Cass Business School
City of London

MODELLING ELECTRICITY PRICE RISK FOR THE VALUATION OF POWER CONTINGENT CLAIMS:

The Case of Nord Pool

Orestes Soldatos

The present thesis is submitted for the degree of Doctor of Philosophy

Faculty of Finance

Supervised by: Dr. Nikos Nomikos and Professor Michael Tamvakis

DECEMBER 2007

CONTENTS

FIGURES

TABLES

"To My whole Family, for their love and great support!"

ACKNOWLEDGMENTS

The completion of the PhD thesis has largely benefited from the support of a few people. First of all I would like to express my gratitude to my first supervisor Dr. Nikos Nomikos. The thesis has largely benefited from the numerous discussions we had, his patience as well as his unfailing guidance.

I am also very grateful to my second supervisor Professor Michael Tamvakis for his support from the begging until the end of the thesis, and always giving the right directions when needed. I would also like to express my deepest appreciation to Professor Costas Grammenos, who gave me the initiative to undertake this thesis and spend a lot of his time to give me guidance, very helpful advices and courage from the day of our first meeting. Moreover, the completion of the thesis is benefited from the comments of my two VIVA examiners, Professor Derek Bunn and Professor John Hatgioannides, which I would also like to thank.

The thesis has also largely benefited from discussions and comments of participants in the 2005 European Energy Conference in Bergen, Norway, the 2005 Cass Business School Energy Risk Management Seminar in London and the International Symposium in Advances in Financial Forecasting in Athens, Greece. I am also grateful to the people working in Viz Risk Management Services AS and in Nord Pool for providing the required data for the thesis, as well as for giving very helpful comments based on their experience in the market.

Moreover I would like to thank my friends John, Kostas, Nasi, Rob and Daphne, and the PhD students at Cass, for their advices and moral support all this time. Finally I would like to express my deepest gratitude to my parents, Alexandros and Danae Soldatos and my girlfriend, Fani Kallianou, for being there for me all these years, believing in me and making this part of my life such a great time.

ABSTRACT

Reconstruction and deregulation in the international power markets has let prices to be determined by the fundamental rules of Supply and Demand, which brought a substitution from *Supply Risk* pre-regulation, to *Price risk*, thus increasing the necessity of hedging using derivatives such as futures and options and therefore brought the issue of pricing these derivatives into focus. However the traditional approaches for the pricing of derivatives are not applicable to electricity due to the unique features of the power market such as the fact that electricity is not storable. Under these circumstances, arbitrage across time and space is limited in the electricity market. As a consequence there is a need for a good model that is able to capture the dynamics of the electricity spot prices for the purposes of Derivatives pricing and Risk Management. In this thesis we propose three different spot models for the Scandinavia electricity market; First, we propose a seasonal affine jump diffusion spike model, which can distinguish the behaviour of electricity spot prices between normal periods and periods when spikes occur. Second, we propose a seasonal affine jump diffusion regime-switching spike, which is an extension of the spike model but contains two separate regimes to distinguish between periods of high and low water levels in the reservoirs, reflecting the availability of hydropower in the market. Third, we propose a seasonal affine jump diffusion three-factor spike model which again extends the spike model but allows the equilibrium level to be stochastic in order to capture the long-run dynamics of the market that are uncovered from the shape of the forward term structure. The performance of our models is compared to that of other models proposed in the literature in terms of fitting the observed term structure, as well as by generating simulated price paths which have the same statistical properties as the actual prices observed in the market. In particular, our models perform well in terms of capturing the spikes and explaining their fast mean reversion as well as in terms of reflecting the seasonal volatility observed in the market. Then we use these models and provide semi-closed form solutions for European option prices and investigate whether the shape of the model implied volatility smile is consistent to the one that is anticipated to be observed in the market. Furthermore, we also perform a sensitivity analysis for Asian option prices which are widely used in the market. Finally, we apply a modified Least Squares Monte Carlo algorithm for the pricing of swing options, and investigate the sensitivity of the incremental swing premium to changes of different parameters used to capture the stochastic behaviour of the power spot prices.

