

Modelling spot and forward prices for energy companies

Dafydd Steele

MSc Stochastics and Financial Mathematics

dafydd.steele@edf-energy.com

August 5, 2010

Abstract

The focus of this thesis is on modelling forward and spot prices for energy companies. The two main ways of modelling power prices are stochastic models and fundamental models. The purpose of the thesis is to investigate the different approaches to modelling and understand which approach is most appropriate for particular applications within an energy company. A stochastic model will be implemented based on recent literature, and applied to multiple markets (coal, gas and power). The fundamental model is a mixed-integer programming stack model. It will be implemented in R as a mixed integer prototype but we will then use the industry standard software Plexos (which has heuristics which enables us to obtain large scale suboptimal solutions). A hybrid model will be implemented which combines both approaches.

Contents

1	Introduction	3
1.1	UK Power Markets	3
1.1.1	History	3
1.1.2	Forward and Spot Market	4
1.1.3	Balancing Market	5
1.1.4	Reserve Market	6
1.2	Gas and Coal markets	6
1.2.1	Gas market	6
1.2.2	Coal market	7
1.3	Why do companies need spot and forward price models?	7
2	Stochastic Models	10
2.1	Mathematical formulation	10
2.2	Traditional models: drawbacks and pitfalls	11
2.2.1	What's wrong with Geometric Brownian Motion?	11
2.3	Risk-neutral probability measures and risk premium	18
2.4	2-Factor Heath-Jarrow-Morton model	20
2.4.1	Volatility Term Structure	22
2.4.2	Extension to multi-commodity simulations	24
2.4.3	Correction for Product Delivery Period	25
2.4.4	Monte Carlo simulation	28
2.4.5	Parameter Estimation	32
2.5	How appropriate is the HJM two-factor model for power, gas and coal?	42

3	Fundamental Models	45
3.1	Background	45
3.2	Mixed Integer formulation	47
3.2.1	Additional constraints	49
3.2.2	Computing a System Marginal Price (SMP)	50
3.2.3	Uplift	50
3.3	Results	51
3.3.1	Daily granularity - backtest against closing Day Ahead prices	51
3.3.2	Half-hourly granularity - backtest against APX prices . . .	52
4	Hybrid Models	61
4.1	Concept	61
4.1.1	Demand modelling	62
4.1.2	Availability modelling	62
4.2	Results	63
4.3	Criticisms and Conclusion	63

Chapter 1

Introduction

1.1 UK Power Markets

1.1.1 History

In the late 1980's the UK Government wanted to restructure the power industry to be more competitive. The problem was achieving fairer prices, whilst also ensuring the stability of the system itself. Many South American countries had deregulated their markets already. In 1990, the UK achieved this aim under the New Electricity Trading Arrangements (NETA) and with mass privatisation. Other countries followed suit. The focus was no longer on how to achieve a more efficient market, and whether that was possible, but on concepts that were already mainstream in other markets - risk management and modelling techniques. How should these new private companies manage risk? Was it possible to manage risk? Was it possible to model power prices?

Black and Scholes had formulated their infamous option pricing formulas 26 years earlier and the power industry was only beginning to attract research interest for the first time. However, in the early 1990's these researchers faced a lack of data. Only towards the end of the 1990's and early 2000's did we start to see research into stochastic processes designed to model power [14, 9].

1.1.2 Forward and Spot Market

The electricity market in the UK is split into half-hour periods. For a given half-hour period, each electricity generator submits a Final Physical Notification (FPN) one hour before the beginning of the half-hour to the Grid Operator specifying its planned generation output. In the lead up to a half-hour of generation, there are 3 separate markets that a generator normally participates in. Initially, a generator will normally build up a position in the long term forward market. This is by entering bilateral agreements with other generators or suppliers through a broker or directly for the physical delivery of electricity. Physical delivery means that electricity must contractually be delivered or consumed. These contracts are agreed in most cases for hedging reasons; most companies would rather reduce price risk by selling (or buying for a net consumer of electricity) some of their final delivery position in advance. Hedging is the reason why it is so important to be able to model the forward price dynamics. The type of products that may be purchased in this way vary. Most of the time they will be standard contracts for a certain period of time. A calendar that is useful in the energy industry is the EFA calendar. An EFA day starts at 23:00 and consists of 6 4-hour blocks, labelled 1-6. The standard contracts are Baseload (24 hours, 7 days per week), Peak (EFA blocks 3-5 on weekdays) and Overnights (EFA blocks 1 and 2). These products are traded for various delivery periods in the future, which can be EFA Seasons (Winter - EFA week 40 - 13 - or Summer - EFA week 14 - 39), EFA Quarters (Q1-4), EFA months, weeks or days. The further into the future a delivery period is, the less liquid it is generally.

There also exist bilateral bespoke forward contracts with different delivery volumes for different half-hour periods within a day (called shape contracts). These must be priced somehow, and it is one of the motivations for obtaining a good model for half-hourly granularity power.

The 'spot' exchange is managed by the APX Group where it is possible to trade individual half-hour blocks (from 49.5 hours prior to the start of delivery). These trades are volume weighted to give a reference price, the APX price. Prices are

Instrument	PDC PDS	High Low	Last	Vol	Drk	Spk	Bids	Market Price	Offers
Win 10 OFFPKS	00.00 33.80						10 37.32 BEPET	37.9	38.47 10 BEPET
Win 10 Blocks 1+2	30.70 26.53						10 30.37 10 31.55 BEPET ICAP	31.6	31.65 10 ICAP
Win 10 BSLD	39.60 39.41	41.00 40.10	41.00 ↑ 0.00				10 40.60 15 40.87 SPEC BEPET*	41	41.12 10 BEPET
Win 10 WD Blocks 1+2	00.00 27.35							28.14	
Win 10 PKS	45.15 49.52	46.10 46.10	46.10 0.00				10 46.00 10 46.25 GFI GFI	46.38	46.50 10 GFI
May-10 BSLD	35.45 35.50	36.00 35.55	35.90 ↓ -0.10				10 35.85 15 35.89 SPEC BEPET*	35.95	36.00 10 SPEC
May-10 PKS	40.55 41.88	41.15 40.85	41.15 ↑ 0.00				10 40.75 10 41.00 PREB SPEC	41.12	41.25 10 SPEC
May-10 Blocks 1+2	25.75 28.53						10 26.00 10 27.00 PREB PREB	27.88	28.75 10 PREB
Jun-10 BSLD	35.50 35.49	36.10 36.00	36.10 ↑ 0.10				10 35.90 10 35.90 ICAP ICAP	36.08	36.25 10 ICAP
Jun-10 PKS	40.75 47.20	41.40 41.00	41.40 0.20				10 40.70 10 41.00 SPEC SPEC	41.25	41.50 10 GFI
Wednesday BSLD	36.90 37.58							32.39	

Figure 1.1: Examples of liquid forward contracts

quoted as £/MWh. Note that technically this is not a spot market in the truest sense, it is merely a very short term forwards market. A classical spot market is not possible because the Grid Operator needs advanced notice to check that the schedule is feasible and within various transmission constraints.

There is also a liquid Day-Ahead market via broker trades. It is more liquid than the exchange products so its closing price will be used as a spot price benchmark unless we have to look at the half-hour level, when we will use the APX exchange half-hourly reference price as a benchmark. Note that these spot price benchmarks are strictly based on forward trades.

1.1.3 Balancing Market

The Grid Operator will ensure that agents that are under or over-contracted must pay/receive money for any imbalance. To facilitate this, the Grid Operator is noti-

fied of all forward contracts in advance. It is then the role of the Grid Operator to ensure that the system is balanced, i.e. that supply meets demand, and this is done using the Balancing Mechanism. This is an optional market where agents may submit bids (to reduce generation/increase consumption) and offers (to increase generation/reduce consumption). The cost of balancing the system in this way is recovered by the Grid Operator through a BSUoS (Balancing Services Use of System) charge, on a £/MWh basis. Agents also pay or are charged for energy imbalances (the difference between their physical metered energy output/input and their contracted output/input) according to some prices specified by the Grid Operator. These prices give agents the financial incentive to match physical delivery with their contracted position.

1.1.4 Reserve Market

Additionally there are reserve contracts where generators are contracted by the Grid Operator to remain on standby. They also receive additional payments if they are utilised. These are known as STOR (Short Term Operating Reserve) contracts. There are also BM Start-up contracts which are paid to warm up certain kinds of generating plants (Oil) which have long start up times and high fuel costs. Conceptually, reserve is often split into two cases, *spinning reserve* and *standing reserve*. Spinning reserve are units that are in an on state and are able to raise or lower their generation as instructed by the Grid Operator. Standing reserve are units that are off, waiting for the instruction to be called on.

1.2 Gas and Coal markets

1.2.1 Gas market

The National Balancing Point (NBP) is a virtual trading location for the sale and purchase of UK natural gas. It is similar to the Henry Hub in the US, except that it is not an actual physical location. The most granular forward contracts for gas in the UK are Day Ahead NBP Prompt Gas contracts, so from the point of view of this thesis, this will be our proxy for spot. These are contracts for delivery of

some amount of gas by the end of the following day. The contracts are priced in pence/therm. The system is operated by National Grid Ltd who are responsible for the actual physical transportation of gas. Trades are made through an anonymous trading service operated by APX-ENDEX where it is possible to post bids or offers for gas contracts. The UK is connected to Europe via several pipelines; in particular the UK-Zeebrugge (Belgium) interconnector, the pipeline to Langeled, Norway and the BBL pipeline between UK and the Netherlands. In the UK, the Gas market is seasonal, with demand being highly dependent on temperature. There is the ability to store Gas to some extent in gas storage facilities.

1.2.2 Coal market

Coal is traded in three main locations in the world. ARA (Amsterdam-Rotterdam-Antwerp) coal, RB (Richards Bay, South Africa) coal and coal traded at Newcastle, Australia. Each of these trading centres have corresponding price benchmarks known as the API indices. API 2 corresponds to ARA coal, API 4 corresponds to RB coal and API 6 to Newcastle coal. API financial contracts are traded both bilaterally and on the European Energy Exchange (EEX)[3]. The forward contracts of shortest granularity are monthly contracts.

1.3 Why do companies need spot and forward price models?

Before attempting to build any models, it is important to understand why they are needed. The common reasons are:

- Investment decisions. This could be a decision to invest in new plant, or whether to accept an offer for part of the company's fleet. The decision to install costly desulphurisation or carbon capture facilities or not relies on good models. Modelling spreads are important - as will be explained later on.
- Pricing contingent claims such as options.

- Pricing of standard forward contracts as well as bespoke ‘shape’ contracts (more complex contracts for different volumes at different times of day).
- Informing hedging decisions (and their timing). How much volume should be hedged and when?
- Assessing risks around market prices. Assuming that a model is well-specified, it gives us accurate distributions of prices which enable us to assess risks using measures such as Value at Risk and Earnings at Risk.

It is important to emphasise how *spreads* are important in the energy industry, as this is one of the main motivations for a hybrid model. Some definitions:

Definition 1.3.1. The *spark / dark spread* is the gross theoretical margin in £ of a gas / coal power plant from selling a MWh of electricity, having bought the fuel required to produce the MWh of electricity.

Definition 1.3.2. The *clean spark / dark spread* is the spark / dark spread minus the cost of the CO₂ emission allowances required to produce the MWh of electricity.

Generally there is a premium on average for power over both gas and coal in the UK. This is because the ‘fuel mix’ is mainly gas and coal; and the gas and coal generators will only run if they make a profit. Therefore, coal/gas power generating assets may be regarded as *real options* on the clean dark/spark spreads. This is because the generator has the option to dispatch or not to dispatch. It will only dispatch when it is profitable to do so, and that is when the price of power is greater than the price of the fuel (plus any other operating costs). Real option theory is the application of derivative pricing theory to the optionality embedded in real assets. The traditional approach to valuing a power station was to run N fuel and electricity simulations, then to sum the discounted expected earnings for each time period. A distribution of earnings would be obtained this way. An analogy from the world of derivatives would be to price European call options by running N simulations of the underlying asset, holding the option until expiry and taking the price to be the discounted expected payout. This generally underestimates the value of the option. An intuitive way of thinking about the option is that it comprises an intrinsic value and a time value (sometimes called extrinsic value) and by

holding the option until expiry, only intrinsic value will ever be realised. The optimal way to extract the value of the option is by delta hedging - ignoring various real-life constraints such as transaction costs, non-continuous price paths and incomplete markets. More on delta hedging can be found in any of the introductory mathematical finance texts - for instance, Hull [13]. However, not so common is the application of delta hedging to pricing real assets. The power generator can also be delta hedged (at least to some extent) with a self-financing replicating portfolio and it is possible to extract the time value of the real option in that way. We will not look further into the subject of delta hedging real options. For an excellent introduction, see [11]. Deng, Johnson and Sogomonian [10] show how the theory can be used to price generation assets under the assumption of Brownian motion and mean-reverting price processes.

Because it is possible to think of the power station as an option on the clean dark/spark spread, which can be delta hedged, the distribution of the spread itself is perhaps the most important aspect of modelling for hedging and investment purposes. This will be our eventual goal - the creation of a hybrid model which is able to provide distributions of spreads. Why is the model called *hybrid* you may ask? It is a blend of modelling approaches. The fuel prices will be modelled stochastically and the fundamental model will use those fuel prices to give a power price.

Chapter 2

Stochastic Models

2.1 Mathematical formulation

Electricity is delivered and consumed continuously. However, the physical ‘spot’ market has a discrete granularity of half an hour. Also, (forward) products with observable prices widely differ regarding their delivery structure (time and length of delivery).

Definition 2.1.1. *Market Granularity* is the smallest delivery length of any type of spot or forward contract. In the UK this is a half-hour.

Most forward contracts in fact have much longer delivery length than market granularity. An example of a liquid forward would be Summer 11 Baseload, which is a contract to deliver a predetermined volume of electricity in each half hour of Summer 2001.

Definition 2.1.2. Any forward product with delivery length being equal to the market granularity is an *unitary product*.

Definition 2.1.3. The set of forward prices $\{F(T) : T = 0, 1, 2, \dots, N\}$ of all unitary products for delivery times T , in units of market granularity, is called the *forward curve*.

Since forwards can be traded at any time given sufficient liquidity, we have a continuous-time forwards market. However, the spot market quote prices on a

half-hourly basis and therefore should strictly be modelled as a time series as opposed to a continuous process as Benth, Benth and Koekebakker [4] point out. They show however that it is valid to introduce an *unobserved* continuous-time stochastic process $\tilde{S}(t)$ which represents the instantaneous spot price of electricity at time t with delivery in the interval $[t, t + dt)$, with respect to a filtration \mathcal{F}_t and to use this to define $S(t) = \tilde{S}(t)$ for discrete t . This is quite intuitive and is in fact assumed without comment in other papers [17]. This allows us to think of the spot price as a continuous process, which makes it easier to obtain the consistency we would like between the forward and the spot dynamics. It is important also to emphasise that the spot price in the UK is neither tradable nor is it observable. We use APX half-hourly prices as a proxy for spot prices when we need to build a model with half-hourly granularity. APX half-hourly prices are in fact volume weighted average traded prices for very short term forwards. Another application will be to build models with daily granularity. We will use broker Day Ahead closing prices in this case.

Assume we have a filtered space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ sufficiently rich such that all models under consideration may be defined on it.

2.2 Traditional models: drawbacks and pitfalls

2.2.1 What's wrong with Geometric Brownian Motion?

In equity markets, there is a long history of modelling prices using Geometric Brownian Motion (GBM). Explicitly in shorthand notation ¹:

$$\frac{dS}{S} = \mu_t dt + \sigma_t dW.$$

Johnson and Barz (1999) [14] were among the first to explicitly discourage the use of this model for pricing electricity. They summarized the properties of electricity prices as follows:

¹All equations between stochastic variables throughout this dissertation are to be understood as almost surely equations under the given probability measure.

- Mean reversion. The cyclical nature of demand causes mean-reversion in the short term, whilst there is long term mean reversion due to the cost of new generating plant.

