


Pricing Swing Options and other Electricity Derivatives



Tino Kluge
St Hugh's College
University of Oxford

Doctor of Philosophy

Hillary 2006

This thesis is dedicated to
my mum and dad
for their love and support.

Acknowledgements

I would like to thank my supervisors Sam Howison and Ben Hambly for their continuous support as well as patience in times when progress was rather slow. A big thank you to Vicky Henderson and David Hobson for many invaluable remarks and also for giving me the great opportunity to visit Princeton and to work with them.

Many thanks also go to my friends Dave Buttle, Antonis Papapantoleon and Dr P Hewlett for very helpful discussions and to Dr Kofi Appiah for his most inspiring random comments. I am deeply indebted to my former boss Steffi Kammer who always leads me by example.

I am grateful for the financial support I have received from the German Academic Exchange Service (DAAD), the British Engineering and Physical Sciences Research Council (EPSRC) and the software company KWI which is now part of Global Energy.

Many thanks to all my other friends in Oxford for making this part of my life such a great time.

Abstract

The deregulation of regional electricity markets has led to more competitive prices but also higher uncertainty in the future electricity price development. Most markets exhibit high volatilities and occasional distinctive price spikes, which results in demand for derivative products which protect the holder against high prices.

A good understanding of the stochastic price dynamics is required for the purposes of risk management and pricing derivatives. In this thesis we examine a simple spot price model which is the exponential of the sum of an Ornstein-Uhlenbeck and an independent pure jump process. We derive the moment generating function as well as various approximations to the probability density function of the logarithm of this spot price process at maturity T . With some restrictions on the set of possible martingale measures we show that the risk neutral dynamics remains within the class of considered models and hence we are able to calibrate the model to the observed forward curve and present semi-analytic formulas for premia of path-independent options as well as approximations to call and put options on forward contracts with and without a delivery period. In order to price path-dependent options with multiple exercise rights like swing contracts a grid method is utilised which in turn uses approximations to the conditional density of the spot process.

Further contributions of this thesis include a short discussion of interpolation methods to generate a continuous forward curve based on the forward contracts with delivery periods observed in the market, and an investigation into optimal martingale measures in incomplete markets. In particular we present known results of q -optimal martingale measures in the setting of a stochastic volatility model and give a first indication of how to determine the q -optimal measure for $q = 0$ in an exponential Ornstein-Uhlenbeck model consistent with a given forward curve.

Contents

1	Introduction	1
2	Electricity markets	3
2.1	The spot market	5
2.1.1	Spot price determination	5
2.1.2	Spot market data	7
2.2	The forward and future market	10
2.2.1	Forwards with and without a delivery period	11
2.2.2	Building a continuous forward curve	12
2.2.3	Call and puts	22
3	Stochastic spot price model	24
3.1	Existing models	24
3.2	A mean-reverting model exhibiting seasonality and spikes	27
3.3	Parameter estimation based on historical data	29
3.4	Properties of the stochastic process	33
3.4.1	The spike process	36
3.4.2	Approximations of the spike process	39
3.4.3	The combined process	44
3.4.4	Approximations of the combined process	46
3.4.5	Conditional expectations	50

4	Option pricing	58
4.1	Utility indifference pricing	58
4.1.1	Pricing without the possibility of hedging	59
4.1.2	Pricing and hedging with a correlated asset	61
4.2	Arbitrage pricing and risk neutral formulation	62
4.3	Pricing path-independent options	65
4.3.1	Pricing call options	66
4.3.2	Pricing options with arbitrary payoff	68
4.3.3	Pricing options on Forwards	69
4.3.4	Pricing options on Forwards with a delivery period	71
4.4	Pricing swing options	76
4.4.1	The grid approach	77
4.4.2	Comparison of algorithms	79
4.4.3	Numerical results	81
4.4.4	Dimension reduction	84
4.4.5	General swing contracts	86
4.5	PIDE formulation	87
4.6	Hedging contingent claims	88
5	Equivalent martingale measures	90
5.1	Introduction	90
5.2	Pricing measures in a stochastic volatility models	92
5.3	Pricing measures in a spot electricity model	101
6	Outlook and extensions	104
A	Elementary probability	106
A.1	Products and sums	106
A.2	Conditional expectations	108

B	The Ornstein-Uhlenbeck process	112
B.1	Solution of the sde	112
B.2	Parameter estimation	114
C	Transform analysis	116
C.1	Moment generating function	116
C.2	Expectation of a call-option payoff	118
D	Option pricing and the inability to hedge with the underlying	120
D.1	Risk neutral valuation	120
D.2	Valuation by a pde approach	121
D.3	Solution for a mean-reverting process	123
D.3.1	Non-uniform grids	124
	Bibliography	125

Chapter 1

Introduction

Historically, electricity prices were generally determined by regulatory authorities controlled by the government of each individual country. Prices were intended to reflect the (marginal) cost of production and did not change very often and even then quite predictably. In the early 1990s, a few countries started to liberalise their electricity markets by leaving the price determination to the market principles of supply and demand. Many countries have since reformed their power sector. One important consequence is the trade of electricity delivery contracts on exchanges, similar to the trade of shares on a stock exchange, for example. The new freedom achieved has brought the drawback of increased uncertainty about the price development and indeed, many markets exhibit very high rates of volatility. Although households do not buy electricity directly from an exchange, many companies with high power consumption do. This creates demand for contracts which protect them against high prices but give the optionality to profit from low prices. Such contracts are called options or derivatives.

For pricing derivative contracts and managing risk, there is now a very comprehensive theory for financial markets which can be utilised. However there are distinct differences between financial and electricity markets which require further investigations. Although the general arbitrage pricing theory can be applied, it is vital to utilise an appropriate stochastic model for the underlying price dynamics. In the literature, two main approaches are considered: the modelling of the spot price dynamics and the entire forward curve, respectively. Forward curve models are very well suited for pricing options on forwards but, as they normally imply a very complex non-Markovian dynamics for the spot price, it is hard to value path dependent options. As one of our main aims is to be able to price swing options – a complex path dependent option giving the holder the opportunity to exercise a certain right repeatedly over a period of time – we exclusively focus on spot price models.

In this thesis we propose and examine in detail a simple mean-reverting process exhibiting price spikes. A distinct feature of electricity markets is the formation of price spikes and are caused by events where the maximum supply is approached by current demand. The occurrence of spikes has far reaching consequences for risk management and pricing purposes which is why we believe it is vital to model this feature appropriately. We do not claim that our proposed model perfectly fits the market but rather recommend it because it reflects some main properties and is analytically tractable.

This thesis is organised as follows: An introduction into electricity markets is given in Chapter 2 which contains a technical description of the NordPool electricity exchange. In addition we propose an interpolation algorithm that will enable us to create a continuous forward curve based on the few forward contracts observed in the market.

We propose a stochastic spot price model in Chapter 3 and examine its properties in detail. In particular the moment generating function of the spot price is given and approximations to the density function are derived which are later used as the basis for a numerical algorithm to price swing options in Chapter 4.

We begin Chapter 4 with a short introduction into utility indifference pricing, we then focus on arbitrage pricing and derive the risk neutral dynamics of the model under a slightly restricted set of equivalent martingale measures and show there is a subset consistent with observed forward prices, i.e. the model can be calibrated to any smooth forward curve. After stating the well known result of pricing path-independent options based the moment generating function we derive approximations to prices of options on forwards with and without a delivery period. This is followed by a section on pricing swing options.

Due to the incompleteness of the electricity market we devote Chapter 5 to an introduction to the choice of optimal martingale measures. We make a short excursion into the setting of a stochastic volatility model in equity markets as theory there is developed further. Our first attempts on finding q -optimal measures in the setting of electricity markets in a special case concludes the chapter.

Chapter 6 proposes model extensions and concludes.

Chapter 2

Electricity markets

The aim of this chapter is to introduce peculiarities of electricity markets. After a discussion of the basic differences between electricity as a commodity and share markets, a detailed description of the electricity spot and derivative markets is given in the subsequent sections. Particular attention is given to the question of how to generate a continuous forward curve based on the few forward contracts observed in the market.

Due to the profound differences between electricity and other financial markets like share markets, classical financial theories cannot be directly applied in electricity markets but modifications and adaptations have to be made. Nevertheless, the absence of arbitrage remains the fundamental principle on which we base the pricing of derivatives. The differences and similarities of the two markets are described below.

Underlying unit: Where in share markets the underlying unit is simply one specific share of a company, in the electricity market it is a specific unit of energy (usually 1 MWh). In an abstract sense and leaving aside the technological restrictions, one could imagine the energy units to be stored in very small objects which lie in a big storehouse. Buying these units as a financial commodity would only involve an electronic money transaction and an assignment of the bought energy units into the buyers portfolio without actual physical delivery. So far, everything sounds identical to share markets.

Production and consumption: In share markets the number of shares basically remain the same over time (unless the company issues new shares) and give the owner codetermination rights. Electrical energy can be produced and consumed and even with the hypothetical ability to store, that has a profound effect on the price per unit. Based on microeconomic considerations, one would expect in the long term the price to revert to the production cost. This is the reason why in commodity markets mean-reverting models are mainly used.

Inability to store: In reality, current technology does not allow electrical energy to be stored efficiently. It is virtually impossible to store the amount of electrical energy a big factory consumes on a single day, let alone the energy of an entire country. Electrical energy is therefore considered to be non-storable as far as the power market is concerned. This has far-reaching consequences.

- Electricity can be described as a pure flow variable (energy per time, measured in MW) and it requires time to transfer a certain amount of energy. In particular, derivative contracts will always specify a delivery period. Also, limitations in the transmission grid can cause congestion.
- Production and consumption have to be in balance all the time. A small imbalance can be absorbed in voltage changes and, for supply excess, dissipation in the grid and generating plants. The supply dropping below the demand could result in a black-out. This real-time balance of demand and supply introduces seasonality of the underlying price as the demand changes over the day, week and year. In addition, inelasticity of demand and supply¹ make electricity prices very sensitive to extreme events like power plant outages. In such an event, the maximum supply could drop to levels near the current demand causing the price to rise considerably. After a short time, however, the power outage could be resolved or spare power stations be activated, normalising the situation and bringing the price down to previous levels. Such price events are called spikes.
- Hedging derivative contracts with the underlying requires the ability to store and therefore cannot be done for electricity derivatives. Hence this market is automatically incomplete, independent of the stochastic process used to model the underlying. In simple terms, the risk neutral probability measure $\mathbb{Q}$ is not unique but can be determined based on market observations of derivative products like forwards.

Having introduced the main theoretical properties of electricity markets, a more detailed and technically oriented description follows.

¹End-users usually receive electricity for a fixed price and would not reduce their consumption if power prices went up on the exchange and power stations need a certain warm-up time until they are ready to produce electricity.

2.1 The spot market

Due to technical limitations in electricity transmission,² markets are localised to specific regions, like individual countries. Each market has their own rules. The following descriptions are based on the specifications of the NordPool market (The Nordic Power Exchange), owing to its long history. Founded in Norway in 1991, NordPool was the world's first international power exchange. Later on, the countries Sweden (1996), Finland (1998) and Denmark (2000) joined this market, resulting in a total generation of almost 400 TWh per year³ serving a total population of about 24 million.

The liberalisation of a power market does not necessarily require the establishment of a power exchange; however, it makes market information more transparent, and improves competition and liquidity. The NordPool spot market (Elspot) operates in direct competition with non-exchange-market trading and had a market share of about 32% in 2002.

In addition to the power generation, the electricity needs to be transferred to its destination through a transmission grid which is operated by transmission system operating companies (TSO). This part of the market is monopolistic and tariffs are set by regulators. Prices should reflect the maintenance cost and the energy loss, as it is the responsibility of the TSOs to buy the energy lost through transmission from the spot market. It is therefore guaranteed that the seller submits and the buyer receives the exact amount of energy as specified in the spot market contracts. As a consequence, the total power procurement cost consists of the spot price, trading fees, transmission charges and eventual imbalance costs based on the real-time market, as will be described below. Despite this complex structure, theory only takes spot prices into account as derivative products are solely based on them.

2.1.1 Spot price determination

Electricity prices per MWh are determined using a bidding system. Everybody with access to the transmission grid and who meets the requirements⁴ set by NordPool can submit bids, which are essentially functions saying how much energy would be bought or sold depending on the price. In simple terms,⁵ the price is then given by the intersection of the aggregate demand and supply curve. Based on this price it is clear how many units of energy each participant sells or buys. To allow for the generators to prepare for delivery,

²Energy gets wasted over long distances due to the inner resistance of the wires.

³About 387 TWh was generated in 2001 which is on average about 44000 MW or 1.84 kW per capita. The trade through the spot market was about 115 TWh and the rest were over the counter trades.

⁴E.g. a security amount in a pledged bank account is needed. The amount depends on the trading activity but is at least NOK 100000.

⁵In reality the price will also be modified if congestion in the grid system is anticipated.




Figure 2.1: Marginal cost merit order chart based on [Noller, 2002].

prices are determined on a day-ahead basis for each individual hour. The average price over the entire day is called the base load price, the average over the most demanding hours (depending on the regional market and day of week) is called peak price and the average over the remaining hours is called off-peak price. Further adjustments to the demand or supply of each participant can be made in the balancing market called Elbas (one hour-ahead) and in the real-time market where prices are set in a way to penalise reduction of supply or increase of demand and to discourage speculation in these markets. As long as the adjustments are within a tolerance level, producers can immediately meet a change of demand. However, stronger adjustments could cause the TSOs to switch off certain consumers to be able to meet demand. The likelihood of such events is supposed to be extremely small. However, events in the past have shown that those blackouts can occur. Finally, the current consumption level is metered and differences to the contractual volume are priced at the real-time market.

Bidding strategies of each participant could be quite complex but one would expect the producers not to bid below their marginal cost of production. Based on this idea and additional assumption on the behaviour of consumers one could use a supply-demand driven model to describe the spot price process. A simplified model is given in [Barlow, 2002] where the demand is assumed to be a stochastic process and independent of the price and the supply an increasing function of the price. The price is then given by that value which matches demand and supply. Figure 2.1 shows the approximate marginal cost of production of the power stations in Germany and could be used as a basis of a more realistic supply-demand model.




Figure 2.2: Daily average prices at NordPool over a ten-year period. The first two small arrows indicate the time region shown in Figure 2.3 and the second two indicate that of Figure 2.4.

2.1.2 Spot market data

Figure 2.2 shows a ten-year history of the NordPool spot market. Unlike in stock markets, prices appear to revert to a mean level and do not seem to behave like exponential Brownian motion. In addition, a pattern of seasonality is clearly visible. Prices generally tend to be higher in winter than in summer which is certainly caused by a higher demand in winter due to the cold climate. An exception is the year 1996 where the price did not go down during summer. In the Scandinavian countries, more than half of the energy generation comes from hydro power plants. In order to satisfy the increased demand during winter months, water from hydro reservoirs is used to generate more electricity. This makes the market sensitive to the rainfall during summer months or the amount of snow-melt during winter months. In addition, as is the case in any power market, the weather also influences the demand side. This might explain part of the deviations from the seasonality patterns. The years 1998, 1999 and 2000 show a particularly similar yearly seasonality pattern.

Analysing the finer structure of the price series reveals further seasonalities which are shown in Figure 2.3. The first graph shows a weekly seasonality with low prices during weekends and the second one shows intraday data with hourly resolution. A reduction of prices overnight is obvious. It also needs to be remarked that the deviations from the daily average price (base load) are very low compared to other markets like the UKPX (UK) and the EEX (Germany).

Another peculiar property of the market data is the occurrence of spikes. There are several apparent spikes in Figure 2.2. It is remarkable how fast prices revert to the previous level




Figure 2.3: Fine-structure of NordPool's electricity prices; red arrows point at Sundays




Figure 2.4: Electricity prices in the Nordic, German and UK market. The spike in the German market goes up to EUR 240 per MWh.

after an upward jump occurred. A closer analysis also reveals downward jumps. The fine-structure of a spike is shown in the third graph of Figure 2.3. The spike suddenly occurred on the 24th of January with a daily average of almost NOK 400 per MWh after about NOK 130 the day before. The price went down on the following day and was back to normal on the day after. The intraday movement is extremely volatile with levels of up to almost NOK 1800 per MWh during high demand hours. Over night, when demand is at a low level, prices revert to nearly normal levels. If a spike is caused by a power plant outage, such a behaviour can be explained by a supply-demand model and keeping in mind the shape of the marginal cost of production graph. Assuming constant high demand, a removal of a part of the supply stack would result in a significant increase of the price whereas the increase would be relatively low if the current demand was at a low level.

As electricity markets are poorly interconnected, regional markets can be very different.

Three different markets are plotted in Figure 2.4, the NordPool, the EEX (Germany) and the UKPX (UK). Whereas weekly seasonality is pronounced in all three markets, the volatility, the speed of mean-reversion and the size and occurrence of spikes are different.

2.2 The forward and future market

As mentioned before, the inability to store electricity makes it a pure flow variable and hence all derivative contracts need to specify a delivery period. Daily averages, i.e. base load contracts, are usually the underlying products. Other averages like peak, off-peak⁶ and block contracts can also be the underlying spot price. The most liquidly traded derivatives are futures and forward contracts which can be bought over the counter (OTC) or from the exchange. They specify the time to maturity, the duration of delivery and the futures or forward price.⁷ Due to the many possible combinations of maturity and duration, only a few of these are listed on exchanges. At NordPool, for instance, only futures with delivery durations of one day, one week and four weeks, with time to maturities of usually less than ten times the delivery period, are traded. In addition, season and year forwards are listed where the delivery periods are specified as follows: January to April (Winter 1), May to September (Summer), October to December (Winter 2) and January to December (Year). Derivative contracts can be physically or financially settled. Assuming financial settlement, as it is generally the case at NordPool, a holder of a forward base load contract would receive or pay the difference between the spot price and the forward price on every day during the delivery period. If electricity is required then it can be purchased on the spot market. Given a constant consumption of electricity, they would pay on average the base load price and receive or pay the difference to the forward price so that the net cost would be equal to the forward price times the delivery period. All futures and forward contracts traded on 1 August 2000 and 1 June 2001 are shown in Figure 2.5 where prices appear to reflect expected seasonality.

Due to the inability to store electricity efficiently it is impossible to hedge futures and forward contracts and hence they cannot be priced based on arbitrage arguments. On the other hand, it does not mean that participants selling forward contracts face non-hedgeable risks. Since power generators do not have to buy energy from the spot market but can produce it to a known price, their risk is even reduced by selling a forward contract because their cost and income is then totally determined.⁸

⁶The definition of peak and off-peak depends on the market. For example the EEX defines it as the average price over 8:00-20:00 and the UKPX as the average over 7:00-19:00.

⁷Under a forward price one understands the strike price of a zero cost forward contract.

⁸Ignoring uncertainty in fuel costs and counterparty risk.




Figure 2.5: Forward curve on two different days.

2.2.1 Forwards with and without a delivery period

In order to relate forwards paying out at one point in time to forwards paying over a time period we need to make some definitions. Without going into too much detail we use standard notation (S_t and $\mathcal{F}_t$ are the spot price and information available at time t , respectively and $\mathbb{Q}$ is the risk neutral measure assumed by the market) and simply make the following definition which is backed by arbitrage arguments.

Definition 2.2.1 (Forward)

The strike price K at time t of a zero-cost forward contract paying $S_T - K$ at time T will be denoted by $F_t^{[T]}$ and given by the risk neutral expectation

$$F_t^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t].$$

If the forward contract does not just pay at time T but pays over a time period $[T_1, T_2]$ the strike of a zero-cost forward depends on the precise specification of when the money is paid. In the Nordpool market the forward pays $(S_t - K)\Delta t$ at time t but alternative contracts, either over the counter or in other regional markets, might specify the payment of the whole amount at the end of the delivery period T_2 . This is called instant settlement and settlement at maturity, respectively.

By definition of a forward contract, the strike K has to be set so that the contract is of zero cost at the time t we enter into it, so for settlement at maturity we have

$$\mathbb{E}^{\mathbb{Q}} \left[\int_{T_1}^{T_2} (S_T - K) dT | \mathcal{F}_t \right] = 0.$$

If K satisfies this equation we find

$$K = \frac{1}{T_2 - T_1} \int_{T_1}^{T_2} \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t] dT.$$

In the case of instant settlement the money received can be invested in a risk-less bank account and so

$$\mathbb{E}^{\mathbb{Q}} \left[\int_{T_1}^{T_2} (S_T - K) e^{r(T_2-T)} dT | \mathcal{F}_t \right] = 0,$$

which leads to

$$K = \frac{r}{e^{r(T_2-T_1)} - 1} \int_{T_1}^{T_2} e^{r(T_2-T)} \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t] dT.$$

The following definition takes both cases into account.

Definition 2.2.2 (Forward with delivery)

We denote the strike price of a zero-cost forward contract with a delivery period $[T_1, T_2]$ at time t by $F_t^{[T_1, T_2]}$ and define it to be the weighted average of all instantaneous forwards in that period, i.e.

$$F_t^{[T_1, T_2]} = \int_{T_1}^{T_2} w(T; T_1, T_2) F_t^{[T]} dT, \quad (2.1)$$

where $w > 0$ and

$$\int_{T_1}^{T_2} w(T; T_1, T_2) dT = 1.$$

Note, for settlement at maturity the factor w is given by $w(T) := \frac{1}{T_2 - T_1}$ and for instant settlement we have $w(T) := \frac{r e^{r(T_2-T)}}{e^{r(T_2-T_1)} - 1} = \frac{r e^{-rT}}{e^{-rT_1} - e^{-rT_2}}$. For small delivery periods we can make the first order approximation

$$\frac{r e^{r(T_2-T)}}{e^{r(T_2-T_1)} - 1} \approx \frac{1}{T_2 - T_1},$$

and so it only makes a small difference whether the money is settled at the end or on a daily basis.

2.2.2 Building a continuous forward curve

No market provides forward prices for any arbitrary period $[T_1, T_2]$. The NordPool market, for example, lists prices for around 30 (partly overlapping) periods of which a sample is

shown in Figure 2.5. On a typical day these listed prices could consist of 5 daily,⁹ 5 weekly, 10 monthly, 7 seasonal and 3 yearly contracts. For the purpose of pricing options or even forwards which are not listed, it is desirable to derive instantaneous forward prices $F_t^{[T]}$ for every maturity T . This is an inverse problem which does not have a unique solution, because we look for a continuous function which satisfies a finite number of integral conditions (2.1).

[Fleten and Lemming, 2003] use a bottom up model (Multiarea Power Scheduling (MPS) model) to make a first prediction of the forward curve. A quadratic optimisation method is then used to minimise squared errors of the proposed forward curve and the results of the MPS model subject to constraints imposed by the observed forward bid and ask prices. A second term in the objective function assures low oscillations. In the setting of interest rates [Hagan and West, 2005] give a survey of a wide range of interpolation methods.

In the following subsections we discuss various simple interpolation methods and their limitations and finally suggest a method which uses seasonality information of the spot price history and some form of spline interpolation to satisfy all the integral conditions.

To simplify the notation we assume $t = 0$ in this subsection, i.e. we seek to find an interpolation $F_0^{[T]}$ to the discrete forward contracts $F_0^{[T_i, \hat{T}_i]}$ given at time 0.

2.2.2.1 Approximation by a set of basis functions

Assuming prices for the periods $[T_1, \hat{T}_1]$, $[T_2, \hat{T}_2]$, $\dots$, $[T_n, \hat{T}_n]$, $T_1 \leq T_i$ and $\hat{T}_i \leq \hat{T}_n$ are given where periods are allowed to overlap. Given a set of basis functions $g_i : [T_1, \hat{T}_n] \rightarrow \mathbb{R}$, $i \in \{1, \dots, k\}$, the function $f : [T_1, \hat{T}_n] \rightarrow \mathbb{R}$ approximating the forward curve can be defined as a linear combination

$$f(T) := \sum_{i=1}^k a_i g_i(T).$$

According to Equation (2.1) the integral conditions are

$$\int_{T_i}^{\hat{T}_i} w(s; T_i, \hat{T}_i) f(s) ds = F_0^{[T_i, \hat{T}_i]},$$

and so

$$\sum_{j=1}^k a_j \int_{T_i}^{\hat{T}_i} w(s; T_i, \hat{T}_i) g_j(s) ds = F_0^{[T_i, \hat{T}_i]}.$$

With $G_{i,j}(s) := \int w(s; T_i, \hat{T}_i) g_j(s) ds$ and $v_i := F_0^{[T_i, \hat{T}_i]}$ we get the following system of equations:

$$\sum_{j=1}^k \left(G_{i,j}(\hat{T}_i) - G_{i,j}(T_i) \right) a_j = v_i, \quad i \in \{1, \dots, n\}.$$

⁹Daily, weekly, etc. indicates the duration $T_2 - T_1$ of delivery.

$G_{i,j}(T)$	$w = \frac{1}{\hat{T}_i - T_i}$	$w = \frac{r e^{-rT}}{e^{-rT_i} - e^{-r\hat{T}_i}}$
$g_j(T) = \sin(c_j T)$	$\frac{-1}{c_j(\hat{T}_i - T_i)} \cos(c_j T)$	$\frac{r e^{-rT}}{(e^{-rT_i} - e^{-r\hat{T}_i})(r^2 + c_j^2)} (-r \sin(c_j T) - c_j \cos(c_j T))$
$g_j(T) = \cos(c_j T)$	$\frac{1}{c_j(\hat{T}_i - T_i)} \sin(c_j T)$	$\frac{r e^{-rT}}{(e^{-rT_i} - e^{-r\hat{T}_i})(r^2 + c_j^2)} (-r \cos(c_j T) + c_j \sin(c_j T))$
$g_j(T) = T$	$\frac{T^2}{2(\hat{T}_i - T_i)}$	$\frac{r e^{-rT}}{e^{-rT_i} - e^{-r\hat{T}_i}} \left(-\frac{T}{r} - \frac{1}{r^2}\right)$
$g_j(T) = T^2$	$\frac{T^3}{3(\hat{T}_i - T_i)}$	$\frac{r e^{-rT}}{e^{-rT_i} - e^{-r\hat{T}_i}} \left(-\frac{T^2}{r} - \frac{2T}{r^2} - \frac{2}{r^3}\right)$

Table 2.1: Values of $G_{i,j}(T)$ given the choice of g_j and w .

For some choices of basis functions the integral values $G_{i,j}(T)$ are given in Table 2.1. If one chooses to use fewer basis function than (non-redundant) integral constraints the equation system will not be solvable in general. In this case, one could still minimise the squared errors

$$\sum_{i=1}^n \left(v_i - \sum_{j=1}^k (G_{i,j}(\hat{T}_i) - G_{i,j}(T_i)) a_j \right)^2 \rightarrow \min,$$

using a least square algorithm. A result of such an approximation is shown in Figure 2.6 where only a small number of sine and cosine basis functions have been used and therefore the approximation is not expected to be very good. Given there are 29 non-redundant contracts in this example (6 daily, 6 weekly, 10 monthly, 6 seasonal and 1 yearly forwards) it would require sine and cosine function of about 15 different frequencies to fulfil all integral conditions, the result of which would be highly oscillating. Another disadvantage of this method is its sensitivity to the first few short term contracts. The volatility of short term forwards close to maturity is much higher than contracts far away from maturity.¹⁰ A method which results in function values $F_0^{[T]}$, $T \gg 0$, being sensitive to given prices $F_0^{[T_i, \hat{T}_i]}$ for $\hat{T}_i$ close to 0 results in too high volatilities for long term forwards which is not desirable.

2.2.2.2 Approximation by a piecewise quadratic polynomial

Another approach is to use piecewise polynomial functions (splines) and request continuity of function values and derivatives as appropriate. We assume non-overlapping contracts without any gaps, i.e. $0 < T_1 < T_2 < \dots < T_{n+1}$ and $F_0^{[T_i, T_{i+1}]}$, $i \in \{1, \dots, n\}$. Furthermore we assume for the moment $w(T; T_i, T_{i+1}) = \frac{1}{T_{i+1} - T_i}$. The interpolating function $f : [T_1, T_{n+1}] \rightarrow \mathbb{R}$ is now given by piecewise quadratic functions $f_i : [T_i, T_{i+1}] \rightarrow \mathbb{R}$ as follows

$$f(t) = f_i(t), \quad \forall t \in [T_i, T_{i+1}].$$

¹⁰This is due to the mean-reverting nature of the underlying and the inability to hedge with it.




Figure 2.6: Approximation of the forward curve by a continuous curve. Both curves are linear combinations of sine and cosine functions, where the red one uses two periodicities (yearly and half yearly) and the blue one uses three. Due to the few parameters and the many integral constraints this interpolation is not expected to fit very well.

It turns out to be favourable to represent f_i as

$$f_i(t) = g_i\left(\frac{t - T_i}{T_{i+1} - T_i}\right), \quad t \in [T_i, T_{i+1}], \quad g_i(t) := a_i t^2 + b_i t + c_i, \quad t \in [0, 1],$$

and so we get

$$\begin{aligned} \int_{T_i}^{T_{i+1}} \frac{f_i(t)}{T_{i+1} - T_i} dt &= \int_0^1 g_i(s) ds = \frac{a_i}{3} + \frac{b_i}{2} + c_i, \\ f_i(T_i) &= g_i(0) = c_i, \\ f_i(T_{i+1}) &= g_i(1) = a_i + b_i + c_i, \\ f'_i(T_i) &= \frac{g'_i(0)}{T_{i+1} - T_i} = \frac{b_i}{\Delta T_i}, \\ f'_i(T_{i+1}) &= \frac{g'_i(1)}{T_{i+1} - T_i} = \frac{2a_i + b_i}{\Delta T_i}. \end{aligned}$$

Hence, the conditions to be satisfied by g_i are

$$\begin{aligned} \frac{a_i}{3} + \frac{b_i}{2} + c_i &= v_i && \text{(integral condition),} \\ a_i + b_i + c_i &= c_{i+1} && \text{(continuity),} \\ \frac{2a_i + b_i}{\Delta T_i} &= \frac{b_{i+1}}{\Delta T_{i+1}} && \text{(smoothness),} \end{aligned}$$

which gives us $3n - 2$ equations and $3n$ unknowns. Two more conditions need to be imposed to obtain a unique solution. First we rearrange the equations in order to reduce the computational effort to solve the equation system. Using the integral and smoothness condition to eliminate a_i and b_i from the continuity condition leads to

$$\begin{aligned} c_i &= v_i - \frac{a_i}{3} - \frac{b_i}{2}, \quad i \in \{1, \dots, n\} && \text{(integral),} \\ a_i &= \frac{1}{2} \left(\frac{\Delta T_i}{\Delta T_{i+1}} b_{i+1} - b_i \right), \quad i \in \{1, \dots, n-1\} && \text{(smoothness),} \\ b_{i-1} + 2 \left(\frac{\Delta T_{i-1}}{\Delta T_i} + 1 \right) b_i + \frac{\Delta T_i}{\Delta T_{i+1}} b_{i+1} &= 6(v_i - v_{i-1}), \quad i \in \{2, \dots, n-1\} && \text{(continuity).} \end{aligned}$$

The third equation is a tridiagonal equation system in b and can be solved in $O(n)$ steps. With the knowledge of the values $b_1, \dots, b_n$ results for a_i and c_i can be obtained using the second and first equation, respectively.

As mentioned before, there is no unique solution to the equation system which leaves us with the choice of imposing two more boundary conditions. For example, one could chose to define the slope of the function on the left and right hand side. Say, d_1 and d_2 are given values of the derivative at T_1 and T_{n+1} , respectively, the following two conditions need to be satisfied:

$$\frac{b_1}{\Delta T_1} = d_1, \quad \frac{2a_n + b_n}{\Delta T_n} = \frac{b_{n+1}}{\Delta T_{n+1}} = d_2.$$

For the right derivative we have introduced an additional but imaginary segment with the index $n + 1$ which only makes the equations more elegant because then we get

$$\begin{aligned} a_n &= \frac{1}{2} (\Delta T_n d_2 - b_n), \\ b_1 &= d_1 \Delta T_1, \\ b_{n-1} + 2 \left(\frac{\Delta T_{n-1}}{\Delta T_n} + 1 \right) b_n &= -\Delta T_n d_2 + 6(v_n - v_{n-1}), \end{aligned}$$

which completes the equation system. A result of the procedure can be seen in Figure 2.7. The advantage of this interpolation method is that it produces a smooth function which satisfies precisely all integral conditions and hence does not introduce arbitrage. However, it will not show a seasonality pattern over a segment which only contains the information of a yearly contract.

Note, by applying cubic spline interpolation to the primitive of f we obtain the same interpolation as above. Define $G(t) := \int_{T_1}^t f(x) dx$ and using the same notation $v_i := F_0^{[T_i, T_{i+1}]}$, the integral conditions become

$$G(T_1) = 0, \quad G(T_{i+1}) - G(T_i) = v_i,$$




Figure 2.7: Interpolation of the forward curve by a piecewise quadratic polynomial.

which we can rewrite to obtain a point-interpolation formulation for G

$$G(T_1) = 0, G(T_2) = v_1, \dots, G(T_i) = \sum_{j=1}^{i-1} v_j.$$

Applying standard cubic spline interpolation and differentiating G results in a piecewise quadratic polynomial satisfying the same continuity and smoothness conditions as above.

