_{T_o}^G \end{pmatrix} \sim N \left(\begin{pmatrix} \tilde{\mu}_C \\ \tilde{\mu}_G \end{pmatrix}, \begin{pmatrix} \tilde{\sigma}_C^2 & \tilde{\rho}_{CG} \tilde{\sigma}_C \tilde{\sigma}_G \\ \tilde{\rho}_{CG} \tilde{\sigma}_C \tilde{\sigma}_G & \tilde{\sigma}_G^2 \end{pmatrix} \right), \tag{5.14}$$

where

$$\begin{aligned}
\tilde{\mu}_C &= \log F_t^C - \frac{1}{2} \tilde{\sigma}_C^2, \\
\tilde{\mu}_G &= \log F_t^G - \frac{1}{2} \tilde{\sigma}_G^2, \\
\tilde{\sigma}_C^2 &= \sigma_3^2(T - t), \\
\tilde{\sigma}_G^2 &= \frac{\sigma_1^2}{2\kappa} \left(1 - e^{-2\kappa(T_o - t)} \right) e^{-2\kappa(T_f - T_o)} + \sigma_2^2(T_o - t) + \frac{2\rho_{12}\sigma_1\sigma_2}{\kappa} \left(1 - e^{-\kappa(T_o - t)} \right) e^{-\kappa(T_f - T_o)}, \\
\tilde{\rho}_{CG} &= \frac{1}{\tilde{\sigma}_C \tilde{\sigma}_G} \left\{ \frac{\rho_{13}\sigma_1\sigma_3}{\kappa} \left(1 - e^{-\kappa(T - t)} \right) e^{-\kappa(T_f - T_o)} + \rho_{23}\sigma_2\sigma_3(T - t) \right\}.
\end{aligned}$$

Hence, for the chosen quantile tolerance level q , we can find the elliptical region (centred on the observed prices $(\log F_{T_o}^C, \log F_{T_o}^G)$) for which we require the approximation (5.13) to hold. While this is similar to the calibration of forwards, the key difference is that we require (5.13) to hold throughout the region, as opposed to the weaker condition that the bid stack approximation itself hold. As $T_f \rightarrow T_o$, these two conditions become identical since the forward price essentially becomes the spot price. Again, for PJM, $T_f - T_o$ is typically about 22 days. Hence we now require the bid stack approximation to hold throughout any small ellipses (corresponding to 22 day fuel price changes) whose means may lie anywhere in a (typically) large elliptical region.

If these conditions hold then for example call options on power forwards $V_t^P(T_o, T_f, K)$ may be priced using

$$V_t^P(T_o, T_f, K) = e^{-r(T_o - t)} \mathbb{E}_t^\mathbb{Q} [\max(F_{T_o}^P - K, 0)]$$

and with equation (5.13) substituted in for $F_{T_o}^P$. Recall that we treat the terms A_i as stochastic processes (functions of D_t, M_t) only for very short maturity forwards, and as seasonal constants (dependent only on maturity T) for medium to long forwards. In the latter case, $V_t^P(T_o, T_f, K)$ can be found

simply by simulation or by numerical integration over the bivariate Gaussian distribution above. In the former case, additional work is required to evaluate all coefficients $A_i = \mathbb{E}_{T_o}^{\mathbb{Q}} [\psi_i \left(\frac{D_{T_f}}{D_{T_f} + M_{T_f}} \right)]$ for all required values of Y_{T_o} and $Z_{T_o}^{\text{OU}}$, before integrating over the distributions of Y_{T_o} , $Z_{T_o}^{\text{OU}}$, $F_{T_o}^C$ and $F_{T_o}^G$ (at least a dimension reduction of one relative to considering X^1 and X^2 separately).




Figure 5.7: Forward power prices plotted against forward coal and forward gas prices, for high expected demand (mean 0.8, left graph) and low expected demand (mean 0.5, right graph). We use a one-year maturity for both cases.

Although (5.13) appears very long and complicated, for much of the (C_t, G_t) space (and common values of $\mu_Y(T_i)$), it is very close to planar, at least for PJM. Figure 5.7 plots F_t^P against F_t^C and F_t^G for both high and low expected demand, and for a time to maturity of one year. As we are modelling peak demand, the dependence on F_t^C is generally low, and becomes particularly low for high $\mu_Y(T_f)$. The tilt and level of this plane can thus vary significantly with maturity month T_f , similarly to the variation in intercept and slope of the linear relationship (4.5) for the one-fuel case earlier. For appropriate regions of the $(F_{T_o}^C, F_{T_o}^G)$ space we can find constants $\omega_0, \omega_1, \omega_2$ such that $F_{T_o}^P \approx \omega_0 + \omega_1 F_{T_o}^C + \omega_2 F_{T_o}^G$. Then the pricing of options on power forwards becomes a classical basket option problem, where

$$V_t^P(T_o, T_f, K) = e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(\omega_0 + \omega_1 F_{T_o}^C + \omega_2 F_{T_o}^G - K, 0)] \quad (5.15)$$

and both $F_{T_o}^C$ and $F_{T_o}^G$ are lognormal.

While no closed form solution exists, many approximations and pricing techniques have been proposed (see e.g., Carmona and Durrleman (2003) for details on spread option pricing), including the simplest approach of approximating $\omega_0 + \omega_1 F_{T_o}^C + \omega_2 F_{T_o}^G$ by a lognormal distribution via a matching of the first two moments. This method is sometimes known as Wakeman's method, following its introduction by Turnbull and Wakeman (1991). An extension of this approach is proposed by Borovkova *et al* (2006), where a family of lognormal distributions is suggested, including shifted and negative (flipped) lognormal random variables. All of these cases use moment matching techniques and lead to convenient closed-form option prices (and Greeks) similar to Black-Scholes. In our case, the simplest moment-matching approach involves solving the following pair of linear equations to

find a lognormal distribution for $F_{T_o}^P$ with parameters μ_P and σ_P^2 :

$$\begin{aligned}\mu_P + \frac{1}{2}\sigma_P^2 &= \log \mathbb{E}_t^Q[F_{T_o}^P] = \log \left(\omega_0 + \omega_1 e^{\tilde{\mu}_C + \frac{1}{2}\tilde{\sigma}_C^2} + \omega_2 e^{\tilde{\mu}_G + \frac{1}{2}\tilde{\sigma}_G^2} \right) \\ &= \log \left(\omega_0 + \omega_1 F_t^C + \omega_2 F_t^G \right)\end{aligned}\tag{5.16}$$

$$\begin{aligned}2\mu_P + 2\sigma_P^2 &= \log \mathbb{E}_t^Q[(F_{T_o}^P)^2] = \log \left(\omega_0^2 + \omega_1^2 e^{2\tilde{\mu}_C + 2\tilde{\sigma}_C^2} + \omega_2^2 e^{2\tilde{\mu}_G + 2\tilde{\sigma}_G^2} \right. \\ &\quad \left. + 2\omega_0\omega_1 e^{\tilde{\mu}_C + \frac{1}{2}\tilde{\sigma}_C^2} + 2\omega_0\omega_2 e^{\tilde{\mu}_G + \frac{1}{2}\tilde{\sigma}_G^2} + 2\omega_1\omega_2 e^{\tilde{\mu}_C + \tilde{\mu}_G + \frac{1}{2}\tilde{\sigma}_C^2 + \frac{1}{2}\tilde{\sigma}_G^2 + \tilde{\rho}_{CG}\tilde{\sigma}_C\tilde{\sigma}_G} \right) \\ &= \log \left(\omega_0^2 + \omega_1^2 (F_t^C)^2 e^{\tilde{\sigma}_C^2} + \omega_2^2 (F_t^G)^2 e^{\tilde{\sigma}_G^2} + 2\omega_0\omega_1 F_t^C + 2\omega_0\omega_2 F_t^G + 2\omega_1\omega_2 F_t^C F_t^G e^{\tilde{\rho}_{CG}\tilde{\sigma}_C\tilde{\sigma}_G} \right)\end{aligned}\tag{5.17}$$

In order to use this approach we require the fit of the plane to our surface $F_{T_o}^P(F_{T_o}^C, F_{T_o}^G)$ to be sufficiently strong over the elliptical region discussed above. For example, we may use a tolerance parameter ϵ^{plane} which sets a limit on the acceptable maximum (or average) percentage error in the fit (over the elliptical region).

For regions $\mathcal{A}_{m,n}$ where this planar approach holds (typically shorter maturities m), the basket option pricing technique discussed above (or any other such techniques) can conveniently be adopted for spark and dark spread options, which are often important contracts for the hedging needs of power generators. As in the one-fuel case, these can be priced similarly since the model indicates that their payoffs are linear in forward prices of fuels. Specifically, dark spread options have payoff

$$\max \left(\omega_0 + (\omega_1 - H^C)F_{T_o}^C + \omega_2 F_{T_o}^G - K, 0 \right),$$

while spark spread options have payoff

$$\max \left(\omega_0 + \omega_1 F_{T_o}^C + (\omega_2 - H^G)F_{T_o}^G - K, 0 \right),$$

where H^C and H^G are the heat rates for coal and gas specified in the contracts. Thus for each of these we could repeat the lognormal fit, finding different values of μ_P and σ_P . Furthermore, another advantage of the planar approach is the possibility of defining implied volatilities and correlations, as we shall discuss in Section 5.7.2.

Finally, we note that while the planar fit provides motivation for the use of the above basket option pricing technique, we could employ the same approach while avoiding fitting the plane. Specifically, the values $\mathbb{E}_t^Q[F_{T_o}^P]$ and $\mathbb{E}_t^Q[(F_{T_o}^P)^2]$ are available in closed form from (5.13). In particular, for any $\eta, \xi \in \mathbb{R}$ and integers $i, j \in \{0, 1, 2\}$ such that $i + j \leq 2$, terms of the form

$$\mathbb{E}_t^Q \left[(\log F_{T_o}^C)^i (\log F_{T_o}^G)^j (F_{T_o}^C)^\eta (F_{T_o}^G)^\xi \right]$$

can be found explicitly, though are rather long to write out. In this way, we avoid the need for a planar fit when approximating $F_{T_o}^P$ by a lognormal random variable. Instead, we require an alternative means of testing when this approximation is sufficiently accurate. For example, we can set a tolerance level ϵ^{skew} on the acceptable maximum percentage error in the third moment (linked to skew) of the fitted distribution. This requires comparing $\exp\{3\mu_P + \frac{3}{2}\sigma_P^2\}$ for the lognormal with $\mathbb{E}_t^Q[(F_{T_o}^P)^3]$ which can be still be found in closed form from (5.13). While improving computation time, a disadvantage of this approach is that we require the same bid stack approximation (i.e., the same number of terms N and parameters $\{\lambda, \delta, \gamma\}$) to be used throughout our elliptical region discussed above. This is not necessary for the planar fit technique.

5.7 Option Pricing Techniques and Results

5.7.1 PJM - Computational Techniques for Options

For each observation date n , we can calculate the range of maturities $m = T_o - t$ for which option contracts can be accurately approximated by the various techniques introduced above. (Note that since all observed PJM options have approximately the same 22 day period for $T_f - T_o$, we do not need to consider this extra dimension.) Each pair (n, m) will fall into one of the following categories:¹⁹

1. The joint distribution of $(F_{T_o}^C, F_{T_o}^G)$ lies within the region for which (5.13) holds (satisfying tolerance parameters q and ϵ) for a fixed bid stack approximation (i.e., fixed $\{N, \lambda, \delta, \gamma\}$) and the third moment of $F_{T_o}^P$ for the fitted lognormal satisfies the tolerance level ϵ^{skew} .
2. The above condition is not satisfied, but the joint distribution of $(F_{T_o}^C, F_{T_o}^G)$ lies within the region for which (5.13) holds (satisfying tolerance parameters q and ϵ) for a flexible bid stack approximation (i.e., let $\{N, \lambda, \delta, \gamma\}$ vary through the region to allow the best fit for each pair $(F_{T_o}^C, F_{T_o}^G)$) and the planar fit to the surface $F_{T_o}^P(F_{T_o}^C, F_{T_o}^G)$ satisfies the tolerance level ϵ^{plane} .
3. The two above conditions are not satisfied, but at least the power forward curve has been accurately calibrated using the bid stack approximation for maturities up to and including T_f (this is implicitly assumed in the above conditions).²⁰

For the first category above, we can price options on forwards very quickly, using the simple moment matching technique directly to approximate $F_{T_o}^P$ by a lognormal random variable. However, this is typically possible only for options very close to their maturity. The second category above is more widely applicable (depending of course on ϵ^{plane} , which we typically set to be 0.05), and still leads to fairly rapid calculation of option prices despite the extra step of the planar fit. However for longer maturities, the planar fit becomes weaker and weaker. Finally, for the third category above, we make fewer assumptions and approximations, but require simulation to price options on forwards. The bid stack approximation is still used within the simulations to find $F_{T_o}^P$ after simulating $F_{T_o}^C$ and $F_{T_o}^G$.

As discussed in detail in Section 5.5.1 for calibrating to forwards, there exist a number of techniques for saving computation time by re-using initial calculations. The same is true for options on forwards. In particular, since $T_f - T_o = 22$ days throughout, we can make use of the following:

- For the chosen q and, calculate the contour value $\theta(q, 22/365)$ (for 22 day forwards) such that $P_t^Q\{f_{C,G}(c, g; m) < \theta(q, m)\} = q$. Again, the symmetry of the Gaussian distribution implies that these ellipses retain their shape and size for different fuel forward prices.
- For a large range of forward coal and gas prices (e.g., $\log F_{T_o}^C \in [2.6, 5.6]$, $\log F_{T_o}^G \in [0.3, 3.3]$, with step size 0.05), determine when and how the bid stack approximation may be used to determine $F_{T_o}^P$. In particular for each point in the space:

¹⁹More generally, we could say for any set $(F_t^C, F_t^G, F_t^P, t, T_f)$, since the three observed forward prices form the key information set we need from observation date n , as well the maturity date and time to forward maturity $m = T_f - t$.

²⁰For maturities $m > M_U$ beyond the range of valid calibration, we require instead high dimensional and computationally expensive inverse problem techniques to calibrate.




Figure 5.8: Surface plots illustrating the optimal choice of parameters λ (top left), γ (top right) and δ (bottom left) for the elliptical region corresponding to the 22-day distribution of forwards centred on $(F_{T_o}^C - \frac{1}{2}\sigma_C^2, F_{T_o}^G - \frac{1}{2}\sigma_G^2)$; Bottom right graph illustrates the minimum and maximum values of $\hat{x}_0/x_0$ over these same regions when using these parameter sets $\{\lambda, \delta, \gamma\}$.

- Determine the parameters $\{\lambda, \delta, \gamma\}$ which best fit $x_0(C_t, G_t)$ over the appropriate elliptical region corresponding to the 22 day joint distribution of coal and gas centred on $(F_{T_o}^C - \frac{1}{2}\sigma_C^2, F_{T_o}^G - \frac{1}{2}\sigma_G^2)$. The first three graphs of Figure 5.8 illustrate the values of these parameters for typical ranges of $F_{T_o}^C$ and $F_{T_o}^G$.
 - Using the same matrices of errors $\text{Err}^{\text{AD}}(C_t, G_t)$ as used for calibration to forwards (for a range of possible ν), check the validity of the approximation (given the tolerance ϵ) over the appropriate elliptical region. Figure 5.8 shows that for typical values of coal and gas forwards, this is not much of an issue, as the maximum over- or underestimation of x_0 of only a few percent, for the small elliptical regions.²¹
 - Using the same matrices of number of required N as used for calibration to forwards (for a range of possible ν), find the number of terms required to avoid bid stack divergence again over the appropriate elliptical region.
- Compute the Cholesky decomposition of the covariance matrix Σ of (X_T^1, X_T^2, X_T^3) as given in (3.6) with $T - t = 22/365$ (to be used in simulations).
 - Finally, for the same large range of forward coal and gas prices as above, create a series of surfaces $F_{T_o}^P(F_{T_o}^C, F_{T_o}^G)$ for different values of log demand mean level $x_j = \mu_Y(T_f) + f(T_f)$ and for seasons $i = 1, 2, 3, 4$, as we did for matrices $\mathbf{A}^{(w)}$ earlier. Within each surface, allow the value of $F_{T_o}^P$ to be determined (whenever the approximation is valid) using the parameters $\{N, \lambda, \delta, \gamma\}$ determined above for that particular point $(F_{T_o}^C, F_{T_o}^G)$.

Note that the last of these is only necessary for the case of fitting a plane to the surface $F_{T_o}^P(F_{T_o}^C, F_{T_o}^G)$ and even then it may not be desirable to have this as an initial calculation. From the perspective of computation speed, it depends how many pairs of dates and maturities (m, n) we are interested in, relative to the number of values x_j (times four seasons) for which we create surfaces. Since the creation of these surfaces only requires algebra (i.e., equation (5.13)), it does not add significant computational expense to the formula in either case, particularly relative to the use of simulations. As with the other techniques described above, it is of course possible to build the appropriate matrices as we go, continually adding rows and columns to saved matrices when necessary (e.g., for wider ranges of C_t and G_t , but this is of course less practical to code.

We describe briefly here the steps to pricing options on forwards by simulation; i.e., for the third category of pairs (m, n) explained above. For each forward maturity $m = T_f - t$, we first calibrate to the forward power curve $F^P(t, T_i) \forall T_i \leq T_f$, thus finding $\mu_Y(T_f)$. We then calculate the Cholesky decomposition $\Sigma = AA^T$ of the covariance matrix of $(X_{T_o}^1, X_{T_o}^2, X_{T_o}^3)$ using (3.6) with $T = T_o$ ($T_o = T_f - 22/365$), and simulate $(X_{T_o}^1, X_{T_o}^2, X_{T_o}^3)$ using the standard approach for multivariate Gaussians (i.e., $X = AZ$ where Z is a vector of independent Gaussians). We find our simulated values (2000 runs used in the results of the next section) for $F_{T_o}^C$ and $F_{T_o}^G$ using (3.9) and (3.2), and then check whether they lie in a region for which the bid stack approximation (5.13) holds. If so, we

²¹Interestingly, it also illustrates that the largest gaps between the maximum and minimum of the ratio $\hat{x}_0/x_0$ occur in the same regions of the fuel-price space which already cause the difficulties for the Taylor Series approach (i.e., low C_t and high G_t or high C_t and low G_t), implying that for longer maturities, these regions become even less accessible to the approximation. However, on the bright side, this also means that the fit is strong for more typical central regions in the space.

obtain $F^P(T_o, T_f)$ using the appropriate parameters $\{N, \lambda, \delta, \gamma\}$ found in the initial calculations. In the rare cases that a simulated pair $(F_{T_o}^C, F_{T_o}^G)$ departs from the region of validity, we find the forward price by returning to the ‘exact’ bid stack $B(\cdot)$ given in (5.2). Of course, this requires averaging over an additional set of simulations (within the main simulation) of $(X_{T_f}^1, X_{T_f}^2, X_{T_f}^3)$ conditional on $(X_{T_f}^1, X_{T_f}^2, X_{T_f}^3)$ as well as numerical inversions to find S_{T_f} in each case. (Here we make use of the Cholesky decomposition for the 22-day Σ calculated earlier.) This method of simulations within simulations is clearly computationally very expensive, but fortunately the bid stack approximation makes this unnecessary in all but a very small portion of the simulations (at least for maturities $m \leq M_U$). Finally, it is also worth noting that the closed-form solutions for power options in the one-fuel case make natural candidates for control variates to use in the two-fuel case and further decrease computation time for simulations.

5.7.2 PJM - Results for Options

We test our model against PJM options on forwards over the same recent NYMEX dataset introduced earlier. As with NEPOOL, we ignore options worth 10 cents or less, and we categorise options by ‘moneyness’ using the same definitions of ‘at-the-money’ (ATM), etc, as in Section 4.4.2. We again begin by considering in Figure 5.9 the movement of two individual option prices over the year 2007. We can firstly observe that all of the model prices seem to underprice both the call and the put throughout, but that the gap narrows significantly as we approach maturity. This suggests that we are either underestimating the volatility of a longer term factor, or failing to capture a longer term risk altogether. The results for the planar fit (category 2 in the previous section) are however encouraging, as these option prices remain close to those using the simulation plus bid stack approximation approach (category 3 in the previous section), which is considered the standard model and labeled simply ‘model’ in graphs. The graph also illustrates that our assumption of slightly higher correlation $\rho_{13} = \rho_{23} = 0.5$ only has a relatively small impact on prices. As mentioned in Section 5.5.2, we shall use these adjusted parameters throughout the rest of the empirical results, and are consequently able to price more options with the bid stack approximation, especially in recent years. The right graph of Figure 5.9 emphasises this point by illustrating the number of options which can be priced using various methods, as well as the total number (priced above 10 cents) observed in the market. As we can see, the use of a planar fit is almost as widely applicable as the standard simulation approach. Note that the line labeled ‘gas opt’ is a crude approximation approach using observed gas option prices, and will be introduced in the coming pages.

The first graph of Figure 5.10 plots for PJM the average relative error by observation date (as for NEPOOL in Figure 4.3), including a comparison with errors for gas options.²² Overall, the model performs worse for PJM than for NEPOOL, however the errors are much more consistent for puts and for calls, as well as following a very similar historical path to those of gas options. This is encouraging as it suggests that the errors in pricing gas options account for a significant portion of the errors in power options, and that the rest could potentially be explained by one missing factor or in-

²²Note that the plots from gas options are not exactly the same as in Figure 4.3 because we price fewer gas options here. In particular, we only consider gas options with time to maturity $m \leq M_U$, where M_U is the longest maturity forward for which the PJM bid stack approximation is valid, as described by (5.12) and Figure 5.3.




Figure 5.9: Left graph shows a sample of PJM option price trajectories, (observed data vs the three possible model prices) for the following contracts: a put with strike \$76.5 on a Apr 09 forward (lower lines); a put with strike \$76 on a Dec 07 forward (upper lines); Right graph shows the number of option contracts in the PJM market and how many can be priced using different methods.

appropriate parameter value. Moving to the case of absolute errors, the second graph of Figure 5.10 plots only call options for all ranges of strikes. Interestingly, both ‘DOTM’ (deep-out-of-the-money) and ‘DITM’ (deep-in-the-money) options are well priced by the model, although we should note that DOTM should always perform fairly well in absolute terms, due to their low prices. Results for puts are similar, as can be seen in the third graph of Figure 5.10, showing remarkably little variation between different lines. The alternative of categorising options by time to maturity instead of strike (in the fourth graph) confirms the more significant underestimation for long maturities, as noticed for the two sample options of Figure 5.9. However, some of this can be explained by a similar issue for gas options, plotted in the next graph. Again, and somewhat surprisingly, we see a closer relationship between error estimates for PJM and gas options than we witnessed in Chapter 4 for NEPOOL.

The final graph of Figure 5.10 represents an ad hoc attempt to improve results by identifying the possible missing risk factor, as we make several rough modifications to test various theories. Clearly we require a modification which widens the distribution of power forwards (in both tails) especially for longer maturities, since we are significantly underestimating observed prices. One possibility is the existence of a longer term component of demand or capacity. While we still expect electricity demand to be fairly rapidly mean-reverting on the whole, it is intuitive to expect some longer-term component as well, corresponding for example to cold or warm spells of weather which continue for several weeks. Either a two-factor model for demand could perhaps be more appropriate, or else a somewhat slower speed of mean reversion than the high parameter estimates of Chapter 3. In order to test this hypothesis without any additional computational or modelling effort, we can make a small adjustment to our simulation approach. We add a simulation of Y_{T_o} and then assume that the distribution of Y_{T_f} has mean equal to Y_{T_o} , instead of succumbing to mean-reversion. In this manner we increase the volatility of 22-day forward contracts, as they now depend significantly on demand levels at T_o . Option prices should thus all rise.