Seasonal effects. There are weekly and yearly periodic fluctuations in prices in the UK market. For example, business days generally have a higher average price than non-business days; Winter has higher prices than Summer.

- Price-dependent volatility. The higher the prices, the more volatile they are generally.
- Occasional price spikes.

They noted that GBM did not capture these properties and gave evidence of this using data from the 90's. We will compile descriptive statistics based on recent data to verify this.

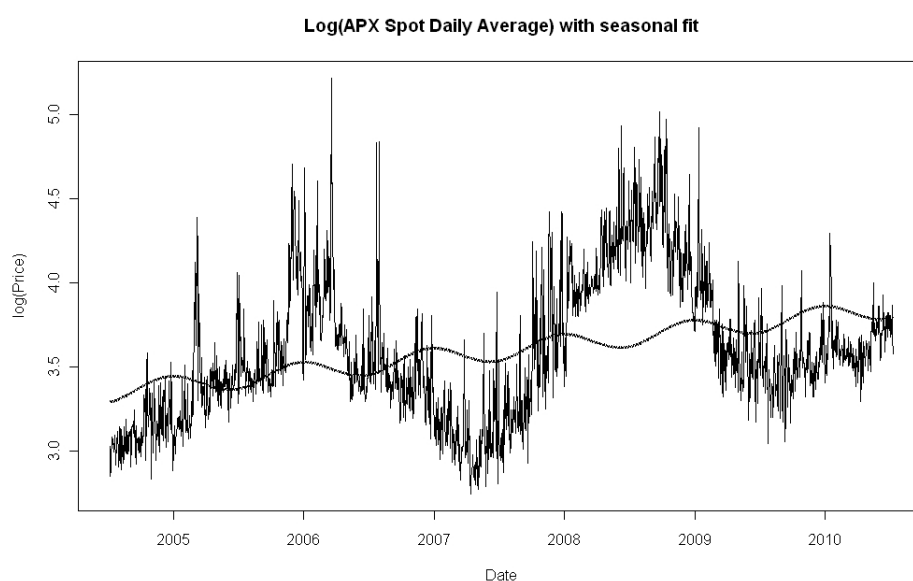


Figure 2.1: Log daily average of APX 'spot' prices from 1st July 2004, with a seasonal fit.

We attempt to remove any seasonality from the data by fitting a seasonal periodic

function to the log-returns. Write

$$\Lambda(t) = a_1 + a_2 t + \sum_{k=1}^2 a_{2k+1} \cos(2\pi l_k(t - a_{2k+2}))$$

for $l_1 = 52$, $l_2 = 1$ (weekly and yearly seasonality). Applying Non-linear Least Squares gives us the following coefficients.

Parameter	Value
a_1	3.3465
a_2	0.0832
a_3	0.0068
a_4	-0.497
a_5	0.059
a_6	-0.5486

Figure 2.2: Seasonal fit parameters.

Then we can subtract the deterministic seasonal effect from the log-prices and gather some descriptive statistics on the log-returns ($\log S_{t+1} - \log S_t$):

Statistic	Value
Mean	0
Standard Deviation	0.173
Skewness	0.262
Kurtosis	5.90

Figure 2.3: Power log-return descriptive statistics.

Observe that the log-returns are leptokurtic compared to what we would expect if they follow a Brownian Motion (where the log-returns should be distributed normally and therefore should have sample kurtosis of 3). Under the Geometric Brownian Motion model we would also expect any tests for stationarity to be rejected because the AR(1) process with a unit root (i.e. non-stationary) is a discrete

approximation of the Brownian Motion. Denote the log of the price in time period t as X_t . Assuming that the model specification of the deseasonalised log-prices is:

$$X_{t+1} = \rho X_t + \epsilon_t ,$$

with

$$H_0 : \rho = 1$$

and

$$H_1 : \rho < 1.$$

We apply the Dickey-Fuller test in R. The null hypothesis is rejected at the 99% level. The Dickey-Fuller statistic obtained was -10.7749 - highly indicative of mean-reversion.

Additionally we examine day-ahead gas and coal prices in the same manner. Coal is not seasonal in nature because it can be stored easily. Therefore we fit the trend function $\Lambda(t) = a_0 + a_1 t$ by Ordinary Least Squares (OLS). Data was only available from 7th March 2006 in this case.

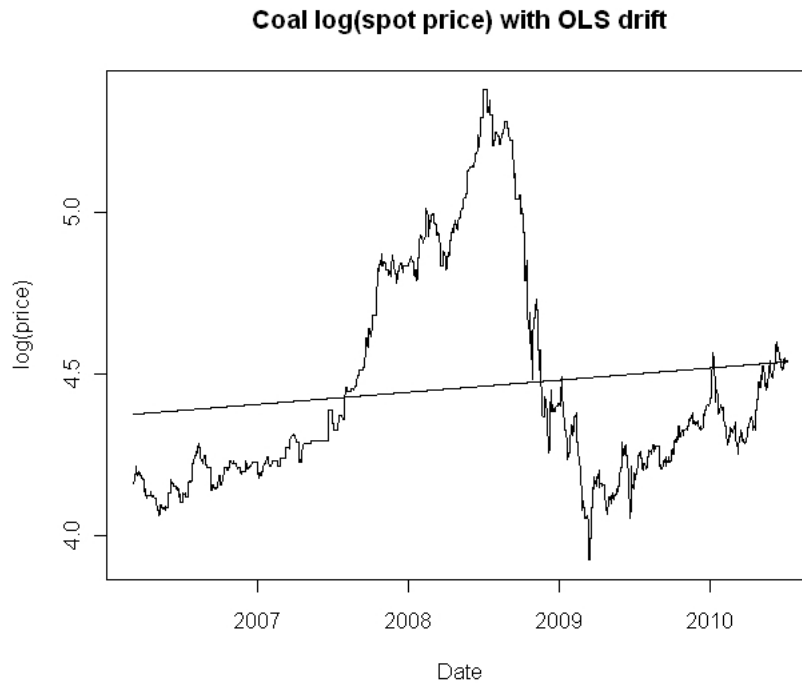


Figure 2.4: Log daily average of Coal ‘spot’ prices from 7th March 2006, with an OLS trendline.

Statistic	Value
Mean	0
Standard Deviation	0.026
Skewness	-0.422
Kurtosis	6.88

Figure 2.5: Coal log-return descriptive statistics.

Observe that the standard deviation (“volatility”) is far less than that of power and once again, the sample kurtosis is much higher than 3, and higher than the sample kurtosis of power. The unit root null hypothesis is rejected at 99.9% level (Dickey Fuller statistic = -25.54) indicating that mean-reversion is present (assuming the model specification above).

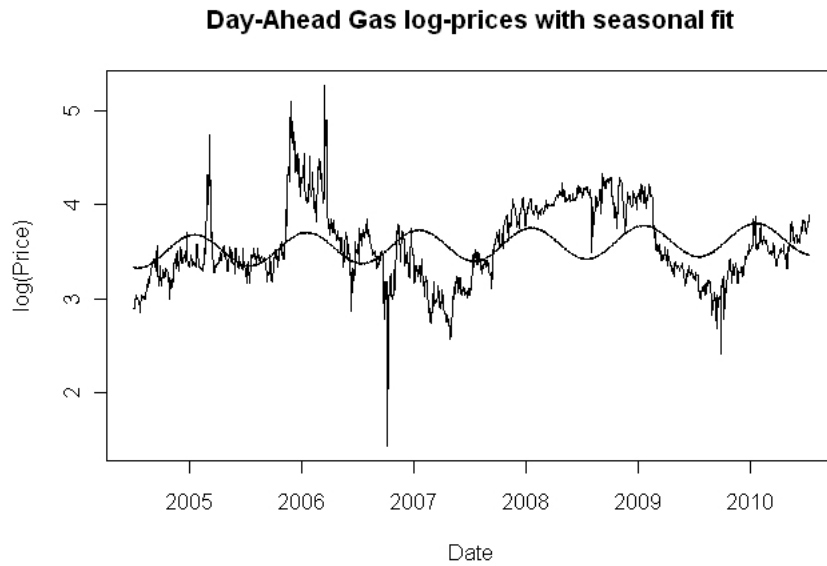


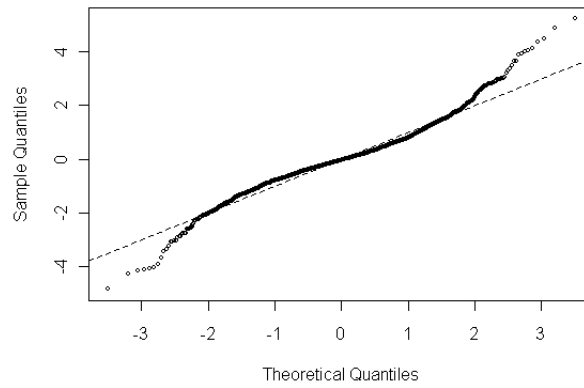
Figure 2.6: Log daily average of Gas “spot” prices from 1st July 2004, with seasonal fit

Statistic	Value
Mean	0
Standard Deviation	0.1
Skewness	0.39
Kurtosis	27.1

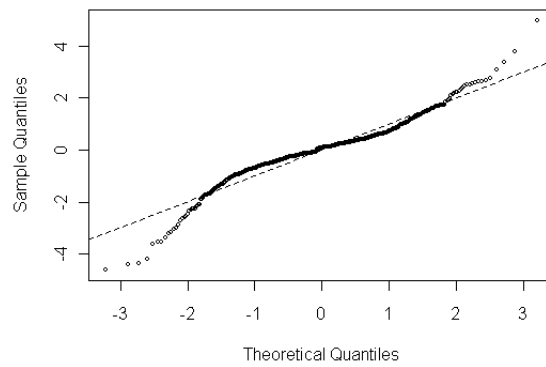
Figure 2.7: Gas log-return descriptive statistics.

Again we have an extremely leptokurtic sample, in part due to the extreme price crash that occurred in October 2006. Note that the sample kurtosis is not a robust estimator and is sensitive to outliers like this. The unit root null hypothesis is rejected at 99.9% level (Dickey Fuller statistic = -38.57) indicating that mean-reversion is present (under the model specification above).

Q-Q plot of power returns against the Normal distribution



Q-Q plot of coal returns against the Normal distribution



Q-Q plot of gas returns against the Normal distribution

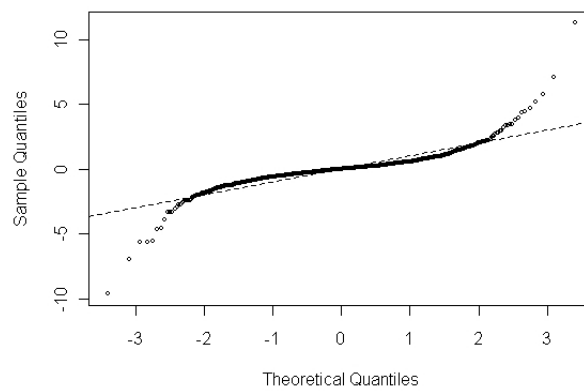


Figure 2.8: Q-Q plots of normalized commodity log-returns against the standard Normal distribution.

The Jarque-Bera test for normality was applied to the deseasonalised log-returns of all three markets and H_0 , the null hypothesis that the samples were from a normal distribution was rejected at the 99.99% level. These unit root tests suggest that we should at least be looking at models that specify mean-reversion. We will build a two-factor mean-reversion model. In fact we should also consider implementing jump-diffusion models which would give us more leptokurtic distributions, but this will be the subject of future work outside this thesis.

2.3 Risk-neutral probability measures and risk premium

The bulk of financial mathematics involves using continuous or càdlàg submartingale processes to model prices using a real-world probability measure \mathbb{P} . The main reason for this is the existence of a set of *risk neutral probability measures*, usually denoted by \mathcal{Q} , under which the submartingale processes become (local) martingales. This is a consequence of Girsanov's theorem (given some assumptions, see for example [19, 15]). We also need to assume that our underlying market satisfies certain properties if we are to substitute our real-world probability measure for a risk-neutral one.

Assumption 2.3.1. The market is arbitrage-free (i.e. arbitrage opportunities do not exist).

Note that this assumption is rarely true in practice in any market and the power market is no exception. However, any arbitrage opportunities that exist will quickly be exploited by market competitors.

We invoke the Fundamental Theorem of Asset Pricing [18]. Since the market is arbitrage-free, the set of risk-neutral measures that are equivalent to \mathbb{P} is non-empty.

Definition 2.3.1. A market is *complete* if every contingent claim admits a unique arbitrage-free price.

Recall that under a risk-neutral probability measure all *tradable* assets are martingales after discounting. The spot asset is not tradable, so we are left with the bank

account, which trivially becomes a martingale under any equivalent measure \mathbb{Q} . It is not possible to establish a unique forward price dynamics based on arbitrage arguments, and the market is incomplete.

Take \mathbb{Q} to be the *forward measure* with the zero-coupon bond as numeraire (for more information regarding this risk adjusted measure see [6]).

Under the forward measure, forward prices are martingales. Defining the spot price S_t as the limit of the forward prices as t approaches T , assuming we can pass limits through the expectation, we have:

$$F(t, T) = \mathbb{E}_{\mathbb{Q}}(S(T)|\mathcal{F}_t),$$

Assumption 2.3.2. Assume that the *rational expectation hypothesis* holds which states that the forward price is the best prediction of the spot price at delivery, i.e.:

$$F(t, T) = \mathbb{E}_{\mathbb{P}}(S(T)|\mathcal{F}_t).$$

Definition 2.3.2. The *risk premium* denoted by $\text{RP}(t, T)$ is given by

$$\text{RP}(t, T) = F(t, T) - \mathbb{E}_{\mathbb{P}}[S(T)|\mathcal{F}_t]$$

Stating that the rational expectation hypothesis holds is equivalent to stating that there is no *risk premium*. The presence of a risk premium would indicate that hedgers would prefer buying forwards over spot or vice versa. This topic in itself has attracted much research, with some arguing that the rational market hypothesis does not hold (in other markets). Evidence from the Pennsylvania-New Jersey-Maryland market shows that the rational market hypothesis does not hold [16], but it is not known whether similar research has been attempted for the UK market. We adopt the current EDF Energy viewpoint that there is no risk premium. This helps us from a modelling point of view - we can choose $\mathbb{Q} = \mathbb{P}$ for the remainder of the chapter.

2.4 2-Factor Heath-Jarrow-Morton model

In bond pricing applications instead of modelling via spot models the forward rates are specified [5]. Clewlow and Strickland [8] propose a similar approach for energy markets which relies on taking a forward curve and simulating how this forward curve will evolve through time. The Stochastic Differential Equation describing the process followed by the unitary forward curve $F(t, T)$ is:

$$\frac{dF(t, T)}{F(t, T)} = \sum_{k=1}^n \sigma_k(t, T) dW_k^{\mathbb{Q}}(t) ,$$

where $W_k, k \in \{1, \dots, n\}$ are possibly correlated Brownian motions under the forward measure $\mathbb{Q}(= \mathbb{P})$. The model assumes that the structure of market prices are described by these Brownian motions. Each factor has its own associated volatility term structure $\sigma_k(t, T)$.

We will now look specifically at a two factor mean-reverting forward curve case with volatility functions as defined by Clewlow and Strickland. That is, $n = 2$. This will be the model that will be implemented.

Define

$$\sigma_1(t, T) = \sigma_s e^{-a(T-t)}$$

and

$$\sigma_2(t, T) = \sigma_l ,$$

where $a > 0$. Implicit in this formulation is an assumption that volatility is not seasonal, which is not necessarily the case [4], but it is adopted for simplicity. The SDE becomes

$$\frac{dF(t, T)}{F(t, T)} = \sigma_s e^{-a(T-t)} dz_{s,t} + \sigma_l dz_{l,t} ,$$

with mean reversion speed a , short-term volatility σ_s and long-term volatility σ_l and where z_s and z_l are Brownian Motions with correlation ρ .