2.2.2.3 Approximation by a piecewise quadratic polynomial: the general case

We now go to the general case and allow any weighting function w . The derivation of the method remains the same but only the integral condition changes:

$$\begin{aligned} \int_{T_i}^{T_{i+1}} w(t; T_i, T_{i+1}) f_i(t) dt &= \int_0^1 w(T_i + (T_{i+1} - T_i)s; T_i, T_{i+1}) (T_{i+1} - T_i) g_i(s) ds \\ &= \alpha_i a_i + \beta_i b_i + c_i, \end{aligned}$$

with

$$\begin{aligned} \alpha_i &:= \int_0^1 w(T_i + (T_{i+1} - T_i)s; T_i, T_{i+1}) (T_{i+1} - T_i) s^2 ds, \\ \beta_i &:= \int_0^1 w(T_i + (T_{i+1} - T_i)s; T_i, T_{i+1}) (T_{i+1} - T_i) s ds. \end{aligned}$$

In the case of $w(T; T_i, T_{i+1}) = \frac{r e^{-rT}}{e^{-rT_i} - e^{-rT_{i+1}}}$ we get

$$\alpha_i = \frac{e^{rT_i} ((r\Delta T_i + 1)^2 + 1) - 2e^{rT_{i+1}}}{r^2 \Delta T_i^2 (e^{rT_i} - e^{rT_{i+1}})},$$

$$\beta_i = \frac{e^{rT_i} (r\Delta T_i + 1) - e^{rT_{i+1}}}{r(e^{rT_i} - e^{rT_{i+1}}) \Delta T_i}.$$

As before, the conditions to be satisfied are

$$\begin{aligned} \alpha_i a_i + \beta_i b_i + c_i &= v_i && \text{(integral condition),} \\ a_i + b_i + c_i &= c_{i+1} && \text{(continuity),} \\ \frac{2a_i + b_i}{\Delta T_i} &= \frac{b_{i+1}}{\Delta T_{i+1}} && \text{(smoothness).} \end{aligned}$$

Rearranging the equations yields

$$\begin{aligned} c_i &= v_i - \alpha_i a_i - \beta_i b_i, \quad i \in \{1, \dots, n\} && \text{(integral),} \\ a_i &= \frac{1}{2} \left(\frac{\Delta T_i}{\Delta T_{i+1}} b_{i+1} - b_i \right), \quad i \in \{1, \dots, n-1\} && \text{(smoothness),} \\ \frac{1 + \alpha_{i-1} - 2\beta_{i-1}}{2} b_{i-1} &&& \\ + \left(\frac{\Delta T_{i-1}}{\Delta T_i} \frac{1 - \alpha_{i-1}}{2} - \frac{\alpha_i}{2} + \beta_i \right) b_i &&& \\ + \frac{\Delta T_i}{\Delta T_{i+1}} \frac{\alpha_i}{2} b_{i+1} &= v_i - v_{i-1}, \quad i \in \{2, \dots, n-1\} && \text{(continuity),} \end{aligned}$$

and for the boundary conditions we get

$$\begin{aligned} a_n &= \frac{1}{2} (\Delta T_n d_2 - b_n), \\ b_1 &= d_1 \Delta T_1, \\ \frac{1 + \alpha_{n-1} - 2\beta_{n-1}}{2} b_{n-1} + \left(\frac{\Delta T_{n-1}}{\Delta T_n} \frac{1 - \alpha_{n-1}}{2} - \frac{\alpha_n}{2} + \beta_n \right) b_n &= v_n - v_{n-1} - \Delta T_n \frac{\alpha_n}{2} d_2. \end{aligned}$$

In the example shown in Figure 2.7 the maximum difference between the interpolating functions using $w = \frac{1}{T_{i+1} - T_i}$ and $w = \frac{r e^{-rT}}{e^{-rT_i} - e^{-rT_{i+1}}}$ is of order 0.1 NOK/MWh and achieved at the end of the delivery period T_{n+1} where an annual interest¹¹ of 5% is assumed.

2.2.2.4 Approximation by a piecewise cubic polynomial

In the context of interpolating points, it is well known that cubic splines provide the smoothest interpolation possible in the sense of Definition 2.2.3, see [de Boor and Lynch, 1966]. However, in this context where integral conditions need to be satisfied rather than function values interpolated this does not seem to be the case as will be demonstrated below.

¹¹ $r = \ln(1.05)$

As before we define

$$f_i(t) = g_i \left(\frac{t - T_i}{T_{i+1} - T_i} \right), \quad t \in [T_i, T_{i+1}], \quad g_i(t) := a_i t^3 + b_i t^2 + c_i t + d_i, \quad t \in [0, 1],$$

and impose integral, continuity, smoothness and in addition curvature conditions and obtain the equation system

$$\begin{aligned} \frac{a_i}{4} + \frac{b_i}{3} + \frac{c_i}{2} + d_i &= v_i && \text{(integral condition),} \\ a_i + b_i + c_i + d_i &= d_{i+1} && \text{(continuity),} \\ \frac{3a_i + 2b_i + c_i}{\Delta T_i} &= \frac{c_{i+1}}{\Delta T_{i+1}} && \text{(smoothness),} \\ \frac{3a_i + b_i}{\Delta T_i^2} &= \frac{b_{i+1}}{\Delta T_{i+1}^2} && \text{(curvature),} \end{aligned}$$

and impose zero curvature boundary conditions:

$$b_1 = 0, \quad 3a_n + b_n = 0.$$

This leaves us with $4n - 1$ equations and $4n$ unknowns so we are free to impose another condition where we have chosen to set the value at the end of the period to be equal to the average integral value:

$$a_n + b_n + c_n + d_n = v_n.$$

We solve this $4n \times 4n$ equation system using a sparse matrix solver. Figure 2.8 shows that the result can be slightly more oscillatory than for quadratic splines.

We further investigate whether quadratic or cubic splines might be the smoothest function possible satisfying the integral constraints.

Definition 2.2.3

Let $f : [a, b] \rightarrow \mathbb{R}$ be a continuously differentiable and piecewise twice continuously differentiable function then we define a measure of smoothness ω by

$$\omega[f] := \int_a^b f''(t)^2 dt.$$

We use a very simple numerical example to illustrate how smooth the interpolations subject to different boundary conditions are. However, it remains inconclusive on whether quadratic or cubic splines are the smoothest possible function satisfying the integral constraints. Consider the function $\sin(\frac{\pi}{3}x)$ on the interval $[0, 3]$ and integral constraints given by the average values of the sin function on $[0, 1]$, $[1, 2]$ and $[2, 3]$, i.e. $v_1 = \frac{3}{2\pi}$, $v_2 = \frac{3}{\pi}$ and $v_3 = \frac{3}{2\pi}$. The smoothness of the sin function is given by

$$\int_a^b (\sin(\alpha x))''^2 = \alpha^4 \left(\frac{b-a}{2} - \frac{\sin(2\alpha b) - \sin(2\alpha a)}{4\alpha} \right),$$




Figure 2.8: Interpolation of the forward curve by a piecewise cubic polynomial.

function f	curvature $\omega[f]$
$\sin(\frac{\pi}{3}x)$	1.8038721
quadratic spline, $f'(0) = f'(3) = \frac{\pi}{3}$	1.857835
quadratic spline, $f''(0) = f''(3) = 0$	2.051754
cubic spline, $f''(0) = f''(3) = 0, f'(0) = \frac{\pi}{3}$	1.808850
cubic spline, $f''(0) = f''(3) = 0, f'''(0) = 0$	2.893934
cubic spline, $f''(0) = f''(3) = 0, f(3) = v_3$	8.293264

Table 2.2: Comparison of the smoothness of functions satisfying the integral conditions. As it turns out the sin function is the smoothest of the given functions.

and approximately 1.8038721 in this particular case. Table 2.2 compares the sin function with various quadratic and cubic splines and as it turns out the sin function is smoother than all the considered spline functions. This counterexample shows that at least the quadratic and cubic splines with the boundary conditions considered do not possess the maximum smoothness property according to the definition of ω . All the functions of Table 2.2 are also plotted in Figure 2.9 and 2.10. In practice this does not play a big role as quadratic splines tend to be very robust and smooth. Cubic splines sometimes exhibit seemingly unnecessary oscillations as in Figure 2.8.

2.2.2.5 Approximation by a seasonal function and spline correction

There is no unique way to infer a continuous forward curve given the forward contracts with delivery periods listed in the market, and so all methods described above which satisfy all




Figure 2.9: Smoothness of the quadratic spline function compared with a sin function. Two cases of boundary conditions to determine the spline are considered: given slope on both ends matching the slope of the sin function and zero curvature on both ends, respectively.




Figure 2.10: Smoothness of cubic spline functions. Three cases of boundary conditions to determine the spline are considered. In all cases the curvature at both ends is set to zero, and for the third condition, the slope on the left is set to match the slope of the sin function, the third derivative on the left is set to zero and the function value on the right is set to the integral value.

the integral conditions represent possibilities of a continuous forward curve consistent with forward market data. Nevertheless, we can identify a few more criteria which seem to be fairly intuitive; they are connected to the dynamics of the curve, seasonality and smoothness. The dynamics of the continuous curve should be similar to the dynamics of the observed prices; in other words an interpolation of today's forward curve should not be too different from tomorrow's curve. Also, as we observe seasonality in the spot price patterns it should be somehow reflected in the forward curve. Finally, as long as the previous conditions are satisfied the interpolated curve should be as smooth as possible.

Our suggestion for building a continuous forward curve consists of two parts. First, a reasonable first approximation $\bar{f}$ of the continuous forward curve needs to be found using other ways, like expert knowledge, historical data, weather forecasts or even additional market data.¹² This first approximation does not need to satisfy the integral conditions and so we make errors

$$e_i := F_0^{[T_i, T_{i+1}]} - \int_{T_i}^{T_{i+1}} w(T; T_i, T_{i+1}) \bar{f}(T) dT.$$

We then interpolate the errors by a quadratic spline $\tilde{f}$ as described above so that

$$\int_{T_i}^{T_{i+1}} w(T; T_i, T_{i+1}) \tilde{f}(T) dT = e_i.$$

The function $\tilde{f}$ can be seen as a smooth correction to the first approximation $\bar{f}$ in order to satisfy all the integral conditions. As the functions $f := \bar{f} + \tilde{f}$ obviously satisfies all integral conditions we choose this to be our continuous forward curve. In Figure 2.11 we have used a sum of sine functions with quarter, half and yearly seasonality as a first approximation $\bar{f}$ to the forward curve. The amplitudes of the sine functions have been chosen to minimise squared errors of $\bar{f}$ to the historical spot price series.

2.2.3 Call and puts

Further commonly traded products are call or put options on futures and forwards. They are mainly OTC traded but a few regional markets, including the NordPool, list these products at their exchanges, too. The contract specifies the exercise price K , the time of maturity T and the delivery period $[T_1, T_2]$ of the underlying forward contract, where normally $T = T_1$. The payoff of a call option is then $(T_2 - T_1)(F_T^{[T_1, T_2]} - K)^+$, payable at time T , where $F_T^{[T_1, T_2]}$ denotes the forward at time T . The buyer could use the payoff and enter straight into a forward contract with exercise price $F_T^{[T_1, T_2]}$ and as $(T_2 - T_1)(F_T^{[T_1, T_2]} - K)^+$ has been

¹²Assuming a liquid option market.




Figure 2.11: Interpolation of the forward curve by a seasonal function and spline correction. Three years worth of spot history data has been used to calibrate a seasonality function which is then used as a first approximation of the forward curve. The difference between the seasonality function and the observed forward prices is then corrected by a piecewise quadratic polynomial as shown in Figure 2.7.

received the exercise price is effectively $\min\{F_T^{[T_1, T_2]}, K\}$.¹³ Such a contract can be useful if a consumer expects increased power consumption during the period of $[T_1, T_2]$ but is not certain at the moment t but will know for sure at time T . They could enter into a forward contract now (t) or when they know for sure at T . In both cases they face market risks, either of not needing the electricity and by fulfilling the forward contract to loose money or by an uncertain forward price at time T . By buying the option they ensure the forward price at time T does not exceed the specified price K and if the energy is not needed then they just take the payoff.

An example of a more complex option which is tailored to the needs of power consumers is a swing option. It normally comes bundled with a base load forward contract and then leaves the consumer freedom to decrease or increase consumption within pre-specified limits. These limits can define minimum and maximum volume per day and overall, allowed dates of volume adjustments and the maximum number of volume adjustments.

¹³Not taking into account interest rate payments.

Chapter 3

Stochastic spot price model

There are two main approaches to construct realistic stochastic processes for the spot price process. One is to understand and model the underlying mechanisms crucial for the price determination. The second approach is to simply observe the market price time series and to construct a stochastic process exhibiting the main properties of the market data. We will follow the second approach and define a Markov process in continuous time. The basic properties of seasonality, mean-reversion and the occurrence of spikes will be reflected by the process. However, it will not be able to emulate the intraday behaviour in a fully satisfactory way, especially when spikes occur. A supply-demand model would then probably be necessary to force prices back to a normal level overnight. Fortunately, most of the derivative contracts are written on daily averages and therefore the following model should be regarded as a continuous model for base load prices.

3.1 Existing models

Regardless of the variety of the models proposed in the literature, they are mainly based on some mean-reverting process, quite often an Ornstein-Uhlenbeck (OU) process. The most basic model, proposed in [Lucia and Schwartz, 2002], is the exponential of an OU process (X_t) ¹ and a seasonal component f . Let W be a standard Brownian motion and let S_t denote the spot price at time t then the model can be formulated as

$$\begin{aligned}dX_t &= -\alpha X_t dt + \sigma dW_t, \\ S_t &= \exp(f(t) + X_t),\end{aligned}\tag{3.1}$$

where σ is a volatility parameter and α the speed of mean-reversion. In this model, we know S_t is log-normally distributed which allows for analytic option price formulae very similar

¹We generally use brackets to indicate that we mean the entire process, i.e. it is an abbreviation for $(X_t)_{t \in [0, T]}$ or $(X_t)_{t \in \mathbb{R}^+}$.

to the formulae in the Black-Scholes model. The formulae are given in Appendix D.3. To allow for a stochastic seasonality, a further component can be inserted into the model and as long as this process has a normal-distribution, the analytical tractability is sustained. Therefore it is suggested in [Lucia and Schwartz, 2002] to consider the model defined by

$$\begin{aligned} dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= \mu dt + \tilde{\sigma} dB_t, \\ S_t &= \exp(f(t) + Y_t + X_t), \end{aligned}$$

where B is a W -independent Brownian motion. The term $f(t) + Y_t$ can be seen as a seasonality with stochastic trend. The main disadvantage of these models are their inability to mimic spikes. To overcome this problem, jumps can be inserted into these models. With (N_t) denoting a Poisson process with intensity λ and J being the jump size, the obvious choice would be to define

$$\begin{aligned} dX_t &= -\alpha X_t dt + \sigma dW_t + J_t dN_t, \\ S_t &= \exp(f(t) + X_t), \end{aligned} \tag{3.2}$$

which is briefly mentioned in [Clewlow and Strickland, 2000, Section 2.8].² Analytic results are given in [Deng, 2000] which are based on transform analysis described in [Duffie et al., 2000]. The issue of calibration to historical data as well as the observed forward curve is discussed in [Cartea and Figueroa, 2005] and practical results for the UK electricity market are given. For these models to exhibit typical spikes it is required that the mean-reversion rate α is extremely high, otherwise jumps do not revert quickly enough. It is suggested in [Benth et al., 2005] to introduce a set of independent pure mean-reverting jump processes of the form

$$S_t = \sum_{i=1}^n w_i Y_t^{(i)}, \quad dY_t^{(i)} = -\alpha_i Y_t^{(i)} dt + \sigma_i dL_t^{(i)}, \quad i = 1, \dots, n,$$

where w_i are some positive weights and $L^{(i)}$ are independent increasing càdlàg pure jump processes. Note, the spot price process is a linear combination of the pure jump processes and as there is no exponential function involved, positivity of the spot is achieved by allowing positive jumps only. The advantage of this formulation is that semi-analytic formulae for option prices on forwards with a delivery period can be derived. However, a full analysis of this class of models still seems to be in its early stages.

An alternative approach is to introduce two different and independent stochastic processes and a Markov switching process, saying which of the processes is active at each time. One

²There the sde is written in terms of S_t by applying Itô's formula.

process can be considered to be the normal regime and the other one the spiky regime. With a Markov switching process m_t with values in $\{0, 1\}$ the model can be described as follows:

$$S_t := X_t^{(m_t)} = \begin{cases} X_t^{(0)} & \text{if } m_t = 0 \\ X_t^{(1)} & \text{if } m_t = 1 \end{cases}.$$

In [de Jong and Huismann, 2002], a time-discrete model is introduced where the normal regime is given by a discrete version of an exponential OU process (3.1) and the spiky regime by a series of independent log-normally distributed random variables. The independence between the two regimes $X^{(0)}$ and $X^{(1)}$ assures the return of the price to a normal level after the occurrence of a spike. The model also allows for analytic formulae for simple options because for the expected value we have $\mathbb{E}[g(S_t)] = \mathbb{E}[g(X_t^{(0)})]P(r_t = 0) + \mathbb{E}[g(X_t^{(1)})]P(r_t = 1)$. However, it does not seem to be obvious how to define an appropriate process for the spiky regime in continuous time. Given we assume the paths of the spike process are continuous there will be dependence between the average sizes of two successive spikes. On the other hand, assuming independence of any two values of the spike process with $t_1 \neq t_2$ will result in some form of white noise. Neither of the two cases would represent reality very well.

Another approach is to model a demand supply equilibrium as described in [Barlow, 2002], where the underlying demand in electricity is assumed to be an OU process (X_t) . The demand is said to be inelastic, i.e. is independent of the current price, but the supply is increasing as prices increase. This principle is based on the fact that the majority of consumers receive electricity at a fixed price and will not reduce consumption if prices on the electricity exchange rise, but more power plants will be happy to generate electricity as the income per MWh goes up, see the marginal cost of production shown in Figure 2.1. Let $u : \mathbb{R}^+ \rightarrow [0, a]$ be the supply function and $u(s)$ the supply of electricity if the price was s . The spot price process is then defined by the equilibrium of supply and demand

$$u(S_t) = X_t.$$

This is an implicit equation for S_t and a solution might not always exist which happens for example if the demand process X_t exceeds the maximum supply $a := \sup_{s \geq 0} u(s)$. If (X_t) is an OU process there is always a positive probability that the value a will be exceeded. To make the process (S_t) well defined one could cap the demand just below the maximum supply, i.e.

$$u(S_t) = \min \{X_t, a - \epsilon\}^+,$$

which is suggested in [Barlow, 2002], or alternatively, one could reflect the process X_t on the maximum supply barrier as soon as it reaches it. This would have the advantage that no

maximum price would need to be imposed as it is implicitly the case by defining a maximum demand $a - \epsilon$. The non-linearity of the supply curve transforms the OU process (X_t) in such a way that one observes price explosions in form of spikes, see Figure 3.1. The disadvantage of the particular demand-supply model illustrated in the figure is that jumps almost always reach $S_{\max}$ and hence the jump size distribution of the real market is not well represented.

3.2 A mean-reverting model exhibiting seasonality and spikes

For the purpose of option pricing one of the main aims of this thesis is to introduce and extensively examine a stochastic model appropriate to describe the spot electricity price with focus on European electricity markets and the Scandinavian market in particular.

In order to keep the model analytically tractable we propose a simple continuous time process (S_t) consisting of three components: a deterministic periodic function f characterising seasonality, an Ornstein-Uhlenbeck (OU) mean-reverting process (X_t) and a mean-reverting process with a jump component to incorporate spikes (Y_t) :

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t, \end{aligned} \tag{3.3}$$

where (N_t) is a Poisson-process with intensity λ and (J_t) is an independent identically distributed (iid) process representing the jump size. Furthermore we require (W_t) , (N_t) and (J_t) to be mutually independent. At this point we make no assumption on the jump size J but will later give examples with exponentially and normally distributed jump sizes.

The model is able to represent typical features of the electricity spot price dynamics like seasonality, mean-reversion and the occasional occurrence of spikes, which in our opinion is crucial for a model to be realistic. However, this model does not claim to fully represent all properties of electricity prices as seen in the market. Historical data indicates a varying volatility over time, see Figure 3.4, and hence would require the introduction of an additional stochastic volatility process. Also, judging from the forward curve dynamics, a further process describing the stochastic component of the seasonality might be needed in order to explain the high volatility of forward contracts maturing in the far future. Finally, it should be pointed out that the risk of a spike occurring is unlikely to be constant over time but rather seasonal dependant. Although, it is not difficult to formulate a stochastic model incorporating all these properties, it would be hard to work with, as far as calibration and option pricing is concerned.




Figure 3.1: A demand supply model. Here the demand is independent of the current price and given by an OU process X_t . The supply depends on the current price and is here simply a deterministic function $u(s) = a - (a - b)e^{-\lambda s}$. The spot price is therefore given by $S_t = g(X_t)$ with $g(x) := \max\{S_{\max}, \frac{1}{\lambda} \ln \frac{a-b}{a-x}\}$ where we truncate the price if $S_{\max}$ is exceeded.

$f(t) =$	$\ln(100) + 0.5 \cos(2\pi t)$	$\sigma =$	1.4	$\beta =$	200
$\alpha =$	7	$\mu_J =$	0.4	$\lambda =$	4
$J_t \sim$	$\exp(1/\mu_J)$				

Table 3.1: Parameters of the sample path.

The only difference of (3.3) to the well studied model (3.2) is the introduction of an independent spike process (Y_t) which allows to choose a different, and indeed higher, mean-reversion rate β in order for the jump to revert much more quickly and so to form a shape similar to a spike. This is crucial for modelling the NordPool market but might not be needed in markets where the speed of mean-reversion is generally very high, like in the UKPX or EEX.

To visualise this process, Figure 3.2 shows a sample path of the processes (X_t) , (Y_t) and the composed process (S_t) . The parameters used are not calibrated to any market but are chosen arbitrarily for the sake of demonstration and given in Table 3.1.

The equations for the spot process (S_t) can be rewritten in order to eliminate the exponential function. Defining $\tilde{X}_t := \exp(X_t)$ and $\tilde{Y}_t := \exp(Y_t)$ and applying Itô's formula yields

$$\begin{aligned} S_t &= \exp(f(t)) \tilde{X}_t \tilde{Y}_t, \\ \frac{d\tilde{X}_t}{\tilde{X}_t} &= \alpha \left(\frac{\sigma^2}{2\alpha} - \ln \tilde{X}_t \right) dt + \sigma dW_t, \\ \frac{d\tilde{Y}_t}{\tilde{Y}_{t-}} &= -\beta \ln \tilde{Y}_{t-} dt + (e^{J_t} - 1) dN_t. \end{aligned}$$

3.3 Parameter estimation based on historical data

As model (3.3) consists of three components $S_t = \exp(f(t) + X_t + Y_t)$ and only S_t is observable, estimating parameters becomes non-trivial. We follow a heuristic approach and first determine the seasonal component. Further assumption about the structure of f need to be made and the obvious choice is to assume some form of yearly and weekly seasonality. Here we define f to be of the form

$$f(t) = a_0 + \sum_{i=1}^6 a_i \cos(2\pi\gamma_i t) + b_i \sin(2\pi\gamma_i t),$$

with $\gamma_1 = 1$, $\gamma_2 = 2$, $\gamma_3 = 4$ for the yearly seasonality and $\gamma_4 = 365/7$, $\gamma_5 = 2 \times 365/7$, $\gamma_6 = 4 \times 365/7$ for the weekly seasonality. The parameters a_i and b_i are chosen to minimise squared errors between observed prices and the seasonal function $\sum_j (\ln S_{t_j} - f(t_j))^2 \rightarrow \min$ and can be solved by a least-square algorithm.




Figure 3.2: Sample paths of X , Y and S .

market	σ	α
NordPool	1.40	6.9
UKPX	2.70	170
EEX	4.15	140

Table 3.2: Estimated parameters of the process (X_t) .

Having determined the seasonal component and removed it from the data we are left with a realisation of the pure stochastic part $\ln S_t - f(t) = X_t + Y_t$. To separate the two processes $(X_t + Y_t)$ we use the fact that the spike process (Y_t) is mainly close to zero and only occasionally assumes big values but then only for a very short time. So we consider the time-series as the realisation of (X_t) occasionally disturbed by big values. We therefore estimate the parameters of the mean-reversion process X_t based on the log-de-seasonalised time-series, see Appendix B.2, knowing that the result is likely to be disturbed by the occurrence of spikes in the data. However, we use the parameters obtained as a first approximation and eliminate all data points likely to be caused by a spike. As we know the conditional distribution of the change in (X_t) ,

$$X_{t+\Delta t} - X_t e^{-\alpha\Delta t} \sim \mathcal{N}\left(0, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\Delta t})\right),$$

we remove all points if they do not fall within a few standard deviations of it. Having removed all likely jumps (i.e. points where (Y_t) is not close to zero) the OU parameters can be estimated again and are now more likely to reflect the parameters of (X_t) . This procedure can be repeated a few times and experiments show that about three iterations seem to be sufficient. One drawback of the algorithm is that it is unable to detect small jumps which are within a few standard deviations of the change in the mean-reverting process. This needs to be taken into account when estimating parameters of the jump size distribution. For the mean-reversion rate β of the spike process (Y_t) we suggest to use some ad-hoc parameter likely to be known by a practitioner. The spike process reverts exponentially ($e^{-\beta\Delta t}$) and experts will have some idea after which time Δt the spike halves ($\Delta t = \frac{\ln 2}{\beta}$) or is decimated ($\Delta t = \frac{\ln 10}{\beta}$). For $\beta = 200$, for example, we have $\Delta t \approx 1.3/365$ and $\Delta t \approx 4.2/365$, respectively.

The result of the parameter estimation can be seen in Figure 3.3, where the market data of three different markets (NordPool, UKPX, EEX) have been used to calibrate the parameters of the model and with which sample paths are generated and plotted. The parameters of the mean-reverting process (X_t) are given in Table 3.2.

Circumstantial evidence suggests that for the NordPool data clusters of high volatility seem to exist, see Figure 3.4. As our parameter estimation procedure is based on the believe that




Figure 3.3: Market data (left) and sample paths of the calibrated stochastic model (right) for three different markets: NordPool, UKPX, EEX.

we have a constant volatility parameter σ it is likely to wrongly identify too many spikes in high volatility regimes. In order to minimise this effect we define a large range in which we assume data-points not to belong to the spike regime, in this case we define the range $[-5, 4]$ times the standard deviation. We still seem to wrongly detect a few jumps which have been caused by high volatility.

Figure 3.4 shows all spikes identified by the algorithm as well as de-seasonalised log-returns, i.e. values $Z_{t+\Delta t} - Z_t$ with $Z_t := \ln S_t - f(t)$. For mean-reverting processes, the distribution of $Z_{t+\Delta t} - Z_t e^{-\alpha\Delta t}$ is also of interest and hence plotted in the same figure as well, but as $e^{-\alpha\Delta t} \approx 0.98$ in this example both graphs look very similar.

Histograms of the log-returns $(Z_{t+\Delta t} - Z_t)$ are plotted in Figure 3.5 and compared to the distribution expected from an OU-process with parameters $\alpha = 7$ and $\sigma = 1.4$. Note, that for an OU-process (X_t) we have

$$X_{t+\Delta t} - X_t = X_t(e^{-\alpha\Delta t} - 1) + \xi_t \sim \mathcal{N}\left(0, \frac{\sigma^2}{2\alpha} ((1 - e^{-2\alpha\Delta t}) + (1 - e^{-\alpha\Delta t})^2)\right),$$

because $\xi_t \sim \mathcal{N}(0, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\Delta t}))$ is independent of X_t and assuming X_t is in its stationary state we have $X_t \sim \mathcal{N}(0, \frac{\sigma^2}{2\alpha})$. The data shows clearly that returns are heavy tailed which we attribute to a non-constant volatility observed in the data and so a normal-distribution does not perfectly fit. We also notice that the maximum-likelihood estimation of the OU-parameters seems to be sensitive to the heavy tails which is why the normal distribution in the figure does not visually fit very well.

Finally, the distribution of the few jumps (24 in total) identified in the NordPool data series between 1995 and 2002 is plotted in Figure 3.6. There are no jumps of a size close to zero as the algorithm is unable to distinguish those from the diffusive mean-reverting process.

3.4 Properties of the stochastic process

The well known properties of an Ornstein-Uhlenbeck (OU) process are given in Appendix B.1. It remains to examine the jump process (Y_t) . For the sake of generality we consider a mean-reverting jump process with a non-zero volatility.

Lemma 3.4.1

Let (Z_t) be a stochastic process satisfying the stochastic differential equation (sde)

$$dZ_t = -\alpha Z_t dt + \sigma dW_t + J_t dN_t,$$




Figure 3.4: Spikes and log-returns in the NordPool market. The upper graph shows the base load price and all spikes identified by the calibration method. Denoting the de-seasonalised log values by $Z_t := \ln S_t - f(t)$, the graph in the middle plots $Z_{t+\Delta t} - Z_t$ and the lower one $Z_{t+\Delta t} - Z_t e^{-\alpha \Delta t}$ for $\alpha := 7$ and $\Delta t = 1/365$.




Figure 3.5: Distribution of log-returns of the NordPool base load. The empirical distribution of $Z_{t+\Delta t} - Z_t$, $Z_t := \ln S_t - f(t)$, is compared to the distribution expected from an OU-process with $\alpha = 7$ and $\sigma = 1.4$, $\Delta t = 1/365$.




Figure 3.6: Histogram of the jump size distribution of the NordPool market. The estimation procedure detects 24 jumps in the period of 1995 – 2002. There are no jumps of a size close to zero (smaller than a few standard deviations of the mean-reverting process X_t) because the method is unable to detect those.

with Z_0 given then the random value at time t is given by

$$Z_t = Z_0 e^{-\alpha t} + \sigma \int_0^t e^{-\alpha(t-s)} dW_s + \sum_{i=1}^{N_t} e^{-\alpha(t-\tau_i)} J_{\tau_i}, \quad (3.4)$$

where τ_i is the random time of the occurrence of the i -th jump.

Proof The random process (Y_t) defined by $Y_t = e^{\alpha t} Z_t$ satisfies $dY_t = e^{\alpha t} (\sigma dW_t + J_t dN_t)$. Integrating this expression and solving for Z_t yields the result. $\square$

3.4.1 The spike process

The last term of 3.4 is a sum over randomly many terms consisting of random jump times in the exponent and random jump sizes and an explicit expression for its distribution does not seem to be known. However, as it turns out, it is possible to obtain an expression for the moment generating function.

Lemma 3.4.2 (Moment generating function of the spike process Y_t)

Let $\{J_1, J_2, \dots\}$ be a series of iid random variables with the moment generating function $\Phi_J(\theta) := \mathbb{E} e^{\theta J}$ being well defined for a subset $\theta \in \Theta \subset \mathbb{C}$. Let furthermore $\{\tau_1, \tau_2, \dots\}$ be the random jump times of a Poisson process (N_t) with intensity λ , then the process (Y_t) with initial condition $Y_0 = 0$ is given by

$$Y_t = \sum_{i=1}^{N_t} e^{-\beta(t-\tau_i)} J_i,$$

and has the moment generating function

$$\Phi_Y(\theta, t) := \mathbb{E} e^{\theta Y_t} = \exp \left(\lambda \int_0^t \Phi_J \left(\theta e^{-\beta s} \right) - 1 ds \right), \quad \forall \theta \in \Theta. \quad (3.5)$$

Furthermore, the first two moments of Y_t are given by

$$\begin{aligned} \mathbb{E}[Y_t] &= \Phi_Y'(0, t) = \frac{\lambda}{\beta} \mathbb{E}[J](1 - e^{-\beta t}), \\ \mathbb{E}[Y_t^2] &= \Phi_Y''(0, t) = \mathbb{E}[Y_t]^2 + \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}), \end{aligned}$$

and in particular we have

$$\text{var}[Y_t] = \frac{\lambda}{2\beta} \mathbb{E}[J^2](1 - e^{-2\beta t}).$$

Proof We prove this lemma by considering the conditional expectation given the first jump, and then by deriving an ode for the moment generating function $\Phi_Y(\theta, t)$ with derivatives in the time variable t .

The mutual independence of J_s , J_t and τ_i (for all $s \neq t$) allows us to write the expectation conditioned on τ_1 as

$$\mathbb{E} \left[e^{\theta Y_t} \mid \tau_1 = s \right] = \mathbb{E} \exp \left(\theta e^{-\beta(t-s)} J \right) \mathbb{E} \exp \left(\theta \sum_{i=2}^{N_t} e^{-\beta(t-\tau_i)} J_{\tau_i} \mid \tau_1 = s \right).$$

Based on properties of the Poisson process, the random sum conditioned on $\tau_1 = s$ has the same distribution as the same sum, starting with the first jump $i = 1$ until time $t - s$ without condition, and hence

$$\begin{aligned} \mathbb{E} \left[e^{\theta Y_t} \mid \tau_1 = s \right] &= \mathbb{E} \exp \left(\theta e^{-\beta(t-s)} J \right) \Phi_Y(\theta, t - s) \\ &= \Phi_J \left(\theta e^{-\beta(t-s)} \right) \Phi_Y(\theta, t - s). \end{aligned} \tag{3.6}$$

Based on these initial considerations we are now able to determine the unconditional expectation. Since the first jump-time is exponentially distributed $\tau_1 \sim \text{Exp}(\lambda)$ we have

$$\Phi_Y(\theta, t) = \mathbb{E} \left[\mathbb{E} \left[e^{\theta Y_t} \mid \tau_1 \right] \right] = \int_0^t \mathbb{E} \left[e^{\theta Y_t} \mid \tau_1 = s \right] \lambda e^{-\lambda s} ds.$$

Inserting the conditional expectation (3.6) yields

$$\begin{aligned} \Phi_Y(\theta, t) &= \int_0^t \Phi_J(\theta e^{-\beta(t-s)}) \Phi_Y(\theta, t - s) \lambda e^{-\lambda s} ds \\ &= \int_0^t \Phi_J(\theta e^{-\beta s}) \Phi_Y(\theta, s) \lambda e^{-\lambda(t-s)} ds, \end{aligned}$$

which is an integral equation but can simply be solved by differentiating with respect to t to give

$$\begin{aligned} \frac{\partial \Phi_Y}{\partial t}(\theta, t) &= \Phi_J(\theta e^{-\beta t}) \Phi_Y(\theta, t) \lambda - \lambda \int_0^t \Phi_J(\theta e^{-\beta s}) \Phi_Y(\theta, s) \lambda e^{-\lambda(t-s)} ds \\ &= \lambda \left(\Phi_J(\theta e^{-\beta t}) - 1 \right) \Phi_Y(\theta, t), \end{aligned}$$

and therefore

$$\Phi_Y(\theta, t) = \exp \left(\lambda \int_0^t \Phi_J(\theta e^{-\beta s}) - 1 ds \right),$$

because $\Phi_Y(\theta, 0) = \mathbb{E} \exp(\theta Z_0) = 1$ as $Z_0 = 0$. □

Example 3.4.3 (Exponentially distributed jump size)

Let the jump size be exponentially distributed with an average jump height μ_J , $J \sim \text{Exp}(1/\mu_J)$, then its moment generating function is $\Phi_J(\theta) = \frac{1}{1-\theta\mu_J}$, $\theta\mu_J < 1$. The integral can then be solved and we obtain

$$\Phi_Y(\theta, t) = \left(\frac{1 - \theta\mu_J e^{-\beta t}}{1 - \theta\mu_J} \right)^{\frac{\lambda}{\beta}}, \quad \theta\mu_J < 1.$$




Figure 3.7: Expectation of the pure spike process $e^{Y_t} = \Phi_Y(1, t)$ with $J \sim \text{Exp}(1/\mu_J)$ and $J \sim \mathcal{N}(\mu_J, \mu_J^2)$, respectively. The parameters coincide with those of the previous example, i.e. $\beta = 200$, $\mu_J = 0.4$ and $\lambda = 4$. For a rough but quick approximation one can use the fact that $\Phi_Y(\theta, t) = 1 + E[Y_t]\theta + \frac{1}{2} E[Y_t^2]\theta^2 + O(\theta^3)$.