NOTATION

Asian Call	Value of an Asian Call option.
Asian Put	Value of an Asian Put option.
$\mathbb{C}$	A complex number.
Call	Price of a European call option.
Cov	Covariance operator.
$d(\tau_j)$	Number of down-swings exercised until τ_j.
DK	Deseasonalised strike price.
E	Expectation operator under the Real measure.
E^*	Expectation operator under the Equivelant Martingale Measure.
$\mathcal{F}$	A filtration or a set of events.
$f(t)$	Deterministic variable or deterministic function of time.
$F(t,T)$	Forward/Future Price at time t, with maturity at time T.
F_{t,T_1,T_2}	Forward price at time t for delivery during the period $[T_1,T_2]$.
$G_{a,b}(y;)$	Price of a security paying e^{aX_T} when at maturity $bX_T \leq y$
$Im[]$	Imaginary part of a complex number.
K	Exercise price.
k	Speed of mean reversion.
l	Jump intensity.
m_i	The ith moment of a distribution.
$N(0, \sigma^2)$	The Normal Distribution with mean 0 and variance σ^2.
$N(d)$	The Normal distribution function.
N^+ or N^-	Total number of up or down swings allowed.
p_i	Probability of being in regime i.
p_{ij}	Transition probability of switching from regime i to regime j.
P_t	Spot price at time t.
Put	Price of a European put option.
q	A Poisson process.
Q	Transition matrix, or variance-covariance matrix in Kalman filter.
r	Continuously compounded interest rate.

t	Time.
T	A maturity time.
τ	Time to maturity.
Δt_R	Refraction time.
u	Error term, or a vector in the Tranform function $\Psi^{\chi}(u, X_t, t, T)$)
$u(\tau_j)$	Number of up-swings exercised until τ_j.
v^+ or v^-	Incremental volume taken or delivered.
V_{cap}	Value of a cap.
V_{floor}	Value of a floor.
W	Brownian Motion.
$\mathbf{X}$	A vector process.
X	Normal short-term factor, or log-price of electricity.
Y	Spike factor or the log-price of generating fuel.
ε	Equilibrium level or long-term factor.
$\tilde{\varepsilon}$	Risk neutral equilibrium level.
λ	Market price of risk.
μ	Trend.
μ_J	Mean jump size.
ρ	Correlation coefficient.
$\mathbb{R}^n$	The n-dimensional real vector space.
σ	Volatility parameter.
${}_{swing}V_{\tau_j}^{u(\tau_j),d(\tau_j)}$	Value of swing option.
$\Psi^{\chi}(u, X_t, t, T)$	The transform function.
Ω	Sample space.
$\{\Omega, \mathcal{F}, P\}$	A probability space.

1. INTRODUCTION AND AIMS OF THE THESIS

1.1 Introduction

The aim of this thesis is to consider the issue of modelling electricity prices in order to encompass the stylised features of the Nordic Market (Norway, Sweden, Finland and Denmark) for the purposes of Risk Management and Derivatives Pricing. More specifically we are developing spot price models that are able to capture the most significant distributional and path characteristics of the electricity market and then use them to explore the implied volatility smiles as well as to price exotic options used widely in the market such as Asian and Swing Options.

The structure of this first chapter is as follows: Section 2 provides a brief outline of the aims of the thesis. After describing the main economic drivers of electricity prices in section 3, section 4 compares the different theories used for the pricing of futures in commodity markets and selects the most appropriate one for electricity. Section 5 discusses why model building for uncertainty is of importance in the power market. Finally section 6 gives an outline of the thesis and section 7 concludes this chapter.