We apply Itô's formula for semi-martingales to the SDE to give an expression for the forward price. We are assuming that F is a continuous semi-martingale.

Define a twice-differentiable function

$$f(x) = \log x.$$

Itô's formula states that for a continuous semi-martingale X :

$$df(X) = f'(X)dX + \frac{1}{2}f''(X)d\langle X \rangle.$$

Then we get

$$d \log F(t, T) = \frac{dF(t, T)}{F(t, T)} - \frac{1}{2F(t, T)^2} d\langle F(t, T) \rangle.$$

We denote $F(t, T)$ as F for conciseness:

$$\begin{aligned} d\langle F \rangle &= d\langle F(e^{-a(T-t)}\sigma_s z_{s,t} + \sigma_l z_{l,t}) \rangle \\ &= F^2(e^{-2a(T-t)}\sigma_s^2 + \sigma_l^2 + 2e^{-a(T-t)}\rho\sigma_s\sigma_l) dt. \end{aligned}$$

Therefore

$$d(\log F) = \frac{dF}{F} - \frac{1}{2}(e^{-2a(T-t)}\sigma_s^2 + \sigma_l^2 + 2e^{-a(T-t)}\rho\sigma_s\sigma_l) dt.$$

If we note

$$V(t_0, t, T) = \int_{t_0}^t \sigma_s^2 e^{-2a(T-u)} du + \int_{t_0}^t \sigma_l^2 du + \int_{t_0}^t 2\rho\sigma_l\sigma_s e^{-a(T-u)} du, \quad (2.1)$$

$$W_s(t_0, t) = \int_{t_0}^t \sigma_s e^{au} dz_{s,u}, \quad (2.2)$$

$$W_l(t_0, t) = \int_{t_0}^t \sigma_l dz_{l,u}. \quad (2.3)$$

Substituting $\frac{dF}{F}$ and rearranging then gives:

$$F(t, T) = F(t_0, T) \exp\left(-\frac{1}{2}V(t_0, t, T) + e^{-aT}W_s(t_0, t) + W_l(t_0, t)\right). \quad (2.4)$$

The spot price can then be defined as $S_t = F(t, t)$ and the spot price dynamics

follow directly:

$$S_t = F(t_0, t) \exp\left(-\frac{1}{2} V(t_0, t, t) + e^{-at} W_S(t_0, t) + W_L(t_0, t)\right). \quad (2.5)$$

2.4.1 Volatility Term Structure

Unlike the Geometric Brownian Motion which has a single volatility value at each point in time, the 2-factor model has a volatility term-structure. This means that the instantaneous volatility of a particular point on the forward curve depends on the valuation date t and the delivery time T . The term structure tells us how volatile the price is depending on how far the product is from delivery.

Definition 2.4.1. We call *instantaneous volatility* of a particular unitary forward product $F(t, T)$ the time-to-maturity varying volatility function $\Sigma_{inst}(t, T)$ obtained while representing the log-return of the product as a unique Brownian Motion process. It is a measure of how volatile the price is with respect to the time-to-delivery.

This definition assumes that we can write the process as:

$$\frac{dF(t, T)}{F(t, T)} = \sigma_s e^{-a(T-t)} dz_{s,t} + \sigma_l dz_{l,t} = \Sigma_{inst}(t, T) d\tilde{Z}_t,$$

for a new Brownian Motion \tilde{Z} .

Lemma 2.4.1. Let z_t^* be a Brownian motion independent to $z_{s,t}$. The linear combination of Brownian Motions $\rho z_{s,t} + \sqrt{1 - \rho^2} z_t^*$ is also a Brownian motion and has correlation ρ with $z_{s,t}$.

Proof. First of all, a linear combination of Brownian motions is obviously a martingale, with initial value 0. Secondly, we calculate the quadratic variation of the linear combination:

$$\langle \rho z_{s,t} + \sqrt{1 - \rho^2} z_t^* \rangle = \rho^2 t + (1 - \rho^2) t = t$$

Therefore, by Lévy's Characterization Theorem it is a Brownian motion. The

covariance is calculated as follows:

$$\begin{aligned}
\text{Cov}(z_{s,t}, \rho z_{s,t} + \sqrt{1 - \rho^2} z_t^*) &= \mathbb{E}[z_{s,t}(\rho z_{s,t} + \sqrt{1 - \rho^2} z_t^*)] \\
&= \mathbb{E}[\rho z_{s,t}^2 + \sqrt{1 - \rho^2} z_{s,t} z_t^*] \\
&= \rho t.
\end{aligned}$$

Therefore we have correlation

$$\text{Corr}(z_{s,t}, \rho z_{s,t} + \sqrt{1 - \rho^2} z_t^*) = \rho t / t = \rho.$$

□

Therefore we can represent two correlated Brownian motions as a linear combination of two independent Brownian motions, which is itself a Brownian motion (the proof of this is based on the fact that the sum of two independent Normal variables are Normally distributed, but we will omit this). Therefore the assumption we make in the definition is valid.

Then we can calculate the quadratic covariation of the left hand side and equate it to the quadratic variation of the right hand side

$$\begin{aligned}
\left\langle \frac{dF}{F} \right\rangle &= \langle \sigma_s e^{-a(T-t)} dz_{s,t} + \sigma_l dz_{l,t} \rangle \\
&= (\sigma_s^2 e^{-2a(T-t)} + \sigma_l^2 + 2\rho\sigma_s\sigma_l e^{-a(T-t)}) dt \\
&= \Sigma_{inst}^2 dt = \langle \Sigma_{inst} d\tilde{Z}_t \rangle.
\end{aligned}$$

This implies that

$$\Sigma_{inst}(t, T) = \sqrt{\sigma_s^2 e^{-2a(T-t)} + \sigma_l^2 + 2\rho\sigma_s\sigma_l e^{-a(T-t)}}.$$

The equation shows that:

- as T becomes large, the instantaneous volatility tends to σ_l and the forward product behaves like a Geometric Brownian Motion,

- as $T - t$ approaches zero, the forward product also behaves like a Geometric Brownian Motion, but with volatility $\sqrt{\sigma_s^2 + \sigma_l^2 + 2\rho\sigma_s\sigma_l}$.

This type of behaviour is observable in the market where the volatility increases as we approach delivery and products with far delivery tend to follow a simple random walk.

Definition 2.4.2. We will call *equivalent volatility*² $\Sigma_{eq}(t_0, t_m, T)$ of the forward product for delivery at time T the average of the instantaneous volatility of $F(t, T)$ between 2 dates t_0 (the quotation date) and t_m (m for maturity).

$$\Sigma_{eq}(t_0, t_m, T) = \sqrt{\frac{1}{t_m - t_0} \int_{t_0}^{t_m} [\sigma_s^2 e^{-2a(T-u)} + \sigma_l^2 + 2\rho\sigma_l\sigma_s e^{-a(T-u)}] du}.$$

This is a very important notion for the following reasons:

- Seen as a function of time-of-delivery T , the equivalent volatility defines the *volatility term-structure* of all forward prices. This term-structure can be observed in the market.
- It can be seen as the volatility to be used when pricing a vanilla option on the forward at delivery T expiring at time t_m via standard methods (namely Black-Scholes closed-form formulas).
- It will be the volatility used for model parameter estimation.

2.4.2 Extension to multi-commodity simulations

Another advantage of this model is that it is also popular for modelling other commodities [8]. In our case, those relevant commodities are coal and gas. The model can be extended to simulate the prices of other commodities, capturing the correlation between them. The spot and forward price dynamics are given in Equations 2.5 and 2.4 when the volatility, correlation and mean-reversion parameters of the underlying single commodity are known. However, markets (power,

²The equivalent volatility is sometimes referred to as marginal volatility or model-implied volatility.

gas and coal in particular) are inter-correlated, so it is necessary to include the correlation between markets if we are going to be modelling, say, distributions of spreads. Stochastic simulations over a set of markets must then be done in a consistent way while drawing random numbers. There will be extra correlation parameters between the risk factors of different markets.

If we have M commodity markets, the number of random variables to draw for each step of a forward curve movement or spot price simulation is $2M$ (short-term and long-term Brownian motion). Denote the correlation coefficient between variables i and j amongst those $2M$ risk factors by ρ_{ij} . The correlation matrix $C = (\rho_{ij})_{i,j=1,\dots,2M}$ describes how the $2M$ risk factors are related. A correlation matrix is always symmetric and positive-semidefinite so it has a Cholesky decomposition.

$$C = LL^T.$$

If we draw a vector of $2M$ independent standard normal variables, it is possible to obtain correlated variables \tilde{Y} by the following operation:

$$\tilde{Y} = LY.$$

This is by well-known properties of the multivariate normal distribution.

2.4.3 Correction for Product Delivery Period

The equations above are only valid for instantaneous product delivery. In reality, there is no such thing in power markets. Power forwards are for a specified delivery period; days, weeks, months, seasons or years. The simulations produced for products with long delivery periods will have volatilities that are too high in reality, because the actual price path within the product is in effect averaged to give the actual product price, and this will affect the volatility of the product itself. It is quite intuitive to expect the volatility of yearly products to be much less than daily products, for example.

To cater for this we adopt the Clewlow and Strickland approach of reducing the volatility before it gets used in the simulations by a certain amount based on the

delivery period of the product. Given an unitary forward curve, from a simple no arbitrage argument, the price of a non-infinitesimal forward product $F_p(t, T, \theta)$ specified by a delivery time T and delivery period θ is given by:

$$F_p(t, T, \theta) = \frac{1}{\theta} \int_T^{T+\theta} F(t, u) du. \quad (2.6)$$

For our simulations, θ will be one day but we keep it general for now to emphasize that it can be used for other delivery lengths. The ‘spot’³ price associated with the same product is:

$$S_p(t, \theta) = \frac{1}{\theta} \int_t^{t+\theta} F(t, u) du.$$

Manipulating Equation 2.6 leads to:

$$F_p(t, T, \theta) - F_p(t_0, T, \theta) = \frac{1}{\theta} \int_T^{T+\theta} (F(t, u) - F(t_0, u)) du.$$

Then from the definition of the model in Equation 2.4 in integral form we have:

$$F_p(t, T, \theta) - F_p(t_0, T, \theta) = \frac{1}{\theta} \int_T^{T+\theta} F(t, u) \left(\int_{k=t_0}^u \sigma_s e^{-a(u-k)} dz_{s,k} + \int_{k=t_0}^u \sigma_l dz_{l,k} \right) du. \quad (2.7)$$

Then the aim is to find a model for the dynamics of the product $F_p(t, T, \theta)$. This means we no longer want the instantaneous-delivery forward prices $F(t, u)$ to appear in the right hand side of the previous equation. A simple approximation involves assuming that $F(t, u)$ can be approximated by $F_p(t, T, \theta)$ for $u \in [T, T + \theta]$. Our justification is that the time period $[T, T + \theta]$ is relatively small because we will be simulating for daily delivery periods. We offer no theoretical basis for this estimate, but it will be possible to validate our assumptions by comparing the volatility of our simulations to the volatility of the empirical data used to calibrate the model. It would also be possible to introduce a seasonal deterministic ‘shape factor’ for longer periods. With slight abuse of notation (differential notation on the left hand side, integral notation on the right):

³The word ‘spot’ should be understood here as ‘instant delivery of a non-instantaneous product’

$$dF_p(t, T, \theta) = \frac{1}{\theta} \int_T^{T+\theta} F_p(t, T, \theta) \left[\int_{k=t_0}^u \sigma_s e^{-a(u-k)} dz_{s,k} + \int_{k=t_0}^u \sigma_l dz_{l,k} \right] du$$

Therefore we can write

$$\begin{aligned} \frac{dF_p(t, T, \theta)}{F_p(t, T, \theta)} &= \frac{1}{\theta} \int_T^{T+\theta} \left(\int_{k=t_0}^u \sigma_s e^{-a(u-k)} dz_{s,k} + \int_{k=t_0}^u \sigma_l dz_{l,k} \right) du \\ &= \frac{1}{\theta} \int_T^{T+\theta} \left(e^{-a(u-T)} \int_{k=t_0}^u \sigma_s e^{-a(T-k)} dz_{s,k} + \int_{k=t_0}^u \sigma_l dz_{l,k} \right) du \\ &= \frac{1}{\theta} \int_T^{T+\theta} \left(e^{-a(u-T)} \int_{k=t_0}^T \sigma_s e^{-a(T-k)} dz_{s,k} + \int_{k=t_0}^T \sigma_l dz_{l,k} \right) du \\ &\quad + \frac{1}{\theta} \int_T^{T+\theta} \left(e^{-a(u-T)} \int_{k=T}^u \sigma_s e^{-a(T-k)} dz_{s,k} + \int_{k=T}^u \sigma_l dz_{l,k} \right) du \end{aligned}$$

The final term in the last statement makes the model non-markovian, which introduces problems in our simulation. It doesn't seem to be easy to simplify either. We must make another approximation. The integrals inside the final term have expectation 0 (property of a Brownian stochastic integral). For relatively small θ the integrals have relatively small variance. Therefore we will assume that the integrals are very small, approximately zero. This gives us the following SDE for the product with delivery period θ . For larger theta the approximation is less valid.

$$\begin{aligned} \frac{dF_p(t, T, \theta)}{F_p(t, T, \theta)} &= \frac{1}{\theta} \int_T^{T+\theta} e^{-a(u-T)} du \sigma_s e^{-a(T-t)} dz_{s,t} + \sigma_l dz_{l,t} \\ &= c_{a,\theta} \sigma_s e^{-a(T-t)} dz_{s,t} + \sigma_l dz_{l,t}, \end{aligned}$$

where we define a new constant

$$c_{a,\theta} := \frac{1}{a\theta} (1 - e^{-a\theta}).$$

Then Itô's formula is applied in exactly the same way as in Equation 2.5 to give new dynamics. This is simple manipulation and the details are omitted.

2.4.4 Monte Carlo simulation

This following section describes how we can generate price simulations using this model in R. We will implement the multi-commodity spot model with non-infinitesimal delivery period. The spot model is essentially the same as the forward model but the observation date is the same as the delivery date, i.e. instant delivery, $T = t$. We will be implementing the model for a daily granularity, with daily delivery period. The constant $c_{a,\theta}$ will be used, as defined in the previous subsection.

Rewriting Equations 2.1 - 2.3 we have:

$$S(t) = F(t_0, t) \exp[-0.5(V_l + V_s + V_{sl}) + e^{-at}W_s + W_l],$$

where

$$\begin{aligned} V_l &= \int_{t_0}^t \sigma_l^2 du = \sigma_l^2(t - t_0), \\ V_s &= \int_{t_0}^t c_{a,\theta}^2 \sigma_s^2 e^{-2a(T-u)} du = \frac{c_{a,\theta}^2 \sigma_s^2}{2a} [e^{-2a(T-t)} - e^{-2a(T-t_0)}], \\ V_{sl} &= \int_{t_0}^t 2c_{a,\theta} \rho \sigma_l \sigma_s e^{-a(T-u)} du = \frac{2\rho c_{a,\theta} \sigma_l \sigma_s}{a} [e^{(T-t)} - e^{(T-t_0)}]. \end{aligned}$$

Note that only two random processes are required to define a complete forward curve. It is straightforward to compute the deterministic terms (V_l , V_s and V_{sl}). Let us examine the short-term random variable $W_s(t_0, t)$.