Setting $\theta = 1$ gives the expectation value of the exponential spike process and is shown in Figure 3.7, where we compare it to the expectation value of the same process but with a normally distributed jump size. In the long term we have $\Phi_Y(\theta, t) \rightarrow (1 - \theta\mu_J)^{-\lambda/\beta}$ for $(t \rightarrow \infty)$. Furthermore we have for the mean and variance of the spike process Y_t with $Y_0 = 0$:

$$\begin{aligned}\mathbb{E}[Y_t] &= \frac{\lambda\mu_J}{\beta}(1 - e^{-\beta t}), \\ \text{var}[Y_t] &= \frac{\lambda\mu_J^2}{\beta}(1 - e^{-2\beta t}).\end{aligned}$$

Remark 3.4.4 (Asymptotics for $\beta \rightarrow \infty$)

To analyse the behaviour of the moment generating function for large β we make the substitution $u = \theta e^{-\beta s}$ in the integrand to obtain

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \int_{\theta e^{-\beta t}}^{\theta} \frac{\Phi_J(u) - 1}{u} du\right).$$

For a fixed θ the following approximation applies

$$\int_0^{\theta e^{-\beta t}} \frac{\Phi_J(u) - 1}{u} du = \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t}),$$

because $\Phi_J(u) = 1 + \mathbb{E}[J]u + O(u^2)$, $(u \rightarrow 0)$ and so

$$\Phi_{Y_t}(\theta) = \exp\left(\frac{\lambda}{\beta} \left(\int_0^{\theta} \frac{\Phi_J(u) - 1}{u} du - \theta e^{-\beta t} \mathbb{E}[J] + O(e^{-2\beta t})\right)\right).$$

3.4.2 Approximations of the spike process

Ideally we would like to know the density of the spike process at expiry Y_T in an explicit form. What we have so far is the moment generating function from which we could obtain the density by using a Laplace inversion method. However, a simple, yet powerful approximation will lead the way for explicit approximate expressions for the probability density function. The basic idea is that for very high mean reversion rates β and small jump intensities λ only the last jump mainly contributes to the jump distribution. We first introduce the approximation, demonstrate that its moment generating function converges to the moment generating function of the spike process for $\beta \rightarrow \infty$ or $\lambda \rightarrow 0$, then derive its probability density function and finally make further approximations.

Lemma 3.4.5

The truncated spike process, defined by

$$\tilde{Y}_t := \begin{cases} J_{N_t} e^{-\beta(t-\tau_{N_t})} & N_t > 0 \\ 0 & N_t = 0 \end{cases} \quad (3.7)$$

is identically distributed as

$$Z_t := \begin{cases} J_1 e^{-\beta\tau_1} & \tau_1 \leq t \\ 0 & \tau_1 > t \end{cases}$$

Proof This lemma follows directly from the reversibility property of Poisson processes. If (N_t) is a Poisson process then $(-N_{-t})$ is also a Poisson process. As τ_{N_t} is the jump time of the last jump before t this translates to the first jump of the reversed process and hence $t - \tau_{N_t}$ and τ_1 are identically distributed, given $N_t > 0$. If $N_t = 0$ then there has been no jump in $[0, t]$ and the same applies for the reversed process and so this is equivalent to $\tau_1 > t$. The rest follows. $\square$

Lemma 3.4.6 (Moment generating function of the approximated spike process)

Let the moment generating function of J be given by $\Phi_J(\theta)$, $\theta \in \Theta$, then the approximated spike process $\tilde{Y}_t$ as defined above has the moment generating function

$$\Phi_{\tilde{Y}_t}(\theta, t) = 1 + \lambda \int_0^t \left(\Phi_J(\theta e^{-\beta s}) - 1 \right) e^{-\lambda s} ds$$

and the first two moments are given by

$$\begin{aligned} \mathbb{E}[\tilde{Y}_t] &= \frac{\lambda}{\beta + \lambda} \mathbb{E}[J] \left(1 - e^{-(\beta + \lambda)t} \right), \\ \mathbb{E}[\tilde{Y}_t^2] &= \frac{\lambda}{2\beta + \lambda} \mathbb{E}[J^2] \left(1 - e^{-(2\beta + \lambda)t} \right). \end{aligned}$$

Proof According to Lemma 3.4.5 we only need to determine the moment generating function of

$$Z_t := J e^{-\beta\tau} \mathbf{1}_{\tau \leq t}, \quad \tau \sim \text{Exp}(\lambda),$$

where $\mathbf{1}_A$ is the indicator function which yields one if the statement A is true and zero otherwise. Given the jump time τ we have

$$\mathbb{E}[e^{\theta Z_t} | \tau = s] = \Phi_J(\theta e^{-\beta s} \mathbf{1}_{s \leq t}),$$

and so

$$\begin{aligned} \mathbb{E}[e^{\theta Z_t}] &= \mathbb{E}[\mathbb{E}[e^{\theta Z_t} | \tau]] \\ &= \int_0^\infty \Phi_J(\theta e^{-\beta s} \mathbf{1}_{s \leq t}) e^{-\lambda s} ds \\ &= \int_0^t \Phi_J(\theta e^{-\beta s}) e^{-\lambda s} ds + e^{-\lambda t}. \end{aligned}$$

The first two moments are given by $\mathbb{E}[\tilde{Y}_t] = \Phi'_{\tilde{Y}_t}(0)$ and $\mathbb{E}[\tilde{Y}_t^2] = \Phi''_{\tilde{Y}_t}(0)$. $\square$

Remark 3.4.7 (Point-wise convergence of the moment generating functions)

The moment generating function of the truncated spike process converges point-wise to the moment generating function of the spike process for $\lambda \rightarrow 0$ and $\beta \rightarrow \infty$. This can easily be seen for $\lambda \rightarrow 0$. Fix all other parameters and set $g(s; \beta, \theta) := \Phi_J(\theta e^{-\beta s}) - 1$, then

$$\begin{aligned} \Phi_{Y_t}(\theta) &= \exp\left(\lambda \int_0^t g(s; \beta, \theta) ds\right) = 1 + \lambda \int_0^t g(s; \beta, \theta) ds + O(\lambda^2), \\ \Phi_{\tilde{Y}_t}(\theta) &= 1 + \lambda \int_0^t g(s; \beta, \theta) e^{-\lambda s} ds = 1 + \lambda \int_0^t g(s; \beta, \theta)(1 + O(\lambda s)) ds. \end{aligned}$$

To see the convergence for $\beta \rightarrow \infty$ one needs to understand the behaviour of $g(s; \beta, \theta)$ which is equal to one at $s = 0$ and quickly tends to zero for big values of β , because we have

$$g(s; \beta, \theta) = \Phi_J(0) - 1 + \Phi'_J(0)\theta e^{-\beta s} + O(\theta^2 e^{-2\beta s}) = \mathbb{E}[J]\theta e^{-\beta s} + O(\theta^2 e^{-2\beta s}),$$

and so the integral of $g(s; \beta, \theta)$ between 0 and t is almost the same as the same integral between 0 and a very small value, say $\epsilon(\beta)$. In this tiny interval one can approximate $e^{-\lambda s}$ by 1. Finally, linearising the exp function around zero will give the same expressions for Φ_{Y_t} and $\Phi_{\tilde{Y}_t}$. More formally we keep all parameters fixed and split the integral

$$\begin{aligned} \int_0^t g(s; \beta, \theta) ds &= \int_0^{\sqrt{\frac{1}{\beta}}} g(s; \beta, \theta) ds + \int_{\sqrt{\frac{1}{\beta}}}^t g(s; \beta, \theta) ds \\ &= O(\sqrt{1/\beta}) + O(e^{-\sqrt{\beta}}) = O(\sqrt{1/\beta}) \end{aligned}$$




Figure 3.8: Moment generating function of Y_t and $\tilde{Y}_t$, denoted by mgf and approximation, respectively. In the left we use $J \sim \text{Exp}(1/\mu_J)$ and in the right $J \sim \mathcal{N}(\mu_J, \mu_J^2)$. Parameters used are given in Table 3.1 and $t = 1$.

and so we have

$$\Phi_{Y_t}(\theta) = 1 + \int_0^t g(s; \beta, \theta) ds + O(1/\beta).$$

Finally, the difference between the generating functions is

$$\begin{aligned} \Phi_{Y_t}(\theta) - \Phi_{\tilde{Y}_t}(\theta) &= \int_0^t g(s; \beta, \theta)(1 - e^{-\lambda s}) ds + O(1/\beta) \\ &= \int_0^{\sqrt{1/\beta}} g(s; \beta, \theta)(1 - e^{-\lambda s}) ds + O(e^{-\sqrt{\beta}}) + O(1/\beta) \\ &= O(1/\beta) + O(e^{-\sqrt{\beta}}) + O(1/\beta) = O(1/\beta). \end{aligned}$$

Note, this approximation does not necessarily show the highest order of convergence. Two examples of the approximated and exact moment generating function using our standard parameters can be seen in Figure 3.8.

Lemma 3.4.8 (Distribution of the approximated spike process)

Let the jump size distribution be given by a density function f_J , then the approximated spike process $\tilde{Y}_t$ as defined above has the cdf

$$F_{\tilde{Y}_t}(x) = e^{-\lambda t} \mathbf{1}_{x \geq 0} + \int_{-\infty}^x f_{\tilde{Y}_t}(y) dy, \quad t \geq 0,$$

with

$$f_{\tilde{Y}_t}(x) = \frac{\lambda}{\beta} \frac{1}{|x|^{1-\frac{\lambda}{\beta}}} \left| \int_x^{x e^{\beta t}} f_J(y) |y|^{-\frac{\lambda}{\beta}} dy \right|, \quad x \neq 0. \quad (3.8)$$

Proof Based on Lemma 3.4.5 it suffices to determine the distribution of

$$\tilde{Y}_t = JZ \mathbf{1}_{\tau \leq t}, \quad Z := e^{-\beta \tau}, \quad \tau \sim \text{Exp}(\lambda).$$

It follows that Z is the $\frac{\beta}{\lambda}$ th power of a uniformly distributed random variable on $[0, 1]$ and its density is given by

$$f_Z(x) = \frac{\lambda}{\beta} x^{-(1-\frac{\lambda}{\beta})} \mathbf{1}_{x \in [0, 1]}.$$

As $\mathbb{P}(\tau > t) = e^{-\lambda t}$ we obtain the cdf of $Z\mathbb{1}_{\tau \leq t}$ as

$$\begin{aligned} F_{Z\mathbb{1}_{\tau \leq t}}(x) &= e^{-\lambda t} \mathbb{1}_{x \geq 0} + \int_{-\infty}^{\infty} f_{Z\mathbb{1}_{\tau \leq t}}(y) dy, \\ f_{Z\mathbb{1}_{\tau \leq t}}(x) &= \frac{\lambda}{\beta} x^{-(1-\frac{\lambda}{\beta})} \mathbb{1}_{x \in [e^{-\beta t}, 1]}, \end{aligned}$$

and because J and $Z\mathbb{1}_{\tau \leq t}$ are independent, Proposition A.1.1 is applicable and so we obtain the distribution of the product as

$$\begin{aligned} F_{JZ\mathbb{1}_{\tau \leq t}}(c) &= e^{-\lambda t} \mathbb{1}_{c \geq 0} + \int_{-\infty}^c f_{JZ\mathbb{1}_{\tau \leq t}}(x) dx, \\ f_{JZ\mathbb{1}_{\tau \leq t}}(c) &= \int_{-\infty}^{\infty} f_{Z\mathbb{1}_{\tau \leq t}}(c/x) \frac{f_J(x)}{|x|} dx. \end{aligned}$$

With

$$\begin{aligned} f_{Z\mathbb{1}_{\tau \leq t}}(c/x) &= \frac{\lambda}{\beta} \frac{1}{c^{1-\frac{\lambda}{\beta}}} \mathbb{1}_{x \in [c, c e^{\beta t}]} x^{1-\frac{\lambda}{\beta}}, \quad c > 0, \\ f_{Z\mathbb{1}_{\tau \leq t}}(c/x) &= \frac{\lambda}{\beta} \frac{1}{|c|^{1-\frac{\lambda}{\beta}}} \mathbb{1}_{x \in [c e^{\beta t}, c]} |x|^{1-\frac{\lambda}{\beta}}, \quad c < 0, \end{aligned}$$

the desired result follows. $\square$

Example 3.4.9 (Exponential jump size distribution)

Let $J \sim \text{Exp}(1/\mu_J)$ be exponentially distributed. Based on the lemma above the distribution of the truncated spike process $\tilde{Y}_t$ is given by Equation (3.8) and hence

$$f_{\tilde{Y}_t}(x) = \frac{\lambda}{\beta \mu_J^{\frac{\lambda}{\beta}}} \frac{\Gamma(1 - \frac{\lambda}{\beta}, \frac{x}{\mu_J}) - \Gamma(1 - \frac{\lambda}{\beta}, \frac{x e^{\beta t}}{\mu_J})}{x^{1-\frac{\lambda}{\beta}}}, \quad x > 0,$$

where $\Gamma(a, x)$ is the incomplete Gamma function defined by

$$\Gamma(a, x) := \int_x^{\infty} t^{a-1} e^{-t} dt.$$

For $\frac{\lambda}{\beta}$ very small and x big we use the approximation

$$\Gamma(1 + \Delta a, x) = \int_x^{\infty} t^{\Delta a} e^{-t} dt \approx e^{-x}, \quad x > 0, \quad \Delta a \rightarrow 0.$$

Inserting this approximation into the density function and re-normalising the factor to satisfy $\int_0^{\infty} f_{\tilde{Y}_t}(x) dx = 1 - e^{-\lambda t}$ we get

$$f_{\tilde{Y}_t}(x) \approx \frac{1}{\Gamma(\frac{\lambda}{\beta}) \mu_J^{\frac{\lambda}{\beta}}} \frac{e^{-\frac{x}{\mu_J}} - e^{-\frac{x e^{\beta t}}{\mu_J}}}{x^{1-\frac{\lambda}{\beta}}}, \quad x > 0, \quad (3.9)$$




Figure 3.9: Distribution of the spike process (Y_t) at T with a jump size of $J \sim \text{Exp}(1/\mu_J)$. We use approximation (3.9) and compare it with the exact density as produced by a Monte-Carlo simulation. The parameters used are given in Table 3.1.

and so the stationary distribution is similar to a Gamma distribution, i.e.

$$f_{\tilde{Y}_t}(x) \approx \frac{1}{\Gamma(\frac{\lambda}{\beta})\mu_J^{\frac{\lambda}{\beta}}} x^{\frac{\lambda}{\beta}-1} e^{-\frac{x}{\mu_J}} \mathbb{1}_{x>0}, \quad t \rightarrow \infty.$$

The approximation fits very well the exact density for typical market parameters as can be seen in Figure 3.9.

Remark 3.4.10 (Further approximations of the tails of $\tilde{Y}_T$)

For many jump size distributions one will not be able to obtain an explicit expression for the pdf (3.8) due to the nature of the integral. Further approximations can be made to allow for an explicit formula. Assuming $|x|$ is big enough and $\frac{\lambda}{\beta}$ is close to zero we can approximate the term $|y|^{-\frac{\lambda}{\beta}}$ in the integral of Equation (3.8) by the constant it assumes on the lower bound x and so we obtain

$$f_{\tilde{Y}_t}(x) \approx \frac{\lambda}{\beta} \frac{F_J(x e^{\beta t}) - F_J(x)}{|x|}, \quad (x \rightarrow \infty). \quad (3.10)$$

Note, the area under the graph of this approximation is not necessarily $1 - e^{-\lambda t}$ which is required in order to make $F_{\tilde{Y}_t}$ a distribution function, as $P(\tilde{Y}_t = 0) = e^{-\lambda t}$.

The integration of (3.8) goes from x to $x e^{\beta t}$ and so for values of $e^{\beta t}$ close to one we can even approximate the entire integrand of Equation (3.8) by the constant it assumes at the point x . Re-normalising to make $F_{\tilde{Y}_t}$ a distribution results in

$$f_{\tilde{Y}_t}(x) \approx (1 - e^{-\lambda t}) f_J(x), \quad (\beta t \rightarrow 0). \quad (3.11)$$

Example 3.4.11 (Normal jump size distribution)

Assuming $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$ is normally distributed we can use Equation (3.10) or (3.11) to approximate the density of the spike process Y_t . Figure 3.10 shows the quality of the approximations for some standard parameters. As expected they perform quite well for big absolute values of Y .

3.4.3 The combined process

Having examined the properties of the spike process Y_t we can conclude properties of the sum $X_t + Y_t$ and consequently $S_t = \exp(f(t) + X_t + Y_t)$. First we derive the moment generating function of $f(t) + X_t + Y_t$ and then the density function of $X_t + Y_t$ for which we also give approximations.

Corollary 3.4.12

Let the spot process (S_t) be defined by (3.3) and let (Z_t) be its natural logarithm, i.e. $Z_t := \ln S_t = f(t) + X_t + Y_t$. The moment generating function of Z_t is then

$$\mathbb{E} e^{\theta Z_t} = \exp \left(\theta f(t) + \theta X_0 e^{-\alpha t} + \theta^2 \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha t}) + \theta Y_0 e^{-\beta t} + \lambda \int_0^t \Phi_J \left(\theta e^{-\beta s} \right) - 1 \, ds \right). \quad (3.12)$$

Proof The processes X and Y are independent so the expectation of the product is the product of the expectations. The moment generating functions of X and Y are given by Equation B.6 and Lemma 3.4.2 which yield the result. $\square$

The expectation value of the spot process S immediately follows by setting $\theta = 1$.

Corollary 3.4.13 (Expectation of the spot price process S_T)

Let the spot price process (S_t) be defined by (3.3), then its expectation value is

$$\mathbb{E}[S_T | X_t, Y_t] = \exp \left(f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J \left(e^{-\beta s} \right) - 1 \, ds \right).$$

The structure of the expectation value is very clear. It consists of four terms: the seasonal component, initial terms, a contribution from the volatility and a jump term.

Example 3.4.14 (Exponentially distributed jump size)

If the jump size is exponentially distributed, $J \sim \text{Exp}(1/\mu_J)$, then the expectation value is

$$\mathbb{E} S_t = \exp \left(f(t) + X_0 e^{-\alpha t} + Y_0 e^{-\beta t} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha t}) \right) \left(\frac{1 - \mu_J e^{-\alpha t}}{1 - \mu_J} \right)^{\frac{\lambda}{\alpha}}, \quad \mu_J < 1.$$

An example plot is shown in Figure 3.11. In the long term we have

$$\mathbb{E} S_t \rightarrow e^{f(t) + \sigma^2/4\alpha} (1 - \mu_J)^{-\frac{\lambda}{\alpha}}, \quad (t \rightarrow \infty).$$




Figure 3.10: Distribution of the spike process (Y_t) at T with a jump size of $J \sim \mathcal{N}(1/\mu_J, 1/\mu_J^2)$. We use approximation (3.10) for $T = 1$ and (3.11) for $T = 1/365$ and compare it with the exact density as produced by a Monte-Carlo simulation of 100 million runs. The parameters used are given in Table 3.1.




Figure 3.11: Expectation of the spot price process with parameters of Table 3.1, but with $X_0 = 0$ and $Y_0 = 0$.

Anticipating the application of pricing plain vanilla call or put options, we state a result from [Duffie et al., 2000]. See also Appendix C for a brief account.

Remark 3.4.15

Let Z be a Markov process satisfying the regularity conditions of Definition C.2.1. The expectation value, defined by

$$G_{\theta,c}(K, T, z_0) := \mathbb{E} \left[e^{\theta Z_T} \mathbf{1}_{cZ_T \leq K} \right],$$

is then given by

$$G_{\theta,c}(K, T, z_0) = \frac{1}{2} \Phi_Z(\theta, T) - \frac{1}{\pi} \int_0^\infty \frac{1}{\nu} \Im \left(\Phi_Z(\theta + i\nu c, T) e^{-i\nu \theta} \right) d\nu,$$

where Φ_Z denotes the complex valued moment generating function and $\Im(z)$ the imaginary part of a complex number $z \in \mathbb{C}$. If the function $G_{\theta,c}$ is known, it is easy to determine the expectation value of a call option payoff. Let $Z_t := \ln S_t$ be the logarithm of S_t , then we have

$$\mathbb{E}[(S_T - K)^+] = \mathbb{E}[S_T \mathbf{1}_{S_T \geq K}] - \mathbb{E}[K \mathbf{1}_{S_T \geq K}] = G_{1,-1}(-\ln K, T, z_0) - K G_{0,-1}(-\ln K, T, z_0).$$

3.4.4 Approximations of the combined process

We use approximations of the density of Y_t to obtain approximations to the density of the sum of both processes $X_t + Y_t$.

Corollary 3.4.16 (Distribution of $X_t + \tilde{Y}_t$)

Let X_t be the Ornstein-Uhlenbeck process defined by our model (3.3) with $X_0 = 0$ and $\tilde{Y}_t$ the approximation to the spike process defined by (3.7) where the jump size J has a density f_J . Then $X_t + \tilde{Y}_t$ has a density and it is given by

$$\begin{aligned} f_{X_t + \tilde{Y}_t}(c) &= e^{-\lambda t} f_{X_t}(c) + \int_{-\infty}^{\infty} f_{\tilde{Y}_t}(c - x) f_{X_t}(x) dx, \\ f_{X_t}(x) &= \frac{1}{\sqrt{2\pi}\sigma_X} \exp\left(-\frac{x^2}{2\sigma_X^2}\right), \\ f_{\tilde{Y}_t}(x) &= \frac{\lambda}{\beta} \frac{1}{|x|^{1-\frac{\lambda}{\beta}}} \left| \int_x^{x e^{\beta t}} f_J(y) |y|^{-\frac{\lambda}{\beta}} dy \right|, \quad x \neq 0, \end{aligned}$$

with $\sigma_X^2 = \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha t})$ being the variance of X_t .

Proof This is a direct result of Proposition A.1.2 and Lemma 3.4.8 and the fact that $X_t \sim \mathcal{N}(0, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha t}))$, see Equation (B.4). $\square$

Remark 3.4.17

Corollary 3.4.16 is the key to an approximation to the density of our spot price process $S_t = \exp(f(t) + X_t + Y_t)$. Note that $\tilde{Y}_t$ is an approximation to the spike process by considering the last jump only.

For non-zero initial conditions and using the notation of Corollary 3.4.16, the density of $\ln S_t$ is then approximately given by

$$f_{\ln S_t}(x) \approx f_{X_t + \tilde{Y}_t}\left(x - f(t) - X_0 e^{-\alpha t} - Y_0 e^{-\beta t}\right).$$

Remark 3.4.18

For very short time horizons so that $\beta t \rightarrow 0$ we can use Approximation (3.11) for the spike process to obtain

$$\begin{aligned} f_{X_t + \tilde{Y}_t}(c) &\approx e^{-\lambda t} f_{X_t}(c) + (1 - e^{-\lambda t}) \int_{-\infty}^{\infty} f_J(c - x) f_{X_t}(x) dx, \\ &= e^{-\lambda t} f_{X_t}(c) + (1 - e^{-\lambda t}) f_{X_t + J}(c), \\ f_{X_t}(x) &= \frac{1}{\sqrt{2\pi}\sigma_X} \exp\left(-\frac{x^2}{2\sigma_X^2}\right), \end{aligned}$$

with $\sigma_X^2 = \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha t})$ being the variance of X_t . This is a very explicit expression. The only function to be determined is $f_{X_t + J}$ which is the density of the sum of a normally distributed random variable and the jump size distribution J .




Figure 3.12: Distribution of $X_t + Y_t$ for $t = 1/365$ and $J \sim \text{Exp}(1/\mu_J)$. Parameters are based on Table 3.1.

Example 3.4.19 (Exponential jump size distribution)

Let $J \sim \text{Exp}(1/\mu_J)$ and let β be big and βt be very small, then according to Remark 3.4.18 and Lemma A.1.3 we get

$$\begin{aligned} f_{X_t+Y_t}(x) &\approx e^{-\lambda t} f_{X_t}(x) + (1 - e^{-\lambda t}) f_{X_t+J}(x), \\ f_{X_t}(x) &= \frac{1}{\sqrt{2\pi}\sigma_X} \exp\left(-\frac{x^2}{2\sigma_X^2}\right), \\ f_{X_t+J}(x) &= \frac{1}{\mu_J} \exp\left(\frac{\sigma_X^2}{2\mu_J^2} - \frac{x}{\mu_J}\right) N\left(\frac{x}{\sigma_X} - \frac{\sigma_X}{\mu_J}\right), \end{aligned}$$

which is illustrated in Figure 3.12. The function $N(x)$ denotes the cumulative distribution of a standard $\mathcal{N}(0, 1)$ normally distributed random variable.

Example 3.4.20 (Normal jump size distribution)

Let $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$ and β be large and βt very small, then according to Remark 3.4.18 we get

$$\begin{aligned} f_{X_t+Y_t}(x) &\approx e^{-\lambda t} f_{X_t}(x) + (1 - e^{-\lambda t}) f_{X_t+J}(x), \\ f_{X_t}(x) &= \frac{1}{\sqrt{2\pi}\sigma_X} \exp\left(-\frac{x^2}{2\sigma_X^2}\right), \\ f_{X_t+J}(x) &= \frac{1}{\sqrt{2\pi(\sigma_X^2 + \sigma_J^2)}} \exp\left(-\frac{(x - \mu_J)^2}{2(\sigma_X^2 + \sigma_J^2)}\right), \end{aligned}$$

which is illustrated in Figure 3.13.




Figure 3.13: Distribution of $X_t + Y_t$ for $t = 1/365$ and $J \sim \mathcal{N}(\mu_J, \mu_J^2)$. Parameters are based on Table 3.1.

Remark 3.4.21 (Longer time horizons)

For longer time horizons t it does not seem to be immediately clear which approximation might yield a simple expression as it is the case for very short time horizons $t \rightarrow 0$. A very rough approximation would be to assume the distribution is close to a normal distribution but where the mean and variance are matched to the values of the combined process $X + Y$. This works fairly well for option pricing as will be demonstrated later in this thesis but does not describe the tail behaviour sufficiently well. Another way would be to assume the same form as before, see Remark 3.4.18, i.e.

$$f_{X_t + \tilde{Y}_t}(x) \approx pf_{X_t}(x) + (1 - p)f_{X_t + J}(x), \quad (3.13)$$

with p yet to be determined. The reasoning behind this approach is as follows. Assuming we are in a stationary situation, i.e. $t \rightarrow \infty$ then the spike process $\tilde{Y}_t$ is approximately given by $\tilde{Y}_t \approx e^{-\beta\tau} J$ and the distribution of $e^{-\beta\tau}$ is, according to the proof of Lemma 3.4.8, the $\frac{\beta}{\lambda}$ th power of a uniform $[0, 1]$ random variable and hence its weight is concentrated around 0. So the idea is to approximate it by

$$f_{e^{-\beta\tau}}(x) \approx \frac{\lambda}{\beta} x^{-(1-\frac{\lambda}{\beta})} \mathbb{1}_{x \in [\delta, 1]},$$

$$\mathbb{P}(e^{-\beta\tau} = 0) = p = \delta^{\frac{\lambda}{\beta}}.$$

This approximation is exact for $e^{-\beta\tau} \mathbb{1}_{\tau \leq t}$ if we chose $\delta = e^{-\beta t}$. However, if δ is chosen too small, approximation (3.11) is not applicable as it is based on setting the integrand of (3.8) to a constant where the integral goes from a value x to $\frac{x}{\delta}$. There, the closer δ is to 1

the better, so there is a tradeoff between the two approximations but for the moment we assume there exists a value δ so that both approximations are fairly accurate. In practice we determine the parameter $p = \delta^{\frac{\lambda}{\beta}}$ so that it matches the first moment, i.e.

$$p \mathbb{E}[X_t] + (1 - p) \mathbb{E}[X_t + J] = \mathbb{E}[X_t + Y_t],$$

and so we get

$$p = 1 - \frac{\mathbb{E}[Y_t]}{\mathbb{E}[J]} = 1 - \frac{\lambda}{\beta}(1 - e^{-\beta t}).$$

Alternatively one could also match the second moment and obtain

$$p = 1 - \frac{\mathbb{E}[Y_t^2] + 2 \mathbb{E}[X_t] \mathbb{E}[Y_t]}{\mathbb{E}[J^2] + 2 \mathbb{E}[X_t] \mathbb{E}[J]}.$$

Matching the second moment seems to give a slightly better tail approximation. Figure 3.14 shows the approximation for exponentially and normally distributed jump sizes. For the parameters used we get $p = 0.98$ and $p = 0.9898$ for matching the first and second moments, respectively. This corresponds to $\delta \approx 0.364$ and $\delta \approx 0.599$, respectively. In the graphs we choose p to match the second moment.

3.4.5 Conditional expectations

In reality we are faced with the situation that only the processes $(X_t + Y_t)$ is observable but not (X_t) and (Y_t) individually. This makes it difficult to determine the initial conditions X_t and Y_t . If a spot price history till that time is known one could use a filtering method to split the sum $(X_t + Y_t)$ into its two components with some little error. One can view $(X_t + Y_t)$ as the spike process (Y_t) being obscured by some noise (X_t) . More heuristically one could simply say $Y_t = 0$ given the last spike was sufficiently far away, say if $e^{-\beta \Delta t} \leq \epsilon$. Otherwise we are in the spike regime and need to know the last observation $X_s + Y_s$ at the time s just before the spike. There we also assume $Y_s = 0$. Knowing X_s we can make a good estimation of the jump size and so the rest follows.

In this subsection we only consider the case where no history is available and then the best estimate will be to use conditional expectations, i.e. if we know $X_0 + Y_0 = c$ one could say $X_0 = \mathbb{E}[\bar{X} | \bar{X} + \bar{Y} = c]$ where $\bar{X}$ and $\bar{Y}$ are stationary distributions of (X_t) and (Y_t) , respectively.

Proposition 3.4.22

Let X have a density f_X and let J have a density f_J and let Y be defined by approximation (3.11), i.e.

$$\begin{aligned} F_Y(x) &= \int_{-\infty}^x f_Y(y) dy + e^{-\lambda t} \mathbb{1}_{x \geq 0}, \\ f_Y(x) &= (1 - e^{-\lambda t}) f_J(x), \end{aligned}$$




Figure 3.14: Distribution of $X_t + Y_t$ for $t = 1$ and $J \sim \text{Exp}(1/\mu_J)$ and $J \sim \mathcal{N}(\mu_J, \mu_J^2)$, respectively. Parameters are based on Table 3.1.

then we have

$$\begin{aligned}\mathbb{E}[X|X+Y=c] &= \frac{\mathbb{E}[X|X+J=c] + ch(c)}{1+h(c)}, \\ h(c) &:= \frac{e^{-\lambda t}}{1-e^{-\lambda t}} \frac{f_X(c)}{f_{X+J}(c)}.\end{aligned}$$

Proof The conditional expectation is given by Lemma A.2.5 which says

$$\mathbb{E}[X|X+Y=c] = \frac{\int_{-\infty}^{\infty} x f_X(x) f_Y(c-x) dx + p c f_X(c)}{f_{X+Y}(c)},$$

where we have $p = e^{-\lambda t}$ and, according to Remark 3.4.18, it is given by

$$f_{X+Y}(c) = p f_X(c) + (1-p) f_{X+J}(c).$$

Dividing by $(1-p) f_{X+J}(c)$ yields

$$\mathbb{E}[X|X+Y=c] = \frac{\frac{\int_{-\infty}^{\infty} x f_X(x) f_J(c-x) dx}{f_{X+J}(c)} + \frac{p}{1-p} \frac{c f_X(c)}{f_{X+J}(c)}}{1 + \frac{p}{1-p} \frac{f_X(c)}{f_{X+J}(c)}}.$$

□

This Proposition gives some good results for the conditional expectation $\mathbb{E}[X_t|X_t+Y_t=c]$ based on approximation (3.11) which is only valid for small t .