Next, we might also suspect that much of the power option pricing errors stem from the model’s




Figure 5.10: Historical analysis of PJM option pricing results using weekly observations: from top left to bottom right, (i) avg relative error of power and gas options, (ii) avg absolute error of power calls by ‘moneyness’, (iii) avg absolute error of power calls and puts by ‘moneyness’, (iv) avg absolute error of power calls and puts by months to maturity, (v) avg absolute error of gas calls and puts by months to maturity, (vi) avg absolute error of power call and puts when testing modelling modifications.

gas option pricing errors, and hence that using observed gas option data should help, as it did for NEPOOL results in Section 4.4.2. However, as we have seen in (5.15), power options on forwards can only be easily related to coal-gas spread options, for which we have no available market data. Instead, we can test the hypothesis by making the following rather crude approximation:

$$\begin{aligned} V_t^P(T_o, T_f, K) &= e^{-r(T_o-t)} \mathbb{E}_t^{\mathbb{Q}} [\max(\omega_0 + \omega_1 F_{T_o}^C + \omega_2 F_{T_o}^G - K, 0)] \\ &\approx e^{-r(T_o-t)} \int_0^\infty \mathbb{E}_t^{\mathbb{Q}} [\max(\omega_0 + \omega_1 F_{T_o}^C + \omega_2 F_{T_o}^G - K, 0) | F_{T_o}^C = x] f_C(x) dx \\ &= \int_0^\infty \omega_2 V_t^G\left(T_o, T_f, \frac{K - \omega_0 - \omega_1 x}{\omega_2}\right) f_C(x) dx \end{aligned}$$

where $f_C(x)$ is the marginal density of forward coal prices. Thus, we assume coal and gas to be independent in order to extract information about the marginal distribution of gas prices from observed gas options (i.e., the implied vol curve for gas). We implement this by using a small sample of our coal price simulations (100 paths), noting that for some options, some paths may lead us to values of $(K - \omega_0 - \omega_1 x)/\omega_2$ beyond the range of available gas strikes, thus reducing the number of options we can price (see Figure 5.9 again).

While both of these modifications are crude or ad hoc, they allow us to make a quick initial investigation into possible error sources or poor parameter estimates in the model. The results are compared to the standard model in the final graph of Figure 5.10, labeled ‘adjust D’ and ‘gas opt’ for the change to demand assumptions and the use of observed gas options respectively. Though results are not spectacular, they do show some significant and consistent error reduction, particularly for the demand adjustment case. The gas options case shows improvement only in the early part of the dataset, but not in more recent years. This could plausibly be related to the recent increase in both correlation and volatility of coal prices, since a greater proportion of the recent error may now be attributable to either errors in capturing the marginal distribution coal or of capturing the coal gas joint distribution. While fairly speculative, this exercise does provide some clues as to where to begin to look for improvements in the fit to market data. Of course many other suggestions are possible, and for example uncertainty in the future values of weights w_1 and $1 - w_1$ could arguably cause the underestimation witnessed in the option prices. Further analysis is required. In particular, the use of a greater variety of market data could be beneficial in this type of investigation, as here all options are on forwards with next month delivery. Options on weekly forwards or even on spot prices would allow us to better understand which underlying factors (or indeed the bid stack itself) are the source of most error.

As in Section 4.4.2 for the one-fuel model, it is useful to analyse observed option prices in terms of implied volatility instead of simply using error measures. For the two-fuel case, this can be achieved by using the planar fit and Wakeman’s option pricing technique for spread options. Again, we allow parameter σ_1 to vary to create an implied volatility surface. Thus we can define

$$\sigma_1^{\text{imp},P}(t, T_f, K) = \sigma \quad \text{such that} \quad V_{\text{obs}}^P(t, T_o, T_f; K) = V_{\text{mod}}^P(t, T_o, T_f; K, \sigma),$$

where, $V_{\text{mod}}^P(t, T_o, T_f; K, \sigma)$ is the standard Black-Scholes solution

$$V_{\text{mod}}^P(t, T_o, T_f; K, \sigma) = e^{-r(T_1-t)} \left[e^{\mu_P + \frac{1}{2}\sigma_P^2} \Phi\left(\frac{-\log(K) + \mu_P + \sigma_P^2}{\sigma_P}\right) - K \Phi\left(\frac{-\log(K) + \mu_P}{\sigma_P}\right) \right],$$

with μ_P and σ_P solving the equations (5.16)-(5.17).

In order to prove the existence and uniqueness of $\sigma_1^{\text{imp},P}$ for any observed power and fuel prices, we observe firstly that fuel forward prices F_t^C and F_t^G are fixed in (5.16)-(5.17) by observed market data. Therefore $\mu_P + \frac{1}{2}\sigma_P^2$ is fixed as we vary σ_G , and the left hand side of (5.17) can be written $2(\mu_P + \frac{1}{2}\sigma_P^2) + \sigma_P^2$. Hence, if the right hand side is strictly increasing in σ_1 , then σ_P must also be increasing in σ_1 and the same can then be said about call or put options as in the standard Black-Scholes model. It is straightforward to show that this is indeed the case, as the right hand side is strictly increasing (and unbounded above) in both $\tilde{\sigma}_C$ and $\tilde{\sigma}_G$, which are in turn strictly increasing (and unbounded) in both σ_1 and σ_2 .

Figure 5.11 illustrates the results for calls and puts over the usual historical period. Clearly the implied volatilities are well above the initial constant parameter estimate of $\sigma_1 = 0.566$, again emphasising the underestimation of option prices. Secondly, there is a general downwards trend similar to the corresponding graphs for NEPOOL and gas in Figure 4.4, though the correlation of these paths does not appear particularly strong otherwise, and is in fact weaker than in the error plots above. Apart from the unusual spikes for DITM puts, the most noticeable feature of these graphs is that for calls there appears to be a general pattern of implied volatility increasing as we move more out of the money, while for puts the reverse is true. A plausible explanation for this is the positioning of gas further up the bid stack, meaning that small changes in σ_1 have less impact on the lower part of the stack and hence on puts. (Consequently, it may be more suitable to measure implied coal volatility for contracts mainly dependent on the lower half of the stack, though the methodology would then be rather inconsistent.)




Figure 5.11: Historical analysis of PJM implied volatility $\sigma_1^{\text{imp},P}$ for calls (left) and puts (right) at different levels of moneyness again.

It is interesting to note that the same approach used to find implied volatility above can be adapted to calculate implied correlation between coal and gas. Again, the right hand side of (5.17) is increasing in both ρ_{13} and ρ_{23} , and hence so are power option prices. This is intuitive both for standard basket options and for the special case of power markets here. For example, a negative correlation between coal and gas would logically decrease the volatility of power prices since bids

from gas and coal generators would tend to swap places while keeping power prices relatively stable (holding other factors constant). The problem with the implied correlation approach is that correlation parameters should be bounded above by 1, and in case of PJM, market option prices are often still underestimated at maximum correlation levels. Furthermore, it should be noted that for both the implied volatility and implied correlation approaches, there is an implicit assumption that the parameters $\{\omega_0, \omega_1, \omega_2\}$ which determine the power forward plane remain fixed as we vary σ_1 . This is also an approximation, so an alternative but more computationally intensive definition of $\sigma_1^{\text{imp,P}}$ also exists, where these are allowed to vary. However, for high values of σ_1 , the relevant region of the fuel-price space will either no longer admit a planar fit or the bid stack approximation itself will fail as we saw earlier, in which cases $\sigma_1^{\text{imp,P}}$ can then not be defined. The simpler approach is therefore a more useful tool for intuitively representing and analysing observed market price data, which is typically the primary goal of any implied volatility calculation.

5.8 Adding Zero Bids (Must-run fuel types)

The methods discussed in the one-fuel and two-fuel cases are useful for markets where generators bid in clusters whose means and standard deviations move with fuel prices. However, as discussed before, some fuel types fall into the category of ‘must-run’ generators, implying that they make bids of zero either in order to avoid shutting down (e.g. nuclear) or simply because they must operate whenever their resource is available (e.g. wind, solar and some hydropower). While we have simply truncated our PJM and NEPOOL bid data to eliminate these less important bids, this may not always be appropriate. Firstly, we may be interested in off-peak hourly prices when the marginal plant in the stack may be very near the bottom (leading to night-time prices of zero or even below zero in some markets). Secondly, a significant proportion of must-run bids may come from generators whose availability in the market is very volatile, and thus the quantity of zero bids is itself stochastic. Though not applicable to nuclear, this is certainly the case for markets with renewables such as wind power, which is highly dependent on wind speed levels. As renewables are growing in importance globally, it is important to be able to adapt our modelling approach to such cases. For example, anecdotal evidence from the German market suggests that wind speed is already an important factor to model.

Thus we propose a simple modification of the two-fuel methodology to incorporate bids at zero.²³ We consider now the case of a two-fuel market (e.g. powered by wind and natural gas) with weight w_1 of must-run generators, and weight $w_2 = 1 - w_1$ of more traditional fuel price dependent plants with logistically distributed bids (with parameters m_1 and s_1). Then the distribution based approach to the bid stack gives us

$$S_t = x \quad \text{such that} \quad B_t^{-1}(x) = \frac{D_t}{D_t + M_t},$$

where

$$B_t^{-1}(x) = w_1 \mathcal{H}(x) + \frac{w_2}{2} + \frac{w_2}{2} \tanh \left(\frac{x - \alpha_0^G - \alpha_1^G G_t}{2(\beta_0^G + \beta_1^G G_t)} \right), \quad (5.18)$$

where $\mathcal{H}$ is the Heaviside step function.

²³Note that while this extension is feasible for a point mass of bids at any price, it faces the same requirement of non-changing merit order as in the two-fuel case, and is therefore most appropriate at price zero.

Since the support of the logistic distribution is the entire real line the model suggests that a few gas bids are negative, though in practice this is not realistic. Instead, we would typically observe a clear separation between wind and gas bids. Therefore, the same bid stack approximation approach as earlier seems reasonable, starting with two separate one-fuel stacks:

$$B_0(x) = \begin{cases} 0 & \text{if } 0 < x < w_1 \\ m_1 + s_1 (\log(x - w_1) - \log(1 - x)) & \text{if } w_1 < x < 1. \end{cases}$$

As before, we fix the middle of the stack by replacing the term $\log(x - w_1)$ with a 3rd order Taylor expansion of $\log x$ about x_0 :

$$B_1(x) = \begin{cases} 0 & \text{if } 0 < x \leq w_1 \\ m_1 + s_1 \left(-\log(1 - x) + \log x_0 + \frac{1}{x_0}(x - w_1 - x_0) - \frac{1}{2x_0^2}(x - w_1 - x_0)^2 + \frac{1}{3x_0^3}(x - w_1 - x_0)^3 \right) & \text{if } w_1 < x < 1. \end{cases}$$

To ensure continuity at $x = w_1$, we now require:

$$\begin{aligned} B_1(w_1^-) &= B_1(w_1^+) \\ 0 &= m_1 + s_1 \left[\log x_0 - 1 - \frac{1}{2} - \frac{1}{3} - \log(1 - w_1) \right] \\ \implies x_0 &= w_1 \exp \left\{ \frac{11}{6} - \frac{m_1}{s_1} \right\}. \end{aligned} \tag{5.19}$$

As before the approximation may require a series of improvements, such as increasing the number of terms or introducing linear approximations for the highest region, particularly for high G_t . On the other hand, it may not be valid for very low G_t , though it should be noted that in this case it may in fact produce a more accurate bid stack (relative to $B^{\text{obs}}(x)$) than the original $B(x)$.

This approach can clearly be viewed as special case of the two-fuel approximation, in the limit as the first distribution's mean and scale parameters both go to zero. Power forward prices in this case also follow from setting m_1 and s_1 equal to zero in (5.8) or (5.9). However, this technique is more than simply a special case of the two-fuel approximation. It allows us firstly to test the impact of a fixed level of zero bids on spot and forward power prices, particularly for off-peak prices which may be set very low in the stack. Secondly it allows us to conveniently extend the modelling framework to capture stochastic factors which drive these zero bids. We consider wind speed as an example. Since changes in wind speed add or remove zero bids in the stack, they correspond to horizontal shifts of the entire bid stack, and can therefore be an important driving factor of prices, even if wind generators never set the market clearing price. Note that x_0 is a linear function of $w_2 = 1 - w_1$, the weight of gas generators relative to wind. If w_2 can be written as a function f^{wind} of wind speed ξ_t , then our approximation to x_0 can be written as $x_0 = \lambda f^{\text{wind}}(\xi_t) G_t^{-\gamma}$. Since wind speed is a weather-related variable similar to temperature (which drives demand), it seems reasonable to consider it independent of fuel prices. Then our coefficients could be written as functions of the form $\mathbb{E}_t^{\mathbb{Q}}[\psi_i(R_T, \xi_T)]$.

It is important to note that while this method is appealing for some situations, it should not be applied in general for any market which includes some must-run generators. In the case that prices

are never set by these bids of zero, the method is inferior to the simple demand adjustment we can make to capture must-run generators (i.e., we subtract this capacity immediately from demand as it is always used). However, for markets or time periods (e.g., off-peak) when prices of zero are observed, the method above can be beneficial to retain the tractability provided throughout this chapter by the bid stack approximation technique.

Chapter 6

Carbon Emissions Allowances

Our discussion and modelling of electricity prices has so far ignored an increasingly important element of many modern energy markets, namely, the price of carbon emissions. While the empirical data studied above are drawn from American markets before the existence of a mandatory US cap-and-trade system for carbon emissions, such markets already exist and in Europe represent a key new factor in setting power prices. In particular, the EU ETS (European Union Emissions Trading Scheme) began trading carbon credits in 2005, with traded volume increasing from \$8 billion in 2005 to \$95 billion in 2008. As the threat of climate change intensifies and gains further attention in the future, the importance of market-based mechanisms to reduce CO₂ emissions is likely to grow substantially. Moreover, a global carbon market may eventually emerge from the regional systems in existence and under development today. In the US, the North-Eastern regions of PJM and NEPOOL are now at least partly covered by the new Regional Greenhouse Gas Initiative (RGGI), which auctioned its first allowances in September 2008. Though prices remain quite low so far (around \$3 per allowance), the establishment of a more widespread national system is likely in the near future, and prices may then be significantly higher. The left graph of Figure 6.1 shows the progress of carbon prices (for 2007 settlement) in the EU ETS during the first trading period 2005-07. Although prices dropped to essentially zero in 2007 when information about emissions levels was released, they had reached a peak of around 30 euros in 2006, prices high enough to significantly impact the behaviour of electricity and other energy prices. The right graph shows that prices in the second trading period 2008-12 reached similar levels before falling again recently. Figure 6.1 shows that through 2005, allowances for 2007 were highly correlated with those of later years, before a detachment occurred in 2006. As we shall discuss, these unusual dynamics are a natural consequence of the cap and trade mechanism under its current structure.

The behaviour of carbon prices presents an interesting and important modelling challenge which has yet to be widely addressed in academic literature. The market's structure, together with complex dependencies between carbon, power and the underlying fuel prices present several obstacles to creating simple but realistic models. From the perspective of the bid stack model for power, carbon allowance prices intuitively factor into generator bids as a production cost, shifting clusters of bids just like coal and gas prices do. However, supply and demand arguments dictate that the price of these allowances depends on the overall rate of emissions (both historical and current or




Figure 6.1: EU ETS emissions allowance prices (EUAs) for 2007 delivery (left graph) and 2008-12 delivery (right graph).

expected), which in turn depends on the bid stack (or merit order) in the power market. In particular, increasing carbon prices move coal generator bids further to the right relative to bids from gas generators, since carbon emissions per MWh are higher from coal. Thus, the portion of power coming from coal is automatically reduced, in turn reducing total emissions and lowering demand for allowances. So, while carbon emissions and allowance prices depend on the merit order of the electricity supply curve, the merit order is itself influenced by carbon prices, creating an intriguing feedback effect. This relationship provides a significant modelling challenge but can ultimately be captured through an equilibrium price framework which reveals interesting and important features of the carbon market. This key new market represents a fascinating new dimension in the behaviour of energy prices, and an area for much further research.

6.1 Introduction to Cap-and-trade markets and models

Before discussing the use of the bid stack model as a tool for understanding carbon prices, we first describe briefly the structure of a standard cap-and-trade system and mention existing price modelling techniques. As the design of a cap-and-trade mechanism is itself a hotly debated topic, we take the EU ETS as our example structure, as it is by far the largest and most important carbon market to date. In this structure, the timeline is divided into trading periods of several years (2005-07, called Phase I, followed by 2008-12, called Phase II), at the end of which emissions allowances typically expire.¹ All traded allowances have a certain year (or ‘vintage’) attached to them, but within trading periods both banking and borrowing from neighbouring years is allowed. Therefore both 2005 and 2007 credits could be used to cover emissions produced in 2006 in the EU ETS. However, 2008 credits cannot be used in 2007, thus explaining the divergence in the price processes observed in Figure 6.1.

¹In the case of EU ETS, there is very limited ‘banking’ of credits allowed between trading periods, subject to nationally-chosen maximum banking levels. Furthermore, the substitution of CDM (Clean Development Mechanism) credits (known as ERUs) for standard traded credits (EUAs) can provide a form of banking between periods, but this is also very limited for now.

Each year, the market administrator (typically European governments) allocates (or auctions) a fixed quantity of credits (emissions allowances) into the market, split amongst the various power generators and other major emitters of CO₂. These companies must then submit sufficient allowances at the end of the year to cover their emissions for the year. Each allowance gives the holder the right to emit one tonne of CO₂. Thus, companies which implement their own abatement measures to reduce emissions are likely to have extra credits to sell into the market, while over-polluters will be obliged to buy. Such a system is intended to reduce overall emissions in the most economically efficient manner possible, as low cost abatement measures will be employed ahead of high cost ones, with the price of allowances (i.e., the price of carbon) determining which are implemented. Any companies which are unable to provide sufficient allowances to cover their annual emissions must pay a penalty per extra tonne of CO₂ emitted. In the EU ETS, this penalty was set at 40 euros in Phase I, and 100 euros in Phase II. Clearly the price of carbon is highly dependent on the both the penalty level and the fixed cap (total supply) which are set in advance.

A number of early models exist which attempt to fit standard stochastic processes or econometric models to historical carbon emissions allowance prices, but we don't discuss these here, particularly in light of their inability to describe the behaviour of Figure 6.1. However, several recent more realistic approaches exist in which emissions levels (or rates) are modelled, and allowance prices derived based on equilibrium arguments and the probability distribution of total emissions at maturity T . In all cases, maturity corresponds to the end of the trading period, where a penalty is due for each tonne of CO₂ not covered by allowances submitted. It is generally argued that a penalty will never be due in early years of a trading period, since borrowing from the following year should be enough to cover any excess emissions. Firstly, Seifert *et al* (2007) propose a simple model using optimal control theory to find the equilibrium allowance price in a continuous time setting. They assume that the emissions rate process X_t (or equivalently, the allowance demand process) is exogenous and follows

$$dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t,$$

and that the expected accumulated emissions process Y_t is given by

$$Y_t = - \int_0^t U_s ds + \mathbb{E}_t^{\mathbb{Q}} \left[\int_0^T X_s ds \right],$$

where U_t is the abatement process, measuring the amount of emissions abatement in the market, which acts as the control. As is typical for carbon markets, the problem of finding an equilibrium spot allowance price A_t is solved from the perspective of a central market planner who chooses the optimal abatement strategy to minimise overall costs. This is shown to be equivalent to solving all individual cost minimisation problems. A_t equals the marginal cost of abatement and can be found by solving the appropriate HJB equation. By assuming that the cost of abatement as a function of quantity abated is simply $C(u) = cu^2$ for constant c , closed-form solutions can be found. While the authors discuss the role of varying gas/coal spreads in driving X_t due to merit order changes, they do not include this as a factor.

An alternative approach is presented by Chesney and Taschini (2008), who argue that abatement strategies (including fuel switching from coal to gas) are inferior to buying emissions allowances because of implementation time, cost and irreversibility. Hence A_t is found by modelling optimal

strategies of two companies trading allowances only (or waiting to see total accumulated demand at T). There is also some asymmetric information due to a lag in when one companies observes the emissions of the other. While neither of these two approaches captures the role of other energy prices in driving carbon allowance prices A_t (or the non-exogenous nature of demand for emissions), they do introduce several key properties of prices in a simplified framework. Most importantly, allowance prices A_t must be martingales after discounting (under $\mathbb{Q}$) and must satisfy the following:

$$A_t = (\text{discount term}) \times (\text{penalty}) \times (\text{conditional probability under } \mathbb{Q} \text{ of an overall shortage at } T).$$

These are both clearly intuitive since allowances are traded assets which can be held like stocks (and unlike electricity spot prices). Furthermore, they have value equal to either 0 or the penalty at T , depending on whether there's an overall shortage or surplus in the market. Essentially, allowances are like digital options on the total accumulated emissions Y_T with strike equal to the market cap (or allocation). Consequently prices must be bounded below by zero and above by the penalty (say, π) in this framework. Note, however, that a key weakness of this framework is that it ignores the fact that paying the penalty does not in reality excuse the company from supplying additional credits at the end of the following year to cover their shortfall. Thus no constant upper bound should exist on A_t , but instead an upper bound equal to the sum of the penalty and the expectation of next period's carbon price, which is in turn bounded by the sum of next period's penalty and the expectation of the following period's carbon price.

As alluded to above, the impact of carbon prices on the merit order of fuels in the bid stack provides a key relationship in describing price behaviour including correlations with other energy prices. However, it greatly complicates modelling methodology as it precludes the assumption of an exogenous emissions rate process, suggesting instead that X_t is a function of coal price C_t , gas price G_t and carbon price A_t , as well as power demand D_t . Carmona *et al* (2008) capture this important feature through a very general and rigorous framework (in discrete time) for emissions markets, where I firms can use J different technologies to produce K different goods. These each have different associated costs and emissions rates, and the equilibrium price of these goods is proven to be given by the typical merit order pricing rule with allowance price times emissions rate added to production costs. Equilibrium here means that supply meets demand, total positions in traded allowances sum to zero, and each firm is maximising its own wealth. As before, the optimisation problem for the central planner is proven to be equivalent to that of the individual firms. Furthermore, a unique allowance price is shown to exist under the condition of no point mass existing in the distribution of future accumulated emissions. The authors go on to discuss in detail various alternative designs for the market (tax, standard cap and trade, auctions, relative allocation scheme), and which perform best for limiting windfall profits to generators, meeting reduction targets, and keeping costs to consumers low. Finally, the paper considers the case of the Texas electricity market with coal and gas generation, using trinomial trees to get some numerical results.

A special case of the above model is presented by two of the authors in Fehr and Hinz (2006), with similar results derived but for the special case of power markets with coal and gas technologies. In their approach, the 'price of fuel switching' F_t is the key price driver as it determines the marginal cost of abatement at each point in time. It is a stochastic process dependent on the constant heat

rates h_c and h_g , and emission rates e_c and e_g , for coal and gas respectively:²

$$F_t = \frac{h_g G_t - h_c C_t}{e_c - e_g} \approx \frac{1.9 G_t - 0.4 C_t}{0.9 - 0.4}$$

Essentially the model leads to abatement happening in bursts (with as much fuel switching as is permitted by some maximum level) when the allowance price satisfies $A_t \leq F_t$.

The work of Fehr and Hinz (2006) and Carmona *et al* (2008) include various simplifying assumptions. For example, they assume constant heat and emissions rates for all generators of a given fuel type, as well as a lack of impact on merit order from strategic bidding. Furthermore, the assumption of a fixed penalty π at maturity T has already been mentioned as a significant approximation, while the impact of non-power sector demand has not been analysed for the power market example. Nonetheless, their approach provides a very useful simplified framework from which to begin to understand the behaviour of carbon prices, especially in light of the observed European carbon price collapse in 2006. As we have emphasised throughout this thesis for power prices, carbon prices are most realistically modelled in conjunction with other energy prices. While many other factors such as macroeconomic conditions, technological innovation and political developments will determine long-term carbon price evolution, much of the short to medium term price dynamics can be well explained through carbon's relationship with coal, gas and power. In particular, changes in merit order and fuel switching are two crucial effects. As we shall now discuss, the bid stack model (or a simplified version) provides a natural means of exploring these complex new dependencies in today's energy markets.