Recalling its definition, with adjustment for the non-infinitesimal delivery period, we have

$$W_s(t_0, t) = \int_{t_0}^t c_{a,\theta} \sigma_s e^{au} dz_{s,u}.$$

This is a Brownian stochastic integral. From the properties of the stochastic inte-

gral, for some arbitrary time for a small increment dt :

$$\begin{aligned} W_s(t_0, t + dt) - W_s(t_0, t) &\sim N\left(0, \int_t^{t+dt} c_{a,\theta}^2 \sigma_s^2 e^{2au} du\right) \\ &\sim N\left(0, \frac{c_{a,\theta}^2 \sigma_s^2}{2a} e^{2at} (e^{2adt} - 1)\right) \\ &\sim \sqrt{\frac{\sigma_s^2}{2a} e^{2at} (e^{2adt} - 1)} \epsilon_s, \end{aligned}$$

where $\epsilon_s \sim N(0, 1)$. Similarly it can be shown that

$$W_l(t_0, t) \sim \sqrt{\sigma_l^2 (t - t_0)} \epsilon_l.$$

Therefore it is possible to approximate the stochastic processes by generating W_s and W_l step-wise. In the single commodity model, we can correlate the approximate Brownian motions at each stage by drawing two independent $N(0,1)$ distributed variables ϵ and $\tilde{\epsilon}$ then

$$\begin{aligned} \epsilon_s &= \epsilon, \\ \epsilon_l &= \rho\epsilon + \sqrt{1 - \rho^2} \tilde{\epsilon}, \end{aligned}$$

where ρ is the correlation required between ϵ_s and ϵ_l . This is extended for multiple commodity markets using the approach suggested in the earlier subsection ‘Extension to multi-commodity simulations’. For the interested reader, the code for the implementation of the spot model is available in the appendix. The antithetic variance reduction technique was used in the Monte Carlo simulations to reduce the number of simulations required.

Below are some examples of simulations for the coal and gas markets using forward curves provided by EDF Energy, using the following parameters, which are typical of the markets (we will investigate this further in the next section on calibration).

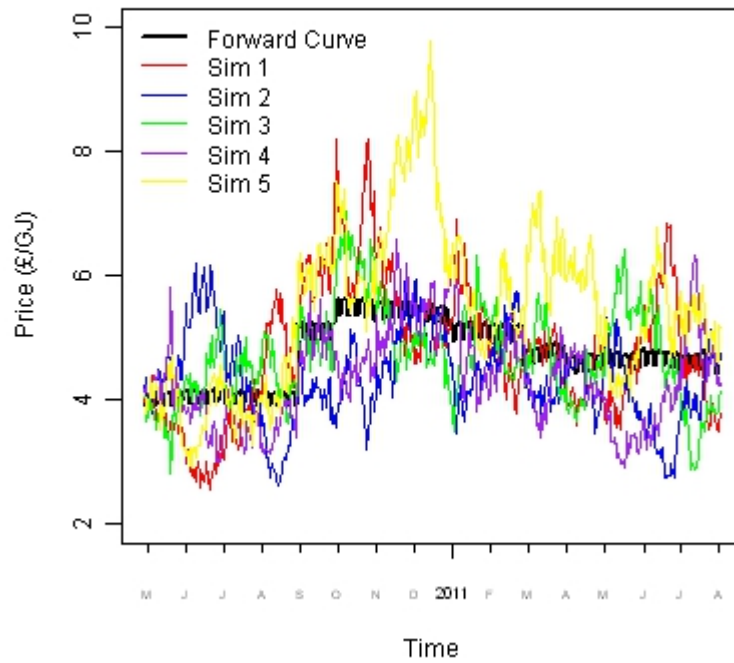
	Gas ST	Gas LT	Coal ST	Coal LT
Gas ST	1	-0.12	0	-0.08
Gas LT	-0.12	1	0	0.56
Coal ST	0	0	1	0
Coal LT	-0.08	0.56	0	1

Figure 2.9: Example correlations between the gas and coal short and long term factors.

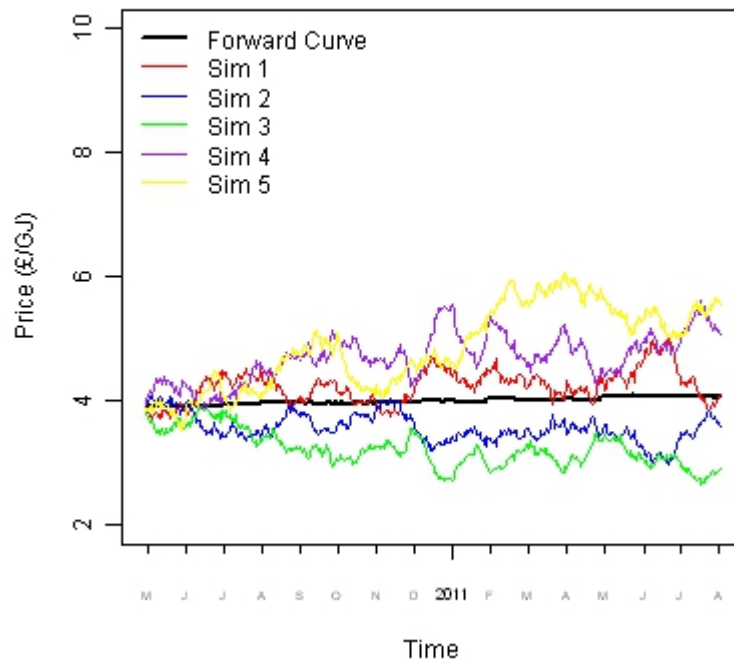
	Gas	Coal
ST Vol	94.4%	0%
LT Vol	38.1%	30.5%
a	14.5	1

Figure 2.10: Example parameters for the coal and gas markets.

Gas Spot Price Simulations (May 10 - Aug 11)



Coal Spot Price Simulations (May 10 - Aug 11)



It is interesting to observe the effect of antithetic sampling between simulations 1 and 2 and simulations 3 and 4.

2.4.5 Parameter Estimation

Two approaches to parameter estimation are described. The first, a linear programming approach, appears to work well initially but it turns out to have drawbacks. An second iterative method is developed and implemented instead which gives better results.

The granularity is set to be equal to 1 day. Recall that the market consists of products with different delivery lengths. The following products have a closing price for each day: Day Ahead, Week Ahead, Week 2, Month 1, M2, M3, M4; Quarter 1, Q2, Q3, Q4, Season 1, S2, S3, S4, S5, and the time periods are consistent with the EFA calendar. There is some product overlap. For example, Season 1, which is either Winter or Summer, will be either Q1 and Q2 or Q2 and Q3. Month 1 may or may not comprise Week 1 and/or Week 2 and so on. As a concrete example, assume today is 21 May 2010. The Day Ahead product is 22 May 2010. Week Ahead is 24-30 May 2010. Week 2 is 31 May - 6 June. Month 1 is 31 May - 4 July (EFA June). Q1 is 5 July - 4 October. Season 1 is 4 Oct 2010 - 3 April 2011, etc.

Define t_0 as the time a product first becomes available to trade in the market, t_m as the time a product is no longer available to trade in the market, and T the start of the delivery period for that product. Each product can have a different delivery period length θ , which as shown previously affects the volatility. Below is a table with the parameters of different electricity products.

Product	$t_m - t_0$	θ	$T - t_m$	Mean Time to Delivery
Day Ahead	1/365	1/365	0	0.5/365
Week Ahead	1/52	1/52	0	0.5/52
W2	1/52	1/52	1/52	1.5/52
M1	1/12	1/12	0	0.5/12
M2	1/12	1/12	1/12	1.5/12
M3	1/12	1/12	2/12	2.5/12
M4	1/12	1/12	3/12	3.5/12
Q1	1/4	1/4	0	0.5/4
Q2	1/4	1/4	1/4	1.5/4
Q3	1/4	1/4	2/4	2.5/4
Q4	1/4	1/4	3/4	3.5/4
S1	1/2	1/2	0	0.5/2
S2	1/2	1/2	1/2	1.5/2
S3	1/2	1/2	2/2	2.5/2
S4	1/2	1/2	3/2	3.5/2
S5	1/2	1/2	4/2	4.5/2

Figure 2.11: Properties of different electricity forward products

Daily price data for the period 1st May 2009 - 30 April 2010 was obtained for these products. The annualized log-returns R_t were computed in the usual manner ($R_t = \sqrt{1/252} \log S_t/S_{t-1}$). The standard deviation of the annualized log-returns of these products give us an estimate of their historical volatilities. It is important to filter out any log-returns that are due to the change in the underlying product itself. For example, on 30 May 2010, the M1 product is June power, but on 31 May, M1 becomes July power, so we would disregard the log-return between these dates because it is misleading. From the definition, with adjustment for the

non-infinitesimal product delivery length, equivalent volatility is given by

$$\begin{aligned}\Sigma_{eq}(t_0, t_m, T) &= \sqrt{\frac{1}{t_m - t_0} \int_{t_0}^{t_m} [c_{a,\theta}^2 \sigma_s^2 e^{-2a(T-u)} + \sigma_l^2 + 2c_{a,\theta} \rho \sigma_l \sigma_s e^{-a(T-u)}] du} \\ &= \sqrt{\frac{1}{t_m - t_0} \left(c_{a,\theta}^2 \frac{\sigma_s^2}{2a} [1 - e^{-2a(t_m-t_0)}] + \sigma_l^2 (t_m - t_0) + \frac{2c_{a,\theta} \rho \sigma_s \sigma_l}{a} e^{-a(T-t_m)} [1 - e^{-a(t_m-t_0)}] \right)}.\end{aligned}$$

The equivalent volatility best fit (relative to historical volatility) should give us our optimal parameters $(\hat{a}, \hat{\sigma}_s, \hat{\sigma}_l, \hat{\rho})$. Initially a linear programming approach was adopted for calibration in a single market (electricity) setting. The following LP was set up and solved.

For products $i \in \{1, \dots, P\}$, minimize the objective function:

$$\sum_{i=1}^p (\Sigma_{hist,i} - \Sigma_{eq,i})^2$$

subject to the constraints:

$$\begin{aligned}|\rho| &\leq 1 \\ \sigma_s &\geq 0 \\ \sigma_l &\geq 0 \\ a &> 0.\end{aligned}$$

The best fit results obtained for the electricity market are noted in the tables below.

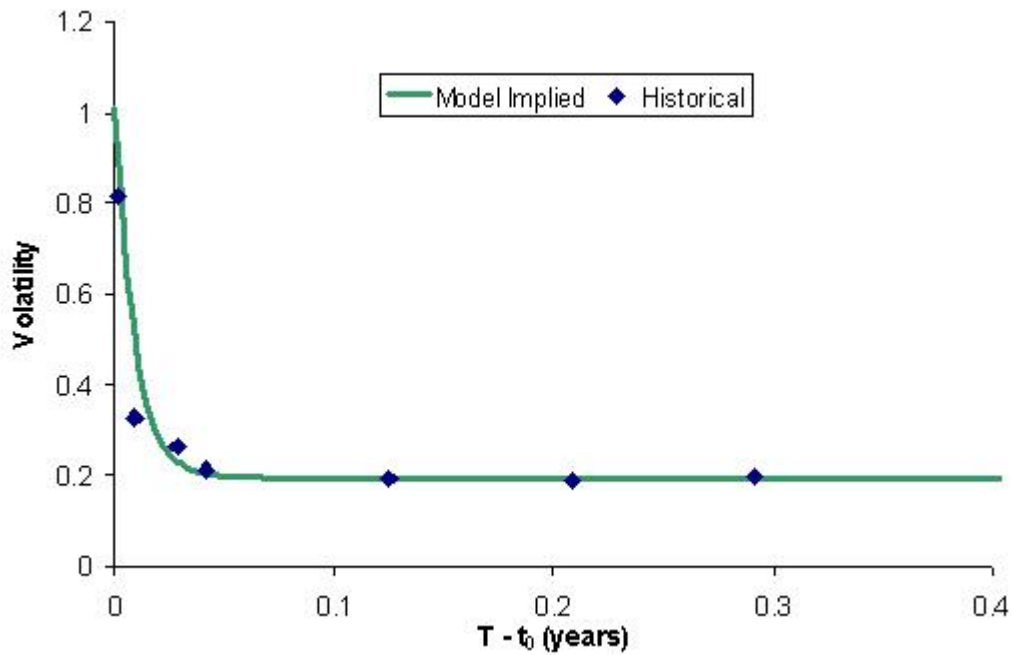
Parameter	Value
σ_s	87.3%
σ_l	19.5%
a	100.3
ρ	0.67

Figure 2.12: Best fit parameters for electricity.

Product	Mean Time to Delivery	Historic Vol. (%)	Equivalent Vol. (%)
Day Ahead	0.001	81.7	81.4
Week Ahead	0.01	32.9	34.8
W2	0.03	26.6	21.3
M1	0.04	21.3	20.5
M2	0.13	19.4	19.51
M3	0.21	19.0	19.51
M4	0.29	19.6	19.51
S2	0.75	18.9	19.51
S3	1.25	18.4	19.51
S4	1.75	17.1	19.51

Figure 2.13: Fitting Equivalent Volatility to Historic Volatility with best fit parameters.

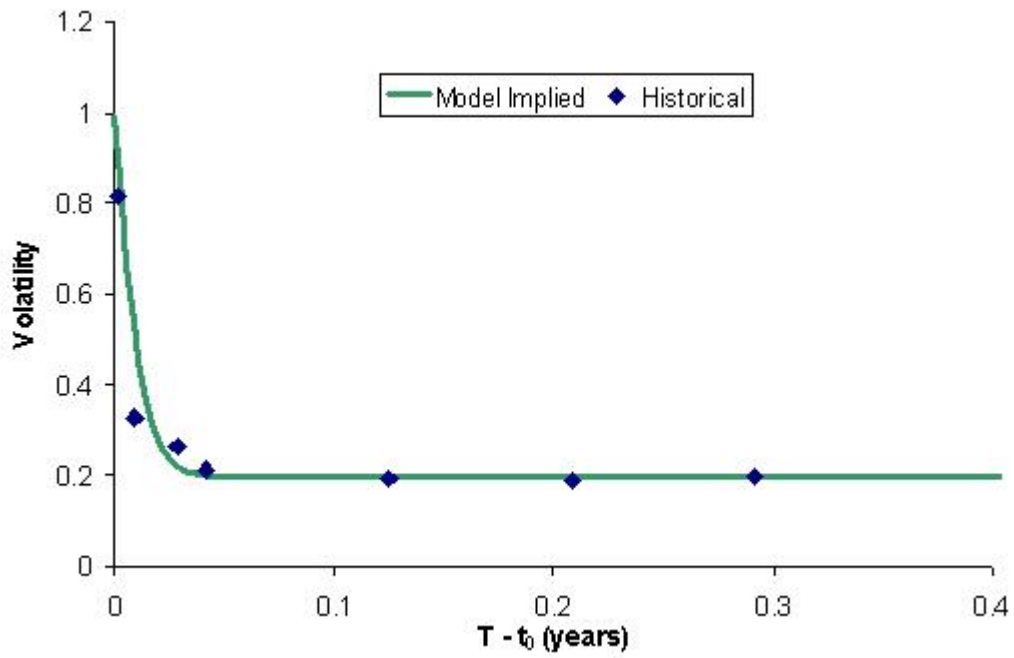
with the following plot of best fit we obtained for the volatility term structure.



Unfortunately however there are some problems with this method. One problem is that of apparent redundancy in the parameters. This was apparent if we fixed $\rho = 0$ (for example) and re-solved the LP to give new parameters.

Parameter	Value
σ_s	97.5%
σ_l	19.7%
a	78.9
ρ	0

Figure 2.14: New model parameters with fixed ρ .



Comparing the graphs of both solutions, there is little difference. We would like to find a way to break this apparent redundancy. Also, this approach does not solve the multiple market calibration problem. Is there more information in the historical data? We develop an iterative parameter estimation procedure which estimates the parameters in a more robust way.

Recall the stochastic differential equation that each product $p \in \{1, \dots, P\}$ satisfies:

$$d(\log F_p) + \frac{1}{2} \Sigma_{inst}^2(t, T, a, \sigma_s, \sigma_l, \rho) dt = \frac{dF_p(t, T)}{F_p(t, T)} = c_{a, \theta_p} \sigma_s e^{-a(T-t)} dz_s + \sigma_l dz_l.$$

Assuming that the model accurately represents historical prices, this implies that the log-returns $R_{i,p}$ of a product p in observation period $i \in \{1, \dots, N\}$ can be written as:

$$R_{i,p} + \frac{1}{2} \int_{t_i}^{t_{i+1}} \Sigma_{inst}^2(s, T, a, \sigma_s, \sigma_l, \rho) ds \approx \sqrt{\frac{c_{a, \theta}^2 \sigma_s^2}{2a}} e^{2a(t_i - T)} (e^{2a(t_{i+1} - t_i)} - 1) \epsilon_{s,i} + \sqrt{\sigma_l^2 (t_{i+1} - t_i)} \epsilon_{l,i}$$

for each product p where $\epsilon_{s,i}, \epsilon_{l,i} \sim N(0, 1)$, assuming that the time between observation periods is small.