Example 3.4.23 (Exponential jump size distribution)

Let $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$ be normally and $J \sim \text{Exp}(1/\mu_J)$ be exponentially distributed. According to Lemma A.2.4 and Lemma A.1.3 we have

$$\begin{aligned}\mathbb{E}[X|X+J=c] &= \mu_X + \frac{\sigma_X^2}{\mu_J} - \sigma_X \frac{\varphi\left(\frac{c-\mu_X-\sigma_X^2/\mu_J}{\sigma_X}\right)}{\text{N}\left(\frac{c-\mu_X-\sigma_X^2/\mu_J}{\sigma_X}\right)}, \\ f_{X+J}(x) &= \frac{1}{\mu_J} \exp\left(\frac{\sigma_X^2}{2\mu_J^2} + \frac{\mu_X - x}{\mu_J}\right) \text{N}\left(\frac{x - \mu_X}{\sigma_X} - \frac{\sigma_X}{\mu_J}\right),\end{aligned}$$

where N and φ are the distribution and density of a $\mathcal{N}(0, 1)$ random variable, respectively, i.e.

$$\varphi(x) := \frac{1}{\sqrt{2\pi}} e^{-x^2/2}, \quad \text{N}(x) := \int_{-\infty}^x \varphi(y) dy.$$

Figure 3.15 shows a plot of the function $\mathbb{E}[X_t|X_t+J=c]$ which cannot be seen as an approximation to $\mathbb{E}[X_t|X_t+Y_t=c]$ but nevertheless seems to provide the right asymptotic behaviour for $|c| \rightarrow \infty$. However, for short time horizons and using the approximation of Proposition 3.4.22 we obtain a very good fit to the function obtained from a Monte-Carlo simulation which is shown in Figure 3.16.




Figure 3.15: Conditional expectations given the knowledge of the sum $X_t + Y_t = c$. An approximation to the function $g(c) := \mathbb{E}[X_t | X_t + Y_t = c]$ is plotted and compared to the result of a Monte-Carlo simulation with 200 million paths. Here we use $\mathbb{E}[X_t | X_t + J = c]$ as the approximation.




Figure 3.16: Conditional expectations given the knowledge of the sum $X_t + Y_t = c$. An approximation to the function $g(c) := \mathbb{E}[X_t | X_t + Y_t = c]$ is plotted and compared to the result of a Monte-Carlo simulation with 200 million paths. Parameters are given in Table 3.1, the jump size is exponentially distributed $J \sim \text{Exp}(1/\mu_J)$ and $t = 1/365$. For these parameters we obtain standard deviations of $\sigma_X \approx 0.0726$ and $\sigma_Y \approx 0.0462$.




Figure 3.17: Conditional expectations given the knowledge of the sum $X_t + Y_t$. An approximation to the function $g(c) := \mathbb{E}[X_t | X_t + Y_t = c]$ is plotted and compared to the result of a Monte-Carlo simulation with 200 million paths. Parameters are given in Table 3.1, the jump size is normally distributed $J \sim \mathcal{N}(\mu_J, \mu_J^2)$ and $t = 1/365$. For these parameters we obtain standard deviations of $\sigma_X \approx 0.0726$ and $\sigma_Y \approx 0.0462$.

Example 3.4.24 (Normal jump size distribution)

Let $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$ and $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$ be normally distributed. According to Lemma A.2.3 we have

$$\mathbb{E}[X | X + J = c] = \frac{\sigma_X^2}{\sigma_X^2 + \sigma_J^2} \left(c - \left(\mu_J - \frac{\sigma_J^2}{\sigma_X^2} \mu_X \right) \right).$$

Figure 3.17 shows how well the approximation of Proposition 3.4.22 works for some typical market parameters and one day to maturity.

Example 3.4.25 (Stationary distribution)

Based on Remark 3.4.21 we can also approximate the conditional expectation for big t or even $t \rightarrow \infty$, i.e. for the stationary case. Because the approximation is only very rough we cannot expect a perfect fit, but at least the shape of the curves is close to the exact functions as shown in Figure 3.18.

Another important issue is to determine the density of the sum of both processes $X_t + Y_t$ given the initial condition $X_0 + Y_0 = c$. This problem turns up when faced with expectations of the form $\mathbb{E}[g(S_t) | S_0 = s]$ where S_t is some function of $X_t + Y_t$. However, this problem is only well defined once we make some assumption of the distribution of X_0 and Y_0 .

Here we compare two approaches, one is to assign X_0 and Y_0 their conditional expectations




Figure 3.18: Conditional expectations given the knowledge of the sum $X_t + Y_t$. An approximation to the function $g(c) := \mathbb{E}[X_t | X_t + Y_t = c]$ is plotted and compared to the result of a Monte-Carlo simulation with 200 million paths. Parameters are given in Table 3.1, the jump size is exponentially and normally distributed, respectively and $t = 1$. For these parameters we obtain standard deviations of $\sigma_X \approx 0.374$ and $\sigma_Y \approx 0.0566$.

based on the knowledge of the sum, i.e.

$$X_0 = c_1 := \mathbb{E}[\bar{X} | \bar{X} + \bar{Y} = c], \quad Y_0 = c_2 := c - X_0,$$

where $\bar{X}$ and $\bar{Y}$ are the independent random variables representing the stationary distribution of (X_t) and (Y_t) . Then the conditional distribution follows directly from the conditional distributions of X_t and Y_t given X_0 and Y_0 :

$$f_{X_t+Y_t|X_0+Y_0=c}(x) = f_{X_t+Y_t|X_0=c_1, Y_0=c_2}(x).$$

Alternatively, one could say X_0 and Y_0 assume the stationary distributions. A comparison of the two approaches is shown in Figure 3.19 where we use the Monte-Carlo method to obtain the densities. As can be seen, both methods result in very similar densities for some values of c but are completely different for others. Generally, the distributions for values of c within about two standard deviations of the process $\bar{X}$ are very similar and for values above about three or four standard deviations are considerably different, i.e. the second approach results in much flatter and heavier tailed distributions. This result is intuitively clear, because if we know $\bar{X} + \bar{Y}$ is below two standard deviations and given $\bar{Y}$ is close to zero most of the time it is very likely that the sum is purely achieved by $\bar{X}$ and $\bar{Y}$ is almost zero. However, if $\bar{X} + \bar{Y}$ is greater than three standard deviations of $\bar{X}$ then it is unlikely $\bar{X}$ could have achieved that value alone and we know a jump has very likely occurred and then it is very hard to give an accurate estimation for $\bar{X}$ as basically we can expect it to be in a range of a few standard deviations of $\bar{X}$ itself. Higher uncertainty in the initial condition will result in a much broader distribution, because if Y_0 is bigger then $X_t + Y_t$ will revert much faster to zero than if Y_0 was smaller and X_0 bigger.

To determine the distribution we denote $Z_t := X_t + Y_t$ and let $f_{Z_t}(x)$ be the density of Z_t given $X_0 = 0$ and $Y_0 = 0$. Recall that for $X_0 = x_0$ and $Y_0 = y_0$ the density is given by $f_{Z_t}(x - x_0 e^{-\alpha t} - y_0 e^{-\beta t})$ and so we state

$$f_{X_t+Y_t|X_0+Y_0=c}(x) = \int_{-\infty}^{\infty} f_{X_t+Y_t}(x - x_0 e^{-\alpha t} - (c - x_0) e^{-\beta t}) f_{\bar{X}|\bar{X}+\bar{Y}=c}(x_0) dx_0.$$

In order to evaluate this expression further we use above approximations which are based on $f_{\tilde{Y}_t}(x) \approx (1 - p)f_J(x)$ and $P(\tilde{Y}_t = 0) = p$:

$$\begin{aligned} f_{X_t+Y_t}(x) &\approx e^{-\lambda t} f_{X_t}(x) + (1 - e^{-\lambda t}) f_J(x), \\ f_{\bar{X}+\bar{Y}} &\approx p f_{\bar{X}}(x) + (1 - p) f_J(x), \\ f_{\bar{X}|\bar{X}+\bar{Y}=c}(x) &= \frac{f_{\bar{X}}(x) f_{\bar{Y}}(c - x)}{f_{\bar{X}+\bar{Y}}(c)}, \\ P(\bar{X} = c | X + Y = c) &\approx p \frac{f_{\bar{X}}(c)}{f_{\bar{X}+\bar{Y}}(c)}, \end{aligned}$$




Figure 3.19: Density of $X_t + Y_t$ given $X_0 + Y_0 = c$ assuming stationary distribution for X_0 and Y_0 . This is plotted by the red curve using Monte-Carlo. The other curves show the distribution obtained if one assumes a particular value for X_0 . Parameters are given in Table 3.1, the jump size is exponentially distributed, $t = 1/365$, and $c = 0.5$ in the left and $c = 1.5$ in the right graph. For the left graph, we have $\mathbb{E}[\bar{X} | \bar{X} + \bar{Y} = 0.5] \approx 0.494$ and the density for $X_0 = 0$ only shows the sensitivity to the choice of the initial condition. For the right graph, we have $\mathbb{E}[\bar{X} | \bar{X} + \bar{Y} = 1.5] \approx 0.793$ but the approximation used in Figure 3.18 yields 0.658.

where p is determined to match one of the moments of the stationary distribution as described in Remark 3.4.21. With that we obtain

$$f_{X_t+Y_t|X_0+Y_0=c}(x) \approx (1-p) \int_{-\infty}^{\infty} f_{X_t+Y_t}(x - x_0 e^{-\alpha t} - (c - x_0) e^{-\beta t}) \frac{f_{\bar{X}}(x_0) f_J(c - x_0)}{f_{\bar{X}+\bar{Y}}(c)} dx_0 \\ + p f_{X_t+Y_t}(x - c e^{-\alpha t}) \frac{f_{\bar{X}}(c)}{f_{\bar{X}+\bar{Y}}(c)}.$$

The second term of the result is the density of $X_t + Y_t$ given $Y_0 = 0$ and $X_0 = c$ scaled by some factor and is actually visible in Figure 3.19 where $c = 1.5$. For a normally distributed jump size J the integral can be solved analytically but for an exponential distribution we fail to solve the integral.

Chapter 4

Option pricing

The electricity market with the model presented in the previous section is obviously incomplete. Not only are we faced with a non-hedgeable jump risk but also can we not use the underlying process (S_t) to hedge derivatives due to inefficiencies in storing electricity. We give a short account of utility based pricing approaches but focus mainly on arbitrage pricing methods. Even though the underlying cannot be used to hedge, derivatives have to satisfy certain consistency conditions. Otherwise arbitrage opportunities could be exploited by setting up a portfolio consisting of different derivatives.

4.1 Utility indifference pricing

Utility theory is based on the belief that each individual agent possesses a utility function and if a decision is to be made the agent acts in such a way which maximises expected utility. For the purpose of this section we do not assume any particular stochastic model for the underlying (S_t) but keep the discussion general.

Say $U : \mathbb{R} \rightarrow \mathbb{R}$ is a utility function which assigns each value of wealth x a number which can be interpreted as the happiness of an agent given they possess x amount of cash. We will only consider monotonic increasing functions U which are called consistent utility functions. Furthermore, we call x the certainty equivalent of a random payoff X if x has the same utility as the expected utility of the random payoff, i.e.

$$U(x) = \mathbb{E}[U(X)].$$

The utility function is called risk averse if the certainty equivalent is less than the expected value of a random payment, $x < \mathbb{E}[X]$, which is the case for all concave utility functions.

To quantify the strength of risk aversion we consider a random payoff X with $\mathbb{E}[X] = 0$ and $\text{var}[X] = 1$ and see by how much the certainty equivalent reduces, i.e.

$$U(x + \delta) = \mathbb{E}[U(x + \epsilon X)].$$

The bigger the ratio $-\delta/\epsilon$ the more risk averse the agent is. Using Taylor series expansion we see

$$U(x) + \delta U'(x) \approx U(x) + \epsilon U'(x) \mathbb{E}[X] + \frac{1}{2} \epsilon^2 U''(x) \mathbb{E}[X^2],$$

and so one defines the risk aversion of U by $-U''/U'$.

For pricing purposes we only consider the exponential utility here, i.e.

$$U(x) = -\gamma e^{-\gamma x},$$

which has a constant risk aversion $-U''/U' = \gamma$ and has the nice property that wealth factors out, i.e. $U(x + y) = U(x) e^{-\gamma y}$.

4.1.1 Pricing without the possibility of hedging

Now, let our market consist of the underlying price process (S_t) and some derivative paying out $g(S_T)$ at maturity T , where the spot price process cannot be used for hedging purposes. Assuming our agent gets the offer to receive c units of the derivative for the price p . The offer will then be accepted by the agent if it increases expected utility or the agent will be indifferent if expected utility remains the same. Such a price p is called the utility indifference price. Assuming initial wealth and the existence of a bank account paying a continuous compounded interest rate r , the utility indifference price p depending on the number of derivatives c is defined by

$$U(e^{rT} x) = \mathbb{E}[U(e^{rT}(x - p(c)) + cg(S_T))],$$

which yields for exponential utility

$$p(c) = -\frac{1}{\gamma} e^{-rT} \ln \left(\mathbb{E} \left[e^{-\gamma c g(S_T)} \right] \right). \quad (4.1)$$

This is the fundamental pricing equation for exponential utility if one cannot hedge with the underlying. Prices are generally non-linear in the quantity c and worse, might not even exist for unbounded functions g and negative values of c , i.e. when the derivative is sold by the agent. In fact, the expectation value does not exist for a call option if S_T is log-normally distributed as it is the case for an exponential Ornstein-Uhlenbeck process. Hence all agents with exponential utility will find it too risky to sell call options on that

underlying. As selling a put option has only a limited downside risk – in the worst case the seller will have to pay the strike price K – it can be priced and is shown in the following example.

Example 4.1.1

Consider an agent with exponential utility and $\gamma = \frac{1}{10}$. To illustrate the risk aversion assume there is a choice of receiving either a guaranteed cash amount of $x = 100$ or $x = d$ with probability q and $x = u$ with probability $(1 - q)$. For $d = 99$ and $u = 101$ the agent would be indifferent only if $q \approx 0.48$, for $d = 90$ and $u = 110$ it will have to be $q \approx 0.27$ and for $d = 0$ and $u = 200$ even $q \approx 0.000045$. To price an option we will have to make further assumptions on the underlying price process (S_t) , say it is an exponential Ornstein-Uhlenbeck process with

$$\begin{aligned} S_t &= S_0 \exp(X_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \end{aligned}$$

and $S_0 = 100$, $\alpha = 7$, $\sigma = 1.4$. Assume further T is big enough for X_T to have the stationary distribution $X_T \sim \mathcal{N}(0, \frac{\sigma^2}{2\alpha})$. The price of a put option is based on Equation (4.1) and given by

$$p(c) = -\frac{1}{\gamma} \ln \left(\int_{-\infty}^{\infty} \exp(-\gamma c(K - S_0 e^x)^+) f_{X_T}(x) dx \right),$$

with f_{X_T} being the pdf of X_T . The unit price $p(c)/c$ for $K = 100$ is shown in Figure 4.1. It is not a coincidence that the expected payoff $\mathbb{E}[(K - S_T)^+]$ is equal to the marginal price $p(c)/c$, $c \rightarrow 0$.

Despite the difficulties with utility indifference pricing, not to mention the problem of agreeing upon one particular utility function, there is a link to arbitrage pricing. As seen in Figure 4.1 the marginal price coincides with the expected value of the payoff which is equal to the price under the real world measure. This is a general result if the payoff function is bounded. We make this plausible by simple linearisation of the utility function. The indifference price $p(c)$ is given by

$$U(x) = \mathbb{E}[U(x - p(c) + e^{-rT} cg(S_T))],$$

and for small c we linearise the utility function at x to obtain

$$U(x) = U(x) + U'(x) (-p(c) + e^{-rT} c \mathbb{E}[g(S_T)]) + O(\dots),$$

indicating that

$$p'(0) = \lim_{c \rightarrow 0} \frac{p(c)}{c} = e^{-rT} \mathbb{E}[g(S_T)].$$

Expectations are all taken under the real world probability measure.




Figure 4.1: Utility indifference price of a put option. The price per option is plotted with respect to the amount of options sold or bought. Exponential utility with $\gamma = \frac{1}{10}$ and a spot price process with log-normally distributed end value S_T is assumed. All parameters are given in Example 4.1.1.

4.1.2 Pricing and hedging with a correlated asset

If a correlated asset is available in the market which can be used to hedge the claim written on the non-tradable or non-storable asset, the solution is not as obvious as in the previous subsection. Let P_t and S_t be the traded and non-traded asset, respectively. Furthermore, let $Z_t^{[\theta]}$ be the value of a self-financing portfolio consisting of P_t and a money market account, where θ_t indicates the cash amount of the portfolio held in the asset P_t . One then defines a value function as the maximum expected utility given initial wealth z , $S_t = s$ and given the agent possesses c units of the derivative:

$$V(t, z, s, c) := \sup_{\theta} \mathbb{E} \left[Z_T^{[\theta]} + cg(S_T) | Z_t = z, S_t = s \right],$$

where the supremum is taken over all admissible and self financing trading strategies θ . It is non-trivial to calculate the value function but stochastic optimal control theory yields an expression for the optimal trading strategy θ^* and a non-linear partial differential equation, the Hamilton-Jacobi-Bellman (HJB) equation, which depends on the precise specification of the dynamics of (S_t) and (P_t) . Once the value function is known the utility indifference price p is given by

$$V(t, z - p, s, c) = V(t, z, s, 0).$$

For more details see [Henderson, 2002]. For the special case of exponential utility and the dynamics

$$\begin{aligned} dS_t &= S_t \mu_1 dt + S_t \sigma_1 dW_t, & (\text{not traded}), \\ dP_t &= P_t \mu_2 dt + P_t \sigma_2 dB_t, & (\text{traded}), \\ dW_t dB_t &= \rho dt, \end{aligned}$$

it is shown that the optimal hedging strategy is to invest the money amount of $\frac{\mu_2 - r}{\sigma_2^2 \gamma} e^{-r(T-t)}$ in the traded asset and the utility indifference prices is given by

$$p(t, s, c) = -\frac{e^{-rT}}{\gamma(1 - \rho^2)} \ln \left(\mathbb{E}^{\mathbb{Q}} \left[e^{-c\gamma(1-\rho^2)g(S_T)} | S_t = s \right] \right),$$

where $\mathbb{Q}$ is the minimal martingale measure of Föllmer and Schweizer making the discounted process of (S_t) a martingale and leaving any orthogonal process of (W_t) unchanged.

For bounded payoff functions it also turns out that the marginal price is linked to the discounted expected payoff under the minimal martingale measure $\mathbb{Q}$:

$$\frac{p(t, s, c)}{c} = e^{-rT} \mathbb{E}^{\mathbb{Q}}[g(S_T) | S_t = s], \quad (c \rightarrow 0).$$

This result is not specific to exponential utility and the particular choice of dynamics. For a more general result see [Davis, 1997]. We henceforth focus on arbitrage pricing.

4.2 Arbitrage pricing and risk neutral formulation

In a market where the underlying cannot be used to replicate¹ derivative products, arbitrage arguments do not immediately lead to a unique price for derivatives. It is only when a market becomes richer in the sense that a broad range of derivatives is liquidly traded in addition to the underlying, that arbitrage theory may provide us with a unique price for a particular derivative we want to sell in that market. The basic idea is to use the traded options to replicate the product we want to sell. See [Björk, 2004, Chapter 15] for a detailed description of arbitrage pricing in incomplete markets. In Appendix D we have also given a short account of it. As it turns out, the market is free of arbitrage if there exists an equivalent measure $\mathbb{Q} \sim \mathbb{P}$ so that the discounted price V of all traded derivatives are martingales under $\mathbb{Q}$, i.e.

$$V_t = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[V_T | \mathcal{F}_t].$$

As the underlying cannot be used to hedge and hence can be considered a non-tradable asset in the sense of arbitrage pricing, the discounted value of the underlying is not necessarily a

¹This is also the case in weather markets where the underlying might be the temperature on one place.

martingale under $\mathbb{Q}$. This implies that the forward curve $F(t, T) := \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t]$ is not of the simple shape of $e^{r(T-t)}$ as it is the case in share or foreign exchange markets, for instance.

If such a measure $\mathbb{Q}$ does not exist then the market offers arbitrage opportunities which we will generally assume not to be the case. If the measure $\mathbb{Q}$ is still not uniquely determined, then the market is called incomplete and it is not possible to replicate any claim on the underlying using the options traded in the market.

We consider the model of the previous section. Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space and $(\mathcal{F}_t)_{t \in \mathbb{R}^+}$ be a filtration generated by (X_t) and (Y_t) . Additionally, let f be a continuous differentiable function, then the dynamics of the spot process is given by

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= -\alpha X_t dt + \sigma dW_t, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t. \end{aligned} \tag{4.2}$$

The processes (W_t) , (N_t) and (J_t) are assumed to be mutually independent.

Measure changes in models with jumps of continuous sizes are subject of ongoing research. [Henderson and Hobson, 2003] give a comprehensive overview of recent papers in the literature and compare option prices in a jump-diffusion model under different equivalent martingale measures. The definition of the measure change is very general and involves point processes and Poisson random measures. Here, we only consider a subset of equivalent measures which leave the structure of the jump process unchanged, i.e. jumps will still be generated by a Poisson process and an independent and stationary jump size distribution under $\mathbb{Q}$. The restriction we impose on the set of possible risk neutral measures might limit the range of arbitrage free prices we will get for certain options. However, it is very common in the literature to restrict the set of possible risk neutral measures in order to obtain a manageable risk neutral spot dynamics. [Merton, 1976] even leaves the jump dynamics unchanged by arguing the jump risk is unpriced. We will later show that the subset of martingale measures as defined below will be sufficient in the sense that a measure $\mathbb{Q}$ can always be found consistent with an observed forward curve.

Define any equivalent measure $\mathbb{Q}$ by $d\mathbb{Q} = \Pi_T d\mathbb{P}$ with the state price process $(\Pi_t)_{t \in [0, T]}$ given by the sde

$$\frac{d\Pi_t}{\Pi_{t-}} = -\lambda\gamma(t) dt - \psi(X_t, t) dW_t + \gamma(t) dN_t, \quad \Pi_0 = 1.$$

We call the function ψ the market price of diffusion risk and $\gamma > -1$ the market price of jump risk. For a more general version see [Henderson and Hobson, 2003, Section 3].

Based on Girsanov's theorem and the fact that J_t , N_t and W_t are mutually independent and based on [Henderson and Hobson, 2003, Remark 3.3] it follows that the dynamics under $\mathbb{Q}$ are given by

$$\begin{aligned} S_t &= \exp(f(t) + X_t + Y_t), \\ dX_t &= (-\alpha X_t - \psi(X_t, t)\sigma) dt + \sigma dW_t^{\mathbb{Q}}, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t^{\mathbb{Q}}, \end{aligned}$$

where $W_t^{\mathbb{Q}}$ is a $\mathbb{Q}$ Brownian motion and $N_t^{\mathbb{Q}}$ is a Poisson process under $\mathbb{Q}$ with intensity $\lambda(1 + \gamma(t))$. In order not to leave the class of models considered in Chapter 3 we further restrict the set of possible risk neutral measures by setting

$$\psi(x, t) = \frac{\hat{\alpha} - \alpha}{\sigma} x - \frac{\hat{\alpha}}{\sigma} \mu(t), \quad \gamma(t) = \frac{\hat{\lambda}}{\lambda} - 1,$$

and so the dynamics become

$$dX_t = \hat{\alpha}(\mu(t) - X_t) dt + \sigma dW_t^{\mathbb{Q}}, \quad dY_t = -\beta Y_{t-} dt + J_t dN_t^{\mathbb{Q}},$$

where $N_t^{\mathbb{Q}}$ has intensity $\hat{\lambda}$. The mean reverting level $\mu(t)$ of the OU process (X_t) can be expressed as an additional term of the seasonality, see Remark B.1.3, to obtain

$$\begin{aligned} S_t &= \exp(f(t) + f_1(t) + X_t + Y_t), \\ dX_t &= -\hat{\alpha} X_t dt + \sigma dW_t^{\mathbb{Q}}, \\ dY_t &= -\beta Y_{t-} dt + J_t dN_t^{\mathbb{Q}}, \end{aligned}$$

where f_1 solves the ordinary differential equation (ode) $f_1'(t) + \hat{\alpha} f_1(t) = \mu(t)$.

Remark 4.2.1 (Risk neutral dynamics)

The risk neutral dynamics of Model 4.2 in an electricity market is given by

$$\begin{aligned} S_t &= \exp(\hat{f}(t) + X_t + Y_t), \\ dX_t &= -\hat{\alpha} X_t dt + \sigma dW_t^{\mathbb{Q}}, \\ dY_t &= -\beta Y_{t-} dt + \hat{J}_t dN_t^{\mathbb{Q}}, \end{aligned} \tag{4.3}$$

where $N_t^{\mathbb{Q}}$ has intensity $\hat{\lambda}$ under $\mathbb{Q}$. The parameters $\hat{\alpha}$, $\hat{\lambda}$, the function $\hat{f}$ as well as the jump size distribution $\hat{J}$ are all determined by the particular choice of measure $\mathbb{Q}$. Only the parameters σ and β remain unchanged by our measure transformation. Note, the drift² of the process (Y_t) under $\mathbb{P}$ could still be different from the drift under $\mathbb{Q}$. This is because the Poisson process is not a martingale and has a drift λdt under $\mathbb{P}$ and $\hat{\lambda} dt$ under $\mathbb{Q}$. Only because we specify our model in terms of the Poisson process and not its compensated version, the term in front of dt remains unchanged by a change of measure.

²In a non-rigorous definition we could say the drift is the process less its martingale part.

Corollary 4.2.2 (Seasonal function consistent with the forward curve)

Let $t = 0$ and $F_0^{[T]}$ be the forward at time 0 maturing at time T , then the risk neutral seasonality function is given by

$$\hat{f}(T) = \ln F_0^{[T]} - X_0 e^{-\hat{\alpha}T} - Y_0 e^{-\beta T} - \frac{\sigma^2}{4\hat{\alpha}}(1 - e^{-2\hat{\alpha}T}) - \hat{\lambda} \int_0^T \Phi_{\hat{J}}(e^{-\beta s}) - 1 ds.$$

Proof The forward price in an arbitrage free market is given by $F_0^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T]$. Based on the risk neutral dynamics (4.3) and the expectation value of S_T given in Corollary 3.4.13 the result follows immediately. $\square$

Remark 4.2.3

The market remains incomplete even with a complete and liquid forward market. The speed of mean-reversion $\hat{\alpha}$, and all the jump parameters $\hat{\lambda}$ and $\hat{J}$ remain undetermined. One way of choosing a measure $\mathbb{Q}$ is to pick the one which is closest to $\mathbb{P}$ in some metrical sense. Section 5.3 gives some ideas on how to find such an optimal $\mathbb{Q}$ and the implications for the parameter $\hat{\alpha}$. Otherwise, if more options are liquidly traded in the market, more constraints on the measure $\mathbb{Q}$ are imposed which might lead to a full determination of all risk neutral parameters. In practice, however, where a liquid option market is still rare, a pragmatic approach is to estimate all parameters from historical data but to calibrate the seasonality function to the observed forward curve.

To simplify notation in this chapter we will assume that the risk neutral dynamics is given by (4.2) but where parameters are calibrated in a consistent way to match observed market prices, in particular we will assume that the model is consistent with the observed forward curve.

4.3 Pricing path-independent options

The purpose of this section is to examine options with a general payoff $g(S_T)$, i.e. only depending on the value of the underlying at maturity T . The arbitrage-free value of the option is given by

$$V = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[g(S_T)|\mathcal{F}_t],$$

and as (S_t) is not a Markov process in our model (4.2) – only the individual processes (X_t) and (Y_t) are Markov – the price depends on the entire history of the spot price process. Assuming the mean-reverting and spike process are individually observable, the option price then only depends on X_t and Y_t and is given by

$$V(x, y, t) = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[g(S_T)|X_t = x, Y_t = y]. \quad (4.4)$$

In practical situations it is fairly obvious when a spike has occurred and so we will assume that both processes are observable. This slight inconsistency with reality will disappear when considering forwards as the underlying process. As it turns out, the forward process of a fixed maturity T is a Markov process (see Equation (4.8)) and hence prices of options on forwards only depend on the initial value of the forward.

For the valuation of option prices we need to be able to calculate the expectation values (4.4). Although we do not have an analytic expression for the probability density of S_T , approximations developed in Section 3.4.3 can be used to obtain approximate pricing formulae. For long time horizons T the distribution of S_T turns out to be similar to a lognormal distribution but with heavier tails. Without the presence of jumps, S_T will be lognormal and pricing formulae are given in Section D.3 and are of the same form as the Black-Scholes equation.

Another method uses the complex valued moment generating function of S_T and the Laplace inversion to derive option prices. For an overview see [Cont and Tankov, 2004, Section 11.1.3]³ or the therein referred papers of [Carr and Madan, 1998] and [Lewis, 2001]. Even [Heston, 1993] has already used this method implicitly to price call options in a stochastic volatility model.

We describe two inversion methods below, one which can only be used to price call and put option and the other one able to deal with general option payoffs.

4.3.1 Pricing call options

Assume the following notation

$$S_t = e^{Z_t},$$

$$\Phi_Z(\theta, t) = \mathbb{E}^{\mathbb{Q}}[e^{\theta Z_t}], \quad \theta \in \Theta \subset \mathbb{C}.$$

Assume today is $t = 0$ and the option payoff is given by $g(S_T)$ then the arbitrage free option price is

$$V(x, y, 0) = e^{-rT} \mathbb{E}^{\mathbb{Q}}[g(e^{Z_T}) | X_0 = x, Y_0 = y].$$

We recall a general result on how to obtain the distribution function from a characteristic function, an extension of which will later be used to price options. See [Stuart and Ord, 1994, Chapter 4] and [Williams, 1991, Chapter 16] for details.

³They describe the method in terms of a complex valued characteristic function and Fourier inversion, but by allowing complex values the method can also be written in terms of Laplace transforms.

Theorem 4.3.1 (Lévy's Inversion Theorem)

Let $\Phi : \Theta \subset \mathbb{C} \rightarrow \mathbb{R}$ be the moment generating function of a random variable Z

$$\Phi(\theta) := \mathbb{E} \left[e^{\theta Z} \right] = \int_{\mathbb{R}} e^{\theta x} dF_Z(x)$$

then the cumulative distribution $F_Z : \mathbb{R} \rightarrow [0, 1]$ is given by

$$F_Z(x) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty \frac{\Im \left(\Phi(0 + i\nu) e^{-i\nu x} \right)}{\nu} d\nu.$$

Proof [Stuart and Ord, 1994, Equation 4.14]. □

Remark 4.3.2

The moment generating function at purely imaginary points is equal to the characteristic function and therefore always exists.

This inversion formula can be generalised to truncated moment generating functions, see also [Duffie et al., 2000, Proposition 2].

Proposition 4.3.3

Let Z be a random variable and its truncated moment generating function be defined by

$$G_\nu(x) := \mathbb{E} \left[e^{\nu Z} \mathbf{1}_{\{Z \leq x\}} \right] = \int_{-\infty}^x e^{\nu y} dF_Z(y).$$

If the moment generating function $\Phi(\nu + i\theta)$ exists for some $\nu \in \mathbb{R}$ and all $\theta \in \mathbb{R}$ then

$$G_\nu(x) = \frac{\Phi(\nu)}{2} - \frac{1}{\pi} \int_0^\infty \frac{\Im \left(\Phi(\nu + i\theta) e^{-i\theta x} \right)}{\theta} d\theta. \quad (4.5)$$

Proof Proposition C.2.3. □

This proposition states that the truncated expectation of $e^{\nu Z}$ is given by some form of an inverse Laplace transform of the moment generating function of Z over θ . We can use this to price put options because we have

$$\mathbb{E}[(K - S_T)^+] = K \mathbb{E}[\mathbf{1}_{S_T \leq K}] - \mathbb{E}[S_T \mathbf{1}_{S_T \leq K}] = KG_0(\ln K) - G_1(\ln K),$$

and the price of a call option can be obtained by put-call parity.

4.3.2 Pricing options with arbitrary payoff

We assume our model (4.2) with zero initial conditions $X_0 = 0$ and $Y_0 = 0$. Let $\bar{Z}_T = f(T) + X_T + Y_T$ given $X_0 = 0$ and $Y_0 = 0$ and $Z_T = f(T) + X_T + Y_T$ given $X_0 = x$ and $Y_0 = y$, as before

$$S_T = \exp(Z_T) = \exp(c + \bar{Z}_T), \quad c = e^{-\alpha T} x + e^{-\beta T} y.$$

Let the option payoff be given by $g(S_T)$. The expected payoff can then be written as the convolution of the payoff and the density function $f_{\bar{Z}_T}$:

$$\begin{aligned} \mathbb{E} \left[g(e^{c+\bar{Z}_T}) \right] &= \int_{\mathbb{R}} g(e^{c+x}) f_{\bar{Z}_T}(x) dx \\ &= \int_{\mathbb{R}} h(c-x) f(x) dx, \end{aligned}$$

with $h(x) := g(e^x)$ and $f(x) = f_{\bar{Z}_T}(-x)$. In the Laplace transformed space a convolution becomes a multiplication, so the idea is to Laplace transform in c , solve the expression in the Laplace transformed space and back-transform.