6.2 A Simple Equilibrium Price Model

We begin by investigating the simple case of a two-fuel power market whose bid stack is formed of two point masses, one for coal at $h_c C_t + e_c A_t$ and one for gas at $h_g G_t + e_g A_t$. Essentially, we assume that all generators of a given fuel type have identical costs and that they make bids equal to their costs. This leads to a very similar model to that of Fehr and Hinz above, although we treat demand and capacity differently and investigate both fuel-switching and exogenous demand from other sectors in a different manner. While the model is clearly an over-simplification considering our results from earlier chapters, it provides some useful insight.

6.2.1 Model Assumptions

We assume that all coal generators (with weight w_1) submit day-ahead bids equal to $h_c C_t + e_c A_t$ and all natural gas generators (with weight $w_2 = 1 - w_1$) submit day-ahead bids equal to $h_g G_t + e_g A_t$. We use a discrete time framework with time steps of length $\Delta t = T/N$, where N is the number of steps covering the trading period $[0, T]$. We assume that electricity demand D_t , defined to correspond to the period $(t, t + \Delta t]$, is not yet observed when bids are made but that the distribution

²The emissions rates $e_c = 0.9$ and $e_g = 0.4$ are in line with the average coal and gas emissions intensity factors (tonnes of CO₂ per MWh) in Europe. In fact, clean spark spread options are typically traded using a factor of 0.411, while the corresponding factor for dark spread options is 2 to 2.5 times larger.

of D_t (conditional on $D_{t-\Delta t}$) is known.³ D_t is then assumed to be observed immediately after bids are submitted at time t . There is a slight difference here with the model of Carmona *et al* (2008), which assumes that demand is first observed and then a production schedule (instead of bid stack) is determined, but the impact of this difference is small.

We retain the notation above for the emissions rate process X_t (or allowance demand process), defined to be the quantity (in tonnes) of CO₂ emissions during $[t, t + \Delta t]$, and hence a function of D_t . Finally, we introduce the accumulated emissions process $Y_t = \sum_{s=0}^{t-1} X_s$, which measures the total emissions so far, before generator bids are submitted at time t . (This is similar to the expected accumulated emissions process of Seifert *et al* (2007), and in fact coincides at maturity T .) We also retain the earlier notation for the fuel switching price $F_t = (h_g G_t - h_c C_t)/(e_c - e_g)$, which can be interpreted as the value of A_t such that all bids from coal and gas generators are equal. For now, we assume that fuel weights are fixed (implying that no generators can switch fuels), and that all demand for carbon allowances stems from the power sector.

Under these simplified assumptions (and with $R_t = \frac{D_t}{D_t + M_t}$ as before), the spot power price is given by

$$S_t = B(R_t) = \begin{cases} (h_c C_t + e_c A_t)1_{R_t \leq w_1} + (h_g G_t + e_g A_t)1_{R_t > w_1} & \text{if } A_t \leq F_t \\ (h_g G_t + e_g A_t)1_{R_t \leq 1-w_1} + (h_c C_t + e_c A_t)1_{R_t > 1-w_1} & \text{if } A_t \geq F_t \end{cases}$$

The emissions X_t over the period $(t, t + \delta t]$ are

$$X_t = \begin{cases} X_t^c = \frac{D_t D_t^{\max}}{R_t} [e_c R_t + (e_g - e_c) \max(R_t - w_1, 0)] & A_t < F_t \\ X_t^g = \frac{D_t D_t^{\max}}{R_t} [e_g R_t + (e_c - e_g) \max(R_t - (1 - w_1), 0)] & A_t > F_t, \end{cases} \quad (6.1)$$

where $D_t^{\max}$ equals the maximum demand (or capacity, like $c^{\max}$ used in Chapter 2) in the power market, measured in MWh. This constant is necessary in the above expression to ensure that emissions X_t is measured in tonnes of CO₂, since D_t is normalised to the size of the power market.⁴ It is also important to understand from (6.1) that the ratio R_t of demand over capacity determines the overall emissions rate per MWh, but that this must be multiplied by the power demand D_t to determine total emissions. Thus, a high value of R_t does not necessarily correspond to a high total emissions, as it may have been caused by a low value of M_t . In fact, rewriting (6.1) as

$$X_t = \begin{cases} X_t^c = e_c D_t^{\max} D_t + (e_g - e_c) D_t^{\max} \max((1 - w_1) D_t - w_1 M_t, 0) & \text{if } A_t < F_t \\ X_t^g = e_g D_t^{\max} D_t + (e_c - e_g) D_t^{\max} \max(w_1 D_t - (1 - w_1) M_t, 0) & \text{if } A_t > F_t, \end{cases}$$

it is easy to check that X_t is increasing in D_t for a fixed M_t , but also non-increasing in M_t for a fixed D_t .

³The notation here is chosen such that prices A_t , C_t and G_t are linked to power demand D_t , emissions rate X_t , and power price S_t (instead of $D_{t+\Delta t}$, $X_{t+\Delta t}$ and $S_{t+\Delta t}$), as they have been in the continuous time framework. Since A_t, C_t, G_t correspond to prices observed at a particular point in time t , while D_t, X_t, S_t cover a time interval $[t, t + \Delta t]$, we should be careful about notation. If Δt equals one day, we should arguably be using the other notation and hence observe a higher correlation between S_t and other energy prices at a one day lag. However, we will typically choose Δt to equal a longer period such as a month in order to ease computation when studying relationships with fuel prices over several years. In this case, our notation is more realistic.

⁴We assume here that no rescaling of demand and capacity has been performed as in Chapter 2. If we have rescaled, then we should of course make a further adjustment to calculate the demand in MWh at time t .

The case $A_t = F_t$ is not covered by the piecewise linear function of demand above, and requires us to make an assumption about the scheduling of power generation when all bids are equal. We can assume that the fraction w_1 of required demand is generated from coal and $1 - w_1$ from gas, irrespective of R_t , such that

$$X_t = D_t D^{\max}(e_c w_1 + e_g(1 - w_1)), \quad \text{if } A_t = F_t. \quad (6.2)$$

Finally, we note that a slightly simpler model version of the model exists in which capacity issues are ignored, so $D_t + M_t = 1$, and demand is the only exogenous factor other than fuel prices. This case is perhaps more intuitive, and we have simply

$$X_t = \begin{cases} X_t^c = e_c D^{\max} D_t + (e_g - e_c) D^{\max} \max(D_t - w_1, 0) & \text{if } A_t < F_t \\ X_t^g = e_g D^{\max} D_t + (e_c - e_g) D^{\max} \max(D_t - (1 - w_1), 0) & \text{if } A_t > F_t. \end{cases}$$

6.2.2 Solving for Carbon Prices

Now, as mentioned earlier, the carbon emissions allowance price A_t is priced as a digital option worth the penalty if the market ends in deficit, and 0 if there is a surplus of credits. Letting π denote the penalty and θ represent the total allocation (or cap) for the trading period,

$$A_t = e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}} [\pi 1_{\{Y_T > \theta\}}] \quad (6.3)$$

$$= e^{-r(T-t)} \pi P_t^{\mathbb{Q}} \left\{ Y_t + \sum_{u=t}^T X_u(D_u, M_u; F_u, A_u) > \theta \right\}. \quad (6.4)$$

Equation (6.4) reveals mathematically the circular relationship in the market and the reason why an equilibrium price emerges. In particular, if $e_c > e_g$ (as should be the case since coal is a dirtier fuel than gas), $X_t^g(d, m; F_t, A_t) < X_t^c(d, m; F_t, A_t) \forall d, m \in (0, 1)$. Therefore, X_t is non-increasing in A_t . Heuristically, as we increase A_t , we expect a larger fraction of future power generation (in the period $[t, T]$) to come from gas, and therefore $\sum_{u=t}^T X_u(D_u; F_u, A_u)$ should decrease. Hence the right-hand side of (6.4) is non-increasing in A_t , while the left-hand side is of course increasing. Figure 6.2 (left graph) illustrates this for the case that $A_t = F_t$ where there is no precise intersection of the diagonal line (left-hand side) with the step function (right-hand side). However, a unique equilibrium price is nonetheless clearly defined, as was proven rigorously to exist by Carmona *et al* (2008).

Note that the assumption in (6.2) above regarding X_t for the case $A_t = F_t$ is not needed here, so any assumption about the ordering of generators for power production will lead to the same result. However, for simulation purposes, it would be needed. The right graph of Figure 6.2 plots X_t as a function of D_t (with $M_t = 1 - D_t$ for simplicity) and illustrates the piecewise linear functions X_t^g (for gas bids lower) and X_t^c (for coal bids lower) as well as the middle case of (6.2) (all bids equal). In theory, for this middle case one can calculate for any D_t , the unique value $X(D_t)$ such that (6.4) holds exactly and then back out the proportion of gas and coal generators which should be used to force this level of emissions. For example, in Figure 6.2 (left graph), the diagonal line comes close to intersecting the lower step (corresponding to using X_t^g), so the required mix of generators would involve more gas than coal. In theory, a clever market administrator could therefore schedule generators to ensure that (6.4) always holds exactly, but this is all extremely hypothetical of course!




Figure 6.2: Left graph illustrates the existence of A_t as the solution of (6.4), plotted for the case that $A_t = F_t$. The right graph illustrates the X_t as a function of D_t (with $M_t = 1 - D_t$) with the lower lines corresponding to X_t^g and the upper ones to X_t^c , while the dotted line comes from (6.2).

Solving for A_t in practice is a dynamic programming problem, involving working backwards from maturity, and solving the above equation at each time step. Even in this simplified framework, we potentially have as many as five dimensions to work back through, since A_t depends on the factors C_t, G_t, Y_t, D_t, M_t . In fact, if G_t is a two-factor process as in Chapters 4-6, then we have six. However, to reduce dimensionality, it is convenient firstly to treat the processes for power demand and margin as serially independent, or similarly, to sample from their invariant distribution. As discussed in Chapters 4-6, we know that this is a reasonable assumption for all but short maturities. So if Δt is about a week or more, we can argue that D_t and M_t are drawn from distributions independent of $D_{t-\Delta t}$ and $M_{t-\Delta t}$. Also, since coal and gas prices only impact the model through F_t , we could also further reduce our number of dimensions by modelling the fuel switching price process directly. Either we use a one-factor model for F_t directly, or assume coal prices C_t are constant and model G_t with a one-factor model.

Now, working backwards from maturity, at time $t = T - \Delta t$ we can first identify three regions of the space (Y_t, F_t) defined by whether coal or gas bids come first in the stack, or whether they are equal. Thus, we find the carbon price A_t as follows (with $r = 0$ to simplify notation):

$$A_t = \begin{cases} \pi P^{\mathbb{Q}}\{X_t^c \geq \theta - Y_t\} & \text{if } F_t > \pi P^{\mathbb{Q}}\{X_t^c \geq \theta - Y_t\} \\ F_t & \text{if } \pi P^{\mathbb{Q}}\{X_t^g \geq \theta - Y_t\} < F_t < \pi P^{\mathbb{Q}}\{X_t^c \geq \theta - Y_t\} \\ \pi P^{\mathbb{Q}}\{X_t^g \geq \theta - Y_t\} & \text{if } F_t < \pi P^{\mathbb{Q}}\{X_t^g \geq \theta - Y_t\}. \end{cases}$$

The probabilities above depend only on the invariant distribution of demand and margin. Let $p^c(y) = P^{\mathbb{Q}}\{X_t^c \geq \theta - y\}$ and $p^g(y) = P^{\mathbb{Q}}\{X_t^g \geq \theta - y\}$. Moving to the previous time-step, since the discounted carbon price is a martingale under $\mathbb{Q}$, for any $t \leq T - 2\Delta t$, we have

$$A_t = \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|F_t, Y_t] = \begin{cases} \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^c] & \text{if } F_t > \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^c] \\ F_t & \text{if } \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^g] < F_t < \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^c] \\ \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^g] & \text{if } F_t < \mathbb{E}_t^{\mathbb{Q}}[A_{t+\Delta t}|X_t = X_t^g]. \end{cases}$$

As before the first case corresponds to coal generators being used first, the next to all generators

being used together and the third to gas generators being used first. For $t = T - 2\Delta t$,

$$\begin{aligned} \mathbb{E}_t^{\mathbb{Q}} [A_{t+\Delta t} | X_t = X_t^c] &= \pi \mathbb{E}_t^{\mathbb{Q}} [p^c(Y_t + X_t^c) 1_{F_{t+\Delta t} > p^c(Y_t + X_t^c)}] + \pi \mathbb{E}_t^{\mathbb{Q}} [F_{t+\Delta t} 1_{p^g(Y_t + X_t^c) < F_{t+\Delta t} < p^c(Y_t + X_t^c)}] + \\ &\quad \pi \mathbb{E}_t^{\mathbb{Q}} [p^g(Y_t + X_t^c) 1_{F_{t+\Delta t} < p^g(Y_t + X_t^c)}]. \end{aligned}$$

We continue this process back through time, always identifying three regions of the space (F_t, Y_t) .

6.2.3 Numerical Methods - Trinomial Trees

In order to implement this numerically, we use a tree-based approach for each of the state variables. In particular, for fuel prices C_t and G_t , we use a recombining trinomial tree for each fuel (under the appropriate risk-neutral measure $\mathbb{Q}$). While for coal we use geometric Brownian Motion, for gas we adapt the tree-building scheme to the two processes considered: geometric Brownian Motion and exponential OU. In both cases, we incorporate coal to gas correlation ρ_{CG} into the structure of the tree, using a variation on the approach of Hull and White (1994a, 1994b). We find that for our purposes (mid-range values of ρ_{CG} and a relatively low number of nodes due to the high dimensionality of the problem), our modified approach may allow us to exactly match correlation throughout the tree in cases when the Hull and White method cannot. Details on the tree-building procedure are discussed in detail in Appendix B. While trinomial trees are simply a special case of explicit finite difference techniques, we favour them here both for their intuitive structure and because they provide us with a grid of carbon price values and transition probabilities which can then easily be used for simulation purposes.

For the other state variables (power demand, margin and emissions levels), the tree structure is more unusual, relative to typical tree-building techniques. In particular, we require a tree for Y_t , the accumulated emissions so far. Firstly, we introduce the shifted processes $\tilde{X}_t$ and $\tilde{Y}_t$, representing emissions and accumulated emissions relative to the market cap. The rate of emissions suggested by the cap is simply θ/N per time period, so processes $\tilde{X}_t$ and $\tilde{Y}_t$ are simply

$$\tilde{X}_t = X_t - \frac{\theta}{N} \quad \text{and} \quad \tilde{Y}_t = Y_t - \frac{\theta}{N} t.$$

These variables are more intuitive as they indicate more clearly how likely an overall shortage is at maturity given the current state and time. The difficulty in building our tree for $\tilde{Y}_t$ originates from the fact that $Y_{t+\Delta t} - Y_t = X_t$ (given by (6.1)) depends on D_t and M_t but also F_t , corresponding to where we are in the C_t and G_t trees. As a result, it may be difficult to create a recombining tree, and one with a regular branching scheme. Firstly we discretise the invariant distribution of demand and margin, such that we only have a finite set of possible values for X_t , repeated throughout the tree. In particular if we choose three possible values for D_t and for M_t ,⁵, then the maximum number of possible values of X_t from (6.1) equals 18; i.e., nine pairs (D_t, M_t) each producing a value of X_t^g (gas below coal in merit order) and of X_t^c (coal below gas). Ideally, we would like to ensure that all of these values of X_t (and hence $\tilde{X}_t$) take us exactly to a node of the tree for $\tilde{Y}_t$ at the next step.

⁵Probabilities can be chosen to match the first three moments of each random variable, as well as their covariance. If we are not concerned about the spacing of the demand and margin values being regular, then we have many free parameters to choose from and multiple ways of matching moments. We also have the option of adding more possible values of demand and margin. We do not discuss the details here.

While this may in general be very difficult, in the case that all possible values $\tilde{X}_t$ are integers, it is simply a case of choosing equally spaced nodes with small enough spacing. It suffices to choose the gap ΔY between nodes to equal the highest common divisor of the differences between 18 possible values. In this case, letting $d_1, \dots, d_{n_D}, m_1, \dots, m_{n_M}$ be our range of possible values of demand and margin respectively, and letting $GCF(x_1, \dots, x_{n_X})$ be the greatest common factor of $x_1, \dots, x_{n_X}$, set

$$\Delta Y = GCF(x_i - x_j : i \neq j, x_i, x_j \in \mathcal{X}),$$

$$\text{where } \mathcal{X} = \left\{ x : x = \tilde{X}_t^c(d_k, m_l) \text{ or } \tilde{X}_t^g(d_k, m_l) \text{ for some } k \in 1, \dots, n_D, l \in 1, \dots, n_M \right\}.$$

Of course, for many cases this value of ΔY may be very small, significantly increasing the computation time for the numerical scheme, while in non-integer cases, it is not defined. For model testing purposes, we can choose round numbers for $d_1, d_2, d_3, m_1, m_2, m_3$, as well as e_c, e_g and w_1 , such that ΔY is not so small. However, in practice these would be either observed values or calibrated to moments of the invariant distributions of demand or margin, so we would have less freedom. An alternative is therefore to simply interpolate linearly between values at neighbouring nodes whenever we require $A_t(C_t, G_t, \tilde{Y}_t)$ for some $\tilde{Y}_t$ not falling on a node. Choosing a sufficiently small value of ΔY , we can keep the resulting error low.

Finally, note that although the calculation of $A_t(C_t, G_t, \tilde{Y}_t)$ at each point requires taking an expectation of $A_{t+\Delta t}$ over nine future values of $\tilde{Y}_{t+\Delta t}$, the size of our grid (or tree) remains manageable due to the bounds on A_t . In particular, let $\tilde{X}^{\max} = \max(\mathcal{X})$ and $\tilde{X}^{\min} = \min(\mathcal{X})$. Then for all $i = 0, \dots, N$, at time $t = i\Delta t$, we know that

$$A_t(C_t, G_t, \tilde{Y}_t) = \pi \quad \forall \tilde{Y}_t > (N - i)\tilde{X}^{\max} \quad (6.5)$$

$$\text{and } A_t(C_t, G_t, \tilde{Y}_t) = 0 \quad \forall \tilde{Y}_t < (N - i)\tilde{X}^{\min}. \quad (6.6)$$

This is a simple consequence of the fact that for $\tilde{Y}_t > (N - i)\tilde{X}^{\max}$, $P_t^{\mathbb{Q}}\{Y_T > \theta\} = P_t^{\mathbb{Q}}\{\tilde{Y}_T > 0\} = 1$, while for $\tilde{Y}_t < (N - i)\tilde{X}^{\min}$, $P_t^{\mathbb{Q}}\{Y_T > \theta\} = P_t^{\mathbb{Q}}\{\tilde{Y}_T > 0\} = 0$. Thus from (6.3), the carbon price must be equal to the penalty (or zero) and remain so until time T since a shortfall (or excess) of credits at maturity is now guaranteed. We make use of this fact to choose a grid for $\tilde{Y}_t$ with step size ΔY , lowest value less than $N\tilde{X}^{\min}$ and highest value greater than $N\tilde{X}^{\max}$. As we work back through time, some nodes near the boundaries of the grid will require us to find future carbon prices outside of this range, but no error will result, since we can make use of (6.5) or (6.6). Note also that although these bounds only hold for a particular choice of discretisation of demand and margin, X_t should also be bounded above by $D^{\max}$ (and below by zero) as long as D_t, R_t and $D_t + M_t \in (0, 1)$. An interesting perspective can be gained by thinking of carbon allowances as barrier options which take value either π or 0 when the accumulated emissions process $\tilde{Y}_T$ hits one of these two barriers. The barriers meet at maturity (at $\tilde{Y}_T = 0$) but become increasingly separated the further we are back in time.

Figure 6.3 illustrates the price of carbon allowances in this simplified framework, as a function of G_t and $\tilde{Y}_t$. We choose parameters $T = 2, N = 80, w_1 = 0.6, e_c = 0.9, e_g = 0.4, h_c = 0.4, h_g = 6, \theta = 32000$ (400 per period), $D^{\max} = 1000, \pi = 40$. Demand and margin distributions are discretised⁶

⁶This choice of discretisation produces a manageable greatest common factor $GCF(x_1, \dots, x_{n_X})$ of emissions rates equal to 10, which can be used as a node spacing to avoid linear interpolation between nodes of $\tilde{Y}_t$.




Figure 6.3: Carbon price at $t = 1/2$ (with $T = 2$) as a function of G_t and $\tilde{Y}_t$ (with C_t held constant): surface plot for case of gas as an exponential OU process (left graph) and contour plot with comparison to the GBM case as well (contours $A_t = 1, 10, 20, 30, 39$ plotted).

using $d_1 = 0.35, d_2 = 0.6, d_3 = 0.85, m_1 = 0.1, m_2 = 0.3, m_3 = 0.5$ with corresponding probabilities 0.3, 0.4, and 0.3 in both cases (D_t and M_t independent for simplicity). Finally, coal is held fixed at $C = \$100$ and the gas tree is centred on $G_0 = \$8$ with $\kappa_G = 1.2, \sigma_G = 0.6$ the mean reversion speed and volatility of the exponential OU process.⁷ The surface plot of $A_t(G_t, \tilde{Y}_t)$ for $t = 1/2$ in Figure 6.3 confirms that carbon price is non-decreasing in $\tilde{Y}_t$ and G_t , and bounded above by 0 and π , but also reveals some rather striking behaviour. There is an interesting region in the middle of the $(\tilde{Y}_t, G_t)$ -space where A_t depends approximately linearly on G_t , with low sensitivity to current accumulated emissions $\tilde{Y}_t$. For this region, $A_t \approx F_t$ (and $A_t = F_t$ for some fraction of the region, depending on grid spacing). On either side of this middle region, A_t depends more and more on $\tilde{Y}_t$ until finally it reaches either the penalty ($\pi = 40$) or zero. We can understand this to mean that for $\tilde{Y}_t$ near zero, the market's internal mechanism of merit order changes is sufficient to partially counteract the influence of changes in emissions demand. What matters is therefore mainly the level of fuel prices and the corresponding carbon price (near F_t) which allows this mechanism of merit order changes to function. The surface plot for the case of gas as a GBM is similar and can best be compared via the contour plot of Figure 6.3. Clearly for higher and lower values of G_t the contours are closer to vertical for the exponential OU case, indicating a weaker dependence on G_t due to the mean reversion. For the remainder of the chapter, we choose the exponential OU case for gas, partly because of evidence discussed in earlier chapters and partly for computational speed since the tree does not continue to expand indefinitely as N increases.

Secondly, we also choose to maintain fixed coal prices for the remainder of the chapter again to improve computation time. While this may seem like a very strong assumption, our goal is to gain an overview of carbon price behaviour as opposed to fitting any data like for power. Therefore, including stochastic coal prices does not provide much additional insight since an increase in coal

⁷To compare with the case of gas as a geometric Brownian Motion (GBM), we choose the volatility of the GBM such that the variances of the two processes are equal at at maturity of one year. For the exponential OU case, we set $y = 2$ and $b = 1 - \sqrt{1 - 1/y}$ to fix the nodes (here, $n = 21$) of the gas tree (see Appendix B for details), and then choose p_g for the GBM case such that the node spacing is the same.

price has a similar impact to a decrease in gas price. More precisely A_t depends on C_t and G_t only through F_t . This is confirmed by Figure 6.4 which shows results A_t as a function of C_t and G_t for a small range of coal and gas prices using the same parameters as for Figure 6.3 except with both coal and gas prices following a geometric Brownian Motion with volatility parameter 0.1. In the general model of Section 6.3, we will no longer be able to use a variable F_t to capture both C_t and G_t , but tests show that the relationship remains similar, with coal and gas movements having similar impacts but in opposite directions.