The left hand side (LHS) of the equation above for a specific product consists of the product's log-return in time period i plus a deterministic term which is a function of t_i, t_{i+1} and T and the parameters a, σ_s, ρ and σ_l . Using the first approximation of the parameters that we obtained from our Least Squares regression, the LHS can be computed for each time period i and for each product p .

Take a 'far-out' product such as Season 4 in the table above. We can assume that the short term volatility is negligible on the right hand side (RHS) of the equation for such a product. This makes it possible for us to compute estimates of the 'long-term shocks' $\sqrt{\sigma_l^2 (t_{i+1} - t_i)} \epsilon_{l,i}$ for each time period i . This implies that the returns of the Day Ahead product has the same long-term shocks. Using the initial OLS parameters, we can calculate LHS for the Day Ahead product and use the derived long-term shocks to give us the short term shocks $\sigma_s \epsilon_{s,i}$. We then measure the correlation between the short term shocks and the long term shocks.

$$\begin{aligned} \text{Corr}(\sigma_l \epsilon_{l,i}, \sigma_s \epsilon_{s,i}) &= \text{Corr}(\epsilon_{l,i}, \epsilon_{s,i}) \\ &\approx \text{Corr}(z_l, z_s) = \rho \end{aligned}$$

We have a new estimate for ρ .

This gives us an iterative procedure, which can be applied across several commodity markets. These are the steps:

- Step 1: Solve the LP for the 1st market to give the initial solution $P_0 = (\rho_0, a_0, \sigma_{s_0}, \sigma_{l_0})$;
- Step 2: Compute estimates of short term and long term shocks using the above procedure using parameters P_i
- Step 3: Compute a new ρ_i based on the correlation between the short term and long term shocks;
- Step 4: Resolve to give new estimates for σ_s , σ_l and a , with ρ fixed ($= P_i$);
- Step 5: Repeat step 2-4 until convergence is attained (e.g. $\rho_{i+1} - \rho_i < \epsilon$);
- Step 6: Repeat for each commodity market;
- Step 7: Calculate the correlations between short and long term shocks of different markets to give us correlation matrix C .

This method ensures that we have correlations across multiple markets that are reasonable and consistent with historical data (no proofs are given regarding the validity of this iterative procedure).

	$\sigma_s(\%)$	$\sigma_l(\%)$	ρ	a	Objective Function
S_0^*	87.3	19.5	0.68	100.3	0.0039
ρ_1			-0.055		
S_1^*	98.4	19.7		77.09	0.005
ρ_2			-0.055		

Figure 2.15: Power iteration scheme results.

Power Parameter	Value
σ_s	98.4%
σ_l	19.7%
a	77.09
ρ	-0.055

Figure 2.16: Power 2-factor model parameters.

We estimate the parameters for the coal and gas markets together. We keep the power market separate for reasons that will become clear later on. A practical point about the coal market is that prices are quoted in US Dollars. For the purposes of this model we simply use historical spot currency rates to convert the prices into UK pounds.

	$\sigma_s(\%)$	$\sigma_l(\%)$	ρ	a	Objective Function
S_0^*	33.00	18.57	-0.223	0.89	0.0000154
ρ_1			-0.124		
S_1^*	31.2	18.1		0.953	0.0000162
ρ_2			-0.124		

Figure 2.17: Coal iteration scheme results.

Coal Parameter	Value
σ_s	31.2%
σ_l	18.1%
a	0.953
ρ	-0.124

Figure 2.18: Coal 2-factor model parameters.

	$\sigma_s(\%)$	$\sigma_l(\%)$	ρ	a	Objective Function
S_0^*	64.7	35	0.71	5.78	0.0457
ρ_1			-0.261		
S_1^*	91.7	35.1		2.92	0.059
ρ_2			-0.261		

Figure 2.19: Gas iteration scheme results.

Gas Parameter	Value
σ_s	91.7%
σ_l	35.1%
a	2.92
ρ	-0.261

Figure 2.20: Gas 2-factor model parameters.

	G_s	G_l	C_s	C_l
G_s	1	-0.26	0.03	0.06
G_l	-0.26	1	0	0.32
C_s	0.03	0	1	-0.12
C_l	0.06	0.32	-0.12	1

Figure 2.21: Correlation between Gas and Coal short-term and long-term factors.

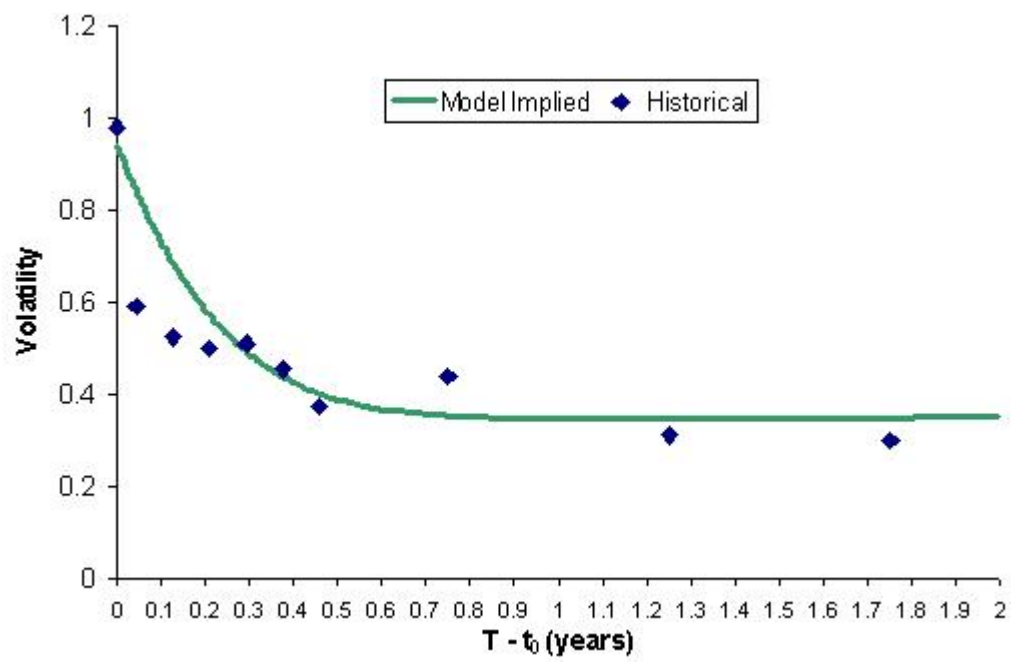


Figure 2.22: Model volatility term structure fit against historical data for Gas.

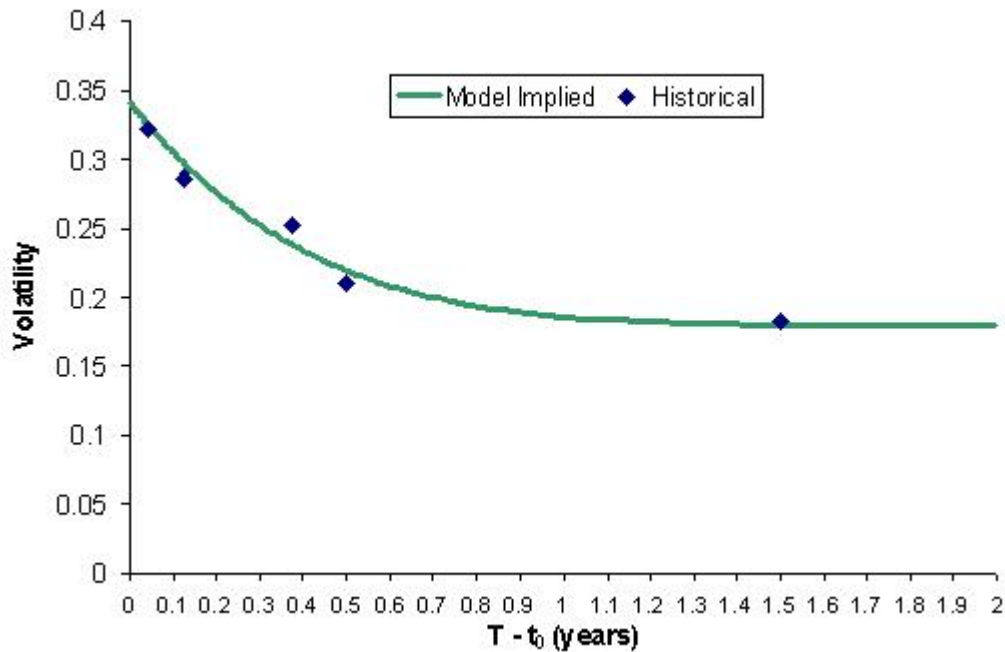


Figure 2.23: Model volatility term structure fit against historical data for Coal.

The fits for Coal and Power look good. The fit for Gas is not quite so good and it is recommended to check the data quality before these parameter estimates are used.

2.5 How appropriate is the HJM two-factor model for power, gas and coal?

Unfortunately the 2-factor HJM model does not capture the excess kurtosis in the power, gas or coal markets. The deseasonalized log-returns of our simulations will necessarily have kurtosis of 3 on average; we have shown in an earlier section that the actual kurtosis is far higher in power, coal and gas markets.

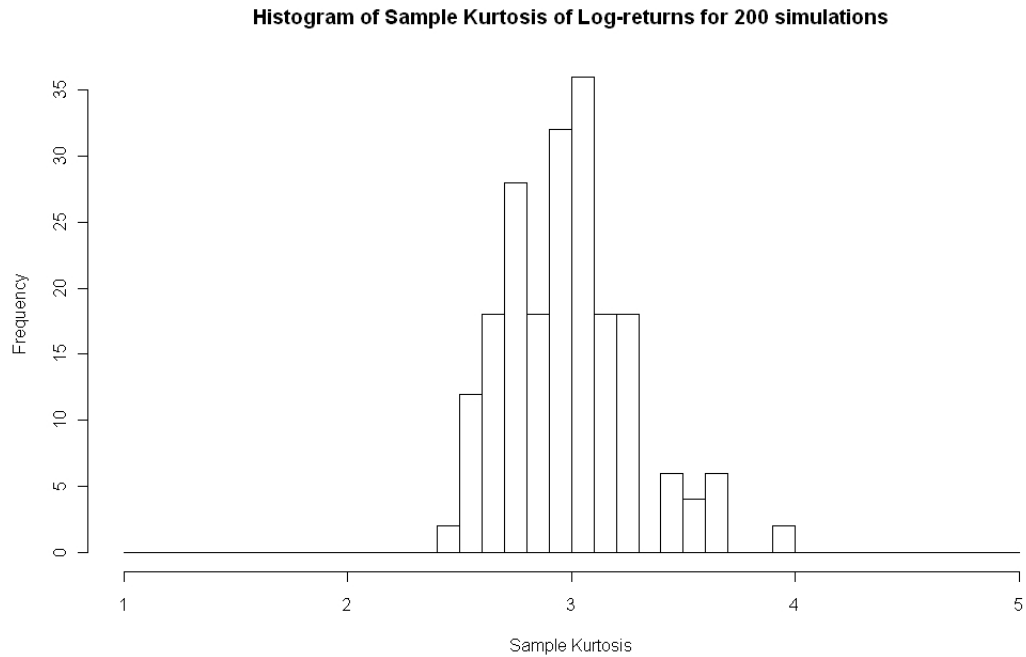


Figure 2.24: Histogram of the sample kurtosis of 200 simulations.

This kurtosis is a result of extreme jumps that occur more often than a normal distribution predicts. There is a great deal of recent literature on stochastic jump-diffusion models in the power market but we shall not go into detail about these in this thesis. For a comprehensive treatment see [4].

The skewness of the log-returns of our simulations will be 0 on average; compare this to the descriptive statistics in the earlier section where we saw that power and gas are slightly positively skewed, and coal is slightly negatively skewed.

There is also a reason why it may be inappropriate to build a 3-market model to compute the distribution of spreads. Recall the definition of clean spark/dark spreads - the price of power minus the price of coal and emission credits required to generate that power. Given the fuel mix of the UK power market, if clean spark and dark spreads both were negative at the same time, then it would not be economical for gas or coal power stations to run. If none of these power stations

ran, then there would not be sufficient power. This effectively forbids the price of these spreads to go negative at the same time. However, our 2-factor model with 3 markets allows negative coal and gas spreads. This would give spurious spread distributions. This is why we investigate the concept of the *fundamental model*.

Chapter 3

Fundamental Models

3.1 Background

The UK has roughly 80 GW of capacity which can convert some fuel (strictly, some energy resource) to power. The fuel split is roughly 20% nuclear, 35% gas, 35% coal, with the remaining capacity being provided by oil and renewables. Not all of this capacity will be available on a given day due to plant outages. Plant outages arise when either maintenance work is required (a planned outage) or when a generator fails (a forced outage). Generators notify the Grid Operator of their available capacity for each half-hour period. This is called a MEL (Maximum Export Limit). We will need a few definitions.

Definition 3.1.1. The *Heat Rate* is a generator's efficiency expressed in GJ/MWh.

Definition 3.1.2. The *marginal cost* of a generator is its Heat Rate (GJ/MWh) * Fuel Cost (£/GJ) + Emissions Rate (Tonne/MWh) * Emissions Cost (£/Tonne).

Each generator of a given fuel type will have a similar marginal cost. Therefore it is possible to construct a generation availability *stack* as shown below.

In a given half-hour period generation must satisfy demand. A simple stack price model assumes that the cheapest available generation will satisfy the required demand, and that no generator will run at a price below its marginal cost. This is the behaviour that would occur in perfectly competitive markets. Of course this is a simplistic assumption. Alternatives to the perfect competition assumption would

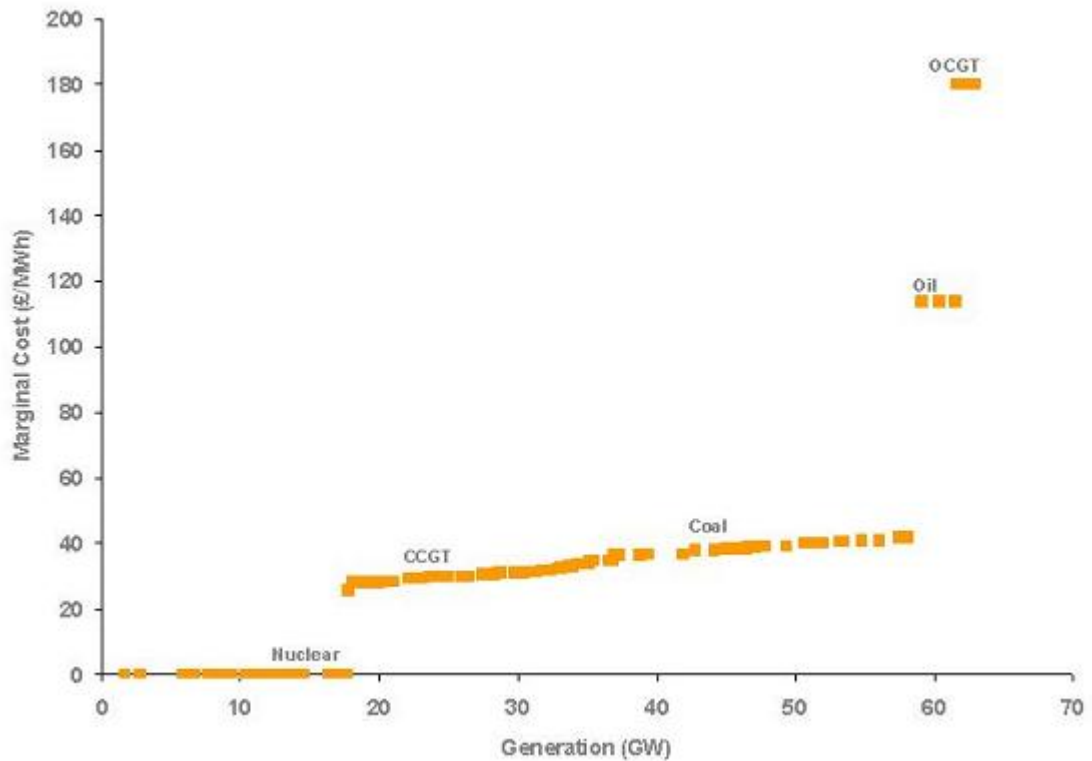


Figure 3.1: UK Availability Stack, Jan 15th 2010 17:00.

be to incorporate Cournot or Bertrand competition [12]. From the curve it is evident that power price jumps can be explained by the extreme convexity of the stack. When the margin (the difference between availability and demand) is tight, then a small change in demand may cause a large change in price. It is an easy problem to solve mathematically - the stack is ordered according to marginal cost then generation is added until demand is met. The system cost will be the maximum marginal cost of the units that are utilised.