Theorem 4.3.4

Let $h : (0, \infty) \rightarrow \mathbb{R}$ be such that its Laplace transform $\hat{h}(\theta)$ exists for $\Re\theta \in I_1$ and $\bar{Z}_T$ be a random variable possessing a probability density and having a moment generating Function $\Phi_{\bar{Z}_T}(\theta)$ defined for all $\Re\theta \in I_2$. Given $I := I_1 \cap -I_2$ is non-empty

$$\mathbb{E}[h(c + \bar{Z}_T)] = \frac{1}{2\pi i} \lim_{R \rightarrow \infty} \int_{\gamma - iR}^{\gamma + iR} \hat{h}(\theta) \Phi_{\bar{Z}_T}(-\theta) e^{-\theta c} d\theta, \quad \forall \gamma \in I. \quad (4.6)$$

Proof Define

$$\begin{aligned} v(x) &:= \mathbb{E} [h(x + \bar{Z}_T)], \\ \hat{v}(\theta) &= \int_{-\infty}^{\infty} e^{\theta x} v(x) dx, \quad \theta \in \mathbb{C}. \end{aligned}$$

The function v can be written as a convolution of h and the density function f of X :

$$v(x) = \int_{\mathbb{R}} h(x-y) f(y) dy, \quad f(y) := f_{\bar{Z}_T}(-y).$$

Note the Laplace transform of f at θ is equal to $\Phi_{\bar{Z}_T}(-\theta)$.

The convolution theorem guarantees existence of $\hat{v}(\theta)$ if the Laplace transform of both convolution terms exists and then it is equal to the product of their Laplace transforms:

$$\hat{v}(\theta) = \hat{h}(\theta) \Phi_{\bar{Z}_T}(-\theta), \quad \forall \theta \in \mathbb{C} : \Re(\theta) \in I.$$

Applying the Laplace inversion theorem obtains the desired result

$$v(x) = \frac{1}{2\pi i} \lim_{R \rightarrow \infty} \int_{\gamma-iR}^{\gamma+iR} \hat{h}(\theta) \Phi_{\bar{Z}_T}(-\theta) e^{-\theta x} d\theta.$$

□

Remark 4.3.5

For a call option payoff $g(x) = (x - K)^+$, the Laplace transformed of $h(x) = g(e^x) = (e^x - K)^+$ is

$$\begin{aligned} \hat{h}(\theta) &= \int_{-\infty}^{\infty} (e^x - K)^+ e^{\theta x} dx \\ &= \int_{\ln K}^{\infty} (e^x - K) e^{\theta x} dx \\ &= \int_{\ln K}^{\infty} e^{(\theta+1)x} dx - K \int_{\ln K}^{\infty} e^{\theta x} dx \\ &= -\frac{1}{\theta+1} e^{(\theta+1)\ln K} + \frac{K}{\theta} e^{\theta \ln K} \\ &= -\frac{K^{\theta+1}}{\theta+1} + \frac{K^{\theta+1}}{\theta} \\ &= \frac{K^{\theta+1}}{\theta(\theta+1)}, \quad \Re \theta < -1. \end{aligned}$$

4.3.3 Pricing options on Forwards

For a forward contract at time t maturing at T we know the strike of a zero-cost forward is given by

$$F_t^{[T]} = \mathbb{E}^{\mathbb{Q}}[S_T | \mathcal{F}_t].$$

The most common options on forwards are puts or calls maturing at the same time as the underlying forward, i.e. the payoff is given by $(F_T^{[T]} - K)^+$ which is equivalent to $(S_T - K)^+$. Therefore we can price these contracts based on the dynamics of the spot and using methods developed in the previous sections. However, by analysing the dynamics of the forward curve implied by the spot price model we will gain further insights and be able to relate the price of an option to the Black-76 formula [Black, 1976], which is still widely used in practice.

Recall the result of Corollary 3.4.13:

$$F_t^{[T]} = \exp \left(f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds \right). \quad (4.7)$$

We fix a maturity T and apply Itô's formula to obtain the dynamics of the forward maturing at T :

$$\frac{dF_t^{[T]}}{F_t^{[T]}} = -\lambda \left(\Phi_J(e^{-\beta(T-t)}) - 1 \right) dt + \sigma e^{-\alpha(T-t)} dW + \left(\exp(J_t e^{-\beta(T-t)}) - 1 \right) dN_t. \quad (4.8)$$

The forward is a martingale by definition, and so the drift term only compensates the jump process. For large time to maturities $T - t$, a jump in the underlying process has only very limited effect on the forward. More precisely, if the relative change in the underlying is $\exp(J_t) - 1$ the forward changes relatively by $\exp(J_t e^{-\beta(T-t)}) - 1$. In addition to the jump component the dynamics follows a deterministic volatility process starting with a low volatility $\sigma e^{-\alpha T}$ at $t = 0$ and ever increasing it approaches σ at maturity. Without the jump component there are clear similarities with the Black-76 model.

For pricing purposes we need to find the distribution of $F_T^{[T]}$ in terms of its initial condition $F_t^{[T]}$ where t is understood to be today. We have

$$\begin{aligned} \ln F_T^{[T]} &= f(T) + X_T + Y_T, \\ \ln F_t^{[T]} &= f(T) + X_t e^{-\alpha(T-t)} + Y_t e^{-\beta(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds. \end{aligned}$$

Subtracting the second equation from the first eliminates the seasonality component $f(T)$, and using the relation

$$\begin{aligned} X_T - X_t e^{-\alpha(T-t)} &= \sigma \int_t^T e^{-\alpha(T-s)} dW_s, \\ Y_T - Y_t e^{-\beta(T-t)} &= \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)}, \end{aligned}$$

we finally get

$$\begin{aligned} \ln F_T^{[T]} - \ln F_t^{[T]} &= \sigma \int_t^T e^{-\alpha(T-s)} dW_s + \sum_{i=N_t}^{N_T} J_{\tau_i} e^{-\beta(T-\tau_i)} \\ &\quad + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds. \end{aligned} \quad (4.9)$$

Without the jump component, $F_T^{[T]}$ would be log-normally distributed. Given a very high mean-reversion rate β of the jump component we make a first approximation by assuming $F_T^{[T]}$ is nearly log-normally distributed, which we can partly justify by the approximations derived in Section 3.4.4. In particular Equation (3.13) says that the density of $\ln S_T - f(T)$ is approximately a weighted sum of a normal density and the density of a random variable related to the jump size distribution. It turns out that for standard market parameters and

medium to long term maturities the weights are almost exclusively on the normal density. Figure 3.14 shows the density of $\ln S_T - f(T) = \ln F_T^{[T]} - f(T)$. Hence, we expect a good fit of option prices of at-the-money calls but due to the heavy tails of the jump size distribution which we completely neglect in the approximation, we expect to underestimate prices of far out of the money calls.

Based on the definition of a forward, $F_t^{[T]}$ is a martingale for a fixed maturity T and in order to keep the same property in our approximation we set

$$\ln F_T^{[T]} \approx \ln F_t^{[T]} + \xi, \quad \xi \sim \mathcal{N}\left(-\frac{1}{2}\hat{\sigma}^2(T-t), \hat{\sigma}^2(T-t)\right),$$

and set $\hat{\sigma}^2(T-t) := \text{var}[\ln F_T^{[T]} | \mathcal{F}_t]$, i.e.

$$\begin{aligned} \hat{\sigma}^2(T-t) &= \text{var} \left[\sigma \int_t^T e^{-\alpha(T-s)} dW_s + \sum_{i=N_t}^{N_T} J_{t_i} e^{-\beta(T-t_i)} \right] \\ &= \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha(T-t)}) + \frac{\lambda}{2\beta} \mathbb{E}[J^2] (1 - e^{-2\beta(T-t)}). \end{aligned}$$

Remark 4.3.6 (Term structure of implied volatility)

Comparing this result with the setting of Black-76 [Black, 1976] where $dF = F\sigma dW$ and so $F_T = F_t \exp(\xi)$ with $\xi \sim \mathcal{N}(-\frac{1}{2}\sigma^2(T-t), \sigma^2(T-t))$ we conclude that $\hat{\sigma}$ is the implied Black-76 volatility and in a first approximation given by

$$\hat{\sigma}^2 \approx \frac{\frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha(T-t)}) + \frac{\lambda}{2\beta} \mathbb{E}[J^2] (1 - e^{-2\beta(T-t)})}{T-t}, \quad (4.10)$$

which is shown in Figure 4.2. It can be seen that the spike process has a much more significant impact on the implied volatility for short maturities rather than for long term maturities. As far as the price of an at the money call is concerned, the additional jump risk adds an almost constant premium to the price to be payed without any jump risk.

Remark 4.3.7 (Implied volatility across strikes)

The approximation does not predict a change of implied volatility across strikes. However, the jump risk introduces a skew as can be seen in Figure 4.2 where the exact solution based on Section 4.3.1 has been used to calculate implied volatilities. The bigger the mean jump size and hence the bigger $\mathbb{E}[J^2]$, the more profound is the skew.

4.3.4 Pricing options on Forwards with a delivery period

As electricity is a flow variable, forwards always specify a delivery period. The results of the previous section can therefore only be seen as an approximation to option prices on forwards with short delivery periods, like one day. We only consider options on forwards




Figure 4.2: Implied volatilities and prices. The left graph shows implied volatilities with respect to time to maturity where approximation (4.10) is used. The three lines correspond to no jumps ($\mu_J = 0$), small jumps ($\mu_J = 0.4$) and big jumps ($\mu_J = 0.8$). In the right graph the corresponding prices of an at the money call are plotted. Parameters are $r = \ln(1.05)$, $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4.0$, $F_0^{[T]} = 100$, $K = 100$. Note, for an exponential distributed jump size J we have $\mathbb{E}[J^2] = 2\mu_J^2$.




Figure 4.3: Implied volatilities across strikes and sample paths. The upper graph shows the implied volatility for one maturity $T = 0.2$ based on the exact solution. The approximate solution (4.10) yields 0.82 and 0.85 for the small and big jumps, respectively. Sample paths of the model with the same parameters are drawn in the lower two graphs, where the left path is generated with a low mean jump size ($\mu_J = 0.4$) and the right with a high mean jump size ($\mu_J = 0.8$). All the other parameters are the same as in Figure 4.2.

maturing at the beginning of the delivery period, i.e. the payoff is given by some function of $F_{T_1}^{[T_1, T_2]}$ at time T_1 . An option on such a forward is conceptually similar to an Asian option in the Black-Scholes world. One method of pricing Asian options is to approximate the distribution of the integral by a log-normal distribution and can be done by matching the first two moments, see [Turnbull and Wakeman, 1991] for example. Once the parameters of the approximate log-normal distribution have been determined, pricing options comes down to pricing in the Black-Scholes or Black-76 setting.

Recalling the relation between the forward with and without a delivery period (2.1)

$$F_t^{[T_1, T_2]} = \int_{T_1}^{T_2} w(T; T_1, T_2) F_t^{[T]} dT,$$

the second moment of $F_{T_1}^{[T_1, T_2]}$ is given by

$$\mathbb{E}^{\mathbb{Q}} \left[\left(\int_{T_1}^{T_2} w(T) F_{T_1}^{[T]} dT \right)^2 \middle| \mathcal{F}_t \right] = \int_{T_1}^{T_2} \int_{T_1}^{T_2} w(T) w(T^*) \mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right] dT dT^*,$$

and the expectation of the product of two individual forwards $\mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} \middle| \mathcal{F}_t \right]$ can be derived using the solution of the forward (4.7) as follows:

$$\begin{aligned} \ln F_{T_1}^{[T]} &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{-\beta(T-T_1)} \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad + \lambda \int_0^{T-T_1} \Phi_J(e^{-\beta s}) - 1 ds - \lambda \int_0^{T-t} \Phi_J(e^{-\beta s}) - 1 ds \\ &= \ln F_t^{[T]} + e^{-\alpha(T-T_1)} \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s + e^{-\beta(T-T_1)} \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) - \lambda \int_0^{T_1-t} \Phi_J(e^{-\beta(T-T_1)} e^{-\beta s}) - 1 ds, \end{aligned}$$

and so

$$\begin{aligned} \ln F_{T_1}^{[T]} + \ln F_{T_1}^{[T^*]} &= \ln F_t^{[T]} + \ln F_t^{[T^*]} \\ &\quad + \left(e^{-\alpha(T-T_1)} + e^{-\alpha(T^*-T_1)} \right) \sigma \int_t^{T_1} e^{-\alpha(T_1-s)} dW_s \\ &\quad + \left(e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)} \right) \sum_{i=N_t}^{N_{T_1}} J_{\tau_i} e^{-\beta(T_1-\tau_i)} \\ &\quad - \frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)}) (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \\ &\quad - \ln \Phi_Y(e^{-\beta(T-T_1)}) - \ln \Phi_Y(e^{-\beta(T^*-T_1)}), \end{aligned}$$




Figure 4.4: Distribution of $F_{T_1}^{[T_1, T_2]}$. Parameters are as before, see Figure 4.2, and $\mu_J = 0.4$, $T_1 = 1$, $T_2 = 1.25$. The red curve is based on a Monte-Carlo simulation with one million sample paths and the green curve is the log-normal approximation matching the first two moments.

which gives

$$\begin{aligned}
 \mathbb{E}^{\mathbb{Q}} \left[F_{T_1}^{[T]} F_{T_1}^{[T^*]} | \mathcal{F}_t \right] &= \mathbb{E}^{\mathbb{Q}} \left[\exp \left(\ln F_{T_1}^{[T]} + \ln F_{T_1}^{[T^*]} \right) | \mathcal{F}_t \right] \\
 &= F_t^{[T]} F_t^{[T^*]} \frac{\Phi_Y(e^{-\beta(T-T_1)} + e^{-\beta(T^*-T_1)})}{\Phi_Y(e^{-\beta(T-T_1)}) \Phi_Y(e^{-\beta(T^*-T_1)})} \\
 &\quad \exp \left(-\frac{\sigma^2}{4\alpha} (1 + e^{-2\alpha(T^*-T)}) (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \right) \\
 &\quad \exp \left(\frac{\sigma^2}{4\alpha} (1 + e^{-\alpha(T^*-T)})^2 (e^{-2\alpha(T-T_1)} - e^{-2\alpha(T-t)}) \right).
 \end{aligned}$$

How well the moment matching procedure works is shown in Figure 4.4 where the density of a forward $F_{T_1}^{[T_1, T_2]}$ is compared to the density obtained by the approximation. The shapes of both densities are similar but differences in values are clearly visible. As a result one might not expect call option prices based on the approximate distribution to be very close to the exact prices for all strikes K but still close enough to be useful. For an at-the-money strike call option, prices for varying maturities are shown in Figures 4.5 and 4.6. As it turns out, the approximations gives very good results for short delivery periods and is still within a 5% range for delivery periods of one year.




Figure 4.5: Value of an at-the-money call option on a forward $F_{T_1}^{[T_1, T_2]}$ depending on the delivery period $T_2 - T_1$. Parameters are as before, see Figure 4.2, and $\mu_J = 0.4$, $T_1 = 1$, $K = 100$. The Monte-Carlo result is based on 100 000 sample paths for each duration. Note, any forward $F_{T_1}^{[T_1, T_2]}$ delivers exactly 1 MWh over the delivery period $[T_1, T_2]$.




Figure 4.6: Same as in Figure 4.5 but where the volume of the forward is proportional to the delivery period, i.e. we assume a constant consumption of 1MWh per year, which is about 114 W. Here, the price is simply $T_2 - T_1$ times the price of a standard call option on $F_{T_1}^{[T_1, T_2]}$.

4.4 Pricing swing options

Swing contracts are a broad class of path dependent options allowing the holder to exercise a certain right multiple times over a specified period but only one right at a time⁴ or per time-interval like a day. Such a right could be the right to receive the payoff of a call option. Other possibilities include the mixture of different payoff functions like calls and puts or calls with different strikes. Another very common feature is to allow the holder to exercise a real valued multiple of a call or put option at once, where the multiple is called volume. This generally involves further restriction on the volume, like upper and lower bounds for each right and for the sum of all trades.

Swing contracts can be seen as an insurance for the holder against excessive rises in electricity prices. Assuming the prices generally revert to a long term mean, even a small number N of exercise opportunities suffices to cover the main risks and hence make the premium of the contract cheaper. Sometimes, swing contracts are bundled with forward contracts. The forward contract then supplies the holder with a constant stream of energy to a fixed pre-determined price. If the strike price of the call options of the swing contract is set to the forward price, the swing contract will allow for flexibility in the volume the customer receives for the fixed price. They can either “swing up” or “swing down” the volume of energy and hence the name swing contract. One cannot assume that the holder always exercises the contract in an optimal way to maximise expected profit but they might also exercise according to their own internal energy demands.

Swing contracts have been around for much longer than academic papers on valuing them based on arbitrage principles. It is only very recent that articles on numerical pricing methods for swing options have appeared in the literature. We can identify a few main approaches all based on dynamic programming principles. A Monte-Carlo method and ideas of duality theory are utilised in [Meinshausen and Hambly, 2004] to derive lower and upper bounds for swing option prices. The main advantages of the method being their flexibility as it can be easily adapted to any stochastic model of the underlying and its ability to produce confidence intervals of the price. Monte-Carlo techniques are also used in [Ibanez, 2004] and [Carmona and Touzi, 2004], where the latter uses the theory of the Snell envelope to determine the optimal exercise boundaries and also utilises the Malliavin calculus. A constructive solution to the perpetual swing case for exponential Brownian motion is also given in [Carmona and Touzi, 2004]. Unfortunately, these methods only work for the most basic versions of swing contracts where at each time only one unit of an option can be exercised.

⁴This will also involve a “refraction period” in which no further right can be exercised.

More general swing contracts with a variable volume per exercise and an overall constrained can be prices with a tree based method introduced in [Jaillet et al., 2004].

In all the above papers a discrete time model for the underlying is used where one time step corresponds to the time frame in which no more than one right can be exercised, i.e. one day in most of the traded contracts. A special case where the number of exercise opportunities is equal to the number of exercise dates is considered in [Howison and Rasmussen, 2002] and a continuous optimal exercise strategy derived which yields a partial integro-differential equation for the option price.

Our method is based on the tree approach of [Jaillet et al., 2004], but instead of using a trinomial tree for the spot price we account for the possibility of the spot price jumping from any value to any other and hence obtain a grid rather than a tree. Besides this modification our approach is identical to the one described in [Jaillet et al., 2004]. To generate the grid we will make use of approximations to the conditional density as given in Section 3.4.2. Although the approximations are not essential for the method to work as we could also use the moment generating function to infer conditional densities, it speeds up the generation of the grid considerably.

4.4.1 The grid approach

The tree method of [Jaillet et al., 2004] requires a discrete time model of the underlying. This is due to the fact that their swing contracts allow the holder to exercise at most one option within a specified time interval, say a day, and this is best modelled if the underlying process has the same time discretisation. Assuming (S_t) is some continuous stochastic process for the spot price we obtain a discrete model by observing it on discrete points in time only, i.e.

$$S_{t_0}, S_{t_1}, S_{t_2}, \dots, S_{t_m},$$

with $t_0 = 0$, $t_{i+1} = t_i + \Delta t$, $t_m = T$ and $\Delta t = \frac{1}{365}$, indicating we can exercise on a daily basis.

Let the maturity date T be fixed and the payoff at time t for simplicity⁵ be given by $(S_t - K)^+$ for some strike price K and we assume it is only allowed to exercise one unit of the underlying at once. Let $V(n, s, t)$ denote the price of such a swing option at time t and spot price s which has n exercise rights left. The general pricing principle is then based on the following equation

$$V(n, s, t) = \max \left\{ e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s], \right. \\ \left. e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n-1, S_{t+\Delta t}, t + \Delta t) | S_t = s] + (s - K)^+ \right\}, \quad (4.11)$$

⁵We could assume any general payoff function.

which basically says, today's value is the maximum of the expected value of the same option tomorrow and the expected value of the same option but with one less exercise opportunity plus the payoff of the option. Given we know the value of all swing options of tomorrow, i.e. we know $V(k, s, t + \Delta t) \forall s, k$, then we can express the conditional expectations as

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s] = \int_{-\infty}^{\infty} V(n, x, t + \Delta t) f_S(x; s) dx,$$

where $f_S(x; s)$ is the density of $S_{t+\Delta t}$ given $S_t = s$. With the final condition known to be $V(n, s, T) = 0, \forall n, s$ and boundary condition $V(0, s, t) = 0, \forall s, t$, the method works backward in time. Discretising the spot variable we approximate

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_j V(n, s_j, t + \Delta t) f_S(s_j; s_i) \Delta s_j.$$

This is only one possible approximation, others might be to use higher order integration rules or using only a few grid points in the sum based on the fact that $f_S(x; s) \rightarrow 0$ for $|s - x|$ big. For a trinomial tree one only uses three grid points, i.e.

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_{j=-1}^1 V(n, s_{i+j}, t + \Delta t) p_{i,i+j},$$

$p_{i,j}$ being the probability of going from node i to node j . However, such a tree approach is not fully suited to our case for two reasons. First, the time step size is determined by the shortest time between two possible exercise dates, which is mainly one day for swing contracts. This limits the accuracy of the algorithm as a refinement of the grid in spot direction will not improve the result. Second, in the presence of jumps, a three point approximation for the conditional density is insufficient due to the heavy tails in the distribution. As a result, we keep our method general and say

$$\mathbb{E}^{\mathbb{Q}} [V(n, S_{t+\Delta t}, t + \Delta t) | S_t = s_i] \approx \sum_j V(n, s_j, t + \Delta t) p_{i,j},$$

where $p_{i,j}$ in some ways resembles the density $f_S(s_j; s_i) \Delta s_j$ but does not have to be the same in order to accommodate higher order integration rules and boundary approximations. With the notation of $V_{i,k}^n := V(n, s_i, t_k)$ we can then write the method as

$$\begin{aligned} V_{i,k}^n &= \max \left\{ e^{-r\Delta t} \sum_j V_{j,k+1}^n p_{i,j}, e^{-r\Delta t} \sum_j V_{j,k+1}^{n-1} p_{i,j} + (s_i - K)^+ \right\}, \\ V_{i,k}^0 &= 0, \\ V_{i,m}^n &= 0. \end{aligned} \tag{4.12}$$

We have not performed a sufficient analysis of different integration methods and non-uniform grids to suggest an optimal method but we can state that the Gaussian quadrature rule as

an integration method and the use of a non-uniform grid improves the result considerably. We use a three point Gaussian quadrature rule, defined by

$$\int_{-1}^1 f(x) dx \approx \sum_{i=-1}^1 w_i f(x_i),$$

with $x_{\pm 1} = \pm \sqrt{3/5}$, $x_0 = 0$, $w_{\pm 1} = 5/9$ and $w_0 = 8/9$. By linear transformation we then get

$$\int_a^b f(x) dx \approx \frac{b-a}{2} \sum_{i=-1}^1 w_i f\left(\frac{b-a}{2}x_i + \frac{a+b}{2}\right).$$

The integration intervals are generated by the method described in Section D.3.1 with a concentration of grid points around zero⁶ and with a 2.5 fold intensity of grid points compared to a uniform grid with the same number of grid points. See Figure D.3 for an illustration.

4.4.2 Comparison of algorithms

We compare the accuracy of our algorithm with results of [Meinshausen and Hambly, 2004] as they are able to calculate upper and lower price bounds based on a Monte-Carlo method [Longstaff and Schwartz, 2001] and duality ideas [Rogers, 2002]. The model they consider is a discrete version of an exponential Ornstein-Uhlenbeck process, which is a special case of our Model (3.3) without the presence of any spikes and no seasonality, i.e.

$$\begin{aligned} dX_t &= -\alpha X_t + \sigma dW_t, \\ S_t &= \exp(X_t). \end{aligned}$$

We observe (S_t) on discrete points in time, only. To compare our results with Table 4.2 of [Meinshausen and Hambly, 2004], we need to match their discrete model with our continuous. They assume a discrete model of the form

$$\tilde{X}_{t+1} = (1 - \tilde{k})\tilde{X}_t + \tilde{\sigma}\xi_t, \quad \xi \sim \mathcal{N}(0, 1),$$

with $\tilde{k} = 0.9$, $\tilde{\sigma} = 0.5$ and $\tilde{X}_0 = 0$. One time-step in their paper corresponds to a time-step of $\Delta t = \frac{1}{365}$ in our context, so

$$\tilde{X}_{t+\Delta t} \sim \mathcal{N}\left((1 - \tilde{k})\tilde{X}_t, \tilde{\sigma}^2\right).$$

Our continuous process $dX_t = -\alpha X_t dt + \sigma dW_t$ has the conditional distribution

$$X_{t+\Delta t} \sim \mathcal{N}\left(e^{-\alpha\Delta t} X_t, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\Delta t})\right).$$

⁶This is because in the numerical examples we calculate the initial conditions are generally $X_0 = 0$ and $Y_0 = 0$.




Figure 4.7: Value of swing contracts with 1–100 exercise rights over 1000 days. The Monte-Carlo bounds are based on [Meinshausen and Hambly, 2004], Table 4.2. We use 42 grid points in spot direction and obtain a result for all options within 2 seconds on an Intel Pentium 4 with 3.4GHz. The grid is sketched in the picture below where each of the 14 intervals indicated by red lines contain three grid points in accordance with the Gaussian quadrature rule.

In order to make both models equivalent, means and variances need to be the same and so

$$\alpha = \frac{\ln(1 - \tilde{k})}{\Delta t},$$

$$\sigma^2 = \frac{2\alpha}{1 - e^{-2\alpha\Delta t}} \tilde{\sigma}^2,$$

and hence we get $\sigma \approx 20.60$ and $\alpha \approx 840.4$.

A swing option contract with 1000 exercise days and up to 100 exercise opportunities is considered in [Meinshausen and Hambly, 2004] with a strike of zero ($K = 0$). We choose parameters of our grid method so that the result is accurate enough to fall within the Monte-Carlo bounds of [Meinshausen and Hambly, 2004]. As it turns out 42 grid points in X direction with a slight grid point concentration⁷ at zero suffice. An important advantage of this method is its speed efficiency of one-dimensional problems like this. The calculation of all swing option prices with 1–100 exercise opportunities only takes two seconds on an Intel Pentium 4 with 3.4GHz. The result is shown in Figure 4.7.

⁷2.5 as dense at zero as in a uniform grid.

4.4.3 Numerical results

We now turn to the model of interest (4.2) which exhibits spikes. Assume that the mean-reversion process (X_t) and the spike process (Y_t) are individually observable and so the value function V of a swing option depends on both variables and the general pricing principle (4.11) becomes

$$V(n, x, y, t) = \max \begin{cases} e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y], \\ e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n-1, X_{t+\Delta t}, Y_{t+\Delta t}, t + \Delta t) | X_t = x, Y_t = y] + (e^{f(t)+x+y} - K)^+. \end{cases}$$

We discretise the spacial variables by generating a grid in x and y direction as previously described in Section 4.4.1. In order to calculate conditional expectations we need to define transitional probabilities. Given one starts at node $(X_t, Y_t) = (x_i, y_j)$ the probability to arrive at node $(X_{t+\Delta t}, Y_{t+\Delta t}) = (x_k, y_l)$ is approximately given by

$$p_{i,j,k,l} \approx f_{X_{t+\Delta t}|X_t=x_i}(x_k) f_{Y_{t+\Delta t}|Y_t=y_j}(y_l) \Delta x \Delta y,$$

because X_t and Y_t are independent. We use a slightly more sophisticated approximation by applying a 3-point Gaussian integration rule within each interval Δx and Δy . The conditional density of the mean-reverting process (X_t) is known as $X_{t+\Delta t}$ given $X_t = x$ is normally distributed with $\mathcal{N}(x e^{-\alpha\Delta t}, \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha\Delta t}))$, see (B.5). As we do not have a closed form expression for the density of the spike process we use approximations developed in Section 3.4.1. For an exponential jump size distribution $J \sim \text{Exp}(1/\mu_J)$ for example we use approximation (3.9) for the spike process at time t given zero initial conditions:

$$f_{\tilde{Y}_t}(x) = \frac{1}{\Gamma(\frac{\lambda}{\beta})\mu_J^{\frac{\lambda}{\beta}}} \frac{e^{-\frac{x}{\mu_J}} - e^{-\frac{x e^{\beta t}}{\mu_J}}}{x^{1-\frac{\lambda}{\beta}}}, \quad x > 0,$$

$$P(\tilde{Y}_t = 0) = e^{-\lambda t},$$

and so the conditional density can be approximated by

$$f_{Y_{t+\Delta t}|Y_t=y}(x) \approx f_{\tilde{Y}_{\Delta t}}(x - y e^{-\beta\Delta t}).$$

The introduction of a second space dimension increases the complexity of the algorithm considerably, essentially by a factor proportional to the square of the number of grid points in the y -direction. To price the swing contract shown in Figure 4.8 which has 365 exercise dates and up to 100 exercise opportunities, our C++ implementation requires about 10 minutes to complete the calculation on an Intel P4, 3.4GHz, and for a grid of 120×60 points in x and y direction, respectively. The same computation but with no spikes and a grid of 120×1 points only takes about one second.




Figure 4.8: Value of a one year swing option per exercise right. Market parameters of the underlying are as before, see Table 3.1, i.e. $\alpha = 7$, $\beta = 200$, $\sigma = 1.4$, $\lambda = 4$, $J \sim \text{Exp}(1/\mu_J)$ with $\mu_J = 0.4$, $f(t) = 0$, $r = 0$, and initial conditions $X_0 = 0$ and $Y_0 = 0$. The swing contract delivers over a time period of one year $T \in [0, 1]$ with up to 100 call rights and a strike price of $K = 1$, where a right can be exercised on any day. As a comparison the price of the same swing option is plotted but where the underlying does not exhibit spikes, i.e. $\lambda = 0$.

Based on Figure 4.8 we make two observations. First, the price per exercise right decreases with the number of exercise rights. This is the correct qualitative behaviour one would expect because n swing options each with one exercise right⁸ only, offer more flexibility than one swing option with n exercise rights.⁹ Second, the premium added due to the spike risk is much more profound for options with small numbers of exercise rights than for large one. This is also intuitively clear, as an option with say 100 exercise rights will mainly be used against high prices caused by the diffusive part and only a very few against spiky price explosions.

In Figure 4.9 we show how sensitive swing option prices are to changes in market parameters. There we consider a swing option with a duration of 60 days and up to 20 exercise opportunities. In each graph we change one parameter by 20% up and down. The most significant change is caused by a change in the volatility parameter σ . Note, the longterm variance of the mean-reverting process (X_t) is $\frac{\sigma^2}{2\alpha}$ and we expect some direct relationship between the long-term variance and the option price. Hence, a change in the mean-reversion parameter α is inversely proportional to the price and quantitatively changes the price less than the volatility σ . The mean-reversion parameter β of the spike process has a similar ef-

⁸This is actually an American option.

⁹The rights of a swing option can only be exercised one at a time.




Figure 4.9: Sensitivity of swing option prices with respect to model parameters. A swing option with 60 exercise dates and up to 20 rights is considered, where the red curve is based on the parameters of Figure 4.8. In each graph one market parameter is shifted up by 20% (green line) and down by 20% (blue line). We always plot the option price divided by the number of exercise rights.

fect on the option price as α has, but where the influence slightly decreases with the number of options. This is consistent with previous observations of the impact of jumps on option prices as seen in Figure 4.8. This effect is much more clearly visible for the other jump parameters λ and μ_J which have the greatest impact on options with only a few exercise rights. For one exercise right, a 20% change in the jump size μ_J has an even greater effect on the price than a 20% change in volatility σ . A possible explanation is that we deal with very heavy tailed jump size distributions.

4.4.4 Dimension reduction

In practical situations, the computational time for pricing options is of crucial importance. It is unlikely that a trader will be happy to wait minutes to get values for an option, especially if it is of a very basic structure as the ones considered here.

The computational time of the grid algorithm is proportional to the square of the number of grid points and so a reduction of grid points will cause a significant improvement in the execution time. Here we document our first attempts to reduce the dimensionality of the problem in order to considerably reduce the number of grid points. Some of the assumptions we make on the way turn out to be quite vague, though intuitive, but as a consequence we cannot necessarily expect the method to produce the right results. It is still hoped that a refinement of the method could be of practical use under a certain range of market parameters. The reduction in computational time is considerable.

The general idea is to use an approximation of the conditional density of $S_{t+\Delta t}$ given X_t and Y_t as derived in Section 3.4.3, in particular see Remark 3.4.18. However, in order to make the problem one-dimensional we need to express $S_{t+\Delta t}$ in terms of S_t and this is where the main difficulty lies. We know the process $S_t = \exp(f(t) + X_t + Y_t)$ is not Markov, despite (X_t) and (Y_t) being Markov processes, and hence the distribution of $S_{t+\Delta t}$ not only depends on S_t but also on the entire history $(S_\tau)_{\tau \in [0, t]}$, given X_t and Y_t are not individually observable. Here we make our intuitive assumption and say if we observe $S_t = s$ then we know $X_t + Y_t = c$ with $c := \ln s - f(t)$ and then we simply assume that X_t and Y_t are given by their conditional expectations $X_t = c_1$ and $Y_t = c_2$ with

$$c_1 = \mathbb{E}^{\mathbb{Q}}[\bar{X} | \bar{X} + \bar{Y} = c], \quad c_2 = \mathbb{E}^{\mathbb{Q}}[\bar{Y} | \bar{X} + \bar{Y} = c] = c - c_1,$$

where $\bar{X}$ and $\bar{Y}$ are the respective stationary distributions. Numerical simulations appear to indicate that this approach produces correct conditional densities for small values of c but rather poor for very big values, see Figure 3.19.




Figure 4.10: Pricing a swing option using our dimension reduction approach. The market parameters as well as the swing option are identical to the those of Figure 4.9. The red line indicates the reference price obtained by the two-dimensional method of the previous section. The prices obtained by our reduced method, indicated by the green line, are between 10%-20% apart from the reference price. The sensitivity of the method towards a different choice of values for X_t and Y_t given $X_t + Y_t = c$ is indicated by the purple line.