Figure 6.4: Surface plot of A_t against C_t and G_t , firstly with $\tilde{Y}_t = 0$ (left graph) and secondly with $\tilde{Y}_t = -1000$ (right graph)

6.2.4 The Impact of Fuel Switching and Exogenous Demand

As seen above, changes in the merit order of coal and gas in the bid stack allow us to solve for an equilibrium A_t and obtain interesting dynamics, where the primary driver of carbon can suddenly change from fuel prices to emissions levels. As we shall now see, a secondary effect known as fuel-switching can also contribute to this unusual behaviour. The term ‘fuel-switching’ is used here to mean the decision by coal generators to temporarily generate electricity from natural gas in order to reduce costs. Although we do not have the data to study the prominence of fuel-switching in recent years, it has been suggested to be one of the primary abatement tools available to the power industry. Fehr and Hinz (2006) categorise fuel switching along with skipping production altogether as the primary short-term abatement tools available, while longer-term measures include installing new scrubbers or investing in new units. However, others such as Chesney and Taschini (2008) argue that switching is not so simple and convenient. Generally, it is unclear what portion of coal generators typically have this ability, what switching costs are involved, and how often a fuel-switch could be made. This is also likely to vary significantly across markets and generators, similarly to the other operational costs and constraints mentioned in Chapter 1. To assess typical impact of fuel switching on A_t , we adopt the simplest approach to analysing fuel switching, assuming that no cost or time constraint is involved and that a certain fraction α of coal generators can immediately switch to gas when desired. As a result, the only change to our model above is a decrease in the emissions rate X_t^g , as some coal generators will now have switched to gas in this case. Thus the

emissions X_t over the period $[t, t + 1]$ are

$$X_t = \begin{cases} X_t^c = \frac{D_t D_t^{\max}}{R_t} [e_c R_t + (e_g - e_c) \max(R_t - w_1, 0)] & A_t < F_t \\ X_t^g = \frac{D_t D_t^{\max}}{R_t} [e_g R_t + (e_c - e_g) \max(R_t - (1 - \alpha w_1), 0)] & A_t > F_t. \end{cases} \quad (6.7)$$




Figure 6.5: Evolution over time (for the cross-section at $G_t = 8$) of the middle and outer regions of the carbon price surface: assessing the impact of fuel switching with $\alpha = 0.5$ (left graph) and of narrowing the invariant distribution of demand and margin (right graph). The narrower distributions of demand and margin correspond to choosing $d_1 = 0.5, d_2 = 0.6, d_3 = 0.7$ and $m_1 = m_2 = m_3 = 0.3$.

While Figure 6.3 provided a snapshot of carbon prices at a fixed time, Figure 6.5 illustrates how the various regions of the price surface progress over the time, with all regions of course converging at the maturity. In these plots, we take a cross-section of the carbon price surface at $G_t = 8$, and identify (for each $t \in (1/2, 2)$) the middle region discussed above, as well as outer regions corresponding to $A_t = 1$ and $A_t = 39$. We determine the middle region where the derivative $\frac{\partial A}{\partial Y}$ suddenly becomes quite low by finding points of maximum and minimum convexity $\frac{\partial^2 A}{\partial Y^2}$. Figure 6.5 illustrates that the regions typically grow linearly as we move backwards in time, but the prominence of the middle region varies for different cases and is significantly affected by coal generators switching to natural gas. Since the inclusion of fuel-switching lowers X_t^g (without impacting X_t^c), the changes in Figure 6.5 (left graph) occur only for higher values of $\tilde{Y}_t$. Essentially, as the jump in X_t at $A_t = F_t$ widens (i.e., the gap between X_t^c and X_t^g widens), the middle region where $A_t \approx F_t$ in our surface of carbon prices becomes more prominent. The right-hand graph illustrates another case in which the middle region widens: a reduction in demand or margin volatility. However, we see clearly that a narrowing of the demand distribution has a much greater impact than a narrowing of the margin distribution, since it plays a larger role in determining X_t through (6.1). Even reducing margin volatility to zero, the width of the jump in X_t is not greatly reduced, so the impact is still low. This suggests that while margin is a key factor for power price modelling, it is not as important for carbon.

Finally, exogenous carbon emissions demand is another factor which may significantly impact the regions described above. In carbon markets where sectors such as transportation and general industry form a large proportion of demand for credits, it is unreasonable to treat power sector demand as the only source of emissions demand. We can easily include another independent demand factor D_t^{exog} in the equation for X_t above, adding it to both X_t^g and X_t^c (and presumably reducing

$D^{\max}$ or increasing the cap θ). As we would expect, the effect here is to reduce the prominence of the middle region of the price surface.

6.3 Generalising the Simple Framework

While the simple approach to the bid stack discussed above provides welcome insight into the dynamics of the carbon markets, it is ultimately a very crude approximation to the electricity bid stack studied in depth in earlier chapters. Moreover, it implies the existence of only (at most) two possible power prices for any given time t , as we must have either $S_t = h_c C_t + e_c A_t$ or $S_t = h_g G_t + e_g A_t$, depending on which fuel sets the market clearing price. The natural remedy to this problem is to revert to our original bid stack model of Chapter 2. However, solving (6.4) becomes highly computationally intensive, owing to the numerical inversion required to obtain the bid stack. An intermediate approach which maintains a reasonable computation time to the above approach while generating a more realistic bid stack and power price process, is to generalise the simple framework above to the case of many point masses of coal and gas bids. This can be viewed as a discretisation of the full bid stack model, which can be made increasingly realistic as the number of bid points increases. Recall that of course real market bid stacks are typically also step functions with many steps.

6.3.1 Model Assumptions

We assume that there are n_C different groups of coal generators, with weights $w_c^1, w_c^2, \dots, w_c^{n_C}$, heat rates $h_c^1, h_c^2, \dots, h_c^{n_C}$, corresponding emissions rates $e_c^1, e_c^2, \dots, e_c^{n_C}$ and fixed costs $f_c^1, f_c^2, \dots, f_c^{n_C}$. We assume also that $h_c^i \leq h_c^j$, $e_c^i \leq e_c^j$ and $f_c^i \leq f_c^j$ for all $i < j$, such that the ordering of point masses of bids in the bid stack is fixed. So coal generator i (or alternatively group i of coal generators) bids $f_c^i + h_c^i C_t + e_c^i A_t$ at time t . Similarly, for gas generators, assume we observe n_G different groups with weights $w_g^1, w_g^2, \dots, w_g^{n_G}$, heat rates $h_g^1, h_g^2, \dots, h_g^{n_G}$, corresponding emissions rates $e_g^1, e_g^2, \dots, e_g^{n_G}$ and fixed costs $f_g^1, f_g^2, \dots, f_g^{n_G}$, with the same inequalities applying as for coal. So gas generator i bids $f_g^i + h_g^i G_t + e_g^i A_t$ at time t . Of course we require $\sum_{i=1}^{n_C} w_c^i + \sum_{i=1}^{n_G} w_g^i = 1$. We also note that if the parameters h_c^i and h_g^i correspond precisely to heat rates (meaning the amount of fuel required per MWh of power), then we should have a close relationship between h_c^i and e_c^i (and between h_g^i and e_g^i) for all i . For example, if we assume all coal generators to be equally polluting per ton of coal but differ only in the amount of coal they need per MWh, then we should set $h_c^i = \beta e_c^i$ for some constant β . However, we could also assume that differences in coal (or gas) bids stem from differences in other costs or bidding strategies which are implicitly captured by parameters either f_c^i or h_c^i , while perhaps e_c^i is constant for all i . Many different scenarios can be considered, but we always require the ordering of bids within each fuel type to remain fixed for all C_t, G_t , and A_t . This condition is not strictly necessary but simplifies notation and computation, as well as appearing quite logical and intuitive.

As in (6.4), the carbon emissions allowance price A_t is priced as a digital option on Y_t (or $\tilde{Y}_t$). However, as we don't have a single 'fuel-switching price' where gas and coal bids meet, we cannot write X_t as a function of F_t , but instead must include both C_t and G_t . Nonetheless, the general

pricing formula still holds, as it does for any bid stack function:

$$A_t = e^{-r(T-t)}\pi P \left\{ Y_t + \sum_{u=t}^T X_u(D_u, M_u; C_u, G_u, A_u) > \theta \right\}. \quad (6.8)$$

As before, this pricing formula implies the existence of a unique carbon allowance price, again corresponding to the intersection of a decreasing step function (the right-hand side) and a continuous increasing function (the left-hand side). Moreover, the step function is clearly bounded between 0 and $e^{-r(T-t)}\pi$ since probabilities are bounded by 0 and 1, which completes the proof of existence and uniqueness. The primary difference from before is that X_t is now a more complicated function, still piecewise linear in D_t and M_t , but determined by stacking all coal and gas bids from lowest to highest, and finding total emissions from all generators needed to match demand. The ordering of coal and gas bids of course depends on C_t , G_t and A_t , and there are a total of $\binom{n_C+n_G}{n_C}$ permutations of the generator bids. Let $\mathcal{P}(C_t, G_t, A_t)$ denote the sequence of power generation conditional on given fuel prices, determined by ranking all bids from lowest to highest. Then let $w_{\mathcal{P}(C_t, G_t, A_t)}^k$, $h_{\mathcal{P}(C_t, G_t, A_t)}^k$, $e_{\mathcal{P}(C_t, G_t, A_t)}^k$ and $f_{\mathcal{P}(C_t, G_t, A_t)}^k$ equal the weight, heat rate, emissions rate and fixed costs for the k -th generator in this sequence (for $k = 1, \dots, n_C + n_G$). Then (suppressing the dependence on (C_t, G_t, A_t) to ease notation),

$$X_t(D_t, M_t; C_t, G_t, A_t) = D^{\max}(D_t + M_t) \left[\sum_{k=1}^{n^*-1} w_{\mathcal{P}}^k e_{\mathcal{P}}^k + e_{\mathcal{P}}^{n^*} \left(R_t - \sum_{k=1}^{n^*-1} w_{\mathcal{P}}^k \right) \right], \quad (6.9)$$

where $D^{\max}$ is again maximum demand (or capacity) in the power market, and

$$n^* = \min \left\{ n : \sum_{k=1}^n w_{\mathcal{P}}^k \geq R_t \right\}.$$

Note also that the power price S_t is of course determined by the bid of the n^* -th generator in $\mathcal{P}(C_t, G_t, A_t)$, and hence

$$S_t(D_t, M_t; C_t, G_t, A_t) = \begin{cases} f_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} + h_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} C_t + e_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} A_t & \text{if the } n^*\text{-th uses coal.} \\ f_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} + h_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} G_t + e_{\mathcal{P}(C_t, G_t, A_t)}^{n^*} A_t & \text{if the } n^*\text{-th uses gas.} \end{cases} \quad (6.10)$$

As in the special case of $n_C = n_G = 1$ from Section 6.2, we have not defined X_t fully, as we have not defined its value when A_t exactly equals a ‘fuel-switching price’, meaning a value for which $f_c^i + h_c^i C_t + e_c^i A_t = f_g^j + h_g^j G_t + e_g^j A_t$ for some $i \in 1, \dots, n_C$ and $j \in 1, \dots, n_G$. For these cases, we can define X_t analogously to (6.2) (using a weighted average of the relevant emissions rates e_c^i and e_g^j), although we only require this definition for simulation purposes, not for finding A_t .

6.3.2 Solving for Carbon Prices

As before, the fuel-switching prices (now several of them) form a key part of the numerical procedure for finding A_t via (6.8). Indeed, for a fixed C_t and G_t , they determine the points at which the right-hand side of (6.8) jumps downwards as we increase A_t . Figure 6.6 plots this function for the case of $n_C = n_G = 2$. The carbon price is determined by finding the intersection of this decreasing step function with the left-hand side of (6.8), which is simply A_t , also shown in the left graph. The




Figure 6.6: Left graph illustrates the existence of A_t as the solution of (6.8) (for $n_C = n_G = 2$, $\pi = 40$) in the two possible cases: $A_t = E_t^k$ for some $k \in \{1, \dots, 5\}$ (blue line) or $A_t = F_t^k$ for some $k \in \{0, \dots, 5\}$ (red line). Here all four fuel-switching prices fall in the range $(0, \pi)$, so $F_t^0 = 0$ and $F_t^5 = \pi$. The right graph illustrates X_t as a function of D_t (with $M_t = 1 - D_t$ for simplification) for various sequences of power generation. The red dotted line corresponds to using the cleaner coal generator first (c1), then cleaner gas (g1), then dirtier gas (g2), then dirtier coal (c2).

blue and red lines (and corresponding change in F_t^3) illustrate the two possible cases, with A_t either exactly equal to a fuel-switching price, or between prices, as we shall discuss further below. Figure 6.6 (right graph) gives an example of emissions X_t as a function of D_t for all possible permutations of the power generation sequence.⁸ The dotted red line illustrates the generation sequence $\{c1, g1, g2, c2\}$ which may correspond to a certain range of carbon prices. However, for a given C_t and G_t , only 5 of these 6 pathways will be possible, since there are four fuel-switching prices. As A_t increases we clearly work our way downwards through this mesh, as cleaner generators move to earlier in the merit order. Each path through the mesh corresponds to one of the steps E_t^k illustrated in the left graph.

Now let F_t equal a matrix of fuel switching prices given by

$$F_t = \begin{pmatrix} \frac{h_g^1 G_t - h_c^1 C_t + f_g^1 - f_c^1}{e_c^1 - e_g^1} & \dots & \frac{h_g^1 G_t - h_c^{n_C} C_t + f_g^1 - f_c^{n_C}}{e_c^{n_C} - e_g^1} \\ \vdots & \ddots & \vdots \\ \frac{h_g^{n_G} G_t - h_c^1 C_t + f_g^{n_G} - f_c^1}{e_c^1 - e_g^{n_G}} & \dots & \frac{h_g^{n_G} G_t - h_c^{n_C} C_t + f_g^{n_G} - f_c^{n_C}}{e_c^{n_C} - e_g^{n_G}} \end{pmatrix}.$$

Thus the (i, j) -th component $F_{t,i,j}$ corresponds to the carbon price which causes the bid from the j -th coal generator (or group of generators) to equal the bid from the i -th gas generator. We are only interested in values $F_{t,i,j}$ in the range $(0, \pi e^{-r(T-t)})$, as these determine the possible sequences of power generation and the ranges of A_t to which they correspond. There are at most $n_C n_G + 1$ such sequences⁹ ranging from all coal bids used first (for low A_t) to all gas bids first (for high A_t).

⁸Note that we plot the simpler case of $D_t = 1 - M_t$ here (so there is perfect capacity availability at all times) which leads to a nice recombining mesh.

⁹There will be exactly this number if $F_{t,i,j} \in (0, \pi e^{-r(T-t)}) \forall (i, j)$ unless two pairs of gas and coal bids coincide simultaneously; i.e., unless $f_c^i + h_c^i C_t + e_c^i A_t = f_g^j + h_g^j G_t + e_g^j A_t$ and $f_c^k + h_c^k C_t + e_c^k A_t = f_g^l + h_g^l G_t + e_g^l A_t$ for some $i, k \in 1, \dots, n_C$ and $j, l \in 1, \dots, n_G$ with $(i, j) \neq (k, l)$.

The ordering of the components of $F_{t,i,j}$ determines the relevant permutations of generators needed to find the equilibrium allowance price A_t . While $F_{t,i,j}$ is always increasing in i and decreasing in j , the overall order of components varies with C_t and G_t . For example, suppose $n_C = n_G = 3$, $\pi = 40$, $r = 0$ and we observe

$$F_t = \begin{pmatrix} 30 & 10 & -5 \\ 45 & 20 & 5 \\ 80 & 70 & 55 \end{pmatrix}.$$

Then we have for 4 fuel switching prices in the range $(0, \pi)$ and hence 5 possible sequences of power generation as follows (where C1,C2,C3,G1,G2,G3 represent the coal and gas generators, from cleanest to dirtiest):

$$\mathcal{P}(C_t, G_t, A_t) = \begin{cases} \text{C1, C2, G1, C3, G2, G3} & \text{for } A_t < 5 \\ \text{C1, C2, G1, G2, C3, G3} & \text{for } 5 < A_t < 10 \\ \text{C1, G1, C2, G2, C3, G3} & \text{for } 10 < A_t < 20 \\ \text{C1, G1, G2, C2, C3, G3} & \text{for } 20 < A_t < 30 \\ \text{G1, C1, G2, C2, C3, G3} & \text{for } A_t > 30. \end{cases}$$

Let $F_t^k(C_t, G_t)$ represent the k -th smallest element of the set $\{f = F_{t,i,j} : f \in (0, \pi e^{-r(T-t)})\}$, containing n_F components (where $n_F \leq n_C n_G$). Define $F_t^0(C_t, G_t) = 0$ and $F_t^{n_F+1}(C_t, G_t) = \pi e^{-r(T-t)}$ to handle the highest and lowest regions of A_t . Then for $k = 1, \dots, n_F + 1$, define the sequence of power generation $\tilde{\mathcal{P}}^k(C_t, G_t)$ to correspond to the case that $A_t \in (F_t^{k-1}(C_t, G_t), F_t^k(C_t, G_t))$. Finally, for $k = 1, \dots, n_F + 1$, set

$$E_t^k(C_t, G_t, \tilde{Y}_t) = \mathbb{E}_t^{\mathbb{Q}} [A_{t+\Delta t} | C_t, G_t, \tilde{Y}_t, \mathcal{P}^k(C_t, G_t)].$$

We can now solve for $A_t(C_t, G_t, \tilde{Y}_t)$ using (6.9) and the following relationship (removing the dependence on (C_t, G_t) for notational simplicity):

$$A_t = \begin{cases} E_t^{k^*}(\tilde{Y}_t) & \text{if } \exists k^* \in \{1, \dots, n_F + 1\} \text{ such that } E_t^{k^*}(\tilde{Y}_t) \in (F_t^{k^*-1}, F_t^{k^*}) \\ F_t^{k^*-1} & \text{if } \exists k^* \in \{1, \dots, n_F + 1\} \text{ such that } E_t^{k^*}(\tilde{Y}_t) < F_t^{k^*-1} \text{ and } E_t^{k^*-1}(\tilde{Y}_t) > F_t^{k^*-1}. \end{cases} \quad (6.11)$$

The two cases above are mutual exclusive and one must always hold. It is easy to see this visually in Figure 6.6. With the slightly higher value of F_t^3 (and blue lines), the diagonal line A_t intersects the blue step function (the right-hand side of (6.8)) at a horizontal section which corresponds to $E_t^3(\tilde{Y}_t)$ in this case (the third step). However, with a small decrease in F_t^3 , the diagonal line ‘intersects’ the red step function at a vertical section, so A_t must equal one of the fuel switching prices, in this case F_t^3 , the third smallest positive one. So the blue step function corresponds to the first case in (6.11) while the red step function corresponds to the second. Note that as we increase n_C and/or n_G , clearly the probability of being in the second case decreases (and converges to zero as $n_C, n_G \rightarrow \infty$). The algorithm for finding A_t proceeds as follows for each node (C_t, G_t) in the fuel price trees discussed in Section 6.2.

1. Find the matrix F_t of fuel switching prices and hence the ordered sequence F_t^k , $k = 1, \dots, n_F$ of prices which lie in the range $(0, \pi e^{-r(T-t)})$.
2. For the corresponding sequences of power generation $\tilde{\mathcal{P}}^k(C_t, G_t)$, $k = 1, \dots, n_F + 1$, calculate $X_t(D_t, M_t)$ using (6.9) for each demand and margin scenario.

3. Beginning with the lowest value of accumulated emissions $\tilde{Y}_t$ in the tree, find $E_t^k(C_t, G_t, \tilde{Y}_t)$ for $k = 1$ (if necessary interpolating linearly in the $\tilde{Y}_t$ dimension between values $A_{t+\Delta t}$).
4. While $E_t^k(C_t, G_t, \tilde{Y}_t) > F_t^k$, let $k = k + 1$ and recalculate $E_t^k(C_t, G_t, \tilde{Y}_t)$ for the next permutation of power generation. Repeat until reaching condition in Step 5.
5. If $E_t^k(C_t, G_t, \tilde{Y}_t) < F_t^k$, stop and check if $E_t^k(C_t, G_t, \tilde{Y}_t) < F_t^{k-1}$ too. If so, we have the second case above, so $A_t = F_t^{k-1}$. If not, $A_t = E_t^k(C_t, G_t, \tilde{Y}_t)$ as in the first case.
6. Move on to the next (higher) value of $\tilde{Y}_t$ and repeat from Step 4 onwards. Note that we do not need to return to the first power generation sequence ($k = 1$) since A_t is increasing in $\tilde{Y}_t$.

6.3.3 The Impact of Fuel Switching

As in the simpler case of Section 6.2, incorporating ‘fuel-switching’ (not to be confused with merit order changes which automatically happen at ‘fuel-switching prices’ $F_t^k(C_t, G_t)$) requires making an assumption about which coal generators can switch to gas and when they can. The simplest approach (analogous to that of Section 6.2) is to assume that for any group of coal generators which bids below or at the highest gas bids, a fraction α switch to burning gas with emissions rate given by the highest gas generators below them in the merit order. Thus for any $k \in \{1 \dots, n_C + n_G\}$ corresponding to the ordered sequence of generation $\mathcal{P}(C_t, G_t, A_t)$, let $I = \{i : i < k \text{ and the } i\text{-th generator uses gas}\}$. Then

$$X_t(D_t, M_t; C_t, G_t, A_t) = D^{\max}(D_t + M_t) \left[\sum_{k=1}^{n^*-1} w_{\mathcal{P}}^k \hat{e}_{\mathcal{P}}^k + e_{\mathcal{P}}^{n^*} \left(R_t - \sum_{k=1}^{n^*-1} w_{\mathcal{P}}^k \right) \right], \quad (6.12)$$

where

$$n^* = \min \left\{ n : \sum_{k=1}^n w_{\mathcal{P}}^k \geq R_t \right\},$$

and

$$\hat{e}_{\mathcal{P}}^k = \begin{cases} \alpha e_{\mathcal{P}}^{\max(I)} + (1 - \alpha) e_{\mathcal{P}}^k & \text{if } I \neq \emptyset \text{ and the } k\text{-th generator uses coal.} \\ e_{\mathcal{P}}^k & \text{otherwise.} \end{cases}$$

Note that fuel-switching under this assumption does not affect the sequences of power generation $\mathcal{P}(C_t, G_t, A_t)$ or their ordering as A_t varies, but only the reduces the amount of emissions produced by some of the sequences. The definition of $\hat{e}_{\mathcal{P}}^k$ above suggests however that a relatively inefficient (or dirty) coal generator could switch to using gas and emit at the rate of a relatively efficient (or clean) gas generator. As a result a more realistic approach for the case that $n_G = n_C$ might be to allow the i -th coal generator to switch to gas if and only if $A_t > F_{t,i,i}$, meaning its bid is higher than that of the i -th gas generator. Let $J_c(i)$ and $J_g(i)$ equal the indices of the i -th coal and i -th gas generators respectively in the sequence $\mathcal{P}(C_t, G_t, A_t)$. Then define for $i = 1, \dots, n_C (= n_G)$,

$$\hat{e}_{\mathcal{P}}^k = \begin{cases} \alpha e_{\mathcal{P}}^{J_g(i)} + (1 - \alpha) e_{\mathcal{P}}^{J_c(i)} & \text{if } k = J_c(i) \text{ for some } i \text{ and } J_g(i) < J_c(i) \\ e_{\mathcal{P}}^k & \text{otherwise.} \end{cases}$$

Clearly any of these assumptions is an approximation to reality, since the decision for a power generator to switch fuels may depend on not just on current prices but factors such as transaction and operation costs, its own forward contracts and expectations of price movements.