Unfortunately generators have substantial *start costs*, and the effect of these is not captured by the simple stack model. Additionally, there are physical constraints which force a generator to be on/off for a period of time, *minimum up time* and *minimum down time*. Suddenly, our simple stack is no longer valid. We need the concept of ON/OFF to reflect these properties and a mathematical formulation. We formulate a mixed integer program to attempt to model these features.

3.2 Mixed Integer formulation

Let I be the number of generators and T be the number of time periods in our problem. Define \mathcal{I} and \mathcal{T} as the index sets $\{1, 2, \dots, I\}$ and $\{1, 2, \dots, T\}$.

Define IT binary variables $\{x_{it} : i \in \mathcal{I}, t \in \mathcal{T}\}$. $x_{it} = 1$ if generator i is ON in time period t , and $x_{it} = 0$ if the generator is OFF. Each generator has a Minimum Stable Level (MSL), m_i , and a Maximum Export Level (MEL), M_i . Each generator may generate between its MSL and its MEL. To allow this we introduce variables y_{it} , which are continuous between 0 and 1, and impose the constraint that $y_{it} \leq x_{it}$. Each generator has a start cost c_i . We introduce variables w_{it} with the constraint $w_{it} \geq x_{it} - x_{i,t-1}$. This is 1 in a period when a generator starts and 0 otherwise. Although w_{it} are integer variables we do not have to impose this condition. It is possible to define it as a continuous variable under a suitable constraint. Define p_{it} as the marginal cost of generator i in time period t and d_t the demand in time period t . We ignore stopping costs and minimum up/down times for now. Using these variables we may define the MIP problem as:

$$\min \sum_{t=1}^T \sum_{i=1}^I (x_{it}m_i + y_{it}(M_i - m_i))p_{it} + w_{it}c_i$$

subject to the constraints:

$$\sum_{i=1}^I x_{it}m_i + y_{it}(M_i - m_i) \geq d_t$$

for $t \in \mathcal{T}$

$$y_{it} \leq x_{it}$$

and

$$w_{it} \geq x_{it} - x_{i,t-1}$$

for $t \in \mathcal{T}$ and $i \in \mathcal{I}$

$$x_{it} \in \{0, 1\}.$$

This is known in the literature as the *unit commitment problem*. We implemented this Mixed Integer program in the R programming language, with the option of

using either the Symphony or the GLPK open-source solvers [2, 1]. Both solvers use some implementation of the branch and bound method. The code as well as example output are documented in the Appendix. To appreciate the power of the branch and bound solvers it is worth considering how many possible combinations of integer variables there are. For a 10 generator system, in an individual time period, our problem has 10 integer variables. Therefore there are $2^{10} - 1$ combinations of those variables. For a 48 period problem, there are $(2^{10} - 1)^{48}$ combinations of integer variables. Also it is worth noting how many variables and constraints we have. We have $3IT$ variables and $48 + I(2T - 1)$ constraints.

In practice then it is not surprising that the MIP formulation takes a long time to run. A test was conducted with 48 time periods and a variable number of generators for a given demand to see how long it would take. These experiments were conducted on a 3.00 GHz Intel Core 2 Duo PC with 3.5 GB of RAM.

Generators	Symphony	GLPK	Matrix Size (MB)
10	1	2	5.7
15	1.76	5.2	12.7
20	5.65	28.2	22.4
25	9.44	89	34.9
30	672	-	50
40	-	-	88.6

The problem is prepared in an $A\mathbf{x} \geq \mathbf{b}$ format and passed to the solver. The matrix has *constraints* multiplied by *variables* elements. Assume that each element uses 4 bytes. For a 10 generator, 48 period problem we have 1440 variables and 998 constraints. The matrix is approximately 5 MB in size. The matrix size increases in proportion to the number of generators squared. It would be better to use a sparse matrix data structure, which stores the position and value of non-zero matrix elements only. This would be useful in our case because our matrix consists of zeros mostly. However these solvers did not have routines to allow us to take advantage of the sparse nature of the matrix. It would be wise to use sparse matri-

ces if there is an option in the solver to use them.

3.2.1 Additional constraints

We will add spinning reserve, standing reserve and minimum up time to the mathematical formulation. Reserve constraints are relatively easy to solve because we will only need to add $3T$ extra constraints, however minimum up time imposes many constraints between time periods (known as temporal constraints) which add great complexity to the problem. Let us define a generator's raise spinning reserve as any remaining capacity a generator has when it is on. Define $SPIN \subseteq \mathcal{I}$ as the index set of generators that may provide raise spinning reserve, and let r_t and l_t be the amount of raise and lower spinning reserve that must be provided in each time period respectively. The additional raise spinning reserve constraint becomes:

$$\sum_{i \in SPIN} (x_{it} - y_{it})(M_{it} - m_{it}) \geq r_t ,$$

and the lower spin reserve constraint is:

$$\sum_{i \in SPIN} y_{it}(M_{it} - m_{it}) \geq l_t$$

for $t \in \mathcal{T}$.

Similarly with standing reserve, define $STAND \subseteq \mathcal{I}$ as the index set of generators which may provide standing reserve, and R_t as the amount of standing reserve that must be provided in each time period. We have:

$$\sum_{i \in STAND} (1 - x_{it})(M_{it} - m_{it}) \geq R_t$$

for $t \in \mathcal{T}$.

Define u_i as the minimum up time of generator i . The uptime constraints are:

$$x_{it} - x_{i,t-1} \leq x_{i,\tau}$$

$$\forall \tau \in [t + 1, \min(t + u_i, T)].$$

This adds much more complexity to the problem: for each generator i we have added approximately $u_i T$ constraints. Unfortunately it also means that the problem is no longer practical as a MIP.

EDF Energy use a commercial software product called Plexos to solve this problem. Plexos uses a heuristic to get around the problem of solving this complex MIP. It works by allowing the integer variables to be continuous, then solving the LP. Then it applies the rounding heuristic - setting the integer variables to either 0 or 1 whilst maintaining the validity of the constraints. Plexos does not give much information about how the product works but they do say that it takes about twice as long as the LP solution. More information about various rounding heuristics can be obtained from Burkard, Kocher and Rudolf [7]; these strategies are however beyond the scope of this thesis.

3.2.2 Computing a System Marginal Price (SMP)

We define the System Marginal Price for a time period t as follows:

$$\text{SMP}_t = \max(p_{it}) \quad \forall i \in \{i : y_{it} > 0\}.$$

The SMP is the maximum marginal cost of all generators that are running in a given period, except the generators that are running at MSL. The generators running at MSL are ignored for price setting purposes because most of the time generators will only run at MSL to avoid start costs in the future. We would not want these generators to be setting price because they would be willing to run at a loss.

3.2.3 Uplift

Consider the generator with the highest marginal price. When it is running (above MSL) then SMP will be its marginal cost. It will receive SMP for running. However, it has not and will not recover its start costs, and will be making a loss. It makes sense to add some amount to the System Marginal Price to compensate

them for their start costs. Plexos has its own algorithm for adding Uplift to SMP to give a Region Price. Developing an uplift model is beyond the scope of this thesis.

3.3 Results

A model of UK generators was set up in Plexos using the EDF Energy plant database. The model was then backtested against historical data. The model price is SMP + Plexos Uplift. Specifically it was possible to obtain historical demand and historical availability, as well as historical fuel spot prices to give us a price for power. We then compared this to the outturn ‘spot’ price of power. What should we use as the benchmark price of spot power? At daily granularity, we use the Day Ahead closing price. The Day Ahead closing price represents the last traded price of the next day’s power. The reason for using Day Ahead closing price is that it is a more liquid market than the APX half-hourly market.

3.3.1 Daily granularity - backtest against closing Day Ahead prices

We also gather descriptive statistics of the Model and Day Ahead prices and their log-returns. The data was for the period 30 March 2009 - 4 April 2010. Spot fuel prices and actual demand were used in the model.

Sample Statistic	Day Ahead	Model	Day Ahead log-returns	Model log-returns
n	371	371	370	370
Mean	34.12	33.2	0	0
Standard Deviation	3.1	3.528	0.13	0.14
Skewness	0.36	0.66	0.16	0.47
Kurtosis	4.85	5.36	3.72	4.44
Min	29.55	26.68	-0.44	-0.45
Max	44.77	50.17	0.39	0.55
75th percentile	36.08	35.98	-	-
99th percentile	44.78	44.20	-	-

This is the model we will use in the next chapter. Notice that the mean is slightly lower in our model than in the Day Ahead market, but otherwise the distribution is reasonably similar.

3.3.2 Half-hourly granularity - backtest against APX prices

Recall that one reason for needing to have price models was to price ‘shape’ contracts. The only market where it is possible to trade half-hourly periods is the APX market, so it is the only benchmark we have. It is less liquid than the Day Ahead market. We must use the APX reference price (half-hourly granularity) instead of the Day Ahead closing prices (daily granularity). Recall that the APX reference price is a weighted average of short term forwards, traded up to 7 days in advance. The raw descriptive statistics show that the half-hourly model results have different properties to the APX prices.

Sample Statistic	APX	Model	APX log-returns	Model log-returns
n	17804	17804	17803	17803
Mean	34.68	33.22	0	0
Standard Deviation	12.41	7.88	0.098	0.086
Skewness	2.72	6.75	0.758	0.34
Kurtosis	21.21	120.4	13.76	194.9
Min	5.27	3.82	-0.91	-2.06
Max	199.29	253.63	1.31	2.09
75th percentile	38.69	36.75	-	-
99th percentile	79.36	51.74	-	-
99.9th percentile	144.83	140.83	-	-

Bear in mind that the APX reference price is a weighted average of forwards. When a forward trade is agreed, neither the buyer nor the seller know the demand or the availability of the generators. Contrast this with the fundamental model which has demand and availability as a deterministic input.

Assumption 3.3.1. The APX reference price comprises 1 day of demand and outage uncertainty (outage uncertainty may also be called availability uncertainty).

Definition 3.3.1. Define $F_{t_i}^A$ as the APX reference price for period t_i and S_{t_i} as the spot. Assuming the rational market hypothesis:

$$F_{t_i}^A = \mathbb{E}(S_{t_i} | \mathcal{F}_{t_i-48}).$$

(Aside: 48 is the number of half-hour periods in one day.)

Consider the half-hourly time series of Demand and Availability values D_{t_i} and A_{t_i} respectively. Define

$$D_{t_i}^* = (D_{t_i} - D_{t_i-48})/D_{t_i-48}, A_{t_i}^* = (A_{t_i} - A_{t_i-48})/A_{t_i-48}.$$

These are the daily percentage changes.

We assume that they are normally distributed because we are restricted to using the Normal distribution in Plexos, with the Standard Deviation as a percentage of the demand value in the model for that half hour. Below are the Q-Q plots for the demand and availability daily percentage change.

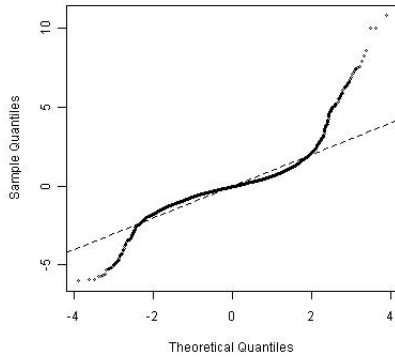


Figure 3.2: Q-Q plot of daily percentage change in Business Day Demand against the Normal distribution.

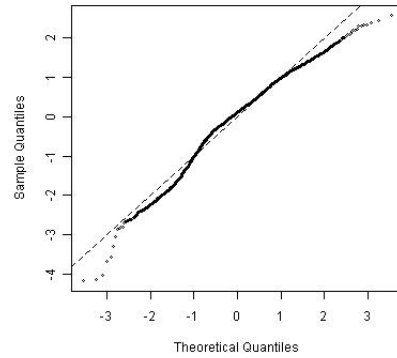


Figure 3.3: Q-Q plot of daily percentage change in Non-Business Day Demand against the Normal distribution.

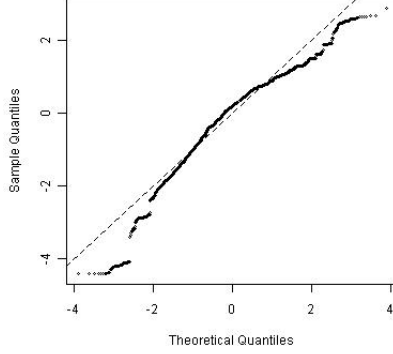


Figure 3.4: Q-Q plot of daily percentage change in Business Day Availability against the Normal distribution.

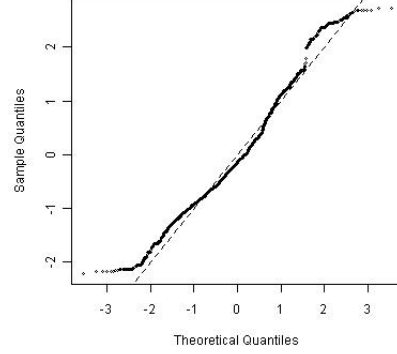


Figure 3.5: Q-Q plot of daily percentage change in Non-Business Day Availability against the Normal distribution.

Unfortunately because of the restrictions imposed on us by the Plexos software, we can only perform Monte Carlo simulations on the Demand and we have to shoehorn the availability simulations into the demand. Otherwise we would model the outages of individual generators. We omit further detail of this approximation. It is also important to note that if we wanted to model the uncertainty in availability in the fundamental model we would need to do so on the generator-level (not total availability). Denote our estimated percentage change standard deviation as σ_{D+A} . We then performed N Monte-Carlo simulations¹ for each half hour to compute the APX price simulation $F_{t_i}^A$ as follows:

$$F_{t_i}^A = \frac{1}{N} \sum_{j=1}^N S(t_i, D_{t_i}(1 + \sigma_{D+A}\epsilon_{j,t_i})).$$

The results below for the half-hourly intra-day shape are very close to the APX price. This appears to validate this approach for approximating the APX price. More work should be done to establish how valid our ‘1-day of uncertainty’ assumption is.

¹60 simulations of 365 days with 48 periods took 11.5 hours

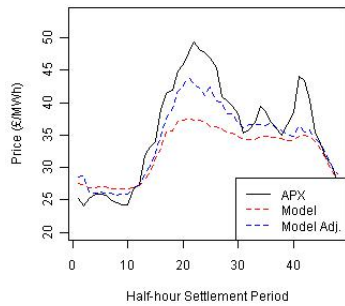


Figure 3.6: April 09 BD.

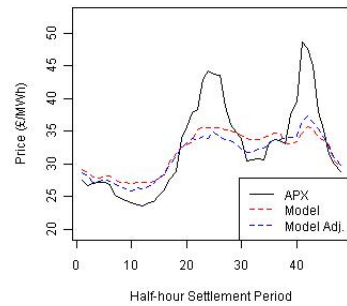


Figure 3.7: April 09 NBD.