We define the method by specifying the conditional densities:

$$\begin{aligned}
 f_{X_{t+\Delta t}+Y_{t+\Delta t}|X_t+Y_t=c}(x) &:= f_{X_{t+\Delta t}+Y_{t+\Delta t}|X_t=0,Y_t=0}(x - c_1 e^{-\alpha\Delta t} - c_2 e^{-\beta\Delta t}), \\
 f_{X_{t+\Delta t}+Y_{t+\Delta t}|X_t=0,Y_t=0}(x) &:= e^{-\lambda\Delta t} f_{X_{t+\Delta t}|X_t=0}(x) + (1 - e^{-\lambda\Delta t}) f_{X_{t+\Delta t}+J|X_t=0}(x), \\
 c_1 &:= \mathbb{E}^{\mathbb{Q}}[\bar{X}|\bar{X} + \bar{Y} = c], \quad c_2 := c - c_1.
 \end{aligned}$$

Numerical simulations show that the method produces qualitatively correct but quantitatively inaccurate results as can be seen in Figure 4.10. As it turns out, one of the main problems this method faces is its sensitivity to the particular choice of values c_1 and c_2 given we know $X_t + Y_t = c$. Using a very poor approximation to the conditional expectation $\mathbb{E}^{\mathbb{Q}}[\bar{X}|\bar{X} + \bar{Y} = c]$ like the one plotted in Figure 3.15 results in completely different values – prices which are by far smaller than the cost of the same options without the premium for the spike risk. This is also intuitively clear as this choice overestimates the relative size of the spike c_2 and underestimates the diffusion part c_1 and as c_2 reverts with a high rate β back to zero, it would appear as if the model was mainly driven by the high rate β and not α .

This sensitivity and the fact that we only know approximations to the conditional expectation $\mathbb{E}^{\mathbb{Q}}[\bar{X}|\bar{X} + \bar{Y} = c]$, see Figure 3.18, might partially explain the errors we observe in the result. Another source of errors might come from the approximate distribution of




Figure 4.11: Valuation of a swing option with ten exercise opportunities per day. The option is the same as in Figure 4.10 but we leave the number of exercise opportunities at 60 and so the duration of the contract becomes 6 instead of 60 days.

$X_{t+\Delta t} + Y_{t+\Delta t}$ given X_t and Y_t which is only valid for very small Δt . Figure 4.11 shows that if the time-step size is reduced to one tenth of a day, the method works much better and errors are generally less than 6% of the reference value. Other situations where the one-dimensional method works also fairly well are those when the mean-reversion rate α gets close to the spike mean-reversion rate β and when the jump intensity λ is close to zero.

4.4.5 General swing contracts

Swing contracts considered so far were of the simplest form. We will now extend it and allow to exercise a certain volume per exercise right and impose overall volume constraints. Assume the payoff per right is given by $g(S_t, t)$ if exercised at time t , as an example we keep in mind a fraction of a call option payoff $g(S_t, t) = \gamma(S_t - K)^+$. Each time an exercise right is executed we can decide how many integer units $u \in \{u_{\min}, u_{\min} + 1, \dots, u_{\max}\}$ of the option g to exercise and so the payoff will be $ug(S_t, t)$. The overall number of units will be limited by $U_{\max}$ and we might impose a penalty if not at least $U_{\min}$ units were exercised over the duration of the contract. Let the value of such a swing contract be given by $V(n, u, s, t)$ if $n \in \mathbb{N}_0$ exercise rights and $u \in \mathbb{N}_0$ units are left and given the price of the

underlying is $S_t = s$, then the dynamic optimisation problem becomes

$$V(n, u, s, t) = \max \left\{ e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n, u, S_{t+\Delta t}, t + \Delta t) | S_t = s], \right. \\ \left. e^{-r\Delta t} \mathbb{E}^{\mathbb{Q}} [V(n-1, u-k, S_{t+\Delta t}, t + \Delta t) | S_t = s] + kg(s, t) : u_{\min} \leq k \leq u_{\max} \right. \\ \left. V(0, u, s, t) = 0, \quad V(n, 0, s, t) = 0, \right. \\ \left. V(n, u, s, T) = 0, \text{ if } u \leq U_{\max} - U_{\min}, \text{ otherwise } V(n, u, s, T) = -\text{penalty}. \right.$$

This is a generalised version of (4.11) and the additional dimension in the volume u increases the complexity of the algorithm roughly in the order of $U_{\max}$.

4.5 PIDE formulation

In order to derive the partial-integro-differential equation (PIDE) the value function of a derivative has to fulfil, we use the ideas of the Feynman-Kac theorem, which states equivalence between expectation values and PIDEs.

Let the dynamics of the spot price in the risk-neutral probability measure $\mathbb{Q}$ be given by (4.2). We rewrite the equations to express the seasonal function f as a time-dependent mean reversion level of the process (X_t) and according to Remark B.1.3 we obtain

$$\begin{aligned} S_t &= \exp(X_t + Y_t), \\ dX_t &= \alpha(\mu(t) - X_t) dt + \sigma dW, \\ dY_t &= -\beta Y_t dt + J_t dN_t, \\ \mu(t) &:= f(t) + \frac{f'(t)}{\alpha}. \end{aligned}$$

The price of a contingent claim with payoff $g(S_T)$ at maturity T in an arbitrage-free market is

$$V(x, y, t) = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[g(S_T) | X_t = x, Y_t = y],$$

and hence $e^{-rt} V(X_t, Y_t, t)$ is a martingale under the filtration created by (X_t) and (Y_t) . From Itô's formula it follows that

$$d(e^{-rt} V(X_t, Y_t, t)) = e^{-rt} (-rV dt + dV),$$

with

$$\begin{aligned} dV(X_t, Y_t, t) &= \left(\frac{\partial V}{\partial t} + \alpha(\mu(t) - X_t) \frac{\partial V}{\partial x} - \beta Y_t \frac{\partial V}{\partial y} + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2} \right) dt \\ &\quad + \sigma \frac{\partial V}{\partial x} dW_t + (V(X_t, Y_t + J_t, t) - V(X_t, Y_t, t)) dN_t. \end{aligned} \tag{4.13}$$

The process $e^{-rt} V$ is a martingale if its drift component minus the compensation of the jump component is zero. The compensation of the jump process of V according to Remark C.1.1 is given by

$$-\lambda \int_0^t \mathbb{E}[V(X_t, Y_t + J_t, t) - V(X_t, Y_t, t) | X_t, Y_t] dt.$$

We finally conclude that $e^{-rt} V$ is a martingale if the following equation is satisfied:

$$\frac{\partial V}{\partial t} + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2} + \alpha(\mu(t) - x) \frac{\partial V}{\partial x} - \beta y \frac{\partial V}{\partial y} + \lambda \mathbb{E}[V(x, y + J, t) - V(x, y, t)] = rV.$$

With a given density function f_J of the jump size, the expectation value can be rewritten and we obtain the integro-differential equation every value function for a contingent claim with payoff $g(S_T)$ has to satisfy:

$$\frac{\partial V}{\partial t} + \frac{\sigma^2}{2} \frac{\partial^2 V}{\partial x^2} + \alpha(\mu(t) - x) \frac{\partial V}{\partial x} - \beta y \frac{\partial V}{\partial y} + \lambda \int_{\mathbb{R}} (V(x, y + z, t) - V(x, y, t)) f_J(z) dz = rV, \quad (4.14)$$

subject to the terminal condition $V(x, y, T) = g(e^{x+y})$.

4.6 Hedging contingent claims

As it is not possible to store electricity, the only way to hedge the risk-exposure of a derivative is to use another derivative. However, it will not be possible to totally eliminate the risk because the model allows for jumps with random sizes. To reduce the risk one could use a hedge minimising the variance of the hedged portfolio or simply hedge the diffusive risk.¹⁰ We illustrate the latter strategy as it is consistent with the particular pricing measure, where expectations over the jumps are the same as under the real world measure.¹¹

Let $V(x, y, t)$ be the value of an option, to be hedged with another option maturing at the same time with value $U(x, y, t)$. We consider the self financing portfolio $\alpha M_t + \delta U_t$ and require its diffusive part to be the same as that of $V(X_t, Y_t, t)$. As it is self-financing, the change in the portfolio-value is $\alpha dM_t + \delta dU_t$. The diffusive parts are, according to (4.13), given by $\delta \frac{\partial U}{\partial x}$ and $\delta \frac{\partial V}{\partial x}$ for the portfolio and the option V , respectively. The hedging strategy is therefore

$$\delta = \frac{\partial V}{\partial x} \bigg/ \frac{\partial U}{\partial x}.$$

¹⁰Terms in front of the Brownian motion in the sde.

¹¹The jump risk is said to be not priced in.

To illustrate this hedging strategy, we choose a forward contract with exercise price K as the hedge U . For obvious arbitrage reasons the value has to be

$$U(x, y, t) = (F(t, T, x, y) - K) e^{-r(T-t)},$$

where $F(t, T, x, y) := \mathbb{E}^{\mathbb{Q}}[S_T | X_t = x, Y_t = y]$ is the forward price. It follows from Corollary 3.4.13 that $\frac{\partial F}{\partial x} = e^{-\alpha(T-t)} F$ and hence $\frac{\partial U}{\partial x} = e^{-\alpha(T-t)} e^{-r(T-t)} F$. For the case where $K = 0$ this simplifies and we obtain the hedging strategy

$$\delta = \frac{e^{\alpha(T-t)}}{U} \frac{\partial V}{\partial x},$$

i.e. the amount of money invested in the hedge is $e^{\alpha(T-t)} \frac{\partial V}{\partial x}$.

We also note that the risks involved in issuing an option also very much depend on whether the issuer actually possesses the ability to produce electricity for some fixed price p , i.e. whether the issuer owns a power plant. Assume a bank sells a call option paying $(S_T - K)^+$ at maturity T . We expect S_T to have a heavy tailed distribution due to the spike risk and so the bank faces the potential risk of high losses if it does not employ an efficient hedging strategy. However, if it is able to supply its own electricity at a fixed cost p it will be able to sell it for S_T in the market and so its profit and loss would be $-(S_T - K)^+ + (S_T - p)^+ = K - p$ for large S_T . In that case the bank is on a safe side even without using an appropriate hedging strategy which in turn relies on model assumptions.

Chapter 5

Equivalent martingale measures

In an arbitrage-free market the price of a derivative is given as the discounted expected payoff with the expectation taken under an equivalent martingale measure $\mathbb{Q}$. In incomplete markets where derivatives cannot be perfectly replicated there is no unique equivalent martingale measure and hence no unique arbitrage-free price. It is the intention of this chapter to give further guidance on which particular measure to pick and why. For stochastic volatility models the theory is quite well established and therefore presented in Section 5.2. This is followed by our first attempts to examine optimal martingale measures in the setting of a simple incomplete electricity market.

5.1 Introduction

If a unique martingale measure $\mathbb{Q}$ cannot be found, one idea of choosing a particular one is to look for the measure $\mathbb{Q}^*$ which is closest to the real world measure $\mathbb{P}$ within the set of all equivalent martingale measures $\mathcal{M}$.

We define “closeness” as some property of the $\mathcal{F}_T$ measurable Radon-Nikodým derivative Π defined by $d\mathbb{Q} = \Pi d\mathbb{P}$. A very well studied class of optimal pricing measures in the literature are the q -optimal measures minimising the q -th moment of the Radon-Nikodým derivative. For $q > 1$ we say

$$\mathbb{E} \left[\left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)^q \right] \rightarrow \min, \quad \mathbb{Q} \in \mathcal{M}.$$

The definition can be extended to make sense for all $q \in \mathbb{R}$.

Definition 5.1.1

Let Π be the Radon-Nikodým derivative so that $d\mathbb{Q} = \Pi d\mathbb{P}$ then define for $q \in \mathbb{R} \setminus \{0, 1\}$

$$H_q(\mathbb{P}, \mathbb{Q}) := \begin{cases} \mathbb{E} \left[\frac{q}{q-1} \Pi^q \right] & \text{if } \mathbb{Q} \ll \mathbb{P} \\ \infty & \text{otherwise,} \end{cases}$$

and for $q \in \{0, 1\}$ define

$$H_q(\mathbb{P}, \mathbb{Q}) = \begin{cases} \mathbb{E} [(-1)^{1+q} \Pi^q \ln \Pi] & \text{if } \mathbb{Q} \ll \mathbb{P} \\ \infty & \text{otherwise.} \end{cases}$$

Any measure $\mathbb{Q} \in \mathcal{M}$ minimising $H_q(\mathbb{P}, \mathbb{Q})$ is then called a q -optimal measure. For $q = 0$ the optimal measure is also called minimal reverse entropy measure, for $p = 1$ it is called minimal relative entropy measure and for $q = 2$ we have the variance optimal measure.

The importance of q -optimal martingale measures is reinforced by the fact that there are strong links with utility indifference pricing. We only give a short sketch of the links, for a more detailed discussion see [Monoyios, 2006, Section 2.2] and [Davis, 1997]. Assume zero interest and let U be a utility function, then we define the maximal expected utility given initial wealth z by

$$u(z) := \sup_{\pi \in \mathcal{A}} \mathbb{E}[U(Z_T^\pi) | Z_0 = z], \quad (5.1)$$

where Z_t^π is the process of a self financing portfolio and $\mathcal{A}$ is the set of all admissible strategies. If the supremum is obtained by an optimal strategy with final value Z_T^* , one can show that under suitable smoothness conditions the marginal utility price $\hat{p}$ of a contingent claim B at time T is given by

$$\hat{p}(z) = \frac{\mathbb{E}[U'(Z_T^*)B | Z_0 = z]}{u'(z)}. \quad (5.2)$$

To get further insights one introduces the dual formulation to (5.1) by

$$v(\eta) := \inf_{\mathbb{Q} \in \mathcal{M}} \mathbb{E} \left[V \left(\eta \frac{d\mathbb{Q}}{d\mathbb{P}} \right) \right], \text{ with } V(\eta) := \sup_x (U(x) - x\eta), \quad \eta > 0, \quad (5.3)$$

where $\mathcal{M}$ is the set of all equivalent martingale measures. One can show that $v(\eta) = \sup_x (u(x) - x\eta)$ and furthermore if $\mathbb{Q}^*$ is the optimal martingale measure of (5.3) then

$$U'(Z_T^*) = u'(z) \frac{d\mathbb{Q}^*}{d\mathbb{P}},$$

which is fundamental as it allows us to rewrite (5.2) to obtain the simple pricing formula for marginal utility prices

$$\hat{p}(z) = \mathbb{E} \left[B \frac{d\mathbb{Q}^*}{d\mathbb{P}} | Z_0 = z \right] = \mathbb{E}^{\mathbb{Q}^*} [B | Z_0 = z],$$

i.e. the marginal utility price of a claim is given by the expected value of the claim under the equivalent martingale measure which minimises the dual problem (5.3). For power utility $U(x) = x^\gamma / \gamma$ we obtain $V(\eta) = -\eta^q / q$ with $q = -\gamma / (1 - \gamma)$, $\gamma < 1$, $\gamma \neq 0$, and hence the optimal measure $\mathbb{Q}^*$ minimises

$$\inf_{\mathbb{Q} \in \mathcal{M}} \mathbb{E} \left[-\frac{1}{q} \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)^q \right],$$

and therefore is a q -optimal measure for $q \in (0, 1)$. On the other hand, exponential utility, $U(x) = -e^{-\alpha x}$, $V(\eta) = \eta/\alpha(\ln(\eta/\alpha) - 1)$, leads to the relative entropy measure ($q = 1$) minimising

$$\inf_{\mathbb{Q} \in \mathcal{M}} \mathbb{E} \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \ln \left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right) \right].$$

See also [Delbaen et al., 2002] for a more detailed account. The case $q = 0$ corresponds to logarithmic utility. Furthermore, the variance optimal measure $q = 2$ is related to mean-variance hedging, see [Föllmer and Sondermann, 1986] and [Duffie and Richardson, 1991].

5.2 Pricing measures in a stochastic volatility models

In the case of stochastic volatility models, q -optimal measures are quite well understood and have been studied extensively for some particular choices of q as well as in general, see [Hobson, 2004]. In this section we state the main results of [Hobson, 2004], compare option prices with respect to q by deriving analytical ordering results as well as by numerical calculations. This has been part of collaborative research with Vicky Henderson, David Hobson and Sam Howison published in [Henderson et al., 2005a].

We assume the class of stochastic volatility models in a zero interest rate world is given by

$$\frac{dS_t}{S_t} = Y_t \lambda(S_t, Y_t, t) dt + Y_t dB_t, \quad dY_t = a(S_t, Y_t, t) dt + b(S_t, Y_t, t) dW_t,$$

where B and W are correlated Brownian motions with coefficient ρ . Introducing a B -independent Brownian motion Z we can rewrite $W_t = \rho B_t + \bar{\rho} Z_t$ with $\bar{\rho} := \sqrt{1 - \rho^2}$. Fixing a maturity T , then an equivalent martingale measure $\mathbb{Q}$ making S into a martingale can be represented by

$$d\mathbb{Q} = \Pi_T d\mathbb{P}, \quad \frac{d\Pi_t}{\Pi_t} = -\lambda(S_t, Y_t, t) dB_t - \psi(S_t, Y_t, t) dZ_t, \quad \Pi_0 = 1,$$

where ψ is a function satisfying certain regularity conditions. If it turns out to be convenient we might use the abbreviation ψ_t for $\psi(S_t, Y_t, t)$, the same applies for λ_t , a_t and b_t . The solution of Π_T is given by the Doléans exponential $\mathcal{E}$

$$\Pi_T = \mathcal{E}(-\lambda B - \psi Z)_T = \exp \left(- \int_0^T \frac{1}{2} (\lambda_t^2 + \psi_t^2) dt - \int_0^T \lambda_t dB_t - \int_0^T \psi_t dZ_t \right).$$

Based on Girsanov's Theorem, the processes $B^\mathbb{Q}$ and $Z^\mathbb{Q}$ are independent Brownian motions under $\mathbb{Q}$ if defined by

$$dB_t^\mathbb{Q} = dB_t + \lambda_t dt, \quad dZ_t^\mathbb{Q} = dZ_t + \psi_t dt,$$

and so the dynamics of the underlying process under $\mathbb{Q}$ is

$$\frac{dS_t}{S_t} = 0 dt + Y_t dB_t^{\mathbb{Q}}, \quad dY_t = (a_t - \rho\lambda_t b_t - \bar{\rho}\psi_t b_t) dt + b_t dW_t^{\mathbb{Q}},$$

with $dW_t^{\mathbb{Q}} = \rho dB_t^{\mathbb{Q}} + \bar{\rho} dZ_t^{\mathbb{Q}}$.

Hobson shows that in the special case of λ being a function of t only, or being adapted to the filtration generated by B , all q -optimal measures coincide with the minimal martingale measure which corresponds to $\psi_t = 0$. A more difficult problem arises if we allow all parameters of the dynamics to be Y dependent, i.e. assume

$$\frac{dS_t}{S_t} = Y_t \lambda(Y_t, t) dt + Y_t dB_t, \quad dY_t = a(Y_t, t) dt + b(Y_t, t) dW_t.$$

Hobson then shows that the q -optimal measure can be obtained by solving a partial differential equation (pde):

$$\begin{aligned} \psi(y, t) &= \bar{\rho} \frac{\partial f}{\partial y}(y, t) b(y, t), \\ \frac{\partial f}{\partial t} &= -\frac{1}{2} b^2 \frac{\partial^2 f}{\partial y^2} + (q\rho b\lambda - a) \frac{\partial f}{\partial y} + \frac{1}{2} R b^2 \left(\frac{\partial f}{\partial y} \right)^2 - \frac{q}{2} \lambda^2, \\ R &= 1 - q\rho^2. \end{aligned} \tag{5.4}$$

In general one might not be able to solve this non-linear pde analytically. However, it is possible to prove certain properties of ψ which will be useful to compare prices under different q -optimal measures. It is one of the main theorems of [Henderson et al., 2005a] which shows that $\psi(y, t) \geq 0$ if $q\lambda(y, t)^2$ is non-decreasing in y and $\psi(y, t) \leq 0$ if $q\lambda(y, t)^2$ is non-increasing in y . The case $\psi(y, t) = 0$ corresponds to the minimal martingale measure and since the main ordering theorem says option prices are decreasing in the market price of volatility risk ψ , we can now compare prices under the q -optimal measures with prices under the minimal martingale measure. The following theorem is based on [Henderson et al., 2005a, Corollary 3].

Theorem 5.2.1

Let $q\lambda(y, t)$ be strictly increasing in y for each $t \in [0, T]$ then option prices under the q -optimal measure are less than those under the minimal martingale measure.

Conversely, if $q\lambda(y, t)$ is strictly decreasing in y for each $t \in [0, T]$ then option prices under the q -optimal measure are greater than those under the minimal martingale measure.

For particular models we can make further statements about prices with respect to q . So would it be desirable to know under which conditions prices are increasing or decreasing in

q . For the remainder of the section we focus on the Heston stochastic volatility model. In its original form, [Heston, 1993], the model is given by

$$\begin{aligned}\frac{dS_t}{S_t} &= \mu dt + \sqrt{v_t} dB_t, \\ dv_t &= \kappa(\theta - v_t) dt + \xi\sqrt{v_t} dW_t, \\ dB_t dW_t &= \rho dt,\end{aligned}$$

with constant parameters, μ being the drift, κ is the rate of mean reversion of the variance process v_t which has a long term mean of θ , and ξ is called the volatility of volatility. The specification of the drift parameter μ of the spot price process did not matter in Heston's analysis as it becomes zero¹ under the risk neutral measure $\mathbb{Q}$. However, our aim is to find risk neutral measures as close as possible to the original measure and then the specification of the drift becomes important. With the choice of $\mu(v, t) = \mu_0\sqrt{v}$ prices under different q -optimal measures turn out to be identical. Below we will examine the case $\mu(v, t) = \mu_0 v$. First we rewrite the dynamics in terms of $Y_t = \sqrt{v_t}$ and obtain

$$\frac{dS_t}{S_t} = \mu_0 Y_t^2 dt + Y_t dB_t, \quad dY_t = \frac{\kappa}{2} \left(\frac{\bar{m}}{Y_t} - Y_t \right) dt + \frac{\xi}{2} dW_t, \quad \bar{m} = \theta - \frac{\xi^2}{4\kappa}. \quad (5.5)$$

In order to obtain a q -optimal measure the pde (5.4) needs to be solved with the functions $\lambda(y, t) = \mu_0 y^2$, $a(y, t) = \frac{\kappa}{2}(\frac{\bar{m}}{y} - y)$ and $b(y, t) = \frac{\xi}{2}$. The solution is given in [Hobson, 2004, Proposition 5.1] and based on the assumption that one can write $f(y, t) = y^2 H(T - t)/2 + G(T - t)$ with $H(0) = G(0) = 0$. Note, the change of measure is then given by $\psi(y, t) = \frac{\bar{\rho}\xi}{2} H(T - t)y$.

Proposition 5.2.2

The dynamics of the particular Heston stochastic volatility model (5.5) is given under the q -optimal measure by

$$\begin{aligned}\frac{dS_t}{S_t} &= 0 dt + Y_t dB_t^{\mathbb{Q}}, \\ dY_t &= \left(\frac{\kappa}{2} \left(\frac{\bar{m}}{Y_t} - Y_t \right) - \frac{\rho\xi\mu_0}{2} Y_t - \frac{\bar{\rho}^2\xi^2}{4} H(T - t) Y_t \right) dt + \frac{\xi}{2} dW_t^{\mathbb{Q}},\end{aligned}$$

where the function H is given on a case by cases basis as follows. Define

$$R = 1 - q\rho^2, \quad A^2 = |R| \frac{\xi^2}{4}, \quad B = \frac{2\kappa + 2q\rho\xi\mu_0}{\xi^2 |R|}, \quad D = q\mu_0^2 + \frac{(\kappa + q\rho\xi\mu_0)^2}{\xi^2 R}, \quad C = \sqrt{|D|}.$$

Case 1: $R = 0$; If $\kappa + q\rho\xi\mu_0 = 0$ then $H(t) = q\mu^2 t$, else

$$H(t) = \frac{q\mu^2}{\kappa + q\rho\xi\mu_0} \left(1 - e^{-(\kappa + q\rho\xi\mu_0)t} \right).$$

¹Or rather r if one considers a non-zero interest rate.

Case 2: $R > 0$;

$$H(t) = \frac{C}{A} \tanh \left(ACt + \tanh^{-1} \left(\frac{AB}{C} \right) \right) - B.$$

Case 3: $R < 0$; If $D > 0$ we have

$$H(t) = \frac{C}{A} \tan \left(ACt - \tan^{-1} \left(\frac{AB}{C} \right) \right) + B,$$

if $D < 0$ it is

$$H(t) = \frac{(A^2 B^2 - C^2)(e^{2ACt} - 1)}{2AC + A(AB + C)(e^{2ACt} - 1)},$$

and finally if $D = 0$ we have

$$H(t) = \frac{A^2 B^2 t}{1 + A^2 B t}.$$

Proof [Hobson, 2004, Proposition 5.1]. □

It follows from [Henderson et al., 2005a, Proposition 6] that the market price of volatility risk ψ is increasing in q for all q where the q -optimal measure is well defined and hence we conclude for the option price.

Proposition 5.2.3

In the stochastic volatility model (5.5) q -optimal prices for European options with convex payoffs are decreasing in q .

Proof [Henderson et al., 2005a, Corollary 7]. □

With this ordering result in mind we now turn to a quantitative analysis on how much option prices are changed by different choices of q . Transforming the volatility process Y back into a variance process v and using log-prices $X_t := \ln S_t$ yields the dynamics under the q -optimal measure

$$\begin{aligned} dX_t &= -\frac{1}{2}v dt + \sqrt{v_t} dB_t^{\mathbb{Q}}, \\ dv_t &= \left(\kappa(\theta - v_t) - \rho\xi\mu_0 v_t - \frac{\bar{\rho}^2 \xi^2}{2} H(T-t)v_t \right) dt + \xi dW_t^{\mathbb{Q}}, \end{aligned} \tag{5.6}$$

and so the variance process follows a mean reverting process with a time depending mean and a time depending speed of mean reversion. Let now the payoff of an option be given by the $\mathcal{F}_T$ measurable random variable O . We then define the value function $U(x, v, T-t)$ of O by

$$U(x, v, T-t) := \mathbb{E}^{\mathbb{Q}}[O | X_t = x, v_t = v],$$

Parameter	Description	Value
θ	long term variance	0.01
κ	rate of mean reversion	2.0
ξ	volatility of volatility	0.2
μ_0	absolute drift on asset $\mu_0 VS$	4.0
v_0	initial variance	0.01
S_0	initial asset price	1.0

Table 5.1: Parameters of the Heston model.

and so $U(X_t, Y_t, t)$ is a martingale under $\mathbb{Q}$. Applying Itô's formula and the fact that the drift has to be zero it follows that the value function satisfies the pde²

$$\begin{aligned}
 U_t &= \frac{1}{2}v (U_{xx} + \xi^2 U_{vv} + 2\rho\xi U_{xv}) + (r_d - r_f - \frac{1}{2}v)U_x + (\kappa(\theta - v) - \gamma(t)v)U_v - r_d U, \\
 \gamma(t) &= \rho\xi\mu_0 + \frac{\bar{\rho}^2 \xi^2}{2} H(t).
 \end{aligned}
 \tag{5.7}$$

The sub-indices of U indicate the derivative of U with respect to that variable. In the case of $\gamma(t) = 0$ and certain boundary conditions, including those corresponding to the valuation of call options, an explicit solution can be found which is based on Fourier inversion, see [Heston, 1993]. In other cases we rely on numerical calculations. We use a finite difference method (Crank-Nicolson) on a non-uniform grid as described in [Kluge, 2002]. Convergence of this method has been tested for call and barrier options ($\gamma = 0$) where an analytical solution was at hand. We use a sample market with parameters given in Table 5.1 representative of the dynamics of currency pairs in the foreign exchange market.

Figure 5.1 uses the above parameter values together with Strike $K = 1$ and $T = 1$. In the upper graph we plot the put price for $\rho \in [-0.5, -0.5]$ and $q \in [-4, 5]$. Over most of this range $qR > 0$ and the q -optimal measure exists for all time, and even for $q = 5$ and $|\rho| = 0.5$ the q -optimal measure exists up to $T = 1$. As anticipated by Proposition 5.2.3, we observe the put price is decreasing in q . The range of the graph represents about 16% difference in prices between the extreme points. If we examine special cases of moving from say $q = 0$ to $q = 2$, the price change is of the order of a couple of percent, depending on the correlation. This is a non-trivial difference, and highlights the fact that the choice of pricing measure has an impact on option prices. In the figure, put option prices are also observed to decrease with correlation. It turns out that this is the case for ‘small’ absolute values of q . Note that in the pricing pde (5.7), there are two drift terms arising from the incompleteness, $\rho\xi\mu_0$ and $\frac{\bar{\rho}^2 \xi^2}{2} H(t)$. In the small q case, the first of these is dominant. If correlation is negative,

²We have reintroduced the interest rates necessary to describe the foreign exchange market, where r_d denotes domestic and r_f foreign interest rate.




Figure 5.1: Price of a 1 year at-the-money put option under the Heston model for various values of ρ with parameters given in Table 5.1.




Figure 5.2: The effect of correlation and volatility of volatility on implied volatilities for a 1 year put option priced under the original Heston model.

this term has a positive effect on the option price, whilst the reverse is true for positive correlation. As the lower part of Figure 5.1 shows, once q is no longer small, the ordering reverses. This is the case as the drift term involving $\bar{\rho}^2 H(t)$ becomes dominant. Outside the range $qR > 0$, the function H may explode to infinity (for sufficiently large T) and prices are small as a result. Similarly, for $q = -5$ and $\rho = 0.5$, H explodes to $-\infty$ and prices are large.

One of the best ways of capturing the effects of a stochastic volatility model is by considering the implied volatilities and the true test of a model is whether it can be calibrated well to market data. By including the adjustment for volatility risk via the q measures, we have a richer class of models which may provide a better fit. We will not focus on calibration here, but rather on the shape of implied volatility smiles from the Heston model under q -optimal measures.

The effects of different parameters on the implied volatilities in the original Heston model (where $\gamma(t) = 0$ in the pricing pde (5.7)) are shown in Figure 5.2. We plot the implied volatilities against the strike K of the put option with parameters of Table 5.1 and $T = 1$ and $S_0 = 1$. The four smiles correspond to different choices of correlation ρ and volatility of volatility ξ . When $\rho = 0$, the smile is symmetric about the at-the-money volatility. Increasing ξ appears to increase the convexity of the smile. Non-zero correlation controls

the smile's asymmetry, important in equity markets. [Hakala and Wystup, 2002b] note the same qualitative effects.

In the final figure we consider the effect of changing the drift parameter μ_0 and the candidate pricing measure parameterised by q . We also consider the effects of varying maturity. In each case the implied volatilities are calculated from the Heston model with correlation $\rho = -0.2$. There are two graphs (one for changing q and one for changing μ_0) for each of three maturities. The key feature that will aid our understanding is that the market price of volatility risk $\psi(Y_t, t) = \frac{\bar{\rho}\xi}{2}H(T-t)Y_t$ is time-inhomogeneous. Since, for each q , $|H(\tau)|$ is increasing in τ , the effects of changing q will become more pronounced as the maturity increases.

We begin with some generic observations which are typical features for stochastic volatility models. The correlation is negative so the smiles are skewed to the left. As maturity increases, the smile becomes flatter or less convex – beware the change in the horizontal scale. By considering the left-hand column we see that as q increases, option prices decrease. This is consistent with Proposition 5.2.3. Conversely, in the right-hand column, we see that as μ_0 increases, the option price increases. This is a consequence of the drift term appearing in the pde (5.7). Since $\rho < 0$, the term $-\rho\xi\mu_0v$ is positive but the second term $\frac{\bar{\rho}^2\xi^2}{2}H(t)$ is negative for $\rho < 0$ and $q > 0$. The two terms will have competing effects, and the overall effect of change in μ_0 will depend on which term dominates. If H is small (q or T small), then the first term dominates and increasing μ_0 shifts the smile upward, as we see in Figure 5.3. However, if H is large then the second term dominates and by increasing μ_0 we expect the smile to shift down (not shown in the figure). Also note, if the correlation is positive both drift terms are positive and so the smile is also expected to shift down whilst μ_0 increases.

A final feature of the graphs that we wish to remark on is the relative magnitude of the implied volatility shifts as maturity changes. For the graphs in the right column, ($q = 0$ and $\mu_0 = 0$ or 4) the change in drift induced by the pricing measure is $-\rho\xi\mu_0v$. The magnitude of the implied volatility changes seems to approximately double each time T increases by a factor of 4. Conversely, on the left hand side, modification to the drift under $\mathbb{Q}$ consists of the two terms $-\rho\xi\mu_0v - \frac{\bar{\rho}^2\xi^2}{2}H(t)$ and as $H(t)$ is positive and increasing in t , the effect of changing q is comparatively larger when the maturity is large.




Figure 5.3: Implied volatility smiles for the Heston model with $\rho = -0.2$ and $T = 0.25$ (top row), $T = 1$ (middle row) and $T = 4$ (bottom row). The solid line in each graph corresponds $q = 0$, $\mu_0 = 4$. The lower line in the graphs in the left column correspond to a higher value of q ($q = 4$, $\mu = 4$) and the lower lines in the graphs on the right correspond to a lower value of μ ($q = 0$, $\mu_0 = 0$). Note that the horizontal scale changes with maturity by a factor of $\sqrt{T}$.