6.4 Carbon in the full Bid Stack Model

Chapter 2 presented our core model for the bid stack, using continuous probability distributions to capture the clusters of bids from different fuel types. The resulting continuous bid stack function $S_t = B(R_t; C_t, G_t)$ provided an intuitive and realistic approximation to the complicated step function observed in the market, and thus a successful model of electricity prices. However, it is now very important to be able to incorporate carbon emissions prices into any fundamental approach to power price modelling. While we do not have historical bid data available from regions with active carbon markets (i.e., Europe), we can suggest their likely impact on generator bids as we did in Section 6.3. We simply allow the parameters of our bid distributions to be linked to emissions allowance prices A_t as we did earlier for fuel prices.

6.4.1 Model Assumptions

For each of the bid distributions considered in Chapter 2, we set the location and shape parameters m_i and s_i to depend linearly on C_t or G_t . Extending this to carbon for the two-fuel case, we let

$$m_1 = \alpha_0^C + \alpha_1^C P_t + \alpha_2^C A_t, \quad m_2 = \alpha_0^G + \alpha_1^G G_t + \alpha_2^G A_t$$

and

$$s_1 = \beta_0^C + \beta_1^C P_t + \beta_2^C A_t, \quad s_2 = \beta_0^G + \beta_1^G G_t + \beta_2^G A_t.$$

Given sufficient historical data, we could of course estimate the new parameters $\{\alpha_2^C, \alpha_2^G, \beta_2^C, \beta_2^G\}$ by linear regression along with the original parameters. In the absence of such data, we can make reasonable assumptions in order to test the model. In particular, the discussion at the end of Chapter 2 indicated that a Gaussian distribution of bids can correspond exactly to a Gaussian distribution of heat rates (or efficiencies) among generators of each fuel type. If all generators are equally polluting per amount of fuel used, then the emissions rates of generators should also be Gaussian (simply a multiple of heat rate). Let γ^C equal the quantity of carbon produced per ton of coal burned, and similarly γ^G for gas. Then for the Gaussian case at least (and by extension for other distributions), we can argue that a sensible choice of parameters would be

$$\alpha_2^C = \gamma^C \alpha_1^C, \quad \alpha_2^G = \gamma^G \alpha_1^G, \quad \beta_2^C = \gamma^C \beta_1^C, \quad \beta_2^G = \gamma^G \beta_1^G.$$

A variety of effects such as variation in pollution rates unrelated to fuel intake and strategic generator bidding may prevent the above relationships from holding, but they serve as a useful approximation.

Under these assumptions and in the Gaussian case, emissions (over the period $[t, t + \Delta t]$) are now given by the following:

$$\begin{aligned} X_t = \frac{D_t D^{\max}}{R_t} & \left[\int_0^{S_t} w_1 \left(\alpha_3^C + \beta_3^C \frac{x - m_1(C_t, A_t)}{s_1(C_t, A_t)} \right) \phi \left(\frac{x - m_1(C_t, A_t)}{s_1(C_t, A_t)} \right) dx \right. \\ & \left. + \int_0^{S_t} w_2 \left(\alpha_3^G + \beta_3^G \frac{x - m_2(G_t, A_t)}{s_2(G_t, A_t)} \right) \phi \left(\frac{x - m_2(G_t, A_t)}{s_2(G_t, A_t)} \right) dx \right], \end{aligned} \quad (6.13)$$

where $w_1 \Phi \left(\frac{S_t - m_1(C_t, A_t)}{s_1(C_t, A_t)} \right) + w_2 \Phi \left(\frac{S_t - m_2(G_t, A_t)}{s_2(G_t, A_t)} \right) = R_t,$

and ϕ and Φ are the density and distribution function for a standard Gaussian. This expression follows from the assumption that the emissions rate of a generator bidding at n standard deviations above (or below) m_i , will emit at a rate $\alpha_3 + \beta_3 \frac{x-m_i}{s_1}$; i.e., n standard deviations above (or below) the average emissions rate.¹⁰ Therefore the intuition above is that we are finding an overall market emissions rate (CO₂ per MWh) by averaging over the emissions rates of all generators used at t , meaning all generators who bid below the market clearing power price S_t .

Under a simpler but less realistic framework, we could instead assume that all coal and gas generators have the same emissions rate e_c or e_g irrespective of their bids. This would suggest that the width of bid distributions should not depend on allowance prices, so we should have

$$\alpha_2^C = e_c, \quad \alpha_2^G = e_g, \quad \beta_2^C = 0, \quad \beta_2^G = 0.$$

The advantage here is a slightly simpler expression for X_t , as

$$X_t = \frac{D_t D^{\max}}{R_t} \left[w_1 e_c \Phi \left(\frac{S_t - m_1(C_t, A_t)}{s_1(C_t)} \right) + w_2 e_g \Phi \left(\frac{S_t - m_2(G_t, A_t)}{s_2(G_t)} \right) \right],$$

where

$$w_1 \Phi \left(\frac{S_t - m_1(C_t, A_t)}{s_1(C_t)} \right) + w_2 \Phi \left(\frac{S_t - m_2(G_t, A_t)}{s_2(G_t)} \right) = \frac{D_t}{D_t + M_t}.$$

Note that in both of the cases above, similar expressions can be written down for other bid stack distributions such as the logistic, Cauchy and Weibull, replacing ϕ and Φ by the appropriate density and distribution functions. In these non-Gaussian cases, the calculation of emissions rates in terms of the number of standard deviations for different points in the bid stack is an approximation but quite reasonable.

6.4.2 Solving for the Carbon Price

In the case of the full bid stack model with a continuous bid stack, (6.8) still holds as before:

$$A_t = e^{-r(T-t)} \pi P \left\{ Y_t + \sum_{u=t}^T X_u(D_u, M_u; C_u, G_u, A_u) > \theta \right\}. \quad (6.14)$$

Moreover, its right-hand side is now continuous, since X_t is now a continuous strictly decreasing function of A_t .

However, the implementation of the model is extremely intensive computationally, even for a small time horizon. This is because at each value $(C_t, G_t, \tilde{Y}_t)$ in our trees, we need to perform a root-finding algorithm to solve (6.8) for A_t . Each step of this algorithm involves finding the electricity price S_t and corresponding emissions X_t for each of the $n_D n_M$ demand and margin scenarios. As we know, finding S_t itself involves a root-finding algorithm, since the two-fuel bid stack can not be written explicitly.

To assess the computation time required, we compare this procedure (performed using the `fzero` function in Matlab), with the simplified framework of Section 6.3, but for increasingly large values

¹⁰For this to hold strictly, we require any variation in the fixed costs of different generators (translating into non-zero values for β_0) to be perfectly correlated with variation in their heat rates and emission rates.

of n_G and n_C . The model of Section 6.3 can be used to approximate the full bid stack model by choosing the fixed costs, heat rates and emissions rates appropriately. For example, we can use a quantile approach to discretising the bid distributions into n_G and n_C point masses of equal weight. We use the Gaussian case described by (6.13), which essentially assumes that a coal generator chosen at random has a bid given by the random variable $Z = f_c + h_c C_t + e_c A_t$ where

$$f_c \sim N(\mu_{fc}, \sigma_{fc}^2), \quad h_c \sim N(\mu_{hc}, \sigma_{hc}^2), \quad e_c = \gamma^C h_c, \quad \text{and } f_c \text{ and } e_c \text{ are perfectly correlated,}$$

implying

$$Z \sim N(\mu_{fc} + \mu_{hc} C_t + \gamma^C \mu_{hc} A_t, (\sigma_{fc} + \sigma_{hc} C_t + \gamma^C \sigma_{hc} A_t)^2),$$

and similarly for gas generators. Under these assumptions the parameters m_1, s_1, m_2, s_2 satisfy

$$\text{For } m_1 : \quad \alpha_0^C = \mu_{fc}, \quad \alpha_1^C = \mu_{hc}, \quad \alpha_2^C = \gamma^C \mu_{hc}$$

$$\text{For } s_1 : \quad \beta_0^C = \sigma_{fc}, \quad \beta_1^C = \sigma_{hc}, \quad \beta_2^C = \gamma^C \sigma_{hc}$$

$$\text{For } m_2 : \quad \alpha_0^G = \mu_{fg}, \quad \alpha_1^G = \mu_{hg}, \quad \alpha_2^G = \gamma^G \mu_{hg}$$

$$\text{For } s_2 : \quad \beta_0^G = \sigma_{fg}, \quad \beta_1^G = \sigma_{hg}, \quad \beta_2^G = \gamma^G \sigma_{hg}.$$




Figure 6.7: Carbon price surfaces at $t = 1/4$ ($T = 1/2$) for the case of $n_C = n_G = 2$ (left graph) and $n_C = n_G = 6$ (right graph).

Now we translate this to a model of point masses of bids by defining the weights $w_c^1, w_c^2, \dots, w_c^{n_C}$, heat rates $h_c^1, h_c^2, \dots, h_c^{n_C}$, corresponding emissions rates $e_c^1, e_c^2, \dots, e_c^{n_C}$ and fixed costs $f_c^1, f_c^2, \dots, f_c^{n_C}$ of the n_C coal generators. For $i = 1, \dots, n_C - 1$, let $q_{fc}(i)$ and $q_{hc}(i)$ represent the i -th n_C -quantile of the Gaussian random variables f_c and h_c respectively (fixed costs and heat rates). Set $q_{fc}(0) = q_{hc}(0) = -\infty$ and $q_{fc}(n_C) = q_{hc}(n_C) = \infty$. Then for $i = 1, \dots, n_C$, let

$$\begin{aligned} w_c^i &= w_1/n_C \\ f_c^i &= n_C \int_{q_{fc}(i-1)}^{q_{fc}(i)} x \phi\left(\frac{x - \mu_{fc}}{\sigma_{fc}}\right) dx \\ h_c^i &= n_C \int_{q_{hc}(i-1)}^{q_{hc}(i)} x \phi\left(\frac{x - \mu_{hc}}{\sigma_{hc}}\right) dx \\ e_c^i &= \gamma^C h_c^i. \end{aligned}$$

We repeat this procedure for the natural gas generators, of course using weights $w_g^i = (1 - w_1)/n_G$ and parameters corresponding to gas.




Figure 6.8: Carbon price as a function of $\tilde{Y}_t$ for $G_t = 7$ as n_C and n_G increase (left); Log-plot of computation time as function of n_C (or n_G) (right).

To test this approach, we now return to the parameter estimates of Chapter 2 for the PJM bid stack (and set $r = 0$). In particular, we use Table 2.2 again (full time period) and find all required parameters as described above. However, we reduce the estimates for μ_{fg} and σ_{fg} (describing the fixed costs for gas generators) by 20 and 10 respectively, in order to increase the occurrence of changes in the ordering of gas and coal generators for $A_t \in (0, \pi)$.¹¹ For all other parameters we retain those from the previous section (for the case of constant coal prices) but now choose $T = 1/2$ as we are no longer comparing with the GBM case. Figure 6.7 shows the resulting carbon price surface for the two cases of $n_C = n_G = 2$ and $n_C = n_G = 6$. In the former, we can still observe two regions where $A_t \approx F_t^k$ for some $k \in \{1, \dots, n_C n_G\}$, and the allowance price is roughly linear in G_t . However, in the latter case, there exist too many ‘fuel-switching prices’ F_t^k and consequently the regions are no longer distinct for our chosen node spacing. The price surface is converging to the smoother surface which we observe using the continuous bid stack and (6.13).

Figure 6.8 illustrates the speed of this convergence as we increase n_C and n_G (with $n_C = n_G$) to better discretise the bid distributions. As in Figure 6.5 earlier, we take a cross-sectional view at $G_t = 7$. For $n_C = n_G = 6$ or more, there is little visible differences between the plots of A_t against $\tilde{Y}_t$. Table 6.1 provides further analysis of this convergence with values of A_t listed for $G_t = 6.1$ or 9.8, and $\tilde{Y}_t = \pm 3000$. Note that $N = 60$ in this example and there are 21 nodes in the gas price tree. The table indicates that the error in the carbon price (relative to the continuous bid stack) is typically less than \$0.10 even for $n_C = n_G = 6$. However, beyond this point the convergence is quite slow while the computation time increases fairly rapidly. Plotting the run time (in MATLAB) against n_C on a log-scale we obtain the scatter plot in Figure 6.8 which seems to have a slope slightly below 2. Thus we conclude that the computation time is approximately of order n_C^2 , which

¹¹Alternatively, we could increase the penalty π to produce this effect, but we keep $\pi = 40$ as before. These required adjustments indicate that much higher carbon prices than those observed in the early stages of RGGI will be necessary to bring about substantial changes in the PJM market’s behaviour.

Table 6.1: Analysis of convergence and computation time as discretising of bid distributions improves (i.e., for increasing n_C and n_G).

n_C (= n_G)	run time (seconds)	G=6.1 Y=3000	G=9.8 Y=3000	G=6.1 Y=-3000	G=9.8 Y=-3000
2	210.7	15.010	39.529	3.117	35.534
4	357.4	13.285	39.255	2.746	34.927
6	609.5	12.689	39.191	2.636	34.777
8	927.8	12.584	39.190	2.640	34.807
10	1346.2	12.415	39.191	2.630	34.756
12	1946.1	12.332	39.165	2.626	34.750
14	2522.3	12.270	39.174	2.611	34.711
16	3203.4	12.268	39.162	2.615	34.776
18	4176.2	12.232	39.185	2.608	34.739
20	5014.7	12.225	39.172	2.646	34.752

is intuitive considering that the number of relevant fuel-switching prices $F_t^k \in (0, \pi)$ should remain roughly proportionally to $n_C n_G$.

Finally, we observe that although the middle regions of our carbon price surfaces appears to disappear in Figure 6.8 as we increase n_C and n_G , this is only true near the end of the trading period. As we move further back in time, we again find clear separations emerging between regions of high sensitivity to G_t and regions of high sensitivity to $\tilde{Y}_t$. Figure 6.9 plots A_t at $t = 1$ with $T = 3$, $N = 60$, $n_C = n_G = 6$, and all other parameters as before (coal constant and gas an exponential OU), and resembles the shape of our original plot in Figure 6.3. As we move further back towards the beginning of the trading period, the derivative $\frac{\partial A}{\partial Y}$ in the middle region decreases, again suggesting that there is still enough time for the market mechanism of merit order changes to offset changes to emission levels. The contour plot of Figure 6.9 illustrates this by showing the changes over time in contour levels $A_t=1, 20$, and 39 . As we can see, for $t = 0$ (3 years before the end of the trading period), all three contours are almost flat for large ranges of $\tilde{Y}_t$. Interestingly, while our simplest initial model of Section 6.2 is too crude near T , tests suggest that it may become a reasonable approximation for the early portion of a trading period.

6.4.3 Simulation Analysis

Further insight into carbon price dynamics can be gained through Monte Carlo simulation analysis. In particular, we can investigate the potential impact of carbon on power price dynamics, as well as correlations of both power and carbon with fuel prices. The surface plots of $A_t(G_t, \tilde{Y}_t)$ highlight the variation we are likely to observe from simulation to simulation, since the dynamics may be quite different depending on which region of the surface our sample path spends most time in.

In order to simulate all prices C_t, G_t, S_t, A_t , as well as emissions Y_t , we first use our tree-based approach to calculate the surface $A_t(C_t, G_t, \tilde{Y}_t)$ for all t in the simulation period. We then simulate forward through the fuel price trees, as well as simulating values of D_t and M_t (from their invari-




Figure 6.9: Carbon price surface as a function of G_t and $\tilde{Y}_t$ for $t = 1$ and $T = 3$ (left graph); Contour plot of carbon price surfaces at $t = 0, 1$ and 2 . (for contours $A_t = 1, 20$ and 39) (right graph).

ant distribution) to calculate the progression of emissions $\tilde{Y}_t$. At each time we also find S_t using (6.10). Note that throughout this section, we are ignoring the impact of risk premiums, since we are using the same probabilities of fuel price movements to simulate dynamics as we did to find A_t . However, we are not comparing results to any market data, but instead analysing general behaviour, so differences between $\mathbb{P}$ and $\mathbb{Q}$ can be safely ignored for our purposes. One advantage of the trinomial tree approach is that we can easily use the resulting grids of values to simulate forwards in time, with the probability structure already set up. Moreover, if we calculate prices of other contracts along with $A_t(C_t, G_t, \tilde{Y}_t)$ as we move backwards through the grid, then these can also be simulated, as well as allowing for the pricing of derivatives on these contracts. For example, we could calculate the forward power price $F^P(t, T)$ at all points $(C_t, G_t, \tilde{Y}_t)$ in the trees by starting with the appropriate payoff S_T at maturity and working backwards, before then pricing options on power forwards $V^P(t, T_o, T)$ by simulating through the tree from t to T_o and evaluating the option payoff at T_o .

For simulation purposes we retain the same parameters as above with a three year trading period and with $n_C = n_G = 6$ as an approximation to the full bid stack. As before, we keep coal fixed for to save computation time but note that the extension to stochastic coal prices is straight-forward. (The impact of coal tested over shorter time periods is similar to that of gas but with the direction reversed and is somewhat weaker due to lower volatility.) Figure 6.10 provides two sets of sample paths from the simulation analysis. In the first of these (plotted in the first row), emissions (and power) demand remains low for the first half of the period, with carbon and power prices also low. As emissions then pick up in the later stages (due in part to gas price rises), carbon prices increase to the almost the penalty level well before the end of the trading period. In the second example, however, emissions levels stay within 1000 tonnes (relative to the cap's schedule) throughout, so the final state of the market (i.e., a shortfall or surplus of credits) at T remains in doubt until almost the very end. As a result, correlation of carbon with gas prices remains high throughout and the two price peaks clearly coincide. Power also follows closely, though with high volatility of course. We can also clearly see in both samples that the existence of a carbon market tends to push high power prices even higher while low prices are less affected, though of course is not always the case as capacity issues alone can be responsible for high carbon prices. Overall, these sample paths appear




Figure 6.10: Each row above corresponds to a sample simulation from $t = 0$ to $T = 3$ and parameters as earlier ($n_C = n_G = 6$). The left graph shows power prices S_t (with and without including the carbon market), the middle graph shows gas prices G_t and carbon prices A_t , while the right graph shows emissions levels $\tilde{Y}_t$ relative to the market cap.

reasonable in light of market structure as well as the historical price behaviour illustrated in Figure 6.1.

In addition to looking at individual sample simulations, we can assess the variation among simulations by observing the distribution of certain statistics such as realised correlations of log returns. Note that these are calculated using only prices up until the point that the carbon price hits 0 or $\pi e^{-r(T-t)}$, as there is clearly zero correlation beyond this point. The first row of Figure 6.10 plots histograms of realised correlations based on 10,000 simulations. The first two graphs illustrate that the inclusion of a carbon market can amplify the correlations of power and gas prices, as the mean of the histogram shifts from about 0.2 to 0.35 when the effect of carbon is added. Of course, we should recall that coal prices are fixed in these simulations, suggesting that with a carbon market we observe a greater sensitivity to gas relative to power demand (and margin). This is intuitive since gas prices now factor in both directly as a production cost and indirectly through carbon. The third graph shows an interesting result, revealing a wide range of possible correlations between carbon and gas. There is almost the occurrence of a double hump, with high correlations near 0.8 very common and mid range correlations near 0.5 very common. This makes sense in light of the distinct regions discussed above. Clearly it is important for power generators and energy companies to understand these unusual dependencies in order to hedge their risks appropriately.




Figure 6.11: The upper row shows histograms of realised sample correlation (of log returns) over the three year period: (i) power against gas, (ii) power against gas (no carbon market), (iii) carbon against gas. The lower row shows histograms of simulated carbon prices at various points in time ($t = 1$, $t = 2$ and $t = 3 (= T)$).

6.5 Discussion

Carbon emissions markets are still a young and rapidly growing field, presenting many challenges and opportunities for policy makers and businesses, as well as researchers from many different disciplines. From a price modelling perspective, we have seen that a relatively simplistic approach to understanding all the key carbon price drivers still leads to a fairly complicated and challenging model to implement. Basing this approach on the bid stack model previously constructed successfully for electricity markets provides a natural link between all the main underlying drivers, as well as giving additional credibility to the framework. Nevertheless, there are still many questions to be answered and improvements to be made, both within this framework and through alternative approaches geared towards different goals.

Firstly, there are a number of interesting questions to explore within the bid stack approach advocated here. Clearly, one of the key advantages of the model is its ability to describe the unusual dependency structure between carbon, power, gas and coal prices, as illustrated for example by the surface plots of Figures 6.3, 6.7 and 6.9 and the simulations of Section 6.4.3. In a multi-fuel market, the dominant feature to capture is the automatic mechanism of merit order changes caused by the incentive to use cleaner fuels when A_t increases. While this has been the main focus of this chapter, it remains to be seen (and analysed) as to how closely the theory will hold in practice. Relatively low carbon prices have so far kept both merit order changes and fuel switching fairly low, though contact with industry (e.g., through conferences) suggests that it is an increasingly debated

topic. As more data becomes available from the EU ETS and an American carbon market begins to develop, there should soon be sufficient historical data to construct calibration techniques and test the assumptions of the framework. One important question is what exactly determines if and when each generator will decide to switch fuels. How significant or costly are the obstacles to switching fuels and how flexible are most generators? For example, it may be appropriate for some generating units to consider fuel switching as an American option problem, if the costs are such that it is irreversible over a trading period. On the other hand, the optionality may more closely resemble a swing option if there are a limited number of times that a switch can be made, or if there is a waiting period between switches. As multiple exercise options are already prominent in energy and electricity markets (e.g., through gas storage problems or standard power swing options), there are perhaps links that can be explored here.

While this chapter has focused solely on the multi-fuel case (and specifically the coal-gas case), it is worth mentioning the behaviour of carbon emission prices in the one-fuel case. In our model, the removal of the interplay between coal and gas bids should lead to a complete detachment of A_t from G_t , with accumulated emissions demand process Y_t now dependent solely on D_t , and A_t a derivative on Y_t as before. Avoiding the complications of fuel price movements, it would thus be interesting to use a one-fuel power market to assess the role of other issues, such as the level of dependence between power S_t and carbon A_t when other industries are also trading credits, and when credits are available via alternative sources such as the Clean Development Mechanism. As carbon markets are typically regional and dependent on a number of more local power markets, new modelling ideas are needed to understand and categorise risk factors according to how widespread their spatial influence is (similarly to identifying clear variation in the time scales of risk factors in this thesis). In addition, political risk is unusually large for carbon markets, owing firstly to its obvious impact on the stringency of future caps κ and penalty levels π , as well as the structure of the market itself, the possibility of additional regulation or policies such as subsidies for renewable energy. All of these issues can translate to carbon price movements, but determining how to best to incorporate political and regulatory risk realistically into a model is a daunting task indeed, even when fuel prices are removed from the equation.

Several other aspects of carbon price dynamics provide more tangible questions to address. In our simplified approach, prices must converge to either zero or to the penalty level depending on the total quantity of emissions at the end of the trading period. While the digital option analogy is appealing, it is lacking in a number of ways. Firstly, the CO_2 being emitted is itself neither tradable nor easily observable to the market as a whole, so typical pricing methods may not be appropriate. This raises the possibility of interesting research related to information flow and uncertainty, as generators can realistically observe their own emissions easily but not necessarily those of the entire market. Furthermore, as mentioned earlier, the end of a trading period can only be loosely equated to an option maturity. Thus, existing approaches are unable to properly handle the behaviour of carbon prices at the end of a trading period, which in turn feeds back to all previous times. Ignoring both the possibility of partial ‘banking’ of credits and also the extra penalty of delivering an additional credit in the following period are both crude approximations to reality. For energy companies planning several years ahead and trying to hedge risks, the correlation between carbon allowances of different maturities is of great importance. Furthermore, the valuation of a

project to construct a power plant nowadays requires some sort of model for carbon emissions prices in the longer-term. While it may be too early to answer many of these questions, the derivative products to hedge these risks are beginning to be traded, and the modelling will need to catch up soon. In these exciting and changing times for global carbon and energy markets, we hope that the bid stack model can make a useful contribution in understanding the modelling challenges ahead.