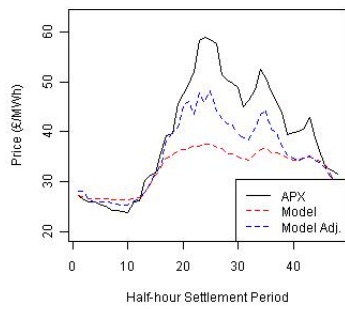


Figure 3.8: May 09 BD.

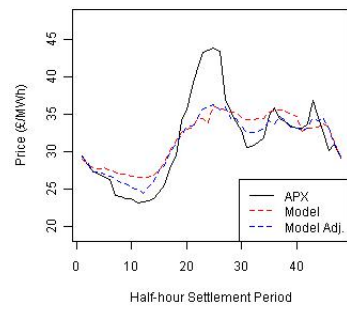


Figure 3.9: May 09 NBD.

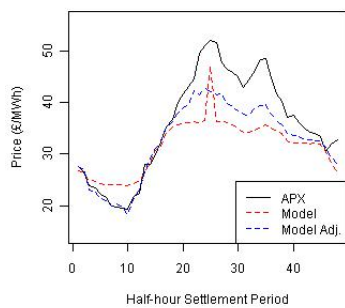


Figure 3.10: June 09 BD.

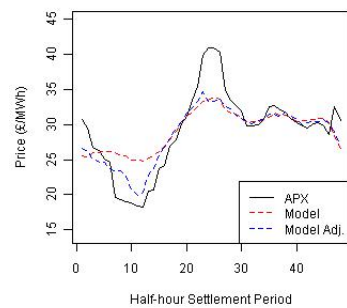


Figure 3.11: June 09 NBD.

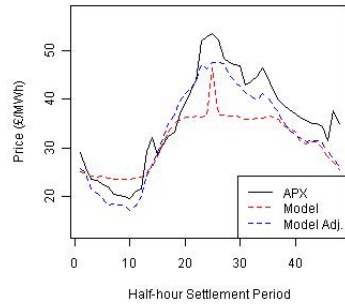


Figure 3.12: July 09 BD.

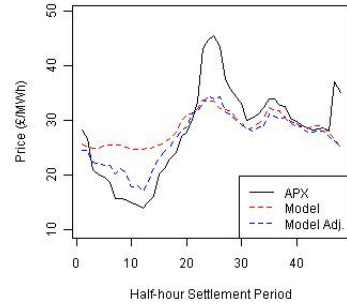


Figure 3.13: July 09 NBD.

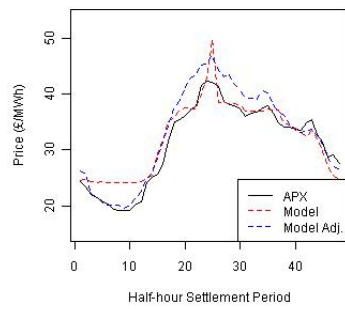


Figure 3.14: August 10 BD.

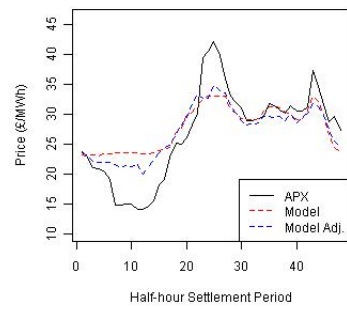


Figure 3.15: August 10 NBD.

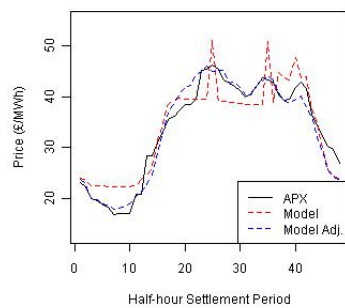


Figure 3.16: September 09 BD.

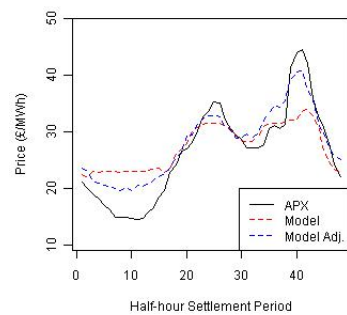


Figure 3.17: September 09 NBD.

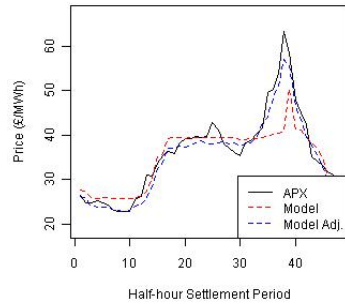


Figure 3.18: October 09 BD.

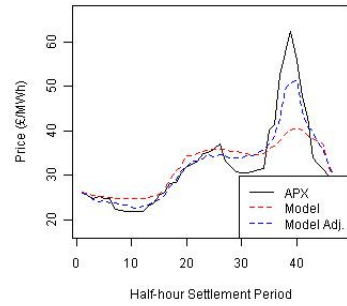


Figure 3.19: October 09 NBD.

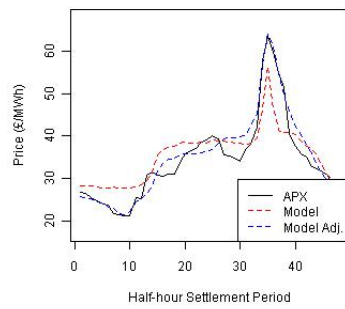


Figure 3.20: November 09 BD.

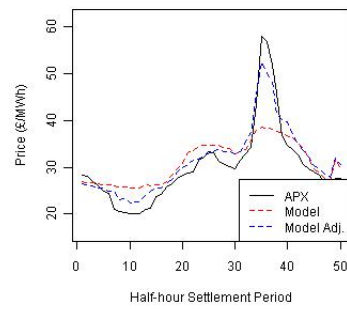


Figure 3.21: November 09 NBD.

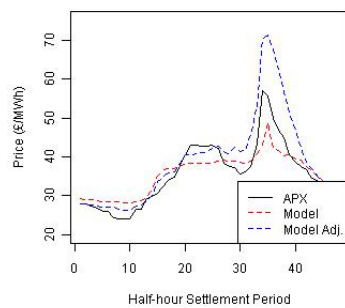


Figure 3.22: December 09 BD.

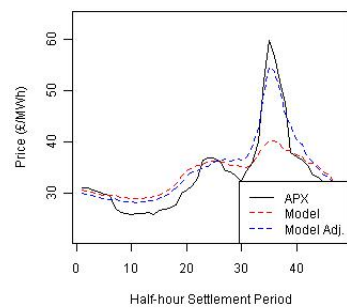


Figure 3.23: December 09 NBD.

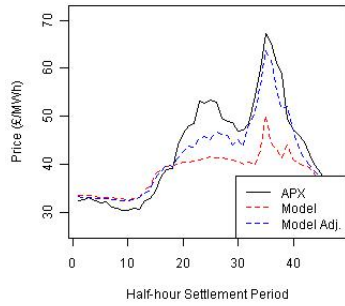


Figure 3.24: Jan 10 BD.

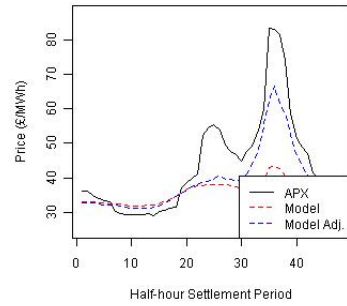


Figure 3.25: Jan 10 NBD.

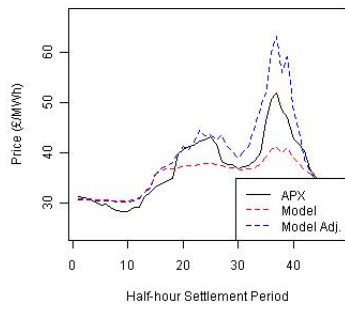


Figure 3.26: Feb 10 BD.

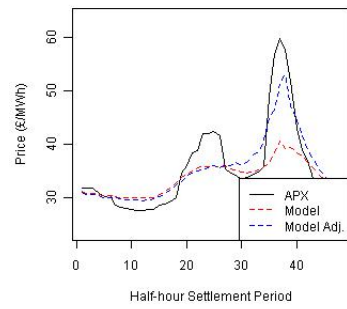


Figure 3.27: Feb 10 NBD.

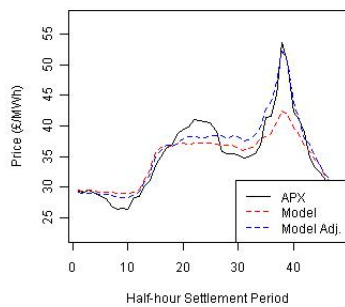


Figure 3.28: March 10 BD.

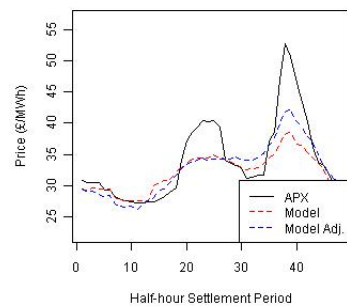


Figure 3.29: March 10 NBD.

The April, May, June and July prices appear to be rather low. One explanation for

this could be that this is when power stations aim to schedule their maintenance, and therefore there is more variability in availability (we assume it has constant variance). One improvement may be to model the seasonal variance in the availability. The descriptive statistics for the Adjusted Model with Simulations are:

Sample Statistic	APX	Model Adj.	APX log-returns	Model Adj. log-returns
n	17804	17804	17803	17803
Mean	34.68	34.34	0	0
Standard Deviation	12.41	9.83	0.098	0.08
Skewness	2.72	1.71	0.758	0.11
Kurtosis	21.21	10.06	13.76	11.43
Min	5.27	5.89	-0.91	-0.815
Max	199.29	119.5	1.31	0.9468
75th percentile	38.69	38.34	-	-
99th percentile	79.36	70.32	-	-
99.9th percentile	144.83	96.11	-	-

The results are much more in line with the APX price distribution, and we conclude that this is a good model for pricing shape contracts. To emphasise this point, and to highlight the differences between our original model and our adjusted model we compare the estimated kernel densities (using the R function density).

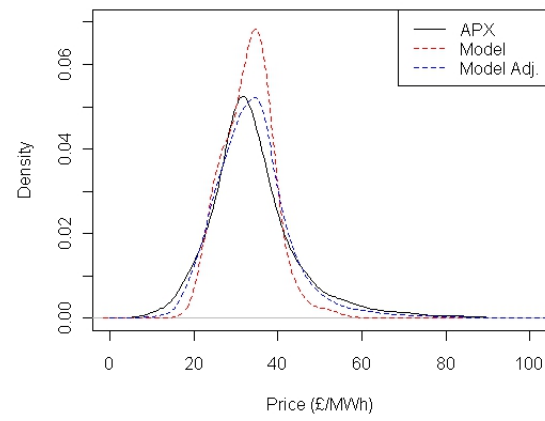


Figure 3.30: Comparison of estimated density functions.

Chapter 4

Hybrid Models

4.1 Concept

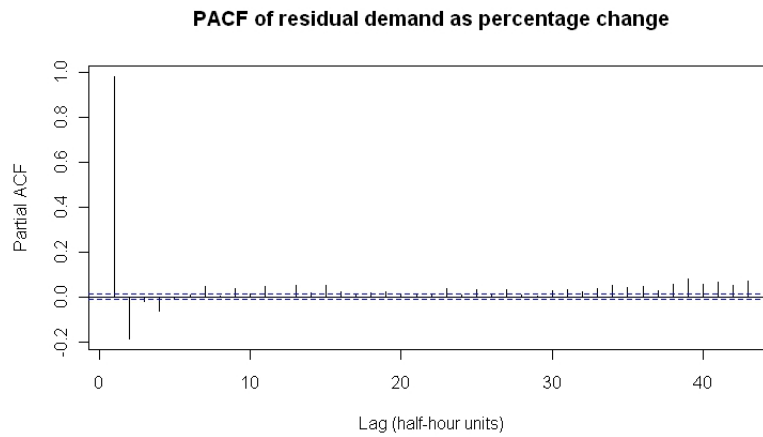
The hybrid model allows us to compute (spot) spread distributions. The concept is as follows. Given that we are interested in simulating prices for a certain time period which comprises discrete steps of time (half-hours).

- Generate N correlated fuel simulations. In our case, the fuels of interest are coal and gas.
- Generate N simulations of demand for each time interval t_i
- Generate N availability/outage schedules for each generator and each discrete time step
- Split up the time period into intervals that the fundamental model can handle. In our tests, daily steps were manageable.
- For each daily step, run the fundamental model N times using the fuel, demand and availability inputs.

This gives us a spot power price based on fundamentals - and a more accurate distribution for the spot spark/dark spreads.

4.1.1 Demand modelling

We developed a method for modelling deseasonalized half-hourly demand. Historical deseasonalized demand was provided by EDF Energy for calibration purposes. EDF Energy were also able to provide a deseasonalized demand forecast which we could use to generate the simulations. Again, we were working within the constraints of Plexos, which forced us to specify standard deviation as a percentage change in the demand. Denote actual historical demand as d_t and seasonal historical demand as d_t^S . We studied the process $d_t^{Res} = \frac{d_t - d_t^S}{d_t^S}$ for autocorrelation between its lags. Below is the partial autocorrelation function for the process d_t^{Res} .



We approximate the percentage change in demand as an AR(2) process with the following parameters:

AR1	AR2	Standard Deviation
1.167	-0.187	0.006

4.1.2 Availability modelling

There are two separate parts to modelling availability of a generator.

- The distribution of its survival time - how long it lasts until an outage occurs
- The distribution of repair times for a generator

- Analysis of outages for nuclear, gas and coal generators was conducted. The analysis is outside the scope of this thesis as it is designed such that it is easily applied in Plexos. If this was not the case, a preferred approach would be to model survival time as a exponentially distributed Markov two-state model, then to draw repair times from an empirical cumulative distribution function based on historical data. Unfortunately, we did not have the flexibility to proceed in this way.

4.2 Results

The model above was implemented with 60 simulations. This took approximately 20 hours to run. The hybrid results revealed a bug in the 3rd party software which is ongoing at the time of writing, so we omit them. However, it is clear from the earlier section ‘Concepts’ that the validity of the hybrid relies only its inputs. Importantly we have developed and implemented a methodology for generating spark and dark spread distributions.

4.3 Criticisms and Conclusion

One of the main criticisms of the stochastic models implemented for power, coal and gas is that jumps are not included. The assumption of gaussian distributed shocks is perhaps too simplistic and jump-diffusion mean-reverting models could be studied as a part of further research. However, it seems that the gaussian assumption is a reasonable estimate at the daily granularity (although certainly not so at the half-hourly level). The stochastic model we developed also has a forward model in-built which we can use to simulate various hedging strategies.

Fundamental models appear to solve the problem of non-gaussianity also - generating power prices with the required jumps due to the non-linear nature of the UK availability stack. The results from the fundamental model were encouraging. We developed a new method of approximating half-hourly prices for the purposes of pricing shape contracts. However, these models have assumptions of their own; most notably the perfect competition assumption. We were also faced with not

being able to solve the mixed integer problem in an acceptable amount of time, forcing us to use a sub-optimal proprietary heuristic. Unfortunately this also restricted some of our other modelling decisions, such as the availability model. With more time, we would have developed our own heuristics for solving this large-scale unit commitment problem.

The hybrid model is the blending of both approaches and it is a good modelling approach for obtaining spreads. The results will provide price signals which will enable companies to improve the timing their hedges or to speculate in the market. The main criticism of the hybrid model is that even with a sub-optimal heuristic for its fundamental component, it takes a long time to run (20 hours for 60 simulations for a time period of 1 year). However, the speed of computer processors increasing and their price decreasing by the day. Another limitation of the current software is that it can only utilise a single computer processor, for all simulations. The Monte Carlo simulations used in the hybrid would appear to be natural candidates for a parallel computing approach. This is because each simulation is independent of all others, which allows each simulation to be run on a separate processor. This implies that even today, the hybrid run time is only limited to approximately the time of the longest simulation, say 20 minutes, given enough computer processors and a more optimal software architecture.