5.3 Pricing measures in a spot electricity model

We assume the underlying spot process cannot be used to hedge derivatives and hence can be considered a non-tradable asset in the sense of the fundamental theorem of option pricing. Hence the risk neutral measure can be any measure $\mathbb{Q}$ equivalent to the real world measure $\mathbb{P}$. Under these assumptions the problem of finding a q -optimal measure becomes trivial as $\mathbb{Q} = \mathbb{P}$ is a possibility and so $\mathbb{P}$ turns out to be the q -optimal risk neutral measure. The problem becomes interesting when we assume a liquid forward market in addition to the dynamics of the underlying. Forwards and futures can be seen as the simplest form of derivatives and the assumption of a liquid electricity forward market is not unrealistic. Assume today is $t = 0$ and we want to price an option maturing at T . Let the delivery price of a forward delivering at a time t be given by $F(t)$. The problem of finding the q -optimal measure is then a constrained optimisation of the form

$$\left\{ \mathbb{E} \left[\left(\frac{d\mathbb{Q}}{d\mathbb{P}} \right)^q \right] \rightarrow \min \right. \quad , \quad \mathcal{M} := \left\{ \mathbb{Q} \sim \mathbb{P} : \mathbb{E}^{\mathbb{Q}}[S_t | \mathcal{F}_0] = F(t), \forall t \in [0, T] \right\}.$$

Although it is very easy to formulate, the problem seems to be hard. In models with traded and non-traded assets³, the constraint on $\mathbb{Q}$ to turn all traded assets into martingales can be parametrised in such a way to obtain an unconstrained optimisation problem. This does not seem to be feasible in our case. Therefore, we only present a solution of a special case. First we focus on $q \in \{0, 1\}$, i.e. the minimal reverse entropy measure and the minimal relative entropy measure, respectively, and later we will further restrict the set of possible equivalent measures $\mathcal{M}$. We define the measure change by

$$d\mathbb{Q} = \Pi_T d\mathbb{P}, \quad \frac{d\Pi_t}{\Pi_t} = -\psi_t dW_t, \quad \Pi_0 = 1.$$

Lemma 5.3.1

The q -optimal measures for $q \in \{0, 1\}$ parameterised by ψ (assuming squared integrability $\mathbb{E}[\psi_t^2]$) are given by minimising

$$\int_0^T \mathbb{E}[\psi_t^2] dt \rightarrow \min, \quad q = 0; \quad \int_0^T \mathbb{E}^{\mathbb{Q}}[\psi_t^2] dt \rightarrow \min, \quad q = 1,$$

each under the constraint of $\mathbb{Q} \in \mathcal{M}$.

Proof According to the definition of q -optimal measures, $q \in \{0, 1\}$, we have

$$\mathbb{E}[-\ln \Pi_T] \rightarrow \min, \quad q = 0; \quad \mathbb{E}[\Pi_T \ln \Pi_T] = \mathbb{E}^{\mathbb{Q}}[\ln \Pi_T] \rightarrow \min, \quad q = 1.$$

³The stochastic volatility model is one example where the spot is traded but not the volatility.

From

$$\begin{aligned} d \ln \Pi_t &= -\frac{1}{2} \psi_t^2 dt - \psi_t dW_t \\ &= \frac{1}{2} \psi_t^2 dt - \psi_t dW_t^{\mathbb{Q}}, \end{aligned}$$

and taking expectations of Π_T the statement follows directly. $\square$

We now specify a model for the dynamics of the electricity price without spikes:

$$\begin{aligned} S_t &= \exp(f(t) + X_t), \\ dX_t &= -\alpha X_t + \sigma dW_t, \end{aligned}$$

and further restrict the set of martingale measures $\mathcal{M}$ so that the risk neutral spot price dynamics is still given by an exponential Ornstein-Uhlenbeck process with a constant speed of mean reversion. Hence, we assume the market price of risk ψ_t is of the form

$$\psi_t = \frac{\hat{\alpha} - \alpha}{\sigma} X_t - \frac{\hat{\alpha}}{\sigma} \mu(t).$$

From Girsanov's theorem we know $dW_t^{\mathbb{Q}} = dW_t + \psi_t dt$ is a $\mathbb{Q}$ -Brownian motion and so the dynamics under $\mathbb{Q}$ is given by

$$\begin{aligned} S_t &= \exp(f(t) + X_t), \\ dX_t &= \hat{\alpha}(\mu(t) - X_t) + \sigma dW_t^{\mathbb{Q}}. \end{aligned}$$

Lemma 5.3.2

Given the above parameterisation of $\psi_t = \frac{\hat{\alpha} - \alpha}{\sigma} X_t - \frac{\hat{\alpha}}{\sigma} \mu(t)$, the constraint

$$\mathbb{E}^{\mathbb{Q}}[S_t | \mathcal{F}_0] = F(t), \quad \forall t \in [0, T],$$

is equivalent to

$$\hat{\alpha} \mu(t) = \hat{\alpha}(\ln F(t) - f(t)) + \frac{\partial}{\partial t}(\ln F(t) - f(t)) - \frac{\sigma^2}{4}(1 + e^{-2\hat{\alpha}t}), \quad \forall t \in [0, T].$$

Proof The dynamics under $\mathbb{Q}$ can also be written as (see Remark B.1.3)

$$\begin{aligned} S_t &= \exp(f(t) + f_1(t) + X_t), & f_1'(t) + \hat{\alpha} f_1(t) &= \hat{\alpha} \mu(t) \\ dX_t &= -\hat{\alpha} X_t + \sigma dW_t^{\mathbb{Q}}, \end{aligned}$$

with the expectation being

$$\mathbb{E}^{\mathbb{Q}}[S_t] = \exp \left(f(t) + f_1(t) + X_0 e^{-\hat{\alpha}t} + \frac{\sigma^2}{4\hat{\alpha}}(1 - e^{-2\hat{\alpha}t}) \right),$$

and so the constraint $\mathbb{E}^{\mathbb{Q}}[S_t] = F(t)$ is equivalent to

$$f_1(t) = \ln F(t) - f(t) - X_0 e^{-\hat{\alpha}t} - \frac{\sigma^2}{4\hat{\alpha}}(1 - e^{-2\hat{\alpha}t}),$$

and as $\hat{\alpha}\mu(t) = f'_1(t) + \hat{\alpha}f_1(t)$ the result follows. $\square$

As the drift μ is completely determined by the forward curve, finding an optimal risk neutral measure is equivalent to an unconstrained optimisation in the speed of mean reversion parameter $\hat{\alpha}$. For the minimal reverse entropy measure we therefore have

$$\int_0^T \mathbb{E}[\psi_t^2] dt = \left(\frac{\hat{\alpha} - \alpha}{\sigma} \right)^2 \int_0^T \mathbb{E}[X_t^2] dt + 2 \frac{\hat{\alpha} - \alpha}{\sigma^2} \int_0^T \hat{\alpha}\mu(t) \mathbb{E}[X_t] dt + \frac{1}{\sigma^2} \int_0^T (\hat{\alpha}\mu(t))^2 dt,$$

and with $\mathbb{E}[X_t] = X_0 e^{-\alpha t}$, $\mathbb{E}[X_t^2] = \frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha t}) + X_0^2 e^{-2\alpha t}$ and inserting the expression for μ we obtain a function with linear terms in $\hat{\alpha}$, $\hat{\alpha}^2 t$, $e^{-2\hat{\alpha}t}$, $e^{-4\hat{\alpha}}$ and cross-terms to be minimised over $\hat{\alpha} > 0$ which can be solved using numerical optimisation methods.

Chapter 6

Outlook and extensions

The stochastic process introduced in this thesis is capable of capturing the apparent properties of electricity spot price time series, that is to say mean-reversion, seasonality and spikes. Despite the additional complication of an independent jump process, the model is still analytically tractable in so far as we can obtain closed form solutions (up to integrals) of expectation values leading to option prices for path-independent options. For call and put options on forwards with delivery periods we are still able to provide approximate solutions. Approximations of the probability distributions of the spike process make it possible to introduce a grid-based method to price path dependent options like swing contracts numerically. The performance of the algorithm is not yet fast enough to be acceptable to a trader and our attempts to introduce a very fast dimensional reduced form did not fully satisfy accuracy requirements. Further investigations into that method are required in order to assess improvements. Extrapolation methods might also be useful to infer the price of an option with many exercise opportunities based on the same option with fewer rights. It is shown in [Howison, 2005] that for an Bermudan option we have the asymptotic expansion $V(n, S, t) = V(\infty, S, t) - W(S, t)/n + O(1/n^{5/2})$. Knowing the price of two different n_1 and n_2 one could determine $V(\infty, S, t)$ and $W(S, t)$ and then extrapolate to any value of n . It is possible that a similar relationship might hold for swing options, too.

We also give guidance on which particular risk neutral measure to choose as the market is incomplete. Since the main constraint of $\mathbb{Q}$ being consistent with the observed forward curve does not suffice to determine a unique measure, we examine q -optimal measures in that setting. However we are only able to give solutions in the very simplest cases. More research needs to be done to find more general solutions.

Finally, extensions to the stochastic spot price model might also be of interest for further research. A natural development would be to include stochastic volatility and a stochastic

seasonality component and would lead to a model like

$$\begin{aligned}
 S_t &= \exp(f(t) + X_t + Y_t + Z_t), \\
 dX_t &= -\alpha X_t dt + \sqrt{V_t} dW_t^{(1)}, \\
 dV_t &= \kappa(\bar{v} - V_t) dt + \xi \sqrt{V_t} dW_t^{(2)}, \\
 dZ_t &= \mu dt + \sigma dB_t, \\
 dY_t &= -\beta Y_{t-} dt + J_t dN_t,
 \end{aligned}$$

with $W^{(1)}$, $W^{(2)}$ and B being Brownian motions with some form of dependency. In addition, the intensity of the Poisson process might be assumed to be time-dependant. Although such a model might be a much better fit to historical data, and to the forward price dynamics, it is likely to be considerably harder to calibrate and to use to price options with it.

Despite the remaining open questions we believe that this work provides a comprehensive overview of a particular stochastic model suitable to describe electricity price dynamics, for the purpose of option pricing.

Appendix A

Elementary probability

Due to the importance of a normally $\mathcal{N}(0, 1)$ distributed random variable we denote its density and distribution by

$$\varphi(x) := \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}},$$

$$N(x) := \int_{-\infty}^x \varphi(y) dy = \frac{1}{2} + \frac{1}{2} \operatorname{erf} \left(\frac{x}{\sqrt{2}} \right).$$

Table A shows some properties of a few distributions.

A.1 Products and sums

Proposition A.1.1 (Product of independent random variables)

Let X and Y be two independent random variables where Y has a density and X has a density except at the point 0 where it has a positive probability $\mathbb{P}(Y = 0) = p$, i.e.

$$F_X(x) = p \mathbb{1}_{x \geq 0} + \int_{-\infty}^x f_X(y) dy,$$

$$F_Y(x) = \int_{-\infty}^x f_Y(y) dy,$$

Distribution	symbol	density	mgf(θ)	$\mathbb{E}$	var
Exponential	$\operatorname{Exp}(\lambda)$	$\lambda e^{-\lambda x} \mathbb{1}_{x \geq 0}$	$\frac{\lambda}{\lambda - \theta}, (\theta < \lambda)$	$\frac{1}{\lambda}$	$\frac{1}{\lambda^2}$
Gamma	$\Gamma(\lambda, \alpha)$	$\frac{1}{\Gamma(\alpha)} \lambda^\alpha x^{\alpha-1} e^{-\lambda x} \mathbb{1}_{x > 0}$	$(\frac{\lambda}{\lambda - \theta})^\alpha, (\theta < \lambda)$	$\frac{\alpha}{\lambda}$	$\frac{\alpha}{\lambda^2}$
Lognormal		$\frac{1}{x} f_{\mathcal{N}}(\ln(x))$		$e^{\mu + \sigma^2/2}$	$e^{2\mu + \sigma^2} (e^{\sigma^2} - 1)$
Normal	$\mathcal{N}(\mu, \sigma^2)$	$\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$	$\exp(\mu\theta + \frac{1}{2}\sigma^2\theta^2)$	μ	σ^2

Table A.1: Some properties of distributions.

then the distribution of the product XY is given by

$$\begin{aligned} F_{XY}(c) &= p\mathbf{1}_{c \geq 0} + \int_{-\infty}^c f_{XY}(x) dx, \\ f_{XY}(c) &= \int_{-\infty}^{\infty} f_X(c/x) \frac{f_Y(x)}{|x|} dx. \end{aligned} \tag{A.1}$$

Proof We have

$$\begin{aligned} F_{XY}(c) &= \mathbb{P}(XY \leq c) = \mathbb{P}(Y \cdot 0 \leq c) \mathbb{P}(X = 0) + \mathbb{P}(XY \leq c \text{ and } X \neq 0) \\ &= p\mathbf{1}_{c \geq 0} + \mathbb{P}(XY \leq c \text{ and } Y \neq 0), \end{aligned}$$

and

$$\begin{aligned} \mathbb{P}(XY \leq c \text{ and } Y \neq 0) &= \int_{\{xy \leq c\}} f_X(x) f_Y(y) dx dy \\ &= \int_{-\infty}^0 \int_{\frac{c}{y}}^{\infty} f_X(x) f_Y(y) dx dy + \int_0^{\infty} \int_{-\infty}^{\frac{c}{y}} f_X(x) f_Y(y) dx dy \\ &= \int_{-\infty}^0 \int_{-\infty}^c -\frac{1}{y} f_X(x/y) f_Y(y) dx dy + \int_0^{\infty} \int_{-\infty}^c \frac{1}{y} f_X(x/y) f_Y(y) dx dy \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^c \frac{1}{|y|} f_X(x/y) f_Y(y) dx dy \\ &= \int_{-\infty}^c \int_{-\infty}^{\infty} f_X(x/y) \frac{f_Y(y)}{|y|} dy dx. \end{aligned}$$

□

Proposition A.1.2 (Sum of independent random variables)

Let X and Y be two independent random variables where X has a density and Y has a density except at the point 0 where it has a positive probability $\mathbb{P}(X = 0) = p$, i.e.

$$\begin{aligned} F_X(x) &= \int_{-\infty}^x f_X(y) dy, \\ F_Y(x) &= p\mathbf{1}_{x \geq 0} + \int_{-\infty}^x f_Y(y) dy, \end{aligned}$$

then $X + Y$ has a density and it is given by

$$f_{X+Y}(c) = pf_X(c) + \int_{-\infty}^{\infty} f_Y(c-x) f_X(x) dx. \tag{A.2}$$

Proof We have

$$\begin{aligned} F_{X+Y}(c) &= \int_{-\infty}^{\infty} F_Y(c-x) f_X(x) dx, \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{c-x} f_Y(y) f_X(x) dy dx + \int_{-\infty}^{\infty} p\mathbf{1}_{x \leq c} f_X(x) dx, \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^c f_Y(y-x) f_X(x) dy dx + pF_X(c), \end{aligned}$$

and changing the order of integration yields the density

$$f_{X+Y}(c) = \int_{-\infty}^{\infty} f_Y(c-x)f_X(x) \, dx + pf_X(c).$$

□

Lemma A.1.3

Let $X \sim \mathcal{N}(\mu, \sigma^2)$ and $Y \sim \text{Exp}(\lambda)$ be a normally and exponentially distributed random variable, respectively. Then the probability density function of the sum $X + Y$ is given by

$$f_{X+Y}(x) = \lambda e^{\frac{1}{2}\lambda^2\sigma^2 + \lambda(\mu-x)} \mathcal{N}\left(\frac{x-\mu}{\sigma} - \lambda\sigma\right).$$

Proof The density of the sum of two random variables is the convolution of their densities, i.e.

$$f_{X+Y}(c) = \int_{-\infty}^{\infty} f_Y(c-x)f_X(x) \, dx,$$

which is

$$\begin{aligned} f_{X+Y}(c) &= \frac{\lambda}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^c e^{-\lambda(c-x)} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \, dx \\ &= \lambda e^{\frac{1}{2}\lambda^2\sigma^2 + \lambda(\mu-c)} \left(1 - \mathcal{N}\left(\lambda\sigma + \frac{\mu-c}{\sigma}\right)\right). \end{aligned}$$

□

A.2 Conditional expectations

Lemma A.2.1

Let $X \sim \mathcal{N}(0, 1)$ be a standard normally distributed random variable then

$$\mathbb{E}[X|X \leq c] = -\frac{e^{-\frac{c^2}{2}}}{\sqrt{2\pi} \mathcal{N}(c)}$$

Proof We know

$$\mathbb{E}[X|X \leq c] = \frac{\mathbb{E}[X\mathbb{1}_{X \leq c}]}{\mathbb{P}(X \leq c)},$$

and with

$$\mathbb{E}[X\mathbb{1}_{X \leq c}] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^c x e^{-\frac{x^2}{2}} \, dx = \frac{1}{\sqrt{2\pi}} \int_{\infty}^{\frac{c^2}{2}} e^{-z} \, dz = -\frac{e^{-\frac{c^2}{2}}}{\sqrt{2\pi}},$$

and

$$\mathbb{P}(X \leq c) = \mathcal{N}(c),$$

we obtain the desired result.

□

Corollary A.2.2

Let $X \sim \mathcal{N}(\mu, \sigma^2)$ be a normally distributed random variable then

$$\mathbb{E}[X|X \leq c] = \mu - \sigma \frac{\varphi\left(\frac{c-\mu}{\sigma}\right)}{\mathcal{N}\left(\frac{c-\mu}{\sigma}\right)}.$$

Proof Define $Y := \frac{1}{\sigma}(X - \mu)$ then $Y \sim \mathcal{N}(0, 1)$ and so

$$\mathbb{E}[X|X \leq c] = \mathbb{E}\left[\mu + \sigma Y | Y \leq \frac{c-\mu}{\sigma}\right] = \mu + \sigma \mathbb{E}\left[Y | Y \leq \frac{c-\mu}{\sigma}\right]$$

□

Lemma A.2.3

Let $X \sim \mathcal{N}(\mu_1, \sigma_1^2)$ and $Y \sim \mathcal{N}(\mu_2, \sigma_2^2)$ two independent normally distributed random variables then

$$\mathbb{E}[X|X + Y = c] = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} \left(c - \left(\mu_2 - \frac{\sigma_2^2}{\sigma_1^2} \mu_1 \right) \right).$$

Proof For any two random variables with a density we have

$$\mathbb{E}[X|X + Y = c] = \int_{\mathbb{R}} x f_{X|X+Y=c}(x) dx = \frac{\int_{\mathbb{R}} x f_X(x) f_Y(c-x) dx}{\int_{\mathbb{R}} f_X(x) f_Y(c-x) dx}.$$

With

$$f_{X|X+Y=c}(x) = f_{X|Y=c-x}(x) = \frac{f_{X,Y}(x, c-x)}{\int_{\mathbb{R}} f_{X,Y}(x, c-x) dx}$$

and given the random variables are independent we obtain

$$\mathbb{E}[X|X + Y = c] = \frac{\int_{\mathbb{R}} x f_X(x) f_Y(c-x) dx}{\int_{\mathbb{R}} f_X(x) f_Y(c-x) dx}.$$

As X and Y are normally distributed we finally obtain

$$\mathbb{E}[X|X + Y = c] = \frac{\sigma_1^2(c - \mu_2) + \sigma_2^2\mu_1}{\sigma_1^2 + \sigma_2^2}.$$

□

Lemma A.2.4

Let $X \sim \mathcal{N}(\mu, \sigma^2)$ and $Y \sim \text{Exp}(\lambda)$ be two independent normally and exponentially distributed random variables, respectively, then with $\tilde{X} := X + \lambda\sigma^2$ we have

$$\mathbb{E}[X|X + Y = c] = \mathbb{E}[\tilde{X}|\tilde{X} \leq c] = \mu + \lambda\sigma^2 - \sigma \frac{\varphi\left(\frac{c-\mu-\lambda\sigma^2}{\sigma}\right)}{\mathcal{N}\left(\frac{c-\mu-\lambda\sigma^2}{\sigma}\right)},$$

and in particular

$$\lim_{c \rightarrow \infty} \mathbb{E}[X|X + Y = c] = \mu + \lambda\sigma^2.$$

Proof Given the densities

$$f_X(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad f_Y(x) = \lambda e^{-\lambda x} \mathbf{1}_{x \geq 0},$$

and by straight forward computation we obtain

$$\begin{aligned} \mathbb{E}[X|X+Y=c] &= \frac{\int_{\mathbb{R}} x f_X(x) f_Y(c-x) dx}{\int_{\mathbb{R}} f_X(x) f_Y(c-x) dx} \\ &= \frac{\int_{-\infty}^c x \exp\left(-\frac{(x-\mu)^2}{2\sigma^2} - \lambda(c-x)\right) dx}{\int_{-\infty}^c \exp\left(-\frac{(x-\mu)^2}{2\sigma^2} - \lambda(c-x)\right) dx} \\ &= \frac{\int_{-\infty}^c x \exp\left(-\frac{(x-\mu-\lambda\sigma^2)^2}{2\sigma^2}\right) dx}{\int_{-\infty}^c \exp\left(-\frac{(x-\mu-\lambda\sigma^2)^2}{2\sigma^2}\right) dx} \\ &= \mathbb{E}[\tilde{X}|\tilde{X} \leq c], \end{aligned}$$

where $\tilde{X} \sim \mathcal{N}(\mu + \lambda\sigma^2, \sigma^2)$ is some normal random variable. $\square$

Lemma A.2.5

Let X and Y be two independent random variables with X having a density and Y having a density except at the point 0 where $P(Y=0) = p$, i.e.

$$F_Y(c) = \int_{-\infty}^c f_Y(x) dx + p \mathbf{1}_{c \geq 0},$$

then the conditional expectation is given by

$$\mathbb{E}[X|X+Y=c] = \frac{\int_{-\infty}^{\infty} x f_X(x) f_Y(c-x) dx + p c f_X(c)}{f_{X+Y}(c)}. \quad (\text{A.3})$$

Proof First note that

$$\begin{aligned} F_{X,X+Y}(a,b) &= \mathbb{P}(X \leq a \text{ and } Y \leq b - X) \\ &= \int_{-\infty}^a \int_{-\infty}^b f_X(x) f_Y(y-x) dy dx + \int_{-\infty}^a p f_X(x) \mathbf{1}_{x \leq b} dx. \end{aligned}$$

For any two random variables X and Z a function g which satisfies

$$\int_A X(\omega) d\mathbb{P}(\omega) = \int_A g(Z(\omega)) d\mathbb{P}(\omega), \quad \forall A \in \mathcal{Y}^{-1}(\mathcal{B}(\mathbb{R})).$$

is a conditional expectation and we write $g(c) = \mathbb{E}[X|Z=c]$. Here, $\mathcal{B}(\mathbb{R})$ denotes the set of all Borel sets in $\mathbb{R}$. This equation is equivalent to

$$\int_a^b \int_{\mathbb{R}} x d\mathbb{P}_{X,Z}(x,z) = \int_a^b g(z) d\mathbb{P}_Z(z), \quad \forall a < b.$$

Now, let $Z := X + Y$ and so the condition becomes

$$\int_a^b \int_{\mathbb{R}} x f_X(x) f_Y(z-x) dx dz + \int_a^b p x f_X(x) dx = \int_a^b g(z) f_{X+Y}(z) dz, \quad \forall a < b,$$

which is satisfied if

$$g(c)f_{X+Y}(c) = \int_{-\infty}^{\infty} x f_X(x) f_Y(c-x) dx + p c f_X(c).$$

□

Lemma A.2.6

Let X be lognormally distributed, i.e. $\log X \sim \mathcal{N}(\mu, \sigma^2)$ then

$$\mathbb{E} \left[(aX + b) \mathbf{1}_{X \in [K, \hat{K}]} \right] = a e^{\mu + \frac{1}{2}\sigma^2} \left(\mathbf{N}(d_1) - \mathbf{N}(\hat{d}_1) \right) + b \left(\mathbf{N}(d_2) - \mathbf{N}(\hat{d}_2) \right),$$

with

$$d_2 = \frac{\mu - \log K}{\sigma}, \quad \hat{d}_2 = \frac{\mu - \log \hat{K}}{\sigma},$$

and $d_1 = d_2 + \sigma$, $\hat{d}_1 = \hat{d}_2 + \sigma$.

Lemma A.2.7

Let X be lognormally distributed, i.e. $\log X \sim \mathcal{N}(\mu, \sigma^2)$ then

$$\begin{aligned} \mathbb{E} \left[(X - K)^+ \right] &= e^{\mu + \frac{1}{2}\sigma^2} \mathbf{N}(d_1) - K \mathbf{N}(d_2) \\ \mathbb{E} \left[(K - X)^+ \right] &= e^{\mu + \frac{1}{2}\sigma^2} (\mathbf{N}(d_1) - 1) - K (\mathbf{N}(d_2) - 1) \end{aligned}$$

with

$$d_2 = \frac{\mu - \log K}{\sigma}, \quad d_1 = d_2 + \sigma.$$

Appendix B

The Ornstein-Uhlenbeck process

A stochastic process X_t which satisfies the stochastic differential equation (sde)

$$dX_t = \alpha(\mu - X_t) + \sigma dW_t \quad (\text{B.1})$$

is called an Ornstein-Uhlenbeck (OU) process with the speed of mean-reversion α , longterm level μ and volatility σ . The parameters can also be functions of time.

B.1 Solution of the sde

Assuming a constant speed of mean-reversion α but allowing for variable longterm level $\mu(t)$ and volatility $\sigma(t)$, the sde can be easily solved by multiplying X_t with $e^{\alpha t}$. Defining $Y_t := e^{\alpha t} X_t$ and applying Itô's formula yields

$$dY_t = e^{\alpha t} (\alpha \mu(t) dt + \sigma(t) dW_t), \quad (\text{B.2})$$

which can be directly integrated. The solution simplifies in the constant coefficient case and is stated in the following Lemma.

Lemma B.1.1 (Solution of an OU process)

Let X_t be a constant coefficient OU process (B.1) then its unique solution is

$$X_t = \mu + (X_0 - \mu) e^{-\alpha t} + \sigma \int_0^t e^{-\alpha(t-s)} dW_s. \quad (\text{B.3})$$

In particular, if $X_0 = x_0$ is not random, X_t is normally distributed with

$$X_t \sim \mathcal{N} \left(\mu + (x_0 - \mu) e^{-\alpha t}, \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha t}) \right). \quad (\text{B.4})$$

Given the knowledge of the state at any time $s < t$, $X_s = x_s$, the conditional distribution is

$$X_t \sim \mathcal{N} \left(\mu + (x_s - \mu) e^{-\alpha(t-s)}, \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha(t-s)}) \right). \quad (\text{B.5})$$

The solution of the OU sde (B.5) also gives an intuitive understanding of the parameters of the process. It says, on average deviations from the mean μ are damped down by a factor of $e^{-\alpha \Delta t}$. A characteristic quantity describing this behaviour is the half-time $\tau_{1/2}$ defined as the time in which the deviation is halved: $\frac{1}{2} = e^{-\alpha \tau_{1/2}}$ or equivalently $\tau_{1/2} = \frac{\ln 2}{\alpha}$. The effect of the volatility σ becomes clear in the long-term behaviour of X_t , i.e. $t \rightarrow \infty$. Then we have

$$X_t \sim \mathcal{N}\left(\mu, \frac{\sigma^2}{2\alpha}\right), \quad (t \rightarrow \infty),$$

and so the standard deviation in the long-term is $\frac{\sigma}{\sqrt{2\alpha}}$.

Remark B.1.2 (Simulation)

Equation (B.5) can be utilised to simulate realisations of the process. For fixed s, t ($s < t$) and an independent, normally distributed random variable $\xi \sim \mathcal{N}(0, 1)$ we can define

$$X_t = \mu + e^{-\alpha(t-s)}(X_s - \mu) + \sqrt{\frac{\sigma^2}{2\alpha}(1 - e^{-2\alpha(t-s)})} \xi.$$

Remark B.1.3 (Relation with seasonality)

Let X be an OU process with zero mean

$$dX_t = -\alpha X_t dt + \sigma dW_t$$

and $f : [0, \infty) \rightarrow \mathbb{R}$ be a twice differentiable function. It is easy to show that the process $Y_t := f(t) + X_t$ is also an OU process but with a time dependent average:

$$dY_t = \alpha(\mu(t) - Y_t) dt + \sigma dW_t, \quad \mu(t) = f(t) + \frac{f'(t)}{\alpha}.$$

Lemma B.1.4 (Moment generating function)

The moment generating function of an OU process X_t with constant coefficients and initial value $X_0 = x_0$ is

$$\Phi_X(\theta, t) := \mathbb{E} e^{\theta X_t} = \exp\left(\theta \mu + \theta(x_0 - \mu) e^{-\alpha t} + \theta^2 \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha t})\right). \quad (\text{B.6})$$

Proof Based on (B.4), X_t is normally distributed with some mean m and standard deviation s . Therefore, the random variable $\exp(\theta X_t)$ is lognormally distributed with mean $\exp(\theta m + \frac{1}{2} \theta^2 s^2)$. $\square$

B.2 Parameter estimation

Given a discrete time series $\{x_0, x_1, \dots, x_n\}$ as a realisation of a stochastic process, a commonly used procedure to estimate unknown parameters of the process is the maximum likelihood method (ML). Thereby, the parameters are chosen to maximise the joint density $f_{X_{t_0}, \dots, X_{t_n}}(x_0, \dots, x_n)$. For Markov processes, the joint density can be represented as a product of transitional densities:

$$f_{X_{t_0}, \dots, X_{t_n}}(x_{t_0}, \dots, x_{t_n}) = f_{X_{t_0}}(x_0) f_{X_{t_1}|X_{t_0}=x_0}(x_1) \cdots f_{X_{t_n}|X_{t_{n-1}}=x_{n-1}}(x_n). \quad (\text{B.7})$$

If transitional densities are not known explicitly, martingale estimation functions can be used to obtain approximate estimations and even if transitional densities are known, the method could provide analytic expressions, whereas the ML method might not. A good overview of the method is given in [Bibby and Sorensen, 1995].

Due to (B.5), we know the transitional density of an OU process. Keeping in mind the density of a $\mathcal{N}(m, \sigma^2)$ distributed random variable, which is $\frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-m)^2}{2\sigma^2}\right)$, the transition density of an OU process with constant coefficients (B.1) is given by

$$f_{X_t|X_s=y}(x) = \sqrt{\frac{\alpha}{\pi\sigma^2(1-e^{-2\alpha(t-s)})}} \exp\left(-\frac{\alpha(x-\mu-e^{-\alpha(t-s)}(y-\mu))^2}{\sigma^2(1-e^{-2\alpha(t-s)})}\right).$$

The log-likelihood function, defined as the logarithm of the joint density (B.7) then is¹

$$L(\mu, \sigma, \alpha) = \frac{1}{2} \sum_{i=1}^n \ln\left(\frac{\alpha}{\pi\sigma^2(1-e^{-2\alpha\Delta t_i})}\right) - \frac{\alpha}{\sigma^2} \sum_{i=1}^n \frac{(x_i - \mu - e^{-\alpha\Delta t_i}(x_{i-1} - \mu))^2}{1 - e^{-2\alpha\Delta t_i}} \quad (\text{B.8})$$

with $\Delta t_i := t_i - t_{i-1}$. The estimated parameters $\hat{\mu}$, $\hat{\sigma}$ and $\hat{\alpha}$ given by the ML method are the solution to the optimisation $L(\mu, \sigma, \alpha) \rightarrow \max$, subject to the constraints $\alpha > 0$ and $\sigma > 0$. Numerical algorithms are capable of solving three-dimensional optimisation problems. However, by checking first order criteria for optimality, further insight can be gained. For a maximum in the interior we require the first derivatives to be zero:

$$\begin{aligned} \frac{\partial L}{\partial \mu} &= \frac{2\alpha}{\sigma^2} \sum_{i=1}^n \frac{x_i - e^{-\alpha\Delta t_i} x_{i-1} - (1 - e^{-\alpha\Delta t_i})\mu}{1 + e^{-\alpha\Delta t_i}} = 0, \\ \frac{\partial L}{\partial \sigma} &= -\frac{n}{\sigma} + \frac{2\alpha}{\sigma^3} \sum_{i=1}^n \frac{(x_i - \mu - e^{-\alpha\Delta t_i}(x_{i-1} - \mu))^2}{1 - e^{-2\alpha\Delta t_i}} = 0, \\ \frac{\partial L}{\partial \alpha} &= 0. \end{aligned}$$

¹We have assumed that the distribution of X_{t_0} is either unknown or $X_{t_0} = x_0$ is deterministic.

The last derivative is quite complex and an analytic solution does not seem to exist, but the first two equations can be solved explicitly, provided α or its estimate is known:

$$\hat{\mu} = \frac{\sum_{i=1}^n \frac{x_i - e^{-\alpha \Delta t_i} x_{i-1}}{1 + e^{-\alpha \Delta t_i}}}{\sum_{i=1}^n \frac{1 - e^{-\alpha \Delta t_i}}{1 + e^{-\alpha \Delta t_i}}},$$

$$\hat{\sigma}^2 = \frac{2\alpha}{n} \sum_{i=1}^n \frac{(x_i - \hat{\mu} - e^{-\alpha \Delta t_i} (x_{i-1} - \hat{\mu}))^2}{1 - e^{-2\alpha \Delta t_i}}.$$

These two estimates can be inserted into $L(\mu, \sigma, \alpha)$ in order to reduce the three-dimensional optimisation into a one-dimensional problem. Further simplifications can be made if μ is known and time-steps are small and equidistant, $\Delta t_i = \Delta t$. The log-likelihood function is then

$$L(\mu, \sigma, \alpha) = n \ln \left(\frac{\alpha}{\pi \sigma^2 (1 - e^{-2\alpha \Delta t})} \right) - \frac{\alpha}{\sigma^2 (1 - e^{-2\alpha \Delta t})} \sum_{i=1}^n (x_i - \mu - e^{-\alpha \Delta t} (x_{i-1} - \mu))^2$$

$$\approx n \ln \left(\frac{1}{2\pi \sigma^2 \Delta t} \right) - \frac{1}{2\sigma^2 \Delta t} \sum_{i=1}^n (x_i - \mu - e^{-\alpha \Delta t} (x_{i-1} - \mu))^2$$

where the approximation $1 - e^{-2\alpha \Delta t} \approx 2\alpha \Delta t$ has been made, and hence

$$\frac{\partial L}{\partial \alpha} \approx -\frac{1}{\sigma^2 \Delta t} \sum_{i=1}^n (x_i - \mu - e^{-\alpha \Delta t} (x_{i-1} - \mu)) \alpha e^{-\alpha \Delta t} (x_{i-1} - \mu),$$

from which the estimation of the mean-reversion rate α follows:

$$\hat{\alpha} = -\frac{1}{\Delta t} \ln \left(\frac{\sum_{i=1}^n (x_i - \mu)(x_{i-1} - \mu)}{\sum_{i=1}^n (x_{i-1} - \mu)^2} \right),$$

$$\hat{\sigma}^2 = \frac{2\alpha}{n(1 - e^{-2\alpha \Delta t})} \sum_{i=1}^n (x_i - \mu - e^{-\alpha \Delta t} (x_{i-1} - \mu))^2.$$

Due to the approximations made for small Δt it is not immediately clear how efficient the parameter estimation is. However, the method based on martingale estimation function also obtains the same expression for $\hat{\alpha}$, see [Bibby and Sorensen, 1995].