Chapter 7

Conclusion

In today's highly interdependent and energy-intensive world, even a quick glance through most international newspapers is often enough to realise the importance of energy prices in our lives. From discussions on petrol price increases, to concerns over energy security and independence, to debates on the cost of constructing wind farms or nuclear power stations, it is clear that views on energy prices shape the decisions of individuals, companies and governments around the world. Against this backdrop, the growing threat of climate change is also bringing together increasing numbers of scientists, economists, industry leaders, and policy makers, with energy markets playing a major role in the search for answers. While much of the economic price forecasting used for decision making may seem very distant from mathematical modelling of electricity option prices, an important contribution can nonetheless be made from the discipline of mathematical finance. A rigorous mathematical and statistical model can often provide new insights into price dependencies and scenarios which may be beyond the reach of more qualitative approaches. Moreover, unlike traditional stock and bond markets, the field of energy price modelling is a prime example of an emergent area which has yet to be fully explored from all perspectives. In light of both environmental concerns and recent record levels and huge volatilities in global oil, gas, coal and power prices, there is undoubtedly a growing need to understand and model the key risks in these markets, as well as their complex interdependencies.

In this thesis, we aimed to make an important contribution to the growing literature on econometric and mathematical modelling of electricity prices, but with a particular focus on links with other energy prices. This focus sets our work apart from most other approaches which treat correlations with energy prices as an optional addition or else a secondary priority to capturing features such as spikes and seasonality. However, as revealed in early plots of monthly natural gas and power prices (e.g., right graph of Figure 2.1), modelling electricity prices without using gas prices equates to missing a key piece of the puzzle, particularly if the model is to then be implemented by companies with multi-commodity exposure, as is almost invariably the case. Yet, despite these observations of correlations with gas prices (or sometimes coal prices), it is well known that for any given hour, electricity demand must be met, and so price is primarily a function of that hour's demand (and available capacity). Only through this link can the short-term volatility and spikes (e.g., left graph of Figure 2.1) be properly explained. Handling this multi-faceted behaviour and

the interaction between these two contrasting price drivers is a significant obstacle, and ultimately one of the major successes of our approach. While a simple fundamental approach such as that of Barlow (2002) ignores the key influence of fuel prices, a more traditional pure spot price model such as that of Deng (1999) treats gas and power together as correlated processes, but ignores the crucial role of demand in setting the marginal price. Successfully integrating both categories of underlying factors into a single model is particularly challenging if we use historical demand, fuel prices and electricity prices as our starting point, but is greatly simplified by basing our approach on market structure. It is from this perspective that the bid stack model of Chapter 2 was born.

The use of observed bid data from PJM and NEPOOL was both the initial motivation and the ultimate justification for the bid stack model created in Section 2.5. In particular, the striking relationships found to exist between the movement of bids in these markets and the movement of coal and gas prices was in itself a remarkable discovery, as well as adding credibility to the pricing results to follow. Though of course it is intuitive that bids should follow generation costs, the distribution-based fitting procedure allows us to notice both the widening and shifting of clusters of bids, and thus to capture the relationship with price intuitively and realistically. It also allows us to assess in more general terms the role of deregulated electricity markets in passing on fuel costs to customers. While in some markets bids may be significantly distorted by active strategic bidding and market power, it appears that the market administrators for PJM and NEPOOL ensure a relatively smooth and fair transition of costs from generators to consumers, which is confirmed also by the existence of PJM's market power monitoring group. Of course this characteristic (as well as the availability and format of bid data) has greatly aided our task of model construction, but this does not imply that the model is not applicable to other markets. On the contrary, the cases of PJM and NEPOOL can be viewed as handy case studies when attempting to adjust the methodology slightly for other power markets. We can also choose to think of the bid stack function not as a description of bid movements themselves but more generally as an innovative and tractable mathematical approach to capturing the transformation from fundamental drivers to power prices.

Indeed, if the first modelling challenge was to understand how to combine the influences of a wide range of underlying variables, then the second modelling challenge was one of tractability. This is important firstly because it sets our approach apart from the common industry optimisation and simulation models, which effectively mimic the role of a market administrator, tallying up individual costs and constraints for hundreds or thousands of generating units. While the distribution-based framework of (2.5) was arguably the key step, the original approaches to modelling demand D_t and M_t in Sections 3.3-3.5 were also a necessary component from a mathematical perspective. These allowed us to represent all of our driving factors ($\{Y_t, Z_t^{\text{OU}}, Z_t^{\text{SP}}\}$ as well as $\{X_t^1, X_t^2, X_t^3\}$ for fuels) as unbounded stochastic processes, and yet capture spikes realistically while avoiding the common obstacle of demand needing to be artificially capped at capacity. Finally, the use of the logistic distribution helped to convert an elegant but computationally intensive framework to a more practical tool for rapid forward and option pricing. This statement is especially true of the one-fuel case presented in Chapter 4, but with the help of the Taylor Series approximation of Chapter 5, was extended to the multi-fuel case for typical regions of the fuel price space (C_t, G_t) .

The expressions for forwards and options in this thesis are appealing not so much from a pure

mathematical perspective, but primarily because they result from an intuitive framework which exploits well-known relationships with underlying observable driving factors in the power markets. Thus, they provide a compromise between the simpler formulas of typical pure spot price models of academia (but based on unobservable, vague or unspecified factors) and the lack of any explicit formulas in more detailed industry models. They also provide a single unified framework to pricing and hedging of single or multi-commodity derivatives, including power and gas derivatives, spark and dark spread options (clean and dirty), demand and weather-dependent contracts, and even the valuation of power stations. Finally, this framework is also flexible enough to incorporate local market information, which may or may not be available, or to use alternative data for calibration. For example, knowledge about anticipated structural changes (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, for example, forward looking capacity or margin data could be incorporated through a time-dependent mean-reversion level μ_Z of the process Z_t^{OU} . Such data have been shown to be useful by Cartea *et al* (2008).

On the other hand, weaknesses of the model include the large number of parameters to be estimated, and the reliance on a large quantity of data. Arguably, however, the complexity of electricity markets justifies a complicated model and the use of more data avoids an over-reliance on historical spot price data. The estimation of risk-neutral dynamics for D_t and M_t has not been fully discussed here and is an extra complication of the model, though calibration of market prices of risk to available derivative data is easily achievable. 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.

As discussed in Section 3.7, another important advantage of the model is its ability to naturally identify and describe the variety of different time scales visible in power market dynamics, varying from very short-term jumps in capacity due to outages, to gradual movements in the long-term levels of coal and gas prices. This useful information allows us to easily determine which derivative contracts are most sensitive to which risk factors and how this varies with maturity. As witnessed in the empirical results in Chapters 4 and 5, options on monthly power forwards have virtually no dependence on current values of demand and margin, but are instead strongly related to options on monthly fuel forwards. Since we treat power prices directly as a function of fuel prices, we can also capture a more complex dependence structure than would be possible in other approaches, for example using correlated Brownian returns or cointegrated price processes. Even in the one-fuel NEPOOL model, where (2.7) leads to a linear dependence on gas, the slope of this relationship is constantly changing due to the seasonality of R_t . Thus we end up with a relationship which mimics cointegration, but is somewhat more involved, helping to explain behaviour of Figure 2.1.

Returning to the two-fuel case, changes in the merit order provide a more fascinating and unusual dependence structure, which other approaches are unable to capture at all. For example, times of high demand typically correspond to higher dependence of power prices on natural gas prices, assuming the typical merit order. However, as gas prices fall very low (or coal prices become very high) then the high demand periods will gradually depend more on coal than gas. This fundamental feature of the model can be illustrated by the following simple graphs using all the factors and parameters estimated for PJM earlier. Figure 7.1 plots the relative sensitivity of a one-year power forward contract $F^P(t, t+1)$ to initial log gas price $\log G_t$ (for a fixed initial price $\log C_t$), by calculating its derivative $\frac{\partial F^P(t, t+1)}{\partial \log G}$ normalised by itself. This is particularly interesting as it reveals clearly the effect described above, with the ordering of the three demand scenarios swapping at some level of gas price. The right graph shows the impact of reducing demand and margin variance to zero, such that we know with certainty the value of demand over capacity, R_{t+1} . As expected this increases the overall sensitivity to fuel prices, but also the difference between sensitivity levels for different values of $\log G_t$. The variation in shape between these three lines is certainly an important strength of the model, providing an additional degree of realism and insight for example as compared with the most similar approach, that of Pirrong and Jermakyan (2008). Under their simpler bid stack model, the relative sensitivities plotted here would remain either strictly increasing or decreasing in G_t (depending on their γ) and very similar to each other in shape, implying a quite different hedging pattern required. Although historical and current gas and coal prices place us to the far right edge of this graph (far from the point of significant merit order change), that may well change in the future, particularly as the extra impact of carbon prices reduces the gap between coal and gas bids. While it is difficult to make a full analysis of the additional benefit of our more complicated model, these simple sensitivity plots emphasise the importance of striving for a more complete understanding of all the interdependencies in energy markets.




Figure 7.1: Relative sensitivity of $F^P(t, t+1)$ to $\log G_t$ for different expected demand scenarios, firstly using all factors and parameters estimated for PJM (left graph) and secondly fixing R_{t+1} by reducing demand and margin variance to zero (right graph)

As investigated in Chapter 6, the bid stack model is also a beneficial tool which can be readily adapted to the new and important topic of carbon emissions markets. While initial findings are interesting and encouraging, it's clearly an area requiring much further study. However, cap-and-trade

markets are just one of many new dimensions and changes occurring in energy markets. As the risk of disastrous climate change begins to dominate political agendas and media headlines, many different ideas are under discussion, most of which are likely to significantly impact energy markets and their price dynamics. While many people agree that a market-based solution involving a cap-and-trade framework for emissions allowances is a key component of any global solution, many also call for much stronger regulation. For example, the recent ambitious EU targets of 20% renewable energy in the power grid by 2020 could massively impact the behaviour of power prices. (As discussed briefly in Chapter 1 and also Section 5.8, renewable power generators typically require different considerations when incorporating them into a fundamental model, as they are at the mercy of often wildly fluctuating resource availability, such as wind speed.) Other policy changes may include subsidies or carbon taxes as an alternative to carbon trading, both of which will clearly impact the merit order in electricity markets. It is interesting to speculate that if these anticipated changes occur with the speed that some are predicting, it is very likely that fundamental modelling approaches such as the one presented here will grow in prominence over pure spot price models. Pure spot price models rely on calibrating long stable price histories and are thus unable to handle significant and recent structural shifts in the market. In contrast, in the bid stack model, all of the historical data for underlying factors would still be as useful and relevant as before, if we are able to represent in the model the cause of the structural change (e.g., a change of fuel type weights). Thus, such approaches may become increasingly popular and important in the coming years and decades, as they allow us to use historical data with more confidence and a better understanding of cause and effect.

Finally, the methodology developed in this thesis can potentially be used to address topics much broader than purely price dynamics. Most notably, the newly formed European Union Emissions Trading Scheme has encountered numerous problems in its first trading period (2005-07), with criticism targeting the lack of auctioning of allowances, windfall profits to power generators, and ultimately an over-allocation of credits leading suddenly to a carbon price collapse (recall Figure 6.1). Clearly there are many lessons to be learned and much work needed in terms of market design and implications for price dynamics, especially in a global setting. While this should certainly draw on a number of different fields of research, the recent work of Carmona *et al* (2008) illustrates the possible contribution of mathematics in the ongoing debate about how to design the optimal cap-and-trade system. Their results from comparing different designs are very interesting and encouraging, but there is room for much more research. Can we provide the right incentives to encourage participation and reduce emissions while avoiding excessive profits and highly volatile prices? Should credits be auctioned at the beginning, given away for free or even allocated dynamically through time? How should the expansion of cap-and-trade into transport and other sectors best be structured? How can allowances be most fairly distributed at a global level, considering huge differences in population, historical emissions and the manufacturing versus consumption balance? Can a single global carbon price emerge from the variety of trading schemes and offsetting programs around the world? While some of these questions have already drawn much research, the literature from a stochastic analysis perspective is very young. To ensure the success of market-based instruments in the fight against climate change, it is important that we adopt a multi-disciplinary approach to understand the challenges and solutions. I therefore believe that the long-term prospects for mathematical research in environmental finance are very promising, and that the energy price model presented here can be a useful stepping stone in this exciting and hugely important direction.

Appendix A

Taylor Series Approximation for the Two-Fuel Bid Stack

A.1 Various Cases of the Bid Stack Approximation

As introduced in Section 5.1, the two-fuel bid stack in the logistic case can be approximated such that the relationship between power forwards and fuel forwards is given by:

$$\begin{aligned}
 F_t^P = & A_1 + A_2^C \log(F_t^C) + A_2^G \log(F_t^G) + A_3^C F_t^C + A_3^G F_t^G + A_4(F_t^C)^{-\delta}(F_t^G)^\gamma + \tilde{A}_4(F_t^C)^{-\tilde{\delta}}(F_t^G)^{\tilde{\gamma}} \\
 & + A_5(F_t^C)^{-2\delta}(F_t^G)^{2\gamma} + \tilde{A}_5(F_t^C)^{-2\tilde{\delta}}(F_t^G)^{2\tilde{\gamma}} + A_6(F_t^C)^{-3\delta}(F_t^G)^{3\gamma} + \tilde{A}_6(F_t^C)^{-3\tilde{\delta}}(F_t^G)^{3\tilde{\gamma}} \\
 & + A_7^C(F_t^C)^{-\delta+1}(F_t^G)^\gamma + \tilde{A}_7^G(F_t^C)^{-\tilde{\delta}}(F_t^G)^{\tilde{\gamma}+1} + A_8^C(F_t^C)^{-2\delta+1}(F_t^G)^{2\gamma} + \tilde{A}_8^G(F_t^C)^{-2\tilde{\delta}}(F_t^G)^{2\tilde{\gamma}+1} \\
 & + A_9^C(F_t^C)^{-3\delta+1}(F_t^G)^{3\gamma} + \tilde{A}_9^G(F_t^C)^{-3\tilde{\delta}}(F_t^G)^{3\tilde{\gamma}+1} \\
 & + A_{10}^C F_t^C \log F_t^C + A_{10}^G F_t^G \log F_t^G + A_{11} F_t^C \log F_t^G + A_{12} F_t^G \log F_t^C
 \end{aligned}$$

Here the coefficients $A_1, A_2^C, \dots$ are found using (5.6) and are given by

$$\begin{aligned}
 A_1 = & \alpha_0^C \mathbb{E}_t^\mathbb{Q} [1_{R_T \leq \tilde{w}_1}] + \alpha_0^G \mathbb{E}_t^\mathbb{Q} [1_{R_T > \tilde{w}_1}] - \beta_0^C \log(w_1) \mathbb{E}_t^\mathbb{Q} [1_{R_T \leq w_A}] + \beta_0^G \log(1 - w_1) \mathbb{E}_t^\mathbb{Q} [1_{R_T > w_B}] \\
 & - \beta_0^C \left(\log(\lambda) - \frac{11}{6} - \frac{\delta \sigma_C^2}{2} + \frac{\gamma \sigma_G^2}{2} \right) \mathbb{E}_t^\mathbb{Q} [1_{w_A < R_T \leq \tilde{w}_1}] + \beta_0^C [\log(R_T) 1_{R_T \leq \tilde{w}_1}] \\
 & + \beta_0^G \left(\log(\tilde{\lambda}) - \frac{11}{6} - \frac{\tilde{\delta} \sigma_C^2}{2} + \frac{\tilde{\gamma} \sigma_G^2}{2} \right) \mathbb{E}_t^\mathbb{Q} [1_{\tilde{w}_1 < R_T \leq w_B}] - \beta_0^G \mathbb{E}_t^\mathbb{Q} [\log(1 - R_T) 1_{R_T > \tilde{w}_1}] \\
 & - \frac{\beta_0^C}{w_A} \left(\log(\lambda) - \frac{11}{6} - \frac{\delta \sigma_C^2}{2} + \frac{\gamma \sigma_G^2}{2} - \log(w_1) \right) \mathbb{E}_t^\mathbb{Q} [R_T 1_{R_T \leq w_A}] \\
 & + \frac{\beta_0^G}{w_B - 1} \left(\log(\tilde{\lambda}) - \frac{11}{6} - \frac{\tilde{\delta} \sigma_C^2}{2} + \frac{\tilde{\gamma} \sigma_G^2}{2} - \log(1 - w_1) \right) \mathbb{E}_t^\mathbb{Q} [(R_T - 1) 1_{R_T > w_B}] \\
 A_2^C = & - \frac{\beta_0^C \delta}{w_A} \mathbb{E}_t^\mathbb{Q} [R_T 1_{R_T \leq w_A}] - \beta_0^C \delta \mathbb{E}_t^\mathbb{Q} [1_{w_A < R_T \leq \tilde{w}_1}] + \beta_0^G \tilde{\delta} \mathbb{E}_t^\mathbb{Q} [1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{\beta_0^G \tilde{\delta}}{w_B - 1} \mathbb{E}_t^\mathbb{Q} [(R_T - 1) 1_{R_T > w_B}] \\
 A_2^G = & \frac{\beta_0^C \gamma}{w_A} \mathbb{E}_t^\mathbb{Q} [R_T 1_{R_T \leq w_A}] + \beta_0^C \gamma \mathbb{E}_t^\mathbb{Q} [1_{w_A < R_T \leq \tilde{w}_1}] - \beta_0^G \tilde{\gamma} \mathbb{E}_t^\mathbb{Q} [1_{\tilde{w}_1 < R_T \leq w_B}] - \frac{\beta_0^G \tilde{\gamma}}{w_B - 1} \mathbb{E}_t^\mathbb{Q} [(R_T - 1) 1_{R_T > w_B}]
 \end{aligned}$$

$$\begin{aligned}
A_3^C &= \alpha_1^C \mathbb{E}_t^{\mathbb{Q}} [1_{R_T \leq \tilde{w}_1}] - \beta_1^C \left(\log(\lambda) - \frac{11}{6} + \frac{\delta \sigma_C^2}{2} + \frac{\gamma \sigma_G^2}{2} - \gamma \rho_{CG} \sigma_C \sigma_G \right) \mathbb{E}_t^{\mathbb{Q}} [1_{w_A < R_T \leq \tilde{w}_1}] \\
&\quad - \beta_1^C \log(w_1) \mathbb{E}_t^{\mathbb{Q}} [1_{R_T \leq w_A}] + \beta_1^C \mathbb{E}_t^{\mathbb{Q}} [\log(R_T) 1_{R_T \leq \tilde{w}_1}] \\
&\quad - \frac{\beta_1^C}{w_A} \left(\log(\lambda) - \frac{11}{6} + \frac{\delta \sigma_C^2}{2} + \frac{\gamma \sigma_G^2}{2} - \gamma \rho_{CG} \sigma_C \sigma_G - \log(w_1) \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T) 1_{R_T \leq w_A}] \\
A_3^G &= \alpha_1^G \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > \tilde{w}_1}] + \beta_1^G \left(\log(\tilde{\lambda}) - \frac{11}{6} - \frac{\tilde{\gamma} \sigma_G^2}{2} - \frac{\tilde{\delta} \sigma_C^2}{2} + \tilde{\delta} \rho_{CG} \sigma_C \sigma_G \right) \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] \\
&\quad + \beta_1^G \log(1 - w_1) \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > w_B}] - \beta_1^G \mathbb{E}_t^{\mathbb{Q}} [\log(1 - R_T) 1_{R_T > \tilde{w}_1}] \\
&\quad + \frac{\beta_1^G}{w_B - 1} \left(\log(\tilde{\lambda}) - \frac{11}{6} - \frac{\tilde{\gamma} \sigma_G^2}{2} - \frac{\tilde{\delta} \sigma_C^2}{2} + \tilde{\delta} \rho_{CG} \sigma_C \sigma_G - \log(1 - w_1) \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}]
\end{aligned}$$

and letting $p(\eta, \xi) = \exp\{\frac{1}{2}(\eta^2 - \eta)\sigma_C^2 + \frac{1}{2}(\xi^2 - \xi)\sigma_G^2 + \eta\xi\rho_{CG}\sigma_C\sigma_G\}$,

$$\begin{aligned}
A_4 &= p(-\delta, \gamma) \left(\frac{3\beta_0^C}{\lambda} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{3\beta_0^C(w_A - w_1)}{\lambda w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_4 &= p(-\tilde{\delta}, \tilde{\gamma}) \left(\frac{3\beta_0^G}{\tilde{\lambda}} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{3\beta_0^G(w_B - w_1)}{\tilde{\lambda}(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_5 &= p(-2\delta, 2\gamma) \left(\frac{3\beta_0^C}{2\lambda^2} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{3\beta_0^C(w_A - w_1)^2}{2\lambda^2 w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_5 &= p(-2\tilde{\delta}, 2\tilde{\gamma}) \left(\frac{-3\beta_0^G}{2\tilde{\lambda}^2} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{-3\beta_0^G(w_B - w_1)^2}{2\tilde{\lambda}^2(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_6 &= p(-3\delta, 3\gamma) \left(\frac{\beta_0^C}{3\lambda^3} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{\beta_0^C(w_A - w_1)^3}{3\lambda^3 w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_6 &= p(-3\tilde{\delta}, 3\tilde{\gamma}) \left(\frac{\beta_0^G}{3\tilde{\lambda}^3} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{\beta_0^G(w_B - w_1)^3}{3\tilde{\lambda}^3(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_7^C &= p(-\delta + 1, \gamma) \left(\frac{3\beta_1^C}{\lambda} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{3\beta_1^C(w_A - w_1)}{\lambda w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_7^G &= p(-\tilde{\delta}, \tilde{\gamma} + 1) \left(\frac{3\beta_1^G}{\tilde{\lambda}} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{3\beta_1^G(w_B - w_1)}{\tilde{\lambda}(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_8^C &= p(-2\delta + 1, 2\gamma) \left(\frac{3\beta_1^C}{2\lambda^2} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{3\beta_1^C(w_A - w_1)^2}{2\lambda^2 w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_8^G &= p(-2\tilde{\delta}, 2\tilde{\gamma} + 1) \left(\frac{-3\beta_1^G}{2\tilde{\lambda}^2} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{-3\beta_1^G(w_B - w_1)^2}{2\tilde{\lambda}^2(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_9^C &= p(-3\delta + 1, 3\gamma) \left(\frac{\beta_1^C}{3\lambda^3} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{\beta_1^C(w_A - w_1)^3}{3\lambda^3 w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_9^G &= p(-3\tilde{\delta}, 3\tilde{\gamma} + 1) \left(\frac{\beta_1^G}{3\tilde{\lambda}^3} \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{\beta_1^G(w_B - w_1)^3}{3\tilde{\lambda}^3(w_B - 1)} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_{10}^C &= -\frac{\beta_1^C \delta}{w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] - \beta_1^C \delta \mathbb{E}_t^{\mathbb{Q}} [1_{w_A < R_T \leq \tilde{w}_1}] \\
A_{10}^G &= -\beta_1^G \tilde{\gamma} \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] - \frac{\beta_1^G \tilde{\gamma}}{w_B - 1} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}]
\end{aligned}$$

$$\begin{aligned}
A_{11} &= \frac{\beta_1^C \gamma}{w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] + \beta_1^C \gamma \mathbb{E}_t^{\mathbb{Q}} [1_{w_A < R_T \leq \tilde{w}_1}] \\
A_{12} &= \beta_1^G \tilde{\delta} \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{\beta_1^G \tilde{\delta}}{w_B - 1} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}].
\end{aligned}$$

The notation is chosen such that all coefficients labeled with a tilde correspond to terms with powers of $\tilde{\delta}$ and $\tilde{\gamma}$, and hence to the portion of the bid stack $x \in (\tilde{w}_1, w_B)$. Moreover, if we choose instead to use the approximation $B_2(x)$ (or equivalently set $\tilde{w}_1 = w_1$), then all tildes can be removed and coefficients $\tilde{A}_4, \tilde{A}_5, \tilde{A}_6$ combine with A_4, A_5, A_6 . Furthermore, the superscripts C and G on the coefficients indicate similar terms but with the roles of coal and gas reversed.