Heath-Jarrow-Morton 2-factor model for multiple commodities

testMultiFuelTwoFactor.R

```
1 source("multitwofactor.R")
2
3 f0<-list()
4 f0[[1]]<-read.csv("gas_curve.csv", header=TRUE)$curve
5 f0[[2]]<-read.csv("coal_curve.csv", header=TRUE)$curve
6 corr<-read.csv("correlation_matrix.csv", header=TRUE)
7
8 vol_s=c(0.94, 0)
9 vol_l=c(0.38, 0.30)
10 a=c(14.5, 1)
11
12 corr<-as.matrix(corr, nrow=4, ncol=4)
13 theta<-1/365
14
15 S<-multiTwoFactorSpotSims(f0, a, vol_s, vol_l, corr, theta)
```

multitwofactor.R

```

1 multiTwoFactorSpotSims<-function(f0, a, vol_s, vol_l, corr, theta){
2   # f0           : a list of forward curves
3   # a           : a vector of mean-reversion parameters
4   # vol_s/vol_l : a vector of short term/long term volatilities
5   # corr        : long and short term correlation matrix of M markets. Size is 2M *
6                   2M
7   #             column/row order: market 1 st, market 1 lt, market 2 st, market
8                   2 lt, etc
9   # theta       : granularity of the product in years
10
11  #number of markets
12  M<-length(vol_s)
13  #cholesky decomposition of the correlation matrix
14  cholesky<-chol(corr)
15
16  days<-length(f0[[1]])
17  t<-(0:(days-1))*theta
18  c<-rep(0, M)
19
20  #number of simulations should be a multiple of 2 for anti-thetic sampling
21  N<-200
22  eps<-0.0001
23
24  V_l<-list(); V_s<-list(); V_sl<-list()
25  for (i in 1:M){
26    #adjustment for non-infinitesimal product delivery period
27    if (a[i] > 0) {
28      c[i]<-(1/(a[i]*theta))*(1-exp(-a[i]*theta))
29    } else {
30      c[i]=1
31    }
32    #deterministic terms
33    V_l[[i]]<-vol_l[i]^2*t
34    V_s[[i]]<-(c[i]^2*vol_s[i]^2/(2*a[i]))*(1-exp(-2*a[i]*t))
35    V_sl[[i]]<-2*c[i]*corr[(i-1)*2+1, (i-1)*2+2]*vol_s[i]*vol_l[i]*(1-exp(-a[i]*t
36    ))/a[i]
37  }
38
39  #create spot price matrix
40  S<-list()
41  for (i in 1:M){
42    S[[i]]<-matrix(nrow=days, ncol=N)
43    S[[i]][1,]<-f0[[i]][1]
44  }
45
46  #the stochastic terms are calculated stepwise, for each simulation
47  W_s<-list(); W_l<-list()

```

```

45 for (i in 1:N) {
46   for (k in 1:M){
47     W_s[[k]]<-rep(0, days)
48     W_l[[k]]<-rep(0, days)
49   }
50   #antithetic sampling
51   if (abs(i/ 2 - floor(i/2))> eps) {
52     #draw correlated numbers
53     iid_norm<-matrix(rnorm(days*2*M), nrow=2*M, ncol=days)
54     corr_norm<-crossprod(cholesky, iid_norm)
55   } else {
56     corr_norm<-corr_norm*(-1)
57   }
58   for (k in 1:M){
59     for (j in 2:days) {
60       W_s[[k]][j]<-W_s[[k]][j-1]*exp(-a[k]*theta)+sqrt((c[k]^2*vol_s[k]^2/(2*a
61         [k]))*(1-exp(-2*a[k]*theta)))*corr_norm[(k-1)*2 + 1, j]
62       W_l[[k]][j]<-W_l[[k]][j-1]+sqrt(vol_l[k]^2*theta)*corr_norm[(k-1)*2 + 2,
63         j]
64       S[[k]][j, i]<-f0[[k]][j]*exp(-0.5*(V_l[[k]][j] + V_s[[k]][j] + V_sl[[k]
65         ][j]) + W_s[[k]][j] + W_l[[k]][j])
66     }
67   }
68 }
return (S)
}

```

Fundamental Unit Commitment model

unitcommitment.R

```
1  # Unit commitment problem using symphony or glpk open source solvers in R
2  # Symphony performs better, so it is the default
3  unitcommitment_solve<-function(demand, generators, solver="SYMPHONY")
4  {
5    if (solver=="SYMPHONY"){
6      library(Rsymphony)
7      cat("Using Symphony solver\n")
8      call_solver<-function(obj, A, constraint_type, rhs, types, max=FALSE){
9        Rsymphony_solve_LP(obj,A, constraint_type, rhs, types=types,max=max)
10     }
11   }
12   else {
13     library(Rglpk)
14     cat("Using GLPK solver\n")
15     call_solver<-function(obj, A, constraint_type, rhs, types, max=FALSE){
16       Rglpk_solve_LP(obj, A, constraint_type, rhs, types=types,max=max)
17     }
18   }
19
20   #number of variables in our problem
21   #x_it - binary unit commitment variable (0=OFF, 1=ON)
22   #y_it - generation between MSL and MEL as a proportion (0 <= y_it <= x_it)
23   #w_it - binary start up variable (0=NO START UP, 1=START UP)
24   # although w_it are binary variables we can write them as a continuous variable
   with a constraint in terms of x
25   # w_it >= x_it - x_{i,t-1}  TODO: solvers have the == option. Does this improve
   performance?
26   vars<-3
27
28   #number of time steps. Normally this will be 48
29   T=length(demand)
```

```

30
31 #time interval length in hours
32 h<-0.5
33
34 #read in demand and generation data
35 names<-generators$names
36 srmc<-generators$srmc
37 startcost<-generators$startcost
38 msl<-generators$msl
39 mel<-generators$mel
40
41 #number of generators
42 I<-length(names)
43 cat("Setting up unit commitment MIP with", formatC(vars*T*I), "variables\n")
44 cat("Building objective function...\n")
45 #build the objective function
46 obj<-rep(0, times=(vars*T*I))
47
48 for (t in 1:T){
49   for (i in 1:I){
50     # objective function coefficient for x_it
51     obj[vars*(t-1)*I + vars*(i-1) + 1]=h*srmc[i]*msl[i]
52     # coefficient for y_it
53     obj[vars*(t-1)*I + vars*(i-1) + 2]=h*(mel[i]-msl[i])*srmc[i]
54     # coefficient for w_it
55     obj[vars*(t-1)*I + vars*(i-1) + 3] =startcost[i]
56   }
57 }
58
59 #build the matrix A
60 #A has dimensions (number of constraints) x (vars * I * T)
61 #we have T demand constraints
62 #TODO: add T spinning reserve constraints
63 #also there are T-1 * I constraints due to converting w from int to continuous
64 constraints<-(T+(T-1)*I)+(I*T)
65 cat("There are", formatC(constraints), "constraints\n")
66 cat("Building solver matrix...\n")
67 a<-rep(0, times=vars*T*I*constraints)
68
69 #build matrix - demand constraint rows
70 for (t in 1:T){
71   for (i in 1:I){
72     # x_it
73     a[(t-1)*T*I*vars + (t-1)*vars*I + (i - 1) * vars + 1]<-msl[i]
74     # y_it
75     a[(t-1)*T*I*vars + (t-1)*vars*I + (i - 1) * vars + 2]<-(mel[i]-msl[i])
76   }
77 }
78 # unit commitment constraint rows
79 row=T*T*vars*I-T*vars*I+(I-1)*vars

```

```

80 for (t in 1:(T-1)){
81   for (i in 1:I){
82     #  $w_{it} - x_{it} + x_{i,t-1} \geq 0$ 
83     row=row + T*vars*I + vars
84     a[row+1] = -1
85     a[row+3]= 1
86     a[row-vars*I+1]=1
87   }
88 }
89
90 row=T*T*vars*I+(T-1)*I*(T*vars*I)-vars - T*vars*I
91 for (t in 1:T){
92   for (i in 1:I){
93     {
94       row=row + T*vars*I + vars
95       #  $x_{it} - y_{it} \geq 0$ 
96       a[row+1]=1
97       a[row+2]=-1
98     }
99   }
100
101 A=t(matrix(a, nrow=T*I*vars))
102 cat("Solver matrix has", formatC(length(A[,1])), "rows and", formatC(length(A
103   [,])), "columns\n")
104
105 #build RHS
106 rhs<-rep(0, times=constraints)
107 for (i in 1:T){
108   rhs[i]<-demand[i]
109 }
110
111 # Define variable types: x_it Binary; y_it, w_it Continuous
112 types<-rep(c("B", "C", "C"), times=constraints)
113 constraint_type<-rep(">=", times=constraints)
114
115 cat("Solving...\n")
116 out=call_solver(obj,A, constraint_type, rhs, types=types,max=FALSE)
117 }

```

main.R

```
1 source("unitcommitment.R")
2
3 main<-function(solver="SYMPHONY")
4 {
5
6   #read generators file
7   cat("Reading generators.csv...\n")
8   generators<-read.table("generators.csv",header=TRUE,sep=",")
9
10  #read demand file
11  cat("Reading demand.csv...\n")
12  demandlist<-read.table("demand.csv", header=TRUE,sep=",")
13  demand<-demandlist$demand
14
15  solverTime<-system.time(out<-unitcommitment_solve(demand,generators, solver))
16  cat("A solution was found. Time elapsed:", formatC(solverTime[[3]]), "seconds\n")
17
18  #analyze the results
19  vars<-3
20  I<-length(generators$names)
21  T<-length(demand)
22  results<-out$solution
23  prices<-rep(0,times=T)
24
25  schedule<-t(matrix(window(results,deltat=vars), nrow=I))
26  for (t in 1:T){
27    maxsrcm=0
28    for (i in 1:I){
29      if (schedule[t,i]==1){
30        if (generators$srcm[i] > maxsrcm){
31          maxsrcm=generators$srcm[i]
32        }
33      }
34    }
35    prices[t]=maxsrcm
36  }
37
38  schedFrame<-data.frame(schedule)
39  names(schedFrame)<-generators$names
40  out<-data.frame(schedFrame, Demand=demand, Price=prices)
41  print(out)
42
43 }
```

Example Output

```
> main()
Reading generators.csv...
Reading demand.csv...
Using Symphony solver
Setting up unit commitment MIP with 1296 variables
Building objective function...
There are 903 constraints
Building solver matrix...
Solver matrix has 903 rows and 1296 columns
Solving...
A solution was found. Time elapsed: 1 seconds
```

	Coal	Gas	Nuclear	Oil	Pumped	CHP	OCGT	Hydro	Wind	Demand	Price
1	1	1	1	0	1	0	0	1	1	1700	60
2	0	1	1	0	0	0	0	1	1	1360	25
3	0	1	1	0	0	0	0	1	1	1100	25
4	0	0	1	0	0	0	0	0	1	500	5
5	0	0	1	0	0	0	0	0	1	400	5
6	0	0	1	0	0	0	1	1	1	800	160
7	0	1	1	0	0	0	0	1	1	900	25
8	0	1	1	0	0	0	0	1	1	1300	25
9	1	1	1	0	0	0	0	1	1	1400	35
10	1	1	1	0	0	0	0	1	1	1500	35
11	1	1	1	0	0	0	0	1	1	1600	35
12	1	1	1	0	1	1	0	1	1	2330	60
13	1	1	1	0	1	1	0	1	1	2310	60
14	1	1	1	0	1	1	0	1	1	2499	60
15	1	1	1	0	0	1	0	1	1	2100	50
16	1	1	1	0	1	1	0	1	1	2400	60
17	1	1	1	0	1	1	0	1	1	2600	60
18	1	1	1	0	1	1	0	1	1	2700	60
19	1	1	1	1	1	1	0	1	1	2850	110
20	1	1	1	0	1	1	1	1	1	2833	160
21	1	1	1	1	1	1	0	1	1	2955	110
22	1	1	1	0	1	1	0	1	1	2700	60
23	1	1	1	0	1	1	0	1	1	2655	60
24	1	1	1	0	1	1	0	1	1	2555	60
25	1	1	1	0	1	1	0	1	1	2601	60
26	1	1	1	0	1	1	0	1	1	2615	60
27	1	1	1	0	1	1	1	1	1	2812	160
28	1	1	1	1	1	1	0	1	1	2850	110
29	1	1	1	1	1	1	0	1	1	2897	110
30	1	1	1	1	1	1	0	1	1	2989	110
31	1	1	1	1	1	1	0	1	1	3120	110
32	1	1	1	1	1	1	0	1	1	3111	110
33	1	1	1	1	1	1	0	1	1	3329	110
34	1	1	1	1	1	1	0	1	1	3400	110
35	1	1	1	1	1	1	0	1	1	3460	110
36	1	1	1	1	1	1	0	1	1	3500	110
37	1	1	1	1	1	1	0	1	1	3200	110
38	1	1	1	1	1	1	0	1	1	3100	110
39	1	1	1	1	1	1	0	1	1	3000	110
40	1	1	1	1	1	1	0	1	1	2900	110
41	1	1	1	0	1	1	1	1	1	2800	160
42	1	1	1	0	1	1	0	1	1	2700	60
43	1	1	1	0	1	1	0	1	1	2600	60
44	1	1	1	0	1	1	0	1	1	2500	60
45	1	1	1	0	1	1	0	1	1	2400	60
46	1	1	1	0	1	1	0	1	1	2300	60
47	1	1	1	0	1	1	0	1	1	2200	60
48	1	1	1	0	0	1	0	1	1	2100	50

Bibliography

- [1] Glpk solver website. <http://www.gnu.org/software/glpk>.
- [2] Symphony solver website. <http://www.coin-or.org/SYMPHONY/index.htm>.
- [3] EEX Product Information Coal. 2006.
- [4] Benth, Benth, and Koekebakker. *Stochastic Modelling of Electricity and Related Markets*. 2009.
- [5] D. Brigo and F. Mercurio. *Interest Rate Models. Theory and Practice*. 2001.
- [6] D. Brigo and F. Mercurio. *Interest Rate Models. Theory and Practice: with Smile, Inflation and Credit*. 2006.
- [7] Rainer E. Burkard, Michael Kocher, and Rüdiger Rudolf. *Rounding Strategies for Mixed Integer Programs Arising from Chemical Production Planning*. 1997.
- [8] L. Clewlow and C. Strickland. *Energy derivatives pricing and risk management*. 2000.
- [9] Shijie Deng. *Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes*. 2000.
- [10] Shijie Deng, Blake Johnson, and Aram Sogomonian. *Exotic electricity options and the valuation of electricity generation and transmission assets*. 2001.
- [11] Julia Frayer and Nazli Z. Uludere. *What Is It Worth? Application of Real Option Theory to the Valuation of Generation Assets*.
- [12] Benjamin F. Hobbs, Carolyn B. Metzler, and Jong-Shi Pang. *Strategic Gaming Analysis for Electric Power Systems: An MPEC Approach*. 2000.

- [13] John C. Hull. *Options, Futures, and Other Derivatives*. 2000.
- [14] Blake Johnson and Graydon Barz. Selecting stochastic processes for modelling electricity prices. 1999.
- [15] Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus*. 1991.
- [16] F. A. Longstaff and A. W. Wang. Electricity Forward Prices: A High-Frequency Empirical Analysis. 2004.
- [17] Julio J. Lucia and Eduardo S. Schwartz. Electricity prices and power derivatives: Evidence from the nordic power exchange. 2001.
- [18] Peter Spreij. *Portfolio Theory Lecture Notes*. 2009.
- [19] Peter Spreij. *Stochastic Integration Lecture Notes*. 2009.