Appendix C

Transform analysis

The following overview about transform analysis of affine jump-diffusion processes is based on [Duffie et al., 2000]. The method allows us to determine the moment generating function of the jump process by solving an ode. Furthermore, a formula is provided leading to the expectation value of a call-option payoff. For the sake of simplicity we only consider processes (Z_t) of the type

$$dZ_t = -\alpha Z_t dt + \sigma dW_t + J_t dN_t, \quad (\text{C.1})$$

keeping in mind that the spot price process is $S_t = \exp(Z_t)$.

C.1 Moment generating function

In short, they define a more general moment generating function which is conditioned on the value at a time t :

$$\Psi_Z(\theta, t, T, x) := \mathbb{E} \left[e^{\theta Z_T} \mid Z_t = x \right].$$

It is easy to show that $\Psi_Z(\theta, t, T, Z_t)$ is a martingale and hence its drift component needs to be zero. This condition yields a pde which can be solved explicitly using the ansatz

$$\Psi(\theta, t, T, x) = \exp(a(t) + b(t)x),$$

given θ and T are fixed, i.e. a and b might also depend on θ and T . Before going into the details we state a general result on how to compensate a jump process to obtain a martingale.

Remark C.1.1 (Compensation of a jump process)

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be some function, $(Z_t)_{t \in \mathbb{R}^+}$ be a left-continuous stochastic process and $(J_t)_{t \in \mathbb{R}^+}$ an iid process (J_s and J_t are iid for all $s \neq t$). Furthermore let $(N_t)_{t \in \mathbb{R}^+}$ denote a

Poisson process with intensity λ , and t_i the random jump times. Under certain regularity conditions, the process

$$\sum_{i=1}^{N_t} f(Z_{t_i-}, J_{t_i}) - \lambda \int_0^t \Gamma_J(Z_s) ds,$$

with $\Gamma_J(z) := \mathbb{E}[f(z, J)]$, is a martingale.

Proof We only give the idea. For the conditional expectation of the jump-part we find

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^{N_T} f(Z_{t_i-}, J_{t_i}) \middle| \mathcal{F}_t \right] - \sum_{i=1}^{N_t} f(Z_{t_i-}, J_{t_i}) &= \mathbb{E} \left[\sum_{t < t_i \leq T} \mathbb{E}[f(Z_{t_i-}, J_{t_i}) | t_i, Z_{t_i-}] \middle| \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\sum_{t < t_i \leq T} \Gamma_J(Z_{t_i-}) \middle| \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\int_t^T \Gamma_J(Z_{s-}) dN_s \middle| \mathcal{F}_t \right] \\ &= \mathbb{E} \left[\int_t^T \Gamma_J(Z_s) \lambda ds \middle| \mathcal{F}_t \right], \end{aligned}$$

and hence

$$\mathbb{E} \left[\sum_{i=1}^{N_T} f(Z_{t_i-}, J_{t_i}) - \int_0^T \Gamma_J(Z_s) \lambda ds \middle| \mathcal{F}_t \right] = \sum_{i=1}^{N_t} f(Z_{t_i-}, J_{t_i}) - \int_0^t \Gamma_J(Z_s) \lambda ds.$$

□

To see which conditions a and b have to satisfy in order to obtain a martingale $\Psi_t := \Psi(\theta, t, T, Z_t) = \exp(a(t) + b(t)Z_t)$ we apply Itô's formula and obtain

$$d\Psi_t = \Psi_t \left(\dot{a}(t) + \frac{1}{2} \sigma^2 b(t)^2 + (\dot{b}(t) - \alpha b(t)) Z_t \right) dt + \Psi_t b(t) \sigma dW_t + \Psi_{t-} (e^{b(t)J} - 1) dN_t.$$

According to Remark C.1.1 (see also [Duffie et al., 2000, Appendix A] for a proof), the jump component minus a drift

$$\sum_{i=0}^{N_t} \Psi_{t_i-} (e^{b(t_i)J} - 1) - \lambda \int_0^t \Psi_s (\Phi_J(b(s)) - 1) ds$$

is a martingale. Therefore, Ψ_t is a martingale if its drift added to the drift component above is zero, i.e.

$$\dot{a}(t) + \frac{1}{2} \sigma^2 b(t)^2 + \lambda (\Phi_J(b(s)) - 1) + (\dot{b}(t) - \alpha) Z_t = 0,$$

which is satisfied if

$$\begin{aligned} \dot{a}(t) &= -\frac{1}{2} \sigma^2 b(t)^2 - \lambda (\Phi_J(b(s)) - 1), \\ \dot{b}(t) &= \alpha b(t), \end{aligned} \tag{C.2}$$

subject to the boundary condition $a(T) = 0$ and $b(T) = \theta$, because $\Psi_T = \mathbb{E}[e^{\theta Z_T} | Z_T] = e^{\theta Z_T}$. Given Φ_J is continuous, the unique solution to the system of ode's is obviously

$$\begin{aligned} a(t) &= \theta^2 \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)}) + \lambda \int_t^T \Phi_J(\theta e^{-\alpha(T-s)}) - 1 \, ds, \\ b(t) &= \theta e^{-\alpha(T-t)}. \end{aligned}$$

The moment generating function of Z therefore is

$$\begin{aligned} \Phi_Z(\theta, t) &= \Psi(\theta, 0, t, z_0) = \exp(a(0) + b(0)z_0) \\ &= \exp\left(\theta^2 \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha t}) + \lambda \int_0^t \Phi_J(\theta e^{-\alpha s}) - 1 \, ds + \theta e^{-\alpha t} z_0\right). \end{aligned}$$

C.2 Expectation of a call-option payoff

In order to formulate an important result which will allow us to determine the expected payoff of a call option, we need to define the well-behavedness property.

Definition C.2.1

We call the process (Z_t) from (C.1) well-behaved at (θ, T) if there exists a unique solution to (C.2) and

1. $\mathbb{E} \left[\int_0^T |\Psi_t(\Phi_J(b(t))) - 1| \, dt \right] < \infty,$
2. $\mathbb{E} \left[\sqrt{\int_0^T (\Psi_t b(t))^2 \, dt} \right] < \infty,$
3. $\mathbb{E}[|\Psi_T|] < \infty.$

Define

$$G_{\theta,c}(K, T, z_0) := \mathbb{E} \left[e^{\theta Z_T} \mathbb{1}_{cZ_T \leq K} \right],$$

then the expected payoff of a call option can be written as

$$\mathbb{E}[S_T - K]^+ = \mathbb{E}[S_T \mathbb{1}_{S_T \geq K}] - \mathbb{E}[K \mathbb{1}_{S_T \geq K}] = G_{1,-1}(-\ln K, T, z_0) - K G_{0,-1}(-\ln K, T, z_0).$$

We cite the result from [Duffie et al., 2000].

Proposition C.2.2

Suppose, for fixed $T > 0$, $\theta \in \mathbb{R}$ and $c \in \mathbb{R}$, that Z is well-behaved at $(\theta + i\nu c, T)$ for any $\nu \in \mathbb{R}$ and that $\int_{\mathbb{R}} |\Phi_Z(\theta + i\nu c, T)| \, d\nu < \infty$. Then $G_{\theta,c}(\cdot, T, z_0)$ is well defined and given by

$$G_{\theta,c}(K, T, z_0) = \frac{1}{2} \Phi_Z(\theta, T) - \frac{1}{\pi} \int_0^\infty \frac{1}{\nu} \Im \left(\Phi_Z(\theta + i\nu c, T) e^{-i\nu \theta} \right) \, d\nu,$$

where Φ_Z denotes the complex valued moment generating function and $\Im(z)$ the imaginary part of any value $z \in \mathbb{C}$.

Finally, we give an alternative proof which is based on the proof of Lévy's Inversion Theorem.

Proposition C.2.3

Let Z be a random variable and its truncated moment generating function be defined by

$$G_\nu(x) := \mathbb{E} [e^{\nu Z} \mathbf{1}_{\{Z \leq x\}}] = \int_{-\infty}^x e^{\nu y} dF_Z(y).$$

If the moment generating function $\Phi(\nu + i\theta)$ exists for some $\nu \in \mathbb{R}$ and all $\theta \in \mathbb{R}$ then

$$G_\nu(x) = \frac{\Phi(\nu)}{2} - \frac{1}{\pi} \int_0^\infty \frac{\Im(\Phi(\nu + i\theta) e^{-i\theta x})}{\theta} d\theta.$$

Proof We show

$$\frac{1}{\pi} \int_0^\infty \frac{\Im(\Phi(\nu + i\theta) e^{-i\theta x})}{\theta} d\theta = \frac{\Phi(\nu)}{2} - G_\nu(x). \quad (\text{C.3})$$

Define

$$\begin{aligned} I_c &:= \frac{1}{\pi} \int_0^c \frac{\Im(\Phi(\nu + i\theta) e^{-i\theta x})}{\theta} d\theta \\ &= \frac{1}{\pi} \int_0^c \frac{\Im(\int_{\mathbb{R}} e^{(\nu+i\theta)y} dF_X(y) e^{-i\theta x})}{\theta} d\theta \\ &= \frac{1}{\pi} \int_0^c \int_{\mathbb{R}} \frac{\Im(e^{(\nu+i\theta)y} e^{-i\theta x})}{\theta} dF_X(y) d\theta \\ &= \frac{1}{\pi} \int_0^c \int_{\mathbb{R}} \frac{e^{\nu y} \sin(\theta(y-x))}{\theta} dF_X(y) d\theta. \end{aligned}$$

The integral is uniformly convergent and so we may change the order of integration and obtain

$$I_c = \frac{1}{\pi} \int_{\mathbb{R}} e^{\nu y} \int_0^c \frac{\sin(\theta(y-x))}{\theta} d\theta dF_X(y).$$

Because the integral with respect to θ is continuous in c and bounded we obtain for the limit

$$\lim_{c \rightarrow \infty} I_c = \frac{1}{\pi} \int_{\mathbb{R}} e^{\nu y} \int_0^\infty \frac{\sin(\theta(y-x))}{\theta} d\theta dF_X(y),$$

where the inner integral has the solution

$$\int_0^\infty \frac{\sin(\theta(y-x))}{\theta} d\theta = \frac{\pi}{2} \text{sign}(y-x) = \frac{\pi}{2} \begin{cases} -1 & y-x < 0 \\ 0 & y-x = 0 \\ 1 & y-x > 0 \end{cases}$$

which yields the desired result

$$\begin{aligned} \lim_{c \rightarrow \infty} I_c &= \frac{1}{2} \int_{\mathbb{R}} e^{\nu y} \text{sign}(y-x) dF_X(y) \\ &= \frac{1}{2} \left(\int_{\{y>x\}} e^{\nu y} dF_X(y) - \int_{\{y<x\}} e^{\nu y} dF_X(y) \right) \\ &= \frac{1}{2} (\Phi(\nu) - G_\nu(x) - G_\nu(x-)). \end{aligned}$$

□

Appendix D

Option pricing and the inability to hedge with the underlying

The inability to store electricity and hence the inability to hedge derivative contracts with the underlying is the main reason why the market is incomplete. Nevertheless, there has to be a consistency between prices of different options in order for the market to be free of arbitrage.

Based on the free-arbitrage principle, the following sections briefly describe how to price options if the underlying is an Itô process with one stochastic source

$$dS_t = \mu(S_t, t) dt + \sigma(S_t, t) dW_t,$$

and a money market account with a constant interest rate r

$$dM_t = M_t r dt.$$

For further details see [Björk, 1998, Chapter 10].

D.1 Risk neutral valuation

According to [Björk, 1998, Proposition 10.3], the prices V_t of all contingent claims are given by the discounted expected final payoff $g(S_T)$ under one equivalent measure $\mathbb{Q}$. The specific choice of the measure $\mathbb{Q}$ is given by the market. Once a measure has been chosen, all derivatives have to be priced under the same measure to avoid arbitrage opportunities.

Proposition D.1.1

Under the absence of arbitrage, the price of an option is given by

$$V_t = e^{-r(T-t)} \mathbb{E}^{\mathbb{Q}}[g(S_T)|S_t],$$

where $\mathbb{Q}$ is some equivalent measure characterised by a density Π_T with $d\mathbb{Q} = \Pi_T d\mathbb{P}$ defined by the sde

$$\frac{d\Pi_t}{\Pi_t} = -\psi(S_t, t) dW_t, \quad \Pi_0 = 1,$$

with some function ψ , called the market price of risk. Under the probability measure $\mathbb{Q}$, the process $W_t^\psi := W_t + \psi t$ is a Brownian motion¹, i.e. the drift under $\mathbb{Q}$ is altered by $-\psi\sigma$:

$$dS_t = (\mu(S_t, t) - \psi(S_t, t)\sigma(S_t, t)) dt + \sigma(S_t, t) dW_t^\psi.$$

Remark D.1.2

In a complete market where hedging with the underlying is possible, the risk neutral measure $\mathbb{Q}$ is uniquely determined. Without going into the details of the derivation, it is required that $e^{-rt} \Pi_t S_t$ is a martingale and hence its drift has to be zero. It follows that $\psi = \frac{\mu - rS_t}{\sigma}$ and therefore the drift of S under $\mathbb{Q}$ is always the same as the money market account, i.e. rS_t .

D.2 Valuation by a pde approach

Although it is not possible to replicate any contingent claim as we are only left with a money market account to invest in, prices of different derivatives have to be consistent. Let us assume (V_t) and (W_t) are the stochastic processes denoting the option prices of two options and the dynamics are given by

$$dV_t = \mu^V(S_t, t) dt + \sigma^V(S_t, t) dW_t,$$

$$dW_t = \mu^W(S_t, t) dt + \sigma^W(S_t, t) dW_t.$$

Now, the idea is to construct a self financing and risk-less portfolio consisting of the derivatives V and W . If no arbitrage exists in the market, the drift of the portfolio value has to be the same as that of the money market account.

If the value of the portfolio at time t is $\alpha_t V_t + \beta_t W_t$, the change in value is then $\alpha_t dV_t + \beta_t dW_t$ because the portfolio is self financing. In order for the portfolio to be risk-less and perform as well as the money market account we require

$$\alpha_t dV_t + \beta_t dW_t = r(\alpha_t V_t + \beta_t W_t) dt.$$

Comparing the terms in front of the Wiener process yields $\alpha = -\frac{\sigma^W}{\sigma^V} \beta$. Using this relation and comparing the drift terms gives $-\sigma^W \mu^V + \sigma^V \mu^W = r(-\sigma^W V + \sigma^V W)$ or equivalently

$$\frac{\mu^W - rW}{\sigma^W} = \frac{\mu^V - rV}{\sigma^V},$$

and the Proposition 10.1 from [Björk, 1998] follows immediately.

¹Girsanov's theorem, see [Karatzas and Shreve, 1991, Section 3.5].




Figure D.1: Risk and return. The slope is determined by the market price of risk ψ .

Proposition D.2.1

Assume that the market for derivatives is free of arbitrage. Then there exists a universal stochastic process ψ such that, with probability 1, and for all t , we have

$$\frac{\mu^V(S_t, t) - rV_t}{\sigma^V(S_t, t)} = \psi_t, \quad (\text{D.1})$$

regardless of the specific choice of the derivative V .

Assuming the price of the option is given by a function, solely depending on the current price of the underlying S_t and time t , expressions for μ^V and σ^V and a pde can be derived. With slight abuse of notation we write $V_t = V(S_t, t)$ and hence

$$dV(S_t, t) = \left(\frac{\partial V}{\partial t} + \frac{\partial V}{\partial s} \mu + \frac{1}{2} \frac{\partial^2 V}{\partial s^2} \sigma^2 \right) dt + \frac{\partial V}{\partial s} \sigma dW_t.$$

Inserting the expressions for μ^V and σ^V into (D.1) yields

$$\frac{\partial V}{\partial t} + (\mu - \psi\sigma) \frac{\partial V}{\partial s} + \frac{1}{2} \sigma^2 \frac{\partial^2 V}{\partial s^2} - rV = 0, \quad (\text{D.2})$$

where the arguments s and t have been suppressed to allow for a compact notation.

Remark D.2.2 (Economic interpretation of the market price of risk)

The market price of risk has a deep economic meaning. To understand this, we need to consider any derivative V as some asset. From equation (D.1) it follows that

$$\frac{\mu^V}{V_t} = r + \psi_t \frac{\sigma^V}{V_t},$$

which means that at a specific time t , the expected instantaneous excess return over r of any asset (which is a derivative on S_T) increases² linearly with its instantaneous risk. The higher the risk (or volatility) of an asset the higher the expected instantaneous excess return which is illustrated in Figure D.1.

Remark D.2.3 (Market price of risk of the risk-neutral and pde approach)

The market price of risk function ψ in both sections coincide. This can be seen by applying the Feynman-Kac formula to convert between expectations and pdes.

²Theoretically it could also decrease but that would not make sense from an economical point of view.

D.3 Solution for a mean-reverting process

We assume the spot price process follows in the risk neutral probability measure $\mathbb{Q}$ the mean reverting process

$$\begin{aligned} dX_t &= -\alpha X_t dt + \sigma dW_t, \\ S_t &= \exp(f(t) + X_t), \end{aligned} \tag{D.3}$$

with constants α and σ . By Itô's formula this it is equivalent to

$$\frac{dS_t}{S_t} = \alpha(\hat{\mu}(t) - \ln S_t) dt + \sigma dW_t, \quad \hat{\mu}(t) = f(t) + \frac{1}{\alpha} \frac{df}{dt}(t) + \frac{\sigma^2}{2\alpha}.$$

According to (B.4), the natural logarithm of S_t is normally distributed with mean $m(T) = f(T) + (X_t - f(t))e^{-\alpha(T-t)}$ and variance $v(t) = \frac{\sigma^2}{2\alpha} (1 - e^{-2\alpha(T-t)})$. The expectation value of S_T (which is equal to the forward price) and the expected payoff of a call option for a lognormally distributed random variable is known to be

$$\begin{aligned} \mathbb{E}[S_T|S_t] &= \exp\left(m(t) + \frac{1}{2}v(t)\right), \\ \mathbb{E}[(S_T - K)^+|S_t] &= \mathbb{E}[S_T \mathbf{1}_{S_T \geq K}|S_t] - K \mathbb{E}[\mathbf{1}_{S_T \geq K}|S_t] \\ &= \mathbb{E}[S_T|S_t] N(d_1) - K N(d_2), \end{aligned}$$

with the cumulative normal distribution $N(x)$ and

$$\begin{aligned} d_1 &:= \frac{-\ln K + m(T) + v(T)}{\sqrt{v(T)}}, \\ d_2 &:= \frac{-\ln K + m(T)}{\sqrt{v(T)}}. \end{aligned}$$

Inserting the expressions for $m(T)$ and $v(T)$, and substituting $f(t)$ by $s(t) = e^{f(t)}$ yields the formulas for the forward price F and the call option value V

$$\begin{aligned} F &= \mathbb{E}[S_T|S_t] \\ &= s(T) \left(\frac{S_t}{s(t)}\right)^{e^{-\alpha(T-t)}} \exp\left(\frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)})\right) \\ V &= e^{-r(T-t)} \mathbb{E}[(S_T - K)^+|S_t] \\ &= e^{-r(T-t)} \left(s(T) \left(\frac{S_t}{s(t)}\right)^{e^{-\alpha(T-t)}} \exp\left(\frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)})\right) N(d_1) - K N(d_2) \right) \end{aligned}$$

with

$$\begin{aligned} d_1 &:= \frac{\ln\left(\frac{s(T)}{K}\right) + \ln\left(\frac{S_t}{s(t)}\right) e^{-\alpha(T-t)} + \frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)})}{\sqrt{\frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)})}}, \\ d_2 &:= \frac{\ln\left(\frac{s(T)}{K}\right) + \ln\left(\frac{S_t}{s(t)}\right) e^{-\alpha(T-t)}}{\sqrt{\frac{\sigma^2}{4\alpha} (1 - e^{-2\alpha(T-t)})}}. \end{aligned}$$




Figure D.2: Option prices of a call option for a mean-reverting process.

To give an intuitive understanding of the pricing formula, figure D.2 shows the prices of a call option for the parameters $\alpha = 1$, $\sigma = 0.7$ and $r = 0$. In the first graph, the initial spot price S_0 changes along the x -axis and the seasonality function is kept constant $f(t) = 1$. It can be seen that for longer term options the initial value of the underlying is almost unimportant, which makes sense as the mean-reverting term of the sde forces deviations back to the mean after a certain time. In the second graph, the seasonal value $s(T)$ changes along the x -axis, whereas the initial value is kept equal to the seasonal value at time zero, i.e. $S_0 = s(0)$. The deviation of the option price from the payoff profile for deep in the money calls can be explained by the fact that the expectation of S_T is not equal to $s(T)$ but $s(T)$ multiplied by a term depending on time to maturity.

We finally state the pde to be satisfied by the value function of the option $V(S, t)$:

$$\frac{\partial V}{\partial t} + \frac{1}{2} S^2 \sigma^2 \frac{\partial^2 V}{\partial S^2} + S \alpha (\hat{\mu}(t) - \ln S) \frac{\partial V}{\partial S} - rV = 0.$$

D.3.1 Non-uniform grids

One approach to generate a non-uniform mesh in one dimension is through a generating function. The idea is to specify a continuously differentiable strictly monotonic increasing function $g : [0, 1] \rightarrow [0, 1]$ which maps a uniform grid in $[0, 1]$ into a non-uniform grid in $[0, 1]$. The resulting grid is then defined by $\{y_i\}_{i=0}^n$ with

$$y_i := g(x_i), \quad x_i := \frac{i}{n}, \quad i = 0, \dots, n.$$

Figure D.3 illustrates this process. One important criteria of choosing a generating function is the rate of which it concentrates grid points at certain positions. This can be expressed in terms of the ration of the distance of two adjacent grid points of the non-uniform grid to the




Figure D.3: Concentration of grid points around $y^* = 0.8$ with 10 fold density $c = 0.1$

uniform grid. The distance in the non-uniform grid at the point $y = g(x)$ is approximately $g'(x)\Delta x$, where $\Delta x = \frac{1}{n}$ is the distance in the uniform grid. So, it is natural to define a distance ratio function r by

$$r(y) = g'(g^{-1}(y)).$$

As a simple example we consider the distance ratio function

$$r(y) := \sqrt{c^2 + p^2(y - y^*)^2}.$$

The parameter y^* can be viewed as the centre of the grid point concentration with c as a measure of the intensity because r assumes its minimum at y^* with $r(y^*) = c$. For big values y the function is almost linear since $r(y) = \sqrt{c^2 + p^2(y - y^*)^2} \approx \sqrt{p^2 y^2} = |py|$. The parameter p has to be set appropriately so that the resulting grid generating function satisfies $g(1) = 1$. By definition of r we have $r(g(x)) = g'(x)$, $g(0) = 0$, which is an ode for g and can be solved explicitly for the example function r we are considering here. The solution is

$$g(x) = y^* + \frac{c}{p} \sinh \left(px + \operatorname{arsinh} \left(-\frac{p}{c} y^* \right) \right). \quad (\text{D.4})$$

The parameter p has to be chosen so that $g(1) = 1$. That can for example be done using the Newton iteration method. With parameters $y^* := 0.8$ and $c := 0.1$ it follows that $p \approx 8.42136$ which Figure D.3 illustrates.

Bibliography

- [ASA, 2002] ASA, N. P. (2002). Derivatives trade at nord pool’s financial market. available online.
- [ASA, 2003] ASA, N. P. (2003). The nordic spot market. available online.
- [Barlow, 2002] Barlow, M. (2002). A diffusion model for electricity prices. *Mathematical Finance*, 12(4):287–298.
- [Barone-Adesi and Gigli, 2002] Barone-Adesi, G. and Gigli, A. (2002). Electricity derivatives. Working paper.
- [Benth et al., 2005] Benth, F., Kallsen, J., and Meyer-Brandis, T. (2005). A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivative pricing. preprint.
- [Bibby and Sorensen, 1995] Bibby, B. and Sorensen, M. (1995). Martingale estimation functions for discretely observed diffusion processes. *Bernoulli*, 1(1/2):17–39.
- [Björk, 1998] Björk, T. (1998). *Arbitrage Theory in Continuous Time*. Oxford University Press, Great Clarendon Street, Oxford OX2 6DP, UK.
- [Björk, 2004] Björk, T. (2004). *Arbitrage Theory in Continuous Time*. Oxford University Press, Great Clarendon Street, Oxford OX2 6DP, UK, 2nd edition.
- [Black, 1976] Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179.
- [Carmona and Touzi, 2004] Carmona, R. and Touzi, N. (2004). Optimal multiple stopping and valuation of swing options. to appear in *Mathematical Finance*.
- [Carr and Madan, 1998] Carr, P. and Madan, D. (1998). Option valuation using the fast Fourier transform. *J. Comput. Fin.*, 2:61–73.

- [Cartea and Figueroa, 2005] Cartea, A. and Figueroa, M. (2005). Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335.
- [Clewlow and Strickland, 2000] Clewlow, L. and Strickland, C. (2000). *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, London, UK.
- [Cont and Tankov, 2004] Cont, R. and Tankov, P. (2004). *Financial Modelling with Jump Processes*. Chapman & Hall/CRC.
- [Davis, 1997] Davis, M. (1997). Option pricing in incomplete markets. In Dempster, M. and Pliska, S., editors, *Mathematics of Derivative Securities*. Cambridge University Press.
- [de Boor and Lynch, 1966] de Boor, C. and Lynch, R. (1966). On splines and their minimum properties. *J. Math. Mech.*, 15:953–969.
- [de Jong and Huismann, 2002] de Jong, C. and Huismann, R. (2002). Option formulas for mean-reverting power prices with spikes. published via ssrn.com.
- [Delbaen et al., 2002] Delbaen, F., Grandits, P., Rheinländer, T., Samperi, D., and Schweizer, M. (2002). Exponential hedging and entropic penalties. *Mathematical Finance*, 12:99–123.
- [Delbaen and Schachermayer, 1994] Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Math. Ann.*, 300:463–520.
- [Delbaen and Schachermayer, 1998] Delbaen, F. and Schachermayer, W. (1998). The fundamental theorem of asset pricing for unbounded stochastic processes. *Math. Ann.*, 312(2):215–250.
- [Deng, 2000] Deng, S. (2000). Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Working paper PWP-073.
- [Duffie et al., 2000] Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.
- [Duffie and Richardson, 1991] Duffie, D. and Richardson, H. (1991). Mean-variance hedging in continuous time. *Ann. Appl. Prob.*, 1:1–15.
- [Eberlein and Stahl, 2003] Eberlein, E. and Stahl, G. (2003). Both sides of the fence: a statistical and regulatory view of electricity risk. *Energy & Power Risk Management*, 8(6):34–38.

- [Eydeland and Wolyniec, 2003] Eydeland, A. and Wolyniec, K. (2003). *Energy and Power Risk Management: New Developments in Modeling, Pricing and Hedging*. John Wiley & Sons Inc.
- [Fleten and Lemming, 2001] Fleten, S. and Lemming, J. (2001). Constructing forward price curves in electricity markets.
- [Fleten and Lemming, 2003] Fleten, S. and Lemming, J. (2003). Constructing forward price curves in electricity markets. *Energy Economics*, 25:409–424.
- [Föllmer and Sondermann, 1986] Föllmer, H. and Sondermann, D. (1986). Hedging of non-redundant contingent claims. In Hildnbrand, W. and Mas-Colell, A., editors, *Contributions to Mathematical Economics*, pages 205–223.
- [Franco, 2003] Franco, J. (2003). Maximum likelihood estimation of mean reverting processes. available online.
- [Hagan and West, 2005] Hagan, P. and West, G. (2005). Interpolation methods for curve construction. to appear in *Applied Mathematical Finance*.
- [Hakala and Wystup, 2002a] Hakala, J. and Wystup, U. (2002a). *Foreign Exchange Risk: Models, Instruments and Strategies*. Risk books.
- [Hakala and Wystup, 2002b] Hakala, J. and Wystup, U. (2002b). Heston’s stochastic volatility model applied to foreign exchange options. In Hakala, J. and Wystup, U., editors, *Foreign Exchange Risk: Models, Instruments and Strategies*. Risk books.
- [Harvey, 1990] Harvey, A. (1990). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- [Henderson, 2002] Henderson, V. (2002). Valuation of claims on nontraded assets using utility maximization. *Mathematical Finance*, 12:351–373.
- [Henderson and Hobson, 2002] Henderson, V. and Hobson, D. (2002). Substitute hedging. *Risk*, 5:71–75.
- [Henderson and Hobson, 2003] Henderson, V. and Hobson, D. (2003). Coupling and option price comparisons in a jump diffusion model. *Stochastics and Stochastics Reports*, 75:79–101.
- [Henderson et al., 2005a] Henderson, V., Hobson, D., Howison, S., and Kluge, T. (2005a). A comparison of option prices under different pricing measures in a stochastic volatility model with correlation. *Review of Derivatives Research*, 8:5–25.

- [Henderson et al., 2005b] Henderson, V., Hobson, D., and Kluge, T. (2005b). Extending figlewski’s option pricing formula. revised version.
- [Heston, 1993] Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343.
- [Hobson, 2004] Hobson, D. (2004). Stochastic volatility models, correlation and the q -optimal measure. *Mathematical Finance*, 14:537–556.
- [Howison, 2005] Howison, S. (2005). A matched asymptotic expansions approach to continuity corrections for discretely sampled options. part 2: Bermudan options. working paper.
- [Howison and Rasmussen, 2002] Howison, S. and Rasmussen, H. (2002). Continuous swing options. working paper.
- [Ibanez, 2004] Ibanez, A. (2004). Valuation by simulation of contingent claims with multiple early exercise opportunities. *Mathematical Finance*, 14(2):223–248.
- [Jaillet et al., 2004] Jaillet, P., Ronn, E., and Tompadis, S. (2004). Valuation of commodity-based swing options. *Management Science*, 50:909–921.
- [Karatzas and Shreve, 1991] Karatzas, I. and Shreve, S. (1991). *Brownian Motion and Stochastic Calculus*. Springer-Verlag.
- [Kluge, 2002] Kluge, T. (2002). Pricing derivatives in stochastic volatility models using the finite difference method. Master’s thesis, TU Chemnitz, Technische Universität Chemnitz, Strasse der Nationen 62, 09111 Chemnitz, Germany.
- [Kluge, 2006] Kluge, T. (2006). *Pricing Swing Options and other Electricity Derivatives*. PhD thesis, University of Oxford.
- [Koekebakker and Ollmar, 2001] Koekebakker, S. and Ollmar, F. (2001). Forward curve dynamics in the nordic electricity market. Working paper.
- [Kwon, 2002] Kwon, O. (2002). A general framework for the construction and the smoothing of forward rate curves. research paper.
- [Lewis, 2001] Lewis, A. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. working paper.

- [Longstaff and Schwartz, 2001] Longstaff, F. and Schwartz, E. (2001). Valuing american options by simulation: a simple least-squares approach. *Review of Financial Studies*, 14:113–147.
- [Lucia and Schwartz, 2002] Lucia, J. and Schwartz, E. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50.
- [Meinshausen and Hambly, 2004] Meinshausen, N. and Hambly, B. (2004). Monte carlo methods for the valuation of options with multiple exercise opportunities. *Mathematical Finance*, 14(4):557–583.
- [Merton, 1976] Merton, R. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3:125–144.
- [Miltersen and Schwartz, 1998] Miltersen, K. and Schwartz, E. (1998). Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates. *Journal of Financial and Quantitative Analysis*, 33(1):33–59.
- [Monoyios, 2006] Monoyios, M. (2006). Characterisation of optimal dual measures via distortion. to appear in *Decisions in Economics and Finance*.
- [Noller, 2002] Noller, J. (2002). Modelling electricity prices. Master’s thesis, University of Oxford.
- [Pilipovic, 1997] Pilipovic, D. (1997). *Energy Risk: Valuing and Managing Energy Derivatives*. McGraw-Hill.
- [Rogers, 2002] Rogers, L. (2002). Monte carlo valuation of american options. *Mathematical Finance*, 12:271–286.
- [Stuart and Ord, 1994] Stuart, A. and Ord, K. (1994). *Kendall’s Advanced Theory of Statistics: Volume 1*. Edward Arnold.
- [Turnbull and Wakeman, 1991] Turnbull, S. M. and Wakeman, L. M. (1991). A quick algorithm for pricing european average options. *Journal of Financial and Quantitative Analysis*, 26(3):377–389.
- [Williams, 1991] Williams, D. (1991). *Probability with Martingales*. Cambridge University Press.