If we require more than the third order Taylor Series expansion (to avoid the divergence near the top and bottom of the bid stack, as discussed earlier), then the following fairly simple modifications are required. Suppose we have an expansion of order $N \geq 4$ both above and below $x = \tilde{w}_1$. Then:

1. In coefficients A_1, A_3^C and A_3^G , replace the fractions $\frac{11}{6}$ by $\sum_{i=1}^N \frac{1}{i}$.
2. In coefficients $A_4, \tilde{A}_4, A_7^C$ and $\tilde{A}_7^G$, replace the constants 3 by N .
3. In coefficients $A_5, \tilde{A}_5, A_8^C$ and $\tilde{A}_8^G$, replace the constants $\frac{3}{2}$ by $\sum_{i=2}^N \frac{1}{i} \binom{i}{i-2}$.
4. In coefficients $A_6, \tilde{A}_6, A_9^C$ and $\tilde{A}_9^G$, replace the constants $\frac{1}{3}$ by $\sum_{i=3}^N \frac{1}{i} \binom{i}{i-3}$.
5. For every integer n such that $4 \leq n \leq N$, add the following four terms to the sum in (5.8):

$$A_{n+12} (F_t^C)^{-n\delta} (F_t^G)^{n\gamma}, \tilde{A}_{n+12} (F_t^C)^{-n\tilde{\delta}} (F_t^G)^{n\tilde{\gamma}}, A_{n+12}^C (F_t^C)^{-n\delta+1} (F_t^G)^{n\gamma}, \tilde{A}_{n+12}^G (F_t^C)^{-n\tilde{\delta}} (F_t^G)^{n\tilde{\gamma}+1},$$

where the coefficients $\{A_{n+12}, \tilde{A}_{n+12}, A_{n+12}^C, \tilde{A}_{n+12}^G\}$ are more general versions of the third order coefficients $\{A_6, \tilde{A}_6, A_9^C, \tilde{A}_9^G\}$. Letting $a_n = \sum_{i=n}^N \frac{1}{i} \binom{i}{i-n}$, then they are given by

$$\begin{aligned}
A_{n+12} &= p(-n\delta, n\gamma) \frac{\beta_0^C a_n}{\lambda^n} \left(\mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^n 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{(w_A - w_1)^n}{w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_{n+12} &= (-1)^{n-1} p(-n\tilde{\delta}, n\tilde{\gamma}) \frac{\beta_0^G a_n}{\tilde{\lambda}^n} \left(\mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^n 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{(w_B - w_1)^n}{w_B - 1} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right) \\
A_{n+12}^C &= p(-n\delta + 1, n\gamma) \frac{\beta_1^C a_n}{\lambda^n} \left(\mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^n 1_{w_A < R_T \leq \tilde{w}_1}] + \frac{(w_A - w_1)^n}{w_A} \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq w_A}] \right) \\
\tilde{A}_{n+12}^G &= (-1)^{n-1} p(-n\tilde{\delta}, n\tilde{\gamma} + 1) \frac{\beta_1^G a_n}{\tilde{\lambda}^n} \left(\mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^n 1_{\tilde{w}_1 < R_T \leq w_B}] + \frac{(w_B - w_1)^n}{w_B - 1} \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \right).
\end{aligned}$$

More generally, when using N terms we have

$$\begin{aligned}
F_t^P &= A_1 + A_2^C \log(F_t^C) + A_2^G \log(F_t^G) + A_3^C F_t^C + A_3^G F_t^G + \sum_{n=1}^N \{ A_{n+12} (F_t^C)^{-n\delta} (F_t^G)^{n\gamma} \\
&\quad + \tilde{A}_{n+12} (F_t^C)^{-n\tilde{\delta}} (F_t^G)^{n\tilde{\gamma}} + A_{n+12}^C (F_t^C)^{-n\delta+1} (F_t^G)^{n\gamma} + \tilde{A}_{n+12}^G (F_t^C)^{-n\tilde{\delta}} (F_t^G)^{n\tilde{\gamma}+1} \} \\
&\quad + A_{10}^C F_t^C \log F_t^C + A_{10}^G F_t^G \log F_t^G + A_{11} F_t^C \log F_t^G + A_{12} F_t^G \log F_t^C
\end{aligned}$$

with $A_1, A_2^C, A_2^G, A_3^C, A_3^G, A_{10}^C, A_{10}^G, A_{11}, A_{12}$ as before but with the modifications to A_1, A_3^C, A_3^G as discussed in the first point above. $A_{n+12}, \tilde{A}_{n+12}, A_{n+12}^C$ and $\tilde{A}_{n+12}^G$ are defined as above but now for all $n \geq 1$.

Finally, for the simpler case of deterministic coal prices (and $N = 3$), the expression for power forwards is:

$$\begin{aligned} F_t^P = & A_1 + A_2 \log(F_t^G) + A_3 F_t^G + A_4 (F_t^G)^\gamma + \tilde{A}_4 (F_t^G)^{\tilde{\gamma}} + A_5 (F_t^G)^{2\gamma} + \tilde{A}_5 (F_t^G)^{2\tilde{\gamma}} \\ & + A_6 (F_t^G)^{3\gamma} + \tilde{A}_6 (F_t^G)^{3\tilde{\gamma}} + \tilde{A}_7 (F_t^G)^{\tilde{\gamma}+1} + \tilde{A}_8 (F_t^G)^{2\tilde{\gamma}+1} + \tilde{A}_9 (F_t^G)^{3\tilde{\gamma}+1} + A_{10} F_t^G \log F_t^G, \end{aligned} \quad (\text{A.1})$$

while coefficients are given by

$$\begin{aligned} A_1 &= m_1 \mathbb{E}_t^{\mathbb{Q}} [1_{R_T \leq \tilde{w}_1}] + \alpha_0^G \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > \tilde{w}_1}] - s_1 \log(w_1) \mathbb{E}_t^{\mathbb{Q}} [1_{R_T \leq w_A}] + \beta_0^G \log(1 - w_1) \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > w_B}] \\ &\quad - s_1 \left(\log(\lambda) - \frac{11}{6} + \frac{\gamma \sigma_G^2}{2} \right) \mathbb{E}_t^{\mathbb{Q}} [1_{w_A < R_T \leq \tilde{w}_1}] - \beta_0^G \left(\log(\tilde{\lambda}) - \frac{11}{6} + \frac{\tilde{\gamma} \sigma_G^2}{2} \right) \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] \\ &\quad + s_1 [\log(R_T) 1_{R_T \leq \tilde{w}_1}] - \beta_0^G \mathbb{E}_t^{\mathbb{Q}} [\log(1 - R_T) 1_{R_T > \tilde{w}_1}] \\ &\quad - \frac{s_1}{w_A} \log \left(\frac{w_1 - w_A}{w_1} \right) \mathbb{E}_t^{\mathbb{Q}} [R_T 1_{R_T \leq \tilde{w}_1}] - \frac{\beta_0^G}{w_B - 1} \log \left(\frac{w_B - w_1}{1 - w_1} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > \tilde{w}_1}] \\ A_2 &= s_1 \gamma \mathbb{E}_t^{\mathbb{Q}} [1_{w_A < R_T \leq \tilde{w}_1}] - \beta_0^G \tilde{\gamma} \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] \\ A_3 &= \alpha_1^G \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > \tilde{w}_1}] + \beta_1^G \left(\log(\tilde{\lambda}) - \frac{11}{6} + \frac{\tilde{\gamma} \sigma_G^2}{2} \right) \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}] + \beta_1^G \log(1 - w_1) \mathbb{E}_t^{\mathbb{Q}} [1_{R_T > w_B}] \\ &\quad - \beta_1^G \mathbb{E}_t^{\mathbb{Q}} [\log(1 - R_T) 1_{R_T > \tilde{w}_1}] + \frac{\beta_1^G}{w_B - 1} \log \left(\frac{w_B - w_1}{1 - w_1} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - 1) 1_{R_T > w_B}] \\ A_4 &= e^{\frac{1}{2}(\gamma^2 - \gamma) \sigma_G^2} \left(\frac{3s_1}{\lambda} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{w_A < R_T \leq \tilde{w}_1}] \\ \tilde{A}_4 &= e^{\frac{1}{2}(\tilde{\gamma}^2 - \tilde{\gamma}) \sigma_G^2} \left(\frac{3\beta_0^G}{\tilde{\lambda}} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{\tilde{w}_1 < R_T \leq w_B}] \\ A_5 &= e^{\frac{1}{2}(4\gamma^2 - 2\gamma) \sigma_G^2} \left(\frac{3s_1}{2\lambda^2} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{w_A < R_T \leq \tilde{w}_1}] \\ \tilde{A}_5 &= e^{\frac{1}{2}(4\tilde{\gamma}^2 - 2\tilde{\gamma}) \sigma_G^2} \left(\frac{-3\beta_0^G}{2\tilde{\lambda}^2} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{\tilde{w}_1 < R_T \leq w_B}] \\ A_6 &= e^{\frac{1}{2}(9\gamma^2 - 3\gamma) \sigma_G^2} \left(\frac{s_1}{3\lambda^3} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{w_A < R_T \leq \tilde{w}_1}] \\ \tilde{A}_6 &= e^{\frac{1}{2}(9\tilde{\gamma}^2 - 3\tilde{\gamma}) \sigma_G^2} \left(\frac{\beta_0^G}{3\tilde{\lambda}^3} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{\tilde{w}_1 < R_T \leq w_B}] \\ \tilde{A}_7 &= e^{\frac{1}{2}((\tilde{\gamma}+1)^2 - (\tilde{\gamma}+1)) \sigma_G^2} \left(\frac{3\beta_1^G}{\tilde{\lambda}} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1) 1_{\tilde{w}_1 < R_T \leq w_B}] \\ \tilde{A}_8 &= e^{\frac{1}{2}((2\tilde{\gamma}+1)^2 - (2\tilde{\gamma}+1)) \sigma_G^2} \left(\frac{-3\beta_1^G}{2\tilde{\lambda}^2} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^2 1_{\tilde{w}_1 < R_T \leq w_B}] \\ \tilde{A}_9 &= e^{\frac{1}{2}((3\tilde{\gamma}+1)^2 - (3\tilde{\gamma}+1)) \sigma_G^2} \left(\frac{\beta_1^G}{3\tilde{\lambda}^3} \right) \mathbb{E}_t^{\mathbb{Q}} [(R_T - w_1)^3 1_{\tilde{w}_1 < R_T \leq w_B}] \\ A_{10} &= -\beta_1^G \tilde{\gamma} \mathbb{E}_t^{\mathbb{Q}} [1_{\tilde{w}_1 < R_T \leq w_B}]. \end{aligned}$$

A.2 Error Analysis for the Bid Stack Approximation

In Section 5.3, we discussed briefly the error resulting from the Taylor Series approximation (in its most general form) to the bid stack. We introduced several possible error measures, $\text{Err}^{\text{AD}}(C_t, G_t)$, $\text{Err}^{\text{MD}}(C_t, G_t)$ and $\text{Err}^{\text{TD}}(C_t, G_t)$, which vary significantly over the fuel price space (C_t, G_t) . We now investigate more specifically how and why the approximation is weaker for some fuel price combinations and when remedies are possible. More generally (considering the use of the model for other markets or other time periods), the error is instead a function of the variables $\{m_1, s_1, m_2, s_2, w_1\}$. However, for simplicity we consider our analysis with the parameters $\{\alpha_0^C, \alpha_1^C, \beta_0^C, \beta_1^C, \alpha_0^G, \alpha_1^G, \beta_0^G, \beta_1^G, w_1\}$ all fixed to be the values estimated for PJM. As the shape of bid distributions in other markets or time periods is likely to be similar, this analysis can be considered fairly representative.

The error in the bid stack approximation stems from four different sources:

1. The choice of parameters λ , δ and γ (and $\tilde{\lambda}$, $\tilde{\delta}$ and $\tilde{\gamma}$) to best fit $x_0(C_t, G_t)$ (and $\tilde{x}_0(C_t, G_t)$).
2. Deviation in the upper and lower regions due to the difference between $B(x)$ and the crude approximation $B_0(x)$ (i.e., the impact of coal bids on the upper stack, and gas bids on the lower stack).
3. Deviation in the upper and lower regions (as we approach w_A from above or w_B from below) due to the divergence of the Taylor Series expansion if an insufficient number of terms is used.
4. Error in the middle of the bid stack due to the method of piecing together the two regions which inevitably produces a kink (discontinuity in the slope) not present in $B(x)$.

We assume for now that the functions x_0 and $\tilde{x}_0$ can be perfectly matched to ensure bid stack continuity, as this is typically a source of relatively little error. However, we shall return to this point as it becomes an issue if the region of (C_t, G_t) we are interested in becomes very large, as it does at long maturities. For now, we investigate the performance of the approximation at particular points (C_t, G_t) . In times of large separation between coal and gas bids (low C_t and high G_t), the second point above leads to very little error, while the third can always be avoided by adding more terms, though a very large number may be required. Thus the remaining error stems primarily from the fourth point above, but can be reduced through a suitable choice for $\tilde{w}_1$ slightly above w_1 . In the case of high C_t and low G_t , the bids of coal and gas generators come together in one large cluster as m_1 approaches m_2 . In this case the error is due primarily to the second point above, and leads to an increasingly significant overestimation of the bid stack below w_1 and underestimation above w_1 . None of the tools discussed above (adjusting the proportion b of the linear approximation in the stack, increasing the number of terms or shifting w_1 to $\tilde{w}_1$) are designed to overcome this, and in fact often make matters worse. Therefore, we stress again that the approximation cannot be used when the merit order of fuels is near to swapping.

Now we shall illustrate the impact of our modifications on the deviation in the middle of the stack, as well as the upper and lower regions. We consider a large range of possible coal and gas prices, letting $\$13 \leq C_t \leq \270 and $\$2 \leq G_t \leq \44 . (More precisely, we use a log space since both coal and gas prices are Gaussian in this space. So $2.6 \leq \log C_t \leq 5.6$ and $0.8 \leq \log G_t \leq 3.8$ with step size 0.05.) Recall from Figure 3.2 that historical spot prices observed in the coal and gas markets imply




Figure A.1: From top left to bottom right: Average absolute deviation, maximum absolute deviation, total deviation, and number of terms required, for the case of $b = 0$ and $\tilde{w}_1 = w_1$.




Figure A.2: Number of terms required, average absolute deviation, and maximum absolute deviation for the case of $b = 0.6$, $\tilde{w}_1 = w_1$ (left column); and number of terms required, shift in $\tilde{w}_1$ required, and maximum absolute deviation for the case of $b = 0.6$, $\tilde{w}_1$ shifted (right column).

that events in the lower right and upper left corners of this space may have very low probabilities of occurring. In particular, parameter estimates from Table 3.1 suggest positive correlation of $\rho_{CG} \approx 0.2$ between the two fuel types, though as discussed, this may be an underestimation. We use this fuel price range for illustrative purposes, but should bear in mind that for many purposes only the middle portion of the space (plus perhaps lower left, upper right) are relevant. High values of error measures in the far left corners of the following plots reflect unavoidable errors caused by point number 2 above, when coal and gas bids merge. Errors from point 3 (caused by Taylor Series divergence) are reflected in the high number of terms in the right corners of the plots (large separation between coal and gas bids). Finally, errors from the middle of the stack (point 4) are visible primarily in the plots of maximum absolute deviation for the near right corners. Since the issue is limited to a narrow section of the stack, the average absolute deviation plots are not visibly affected by it.




Figure A.3: Analysis (over the sample range $\$90 < C_t < \150 , $\$10 < G_t < \16) of the impact of varying b and $\tilde{w}_1$ on the maximum and average of $\text{Err}^{\text{MD}}(C, G)$ (top left and top right), the maximum of $\text{Err}^{\text{AD}}(C, G)$ (bottom left) and on the portion of the sample range for which a three-term Taylor expansion can be safely used. (In axis labels above, $ww = \tilde{w}_1$.)

We begin with the case of $b = 0$ and $\tilde{w}_1 = w_1$, corresponding to the simpler approximation $B_1(x)$ for the bid stack. The plots of Figure A.1 (mostly repeated from Chapter 5) illustrate the various error measures introduced above, as functions of C_t and G_t , as well as the number of terms

required such that the deviation at top of the bid stack (at $x = 0.995$) remains below \$5. Next consider the impact of increasing the percentage of linear cutoff to $b = 0.6$, given by the plots in the left column of Figure A.2. While the overall error is slightly increased relative to Figure A.1, the number of terms needed in the region of high G_t and low C_t is significantly reduced. Finally, consider the impact of slightly increasing $\tilde{w}_1$, as shown by the right column of Figure A.2. We let the value of $\tilde{w}_1$ vary with coal and gas prices, increasing it by as much as is needed in each case to reduce the maximum deviation in the middle of the stack to below \$5. While a plot of average absolute deviation would show little impact, the maximum absolute deviation (bottom right) is greatly reduced in the region of high G_t and low C_t . This was of the course the motivation for introducing the modification: to improve the small range of large deviations near $x = w_1$ in these cases.

The surface plots of Figure A.3 show in more detail the impact of varying b and $\tilde{w}$ over the smaller range $10 \leq G_t \leq 16$ and $90 \leq C_t \leq 150$. For each pair $(b, \tilde{w}_1)$, we have calculated new error measures by finding either the average or the maximum over C_t and G_t of the standard error functions introduced in Section 5.3. The lower left plot of Figure A.3 illustrates $\max_{C,G} \{\text{Err}^{\text{AD}}(C, G)\}$ and is particularly relevant as we use this approach soon when finding valid regions for forward curve calibration. No matter which error measure we use, the plot is always increasing in b , emphasising again that this tool (a non-zero choice for b) is a means of reducing the number of terms in the expansion (e.g., see lower right graph) at the expense of increased error. In contrast, there is typically an optimal choice of $\tilde{w}_1$ slightly about w_1 which minimises the error for the chosen range of fuel prices. However, the lower right plot shows that as $\tilde{w}_1$ increases, the number of terms required also increases as a side-effect.

Finally, one should not forget the error induced by a poor fit of the function $x_0(C_t, G_t) = \lambda C_t^\delta G_t^{-\gamma}$ over certain ranges of coal and gas prices. The fit tends to be very strong for smaller regions of the (C, G) space, but sometimes weak for large regions. The graphs here have all been produced assuming that the fit $\hat{x}_0$ is exactly equal to x_0 as given by (5.4) or (5.5) to ensure bid stack continuity. However, we can produce corresponding graphs for the case that $\hat{x}_0 = \nu x_0$ for some $\nu \neq 1$. Though these are not plotted here, Figure 5.5 provides a sufficient illustration of the impact on both number of terms required and Err^{AD} for the cases of $\nu = 0.8$ and $\nu = 1.4$, representing an 20% underestimation or a 40% overestimation. We can observe that only an underestimation of x_0 increases the number of terms required in the expansion (since the divergence of $B_1(x)$ from $B_0(x)$ happens sooner), while both under- and overestimation can increase the average deviation Err^{AD} due to a larger region of error in the middle of the stack.

Appendix B

Trinomial Tree Techniques for Coal and Gas Prices

In Section 6.2.3, we proposed a numerical scheme for calculating carbon prices backwards through time via a dynamic programming approach. The processes for correlated coal and gas prices are discretised using a variation of the Hull and White trinomial trees, which we now discuss in detail. Firstly, under the assumption of Geometric Brownian Motion (with no drift), the trees are constructed using up and down steps corresponding to a fixed log return, and a middle step corresponding to no change:¹

$$\log C_{t+\Delta t} = \begin{cases} \log C_t + \sigma_C \sqrt{x\Delta t} & \text{up step} \\ \log C_t & \text{mid step} \\ \log C_t - \sigma_C \sqrt{x\Delta t} & \text{down step} \end{cases} \quad \log G_{t+\Delta t} = \begin{cases} \log G_t + \sigma_G \sqrt{y\Delta t} & \text{up step} \\ \log G_t & \text{mid step} \\ \log G_t - \sigma_G \sqrt{y\Delta t} & \text{down step} \end{cases},$$

where σ_C and σ_G correspond to the volatilities of the coal and gas processes (and not to previous notation in Chapters 4-5). In order to match the mean and standard deviation of the Gaussian distributions (as well as the skewness of zero), we choose up, middle and down steps to have probabilities p_c , $1 - 2p_c$ and p_c for the coal process, and p_g , $1 - 2p_g$ and p_g for the gas process, where $p_c = \frac{1}{2x}$ and $p_g = \frac{1}{2y}$. We therefore still have one free parameter to choose for each process, with the restriction that $p_c, p_g \in (0, 1/2]$. The well-known trinomial tree techniques of Hull and White (1994a) recommend the choice $x = y = 3$ above (for error minimization reasons), leading to probabilities $1/6$, $2/3$ and $1/6$ respectively. However, as discussed in Chapter 3, it may be important to incorporate a correlation between gas and coal prices to reflect overall co-movement in commodity prices. In this case, we propose an alternative method of adjusting probabilities similar to that of Hull and White (1994b). Define the three-by-three matrix $P_{C,G}$ of joint probabilities such that the rows (in order) correspond to up, flat and down branches for coal while the columns apply in the same way to gas movements. Then we adjust the corner probabilities by z , such that

$$P_{C,G} = \begin{pmatrix} p_c p_g + z & p_c(1 - 2p_g) & p_c p_g - z \\ (1 - 2p_c)p_g & (1 - 2p_c)(1 - 2p_g) & (1 - 2p_c)p_g \\ p_c p_g - z & p_c(1 - 2p_g) & p_c p_g + z \end{pmatrix}.$$

¹While a binomial tree is sufficient for this case, a trinomial tree takes account of a greater range of pairs (C_t, G_t) , which can be useful to better capture the probabilities of merit order changes and fuel switching effects which lead to discontinuities in $X_t(C_t, G_t)$ as discussed in Chapter 6.

Clearly choosing $z > 0$ induces positive correlation, while $z < 0$ induces negative correlation, while means, standard deviations and skews of $\log C_{t+\Delta t}$ and $\log G_{t+\Delta t}$ remain unchanged. In order to match the covariance of coal and gas, we solve

$$\begin{aligned}
\rho_{CG}\sigma_C\sigma_G\Delta t &= \text{Cov}_t(\log C_{t+\Delta t}, \log G_{t+\Delta t}) \\
&= \begin{pmatrix} \sigma_C\sqrt{x\Delta t} & 0 & -\sigma_C\sqrt{x\Delta t} \end{pmatrix} P_{C,G} \begin{pmatrix} \sigma_G\sqrt{y\Delta t} & 0 & -\sigma_G\sqrt{y\Delta t} \end{pmatrix}^T \\
&= 4(\sigma_C\sqrt{x\Delta t})(\sigma_G\sqrt{y\Delta t})z \\
\implies z &= \frac{\rho_{CG}}{4\sqrt{xy}} = \frac{\rho_{CG}\sqrt{p_cp_g}}{2},
\end{aligned}$$

where ρ_{CG} corresponds to the correlation of the coal and gas processes (and again, not to previous notation in Chapter 5). This matrix $P_{C,G}$ above is valid only if it has no negative probabilities which is guaranteed if $z \leq p_cp_g$ or equivalently $\rho_{CG} \leq 2\sqrt{p_cp_g}$. We can therefore use the probabilities $p_c = p_g = 1/6$ only for low correlations $-1/3 \leq \rho_{CG} \leq 1/3$. Otherwise we increase both p_c and p_g (keeping $p_c = p_g$ for symmetry) until the correlation can be captured by the approach. Note that $p_c = p_g = 1/2$ is the maximum value of the probabilities (reducing our trinomial tree to a binomial tree) and is also sufficient to capture the hypothetical case of perfect correlation $\rho = \pm 1$.

In light of results from Chapter 3, we also consider an exponential OU process for G_t , in which case we can keep the same equally-spaced nodes as in our GBM tree above, but modify the branching scheme and probabilities through an approach similar to that of Hull and White (1994a,1994b). That is, above and below certain thresholds in the tree, the branching scheme switches to be non-centred to reflect mean-reversion (i.e., the fuel price can jump two nodes upwards but none downwards if it is very low or two nodes downwards but none upwards if very high). As before, we fix $y = 3$ only for cases of low correlation, but reduce it if necessary to match high correlation estimates. The more general version of the Hull and White probabilities is given below, where the index j corresponds to being in the state $\log G_t = \log G_0 + j\sigma_G\sqrt{y\Delta t}$, where $\log G_0$ is the mean-reversion level of the process, and midpoint of the tree. Note that σ_G now corresponds to the volatility parameter in the OU process, and κ to the speed of mean reversion.

For the centred branching scheme,

$$p_{\text{up}} = \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 - j\kappa\Delta t}{2}, \quad p_{\text{mid}} = 1 - \frac{1}{y} - (j\kappa\Delta t)^2, \quad p_{\text{down}} = \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 + j\kappa\Delta t}{2}.$$

For the scheme where G_t can no longer increase (for large positive j),

$$p_{\text{up}} = 1 + \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 - 3j\kappa\Delta t}{2}, \quad p_{\text{mid}} = -\frac{1}{y} - (j\kappa\Delta t)^2 + 2j\kappa\Delta t, \quad p_{\text{down}} = \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 - j\kappa\Delta t}{2}.$$

For the scheme where G_t can no longer decrease (for large negative j),

$$p_{\text{up}} = \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 + j\kappa\Delta t}{2}, \quad p_{\text{mid}} = -\frac{1}{y} - (j\kappa\Delta t)^2 - 2j\kappa\Delta t, \quad p_{\text{down}} = 1 + \frac{1}{2y} + \frac{(j\kappa\Delta t)^2 + 3j\kappa\Delta t}{2}.$$

These are obtained by matching the mean and standard deviations of $\log G_{t+\Delta t}$ under the assumption of step sizes of $(\sigma_G\sqrt{y\Delta t} \ 0 \ -\sigma_G\sqrt{y\Delta t})$, $(0 \ -\sigma_G\sqrt{y\Delta t} \ -2\sigma_G\sqrt{y\Delta t})$ or $(2\sigma_G\sqrt{y\Delta t} \ \sigma_G\sqrt{y\Delta t} \ 0)$. The choice of which branching scheme to use is determined by a threshold b such that for $-\frac{b}{\kappa\Delta t} < j < \frac{b}{\kappa\Delta t}$, we use the first, for $j > \frac{b}{\kappa\Delta t}$ the second, and for $j < -\frac{b}{\kappa\Delta t}$ the third. Note that only one

value of j will employ the second and third schemes, as they represent the boundaries of the tree itself. By solving quadratic inequalities in $j\kappa\Delta t$ we find the values of b and y for which the branching scheme is valid (i.e., no negative probabilities are generated). We firstly require $1 < y < 4$ for the centred branching scheme to hold. (In fact, the lower bound here can be increased to $4/3$, since the condition on b below cannot be satisfied otherwise.) Given a choice of y , we then require

$$1 - \sqrt{1 - \frac{1}{y}} \leq b \leq \sqrt{1 - \frac{1}{y}}. \quad (\text{B.1})$$

The left-hand side of this inequality ensures that $p_{\text{mid}} > 0$ for the non-centred branching schemes, while the right-hand side ensures that $p_{\text{mid}} > 0$ for the centred scheme. Like Hull and White (using $y = 3$ only), we find it most efficient to choose b equal to its lower bound. Note that since the midpoint of the range $(1 - \sqrt{1 - 1/y}, \sqrt{1 - 1/y})$ is $1/2$ for all y , we always have $b < 1/2$.

Finally, for the case of correlated coal prices, we again adjust the joint probabilities of the coal and gas trees to match observed correlation. As before, we adjust only the corners of the matrix $P_{C,G}$ and by an amount z , so

$$P_{C,G} = \begin{pmatrix} p_c p_{\text{up}} + z & p_c p_{\text{mid}} & p_c p_{\text{down}} - z \\ (1 - 2p_c)p_{\text{up}} & (1 - 2p_c)p_{\text{mid}} & (1 - 2p_c)p_{\text{down}} \\ p_c p_{\text{up}} - z & p_c p_{\text{mid}} & p_c p_{\text{down}} + z \end{pmatrix}. \quad (\text{B.2})$$

Here p_c is fixed but the gas probabilities vary throughout the tree. Nonetheless, the equation

$$\text{Cov}_t(\log C_{t+\Delta t}, \log G_{t+\Delta t}) = 4(\sigma_C \sqrt{x\Delta t})(\sigma_G \sqrt{y\Delta t})z$$

still holds throughout the tree (for all branching schemes), although the exact expression for the covariance of the Gaussian processes changes. We solve for z similarly to before:

$$\begin{aligned} \frac{\rho_{CG}\sigma_C\sigma_G}{\kappa} (1 - e^{-\kappa\Delta t}) &= \text{Cov}_t(\log C_{t+\Delta t}, \log G_{t+\Delta t}) \\ &= 4(\sigma_C \sqrt{x\Delta t})(\sigma_G \sqrt{y\Delta t})z \\ \Rightarrow z &= \frac{\rho_{CG}}{4\sqrt{y/(2p_c)}} \frac{1 - e^{-\kappa\Delta t}}{\kappa\Delta t}. \end{aligned}$$

We now require $z \leq p_c \min(p_{\text{up}}, p_{\text{down}})$ throughout the two-dimensional tree. Therefore, we require that for the largest positive integer $j < \frac{b}{\kappa\Delta t}$,

$$\rho_{CG} \leq 4 \left(\sqrt{\frac{yp_c}{2}} \right) \frac{\kappa\Delta t}{1 - e^{-\kappa\Delta t}} p_{\text{up}} = 4 \left(\sqrt{\frac{yp_c}{2}} \right) \frac{\kappa\Delta t}{1 - e^{-\kappa\Delta t}} \left(\frac{1}{2y} + \frac{(j\kappa\Delta t)^2 - j\kappa\Delta t}{2} \right), \quad (\text{B.3})$$

and that for the smallest positive integer $i > \frac{b}{\kappa\Delta t}$ (i.e., $i = j + 1$),

$$\rho_{CG} \leq 4 \left(\sqrt{\frac{yp_c}{2}} \right) \frac{\kappa\Delta t}{1 - e^{-\kappa\Delta t}} p_{\text{down}} = 4 \left(\sqrt{\frac{yp_c}{2}} \right) \frac{\kappa\Delta t}{1 - e^{-\kappa\Delta t}} \left(\frac{1}{2y} + \frac{(i\kappa\Delta t)^2 - i\kappa\Delta t}{2} \right). \quad (\text{B.4})$$

These conditions follow firstly from the observation that for $j > 0$, $p_{\text{up}} < p_{\text{down}}$ and p_{up} is strictly decreasing in j for the centred branching scheme if $b < 1/2$, since p_{up} reaches a minimum at $j = 1/(2\kappa\Delta t) > b\kappa\Delta t$ (beyond the range of the centred scheme). Secondly, it is also easy to show that $p_{\text{down}} < p_{\text{up}}$ for the non-centred scheme.² Finally, we observe that in our framework the tree is

²As mentioned before, there is only one value of j (i.e., the smallest positive integer $j > \frac{b}{\kappa\Delta t}$) that we need to consider for the non-centred scheme.

perfectly symmetric, so probabilities at j are identical to those at $-j$ but with their order reversed. Next, we note that by choosing $b = 1 - \sqrt{1 - 1/y}$, the second condition above is stronger than the first condition in almost all cases, and in particular for small enough Δt it can be guaranteed.³

We can now use both free parameters y and p_c to increase the correlation levels which the tree can capture. We can first increase p_c to as much as $1/2$, as the bound on ρ_{CG} above is clearly strictly increasing in p_c . We can then decrease y to as low as $4/3$, which narrows the spacing of nodes, consequently increasing up and down probabilities (at the expense of middle probabilities) and thus decreasing the bound on ρ_{CG} . This can also be shown to hold strictly true for all $j\kappa\Delta t$ consistent with the branching scheme of the tree. Therefore, for small Δt , an approximate bound on ρ_{CG} can be found by setting $b = j\kappa\Delta t$ in (B.3) or $b = (j + 1)\kappa\Delta t$ in (B.4). As $\Delta t \rightarrow 0$, both bounds converges to this expression (since $b - j\kappa\Delta t \rightarrow 0$ and $b - (j + 1)\kappa\Delta t \rightarrow 0$), and we have (with $b = 1 - \sqrt{1 - 1/y}$)

$$\begin{aligned}\rho_{CG} &\leq 4 \left(\sqrt{\frac{yp_c}{2}} \right) \left(\frac{1}{2y} + \frac{b^2 - b}{2} \right) \\ &= 2 \left(\sqrt{\frac{yp_c}{2}} \right) \left(1 - \sqrt{1 - \frac{1}{y}} \right).\end{aligned}$$

Choosing $p_c = 1/2$ and $y = 4/3$, we find that the bound on the correlation parameter in this limiting case is $\rho_{CG} \leq 1/\sqrt{3} \approx 0.577$. (By symmetry the lower bound is -0.577 .) The left graph of Figure B.1 shows the variation of the bound as we vary y (between 1.4 and 3.9), as well as the impact of changing Δt and p_c . Clearly there is a significant drop in the attainable correlation if we allow the coal tree to have a high probability of a middle step (e.g., if $p_c = 1/6$). Also, for larger values of Δt , there is some jaggedness to the plot due to the fluctuation in the difference between the value $\frac{b}{\kappa\Delta t}$ and the smallest integer above it.




Figure B.1: Maximum correlation as node spacing parameter y varies (left graph), and comparison with Hull and White approach across all nodes in the tree (right graph). $\kappa = 1.2$ throughout and n is the number of nodes in the tree.

³Only for values of y near $4/3$ (and large Δt) might we find that $j < \frac{b}{\kappa\Delta t}$ but $j + 1 > \frac{1}{2\kappa\Delta t}$, in which case the first condition may be stronger. In these cases, reducing Δt will make the second condition stronger again.

An obvious means of evaluating this technique is by comparison with the original method of Hull and White (1994b). Applying their approach to incorporating positive correlation between two processes, we adjust the probability matrix as follows:

$$P_{C,G} = \begin{pmatrix} p_c p_{\text{up}} + 5\epsilon & p_c p_{\text{mid}} - 4\epsilon & p_c p_{\text{down}} - \epsilon \\ (1 - 2p_c)p_{\text{up}} - 4\epsilon & (1 - 2p_c)p_{\text{mid}} + 8\epsilon & (1 - 2p_c)p_{\text{down}} - 4\epsilon \\ p_c p_{\text{up}} - \epsilon & p_c p_{\text{mid}} - 4\epsilon & p_c p_{\text{down}} + 5\epsilon \end{pmatrix}, \quad (\text{B.5})$$

where $\epsilon = \rho_{CG}/36$. One advantage of their approach is that it allows us to capture very high positive correlation near the middle of the gas price tree. Moreover, for $\rho_{CG} = 1$, as $\Delta t \rightarrow 0$, the probabilities using the centred branching scheme converge to $1/6$, $2/3$, $1/6$ on the diagonal (if $x = y = 3$), with zeros elsewhere. For very small Δt , the edges of the tree (and hence the other branching schemes) become less important, as the number of branches is very large. However, for more typical choices of Δt , the boundaries of the tree are relevant, and we would therefore like all three branching schemes to capture sufficiently high values of ρ_{CG} . Here our alternative approach to shifting the probabilities provides an advantage. In particular, the matrix $P_{C,G}$ in (B.2) (unlike in (B.5)) requires no change to probabilities involving $(1 - 2p_c)$ or p_{mid} , the middle step probabilities. Therefore letting p_c approach $1/2$ causes no problem, while the small values of p_{mid} produced near the boundaries of the gas tree also have no impact on the maximum correlation we can capture. Recall that the range of b allowed by (B.1) was limited precisely because of negative values of p_{mid} produced at and just before the tree's boundaries. The non-centred branching schemes lead to especially low values of p_{mid} , causing the maximum ρ_{CG} possible under the Hull-White scheme to drop very near 0.

The right graph of Figure B.1 illustrates this comparison by plotting the maximum correlation achievable throughout the gas price tree, rescaled to $(0,1)$.⁴ Comparing the first two lines listed in the legend, the Hull and White scheme (with their suggested $b = 0.184 \approx 1 - \sqrt{1 - 1/3}$) outperforms our alternative for the majority of the tree, but performs very poorly at the edges. The performance at the edges can be somewhat improved by increasing b to 0.32, but this is at the expense of the performance throughout the rest of the tree. Note also the effect of either decreasing p_c or increasing y in our approach, leading to a downwards shift which is parallel in the case of p_c . Overall, we conclude that our approach is a useful alternative for the case of relatively large $\kappa\Delta t$ (such as $\kappa\Delta t \approx 1/50$) and for correlation levels around $\rho \in (0, 0.5)$, as we observe for coal and gas prices. In these cases, it allows us to exactly match the correlation throughout the tree without introducing bias by reducing the correlation at certain nodes, as suggested by Hull and White.

⁴Here 0.5 corresponds to the node at $j = 0$ while 0 and 1 corresponds to the boundary nodes (with non-centred branching) where $j = \pm j^{\text{max}}$, the smallest positive integer above $\frac{b}{\kappa\Delta t}$. Note also that the number of nodes in the tree is given by $n = 2j^{\text{max}} + 1$ which varies for the different schemes and for different values of y . Hence, although plotted on the same range $(0, 1)$, the trees are of different sizes for different cases.

Bibliography

- [1] C. Albanese, H. Lo, S. Tompaidis, *A Numerical Method for Pricing Electricity Derivatives Based on Continuous Time Lattices*
- [2] C.L. Anderson, M. Davison (2008), *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.
- [3] N. Audet, P. Heiskanen, J. Keppo, I. Vehvilainen, (2004) *Modeling Electricity Forward Curve Dynamics in the Nordic Market*, In: Bunn, D. W. (Ed.), *Modelling Prices in Competitive Electricity Markets*. Wiley, pp. 251-65
- [4] M.T. Barlow (2002), *A diffusion model for electricity prices*, Mathematical Finance **12**, 287-98.
- [5] H. Bessembinder, M. Lemmon (2002), *Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets*, Journal of Finance, **57(3)**, 1347-1382.
- [6] S. Borovkova, F.J. Permana, H. v.d. Weide (2006), *A closed form approach to valuing and hedging basket options* working paper
- [7] A. Boogert, D. Dupont (2008), *When supply meets demand: the case of hourly spot electricity prices*, IEEE Transactions on Power Systems **23(2)**, 389-398.
- [8] D. Bunn, N. Karakatsani (2005) *Diurnal Reversals of Electricity Forward Premia*, working paper, London Business School
- [9] M. Burger, B. Klar, A. Muller, G. Schindlmayr (2004), *A spot market model for pricing derivatives in electricity markets*, Quantitative Finance **4**, 109-122.
- [10] M. Burger, B. Graeber, G. Schindlmayr (2007), *Managing Energy Risk: An Integrated View on Power and Other Energy Markets* Wiley Finance, Chichester.
- [11] P. Carr, D. Madan (1999), *Option Pricing and the fast Fourier transform*, Journal of Computational Finance **2**, 61-73.
- [12] R. Carmona, M. Ludkovski (2006), *Energy Trading* Siam Report, June 12, 2006. (available at <http://www.siam.org/news/news.php?id=949>)
- [13] R. Carmona, M. Fehr, J. Hinz, A. Porchet, (2008) *Market design for emission trading schemes*, working paper

- [14] R. Carmona, V. Durrleman, (2003) *Pricing And Hedging Spread Options* Siam Rev, **45(4)**, 627-685
- [15] A. Cartea, M.G. Figueroa (2005), *Pricing in Electricity Markets: A mean reverting jump diffusion model with seasonality*, Applied Mathematical Finance, **12(4)**, 313-335.
- [16] A. Cartea, M.G. Figueroa, H. Geman (2008), *Modelling Electricity Prices with Forward Looking Capacity Constraints*, Applied Mathematical Finance, forthcoming.
- [17] A. Cartea, P. Villaplana (2008), *Spot Price Modelling and the Valuation of Electricity Forward Contracts: The Role of Demand and Capacity*, Journal of Banking and Finance 32(12), 2502-19.
- [18] L. Clewlow, C. Strickland (2000), *Energy Derivatives: Pricing and Risk Management*, London: Lacima Productions.
- [19] L. Clewlow, C. Strickland (1999), *Valuing Energy Options in a One Factor Model Fitted to Forward Prices* working paper
- [20] M. Culot (2003), *An integrated affine jump-diffusion framework to manage power portfolios in a deregulated market* unpublished PhD thesis, Department of Industrial Engineering, Universite Catholique de Louvain
- [21] M. Culot, V. Goffin, S. Lawford, S. de Menten, Y. Smeers (2006), *An Affine Jump Diffusion Model for Electricity* working paper, Catholic University of Leuven.
- [22] M. Davison, C.L. Anderson, B. Marcus, K. Anderson (2002), *Development of a hybrid model for electrical power spot prices*, IEEE Transactions on Power Systems **17(2)**, 257-64.
- [23] C. de Jong (2005), *The Nature of Power Spikes: a regime switch approach*, ERIM Report Series
- [24] C. De Jong, R. Huisman (2003), *Option Pricing for Power Prices with Spikes*, Energy Power Risk Management, **7**, 12-16.
- [25] S. Deng (1999), *Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes*, working paper, University of California Energy Institute.
- [26] D. Duffie, J. Pan, K. Singleton (2000), *Transform Analysis and Asset Pricing for affine jump-diffusions*, Econometrica **68(6)**, 1343-76.
- [27] G. Emery, Q. Liu (2002), *An Analysis of the Relationship between Electricity and Natural-Gas Futures Prices*, The Journal of Futures Markets **22(2)**, 95-122.
- [28] A. Escibano, J. Pena, P. Villaplana (2002), *Modeling Electricity Prices: International Evidence*, working paper
- [29] A. Eydeland, H. Geman (1999), *Fundamentals of Electricity Derivatives*, Energy Modelling & the Management of Uncertainty, Risk Books.
- [30] A. Eydeland, K. Wolyniec (2003), *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging* Wiley Finance, New York.

- [31] M. Fehr, J. Hinz (2006), *A quantitative approach to carbon price risk modeling*, working paper, ETH Zentrum.
- [32] H. Geman, A. Roncoroni (2006), *Understanding the Fine Structure of Electricity Prices*, Journal of Business **79**(3), 1225-1261.
- [33] Hortacsu, Puller (2007), *Understanding Strategic Bidding in Multi-Unit Auctions: A Case Study of the Texas Electricity Market*, The RAND Journal of Economics **39**(1), 86-114.
- [34] J. Hull, A. White (1994a), *Numerical procedures for implementing term structure models I: Single-Factor Models*, Journal of Derivatives, **2**(1), 7-16.
- [35] J. Hull, A. White (1994b), *Numerical procedures for implementing term structure models II: Two-Factor Models*, Journal of Derivatives, **2**(2), 37-48.
- [36] R. Huisman (2008), *The influence of temperature on spike probability in day-ahead power prices*, Energy Economics, **30**, 2697-2704.
- [37] R. Huisman, R. Mahieu (2001), *Regime Jumps in electricity prices*, Energy and Power Risk Management, Risk Publications
- [38] T. Kanamura, K. Ohashi (2007), *A structural model for electricity prices with spikes: measurement of spike risk and optimal policies for hydropower plant operation*, Energy Economics, 29(5), 1010-1032.
- [39] N. Karakatsani, D. Bunn (2008), *Intra-day and regime-switching dynamics in electricity price formation*, Energy Economics, **30**, 1776-1797.
- [40] V. Kholodnyi (2001), *A Non-Markovian Process for Power Prices with Spikes and Valuation of European Contingent Claims on Power*,
- [41] T. Kluge (2006), *Pricing Swing Options and Other Electricity Derivatives*, DPhil Thesis, University of Oxford.
- [42] S. Koekebakker, F. Ollmar (2005), *Forward Curve Dynamics in the Nordic Electricity Market*, Managerial Finance, **31**(6), 74-95.
- [43] J.J. Lucia, E.S. Schwartz (2002), *Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange*, Review of Derivatives Research **5**, 5-50.
- [44] T.D. Mount, Y. Ning, X. Cai (2006), *Predicting Price Spikes in Electricity Markets using a Regime-Switching Model with Time-Varying Parameters*, Energy Economics **28**, 62-80.
- [45] D. Pilipovic (1997), *Energy Risk: Valuing and Managing Energy Derivatives*, McGraw-Hill
- [46] R. Pindyck (1999), *The long-run evolution of energy commodity prices*, Energy Journal, April, IAEE
- [47] C. Pirrong, M. Jermakyan (2008), *The Price of Power: The Valuation of Power and Weather Derivatives*, Journal of Banking and Finance, **32**, 2520-2529.

- [48] C. Pirrong (2006), *The Valuation of Power Options in a Pirrong-Jermakyan Model*, working paper, University of Houston.
- [49] C. Ruibal, M. Mazumdar (2008), *Forecasting the Mean and the Variance of Electricity Prices in Deregulated Markets*, IEEE Transactions on Power Systems, **23(1)**, 25-32.
- [50] P.A. Samuelson (1965), *Proof that properly anticipated prices fluctuate randomly*, Industrial Management Review, **6**, 41-49.
- [51] E.S. Schwartz (1997), *The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging*, The Journal of Finance **3**, 923-973.
- [52] E. Schwartz, J. Smith (2000), *Short-Term Variations and Long-Term Dynamics in Commodity Prices*, Management Science **46** 893-911.
- [53] P. Skantze, M. Ilic, J. Chapman (2000a), *Stochastic Modeling of Electric Power Prices in a Multi-Market Environment*, IEEE Power Engineering Society Winter Meeting, Volume 2, 23-27 Jan. 2000, 1109-1114.
- [54] P. Skantze, M. Ilic, A. Gubina (2000b), *Bid-based Stochastic Model for Electricity Prices: Impact of Fundamental Drivers on Market Dynamics*, MIT Energy Laboratory Publication.
- [55] C. Supatgiat, R. Zheng, J. Birge (2001), *Equilibrium Values in a Competitive Power Exchange Market* Computational Economics **17** 93-121.
- [56] S. Turnbull, L. Wakeman (1991), *A quick algorithm for pricing european average options*, The journal of financial and quantitative analysis, **26(3)**, 377-389.
- [57] I. Vehvilainen, T. Pyykkonen (2004), *Stochastic factor model for electricity spot price - the case of the Nordic Market*, Energy Economics, **27**, 351-367.
- [58] P. Villaplana (2004), *Valuation of Electricity Forward Contracts: The Role of Demand and Capacity*, working paper
- [59] P. Villaplana (2003), *Pricing Power Derivatives: A Two-Factor Jump-Diffusion Approach* working paper
- [60] R. Weron (2008), *Market price of risk implied by Asian-style electricity options*, Energy Economics, **30(3)**, 1098-1115
- [61] R. Weron, M. Bierbrauer, S. Truck (2004), *Modeling Electricity Prices: jump diffusion and regime switching* Physica A 336 39-48.
- [62] PJM bid data available at: <http://www.pjm.com/markets-and-operations/energy/real-time/historical-bid-data.aspx>
- [63] NE bid data available at: http://www.iso-ne.com/markets/hstdata/mkt_offer_bid/index.html
- [64] NYMEX daily fact sheets available at: <http://www.cmegroup.com/trading/energy-metals/nymex-daily-reports.html>