1.2 Aim of the Thesis

Over the last 20 years radical changes have taken place in the structure of electricity markets around the world. Prior to 1980s the electricity industry was a natural monopoly and strong vertical integration was the ideal economical model for electricity utilities. Deregulation and reconstruction however, made it possible to operate power generation and retail supply as different competitive market segments. Also, technological advances in electricity transmission equipment have made possible the economic transmission of power over long distances so that customers can now be more selective in choosing an electricity supplier, which result in greater competition among operators.

Reconstruction was first applied in the UK in 1990, when the then government argued that reconstruction was necessary, and hence privatised the industry in order to remove inefficiencies caused by bureaucracy, bloated workforce and procurement procedures. The reconstruction experiment in the UK proved successful by creating revenue in a time of tax cut, and also persuaded other countries to follow the same path, as stated by the *Managing Energy Price Risk* magazine in 1998.

Reconstruction and deregulation however, let prices to be determined by the fundamental rules of Supply and Demand, which brought a substitution from *Supply Risk* pre-regulation, to *Price risk*, thus increasing the necessity of hedging using derivatives such as futures and options. Hence, the issue of pricing these derivatives came into focus. However the traditional approaches for the pricing of derivatives are not applicable to electricity due to the unique features of the power market. One of these features is the fact that electricity is not storable and this breaks down the relationship between spot and forward prices, which is mainly characterised by the notions of storage and convenience yield in storable commodities such as oil and metals. Under these circumstances, arbitrage across time and space is limited in the electricity market. The problem of non-storability also creates extreme movements in the spot price as in high demand periods, forced outages or transmission constraints may cause extreme movements in price, called spikes. Under these circumstances, the task of identifying an appropriate pricing model that is tractable, parsimonious, capturing the most significant risks, and is easy to apply for the pricing of derivative instruments becomes of paramount importance.

Following the above discussion, the first aim of the thesis is to introduce different stochastic models that are able to accurately describe the behaviour of electricity price movements and at the same time fit the historical forward term structure, for the purposes of derivatives pricing and risk management. The reason why the pricing model needs to fit the forward term structure stems from the fact that even though the electricity spot is not a storable asset, forward and futures contracts are liquid traded financial assets and can be used for hedging strategies of other financial contingent claims, such as options. In this thesis, we examine the Scandinavian electricity market (Nord Pool) which has been the most liquid financial market for electricity worldwide, for the time period we are investigating.

After having introduced the different spot models and examined their properties, our next step is to apply them in options pricing. As a first step, we provide closed-form solutions for European options, using the different proposed pricing models. Then, we examine the properties of the implied volatility shape produced by each model, using as a benchmark the traditional Black-Scholes-Merton model. In this way we are able to investigate whether our proposed models can capture the volatility smiles and smirks that are anticipated in the market. This analysis will assist market participants who make their decisions based on the implied volatility from traded options in the market using the Black and Scholes (1973) model. In addition, we perform a sensitivity analysis in order to identify the main drivers of Asian option prices, which are widely used in the OTC market to hedge long-term price risks.

Finally, we examine the issue of swing options pricing, which are contracts that are widely used in the energy markets, to hedge simultaneously against demand and price risk. Swing options pricing is a challenging topic, especially due to their early exercise features. After discussing the theoretical framework for pricing Swing options and identifying the most suitable numerical method for pricing, we conduct a sensitivity analysis on the incremental premium that is paid relative to a strip of European options. In this way market practitioners can have a benchmark on how the incremental premium paid for early exercise is affected by changes in different underlying parameters, and thus identify potential pitfalls when hedging Swing options with European options, which is generally the case.

Having discussed the aims of the thesis our next step is to understand in depth how electricity prices move and their main economic driving forces, in order to develop a pricing model for derivatives pricing in the power market. Thus the next section focuses on the price determinants of the electricity prices.

1.3 Electricity and its Price Determinants

1.3.1 Power Market Features and Price determinants

Electricity may be considered as a *flow commodity* (see Lucia and Schwartz, 2002) characterised by its limited storability and transportability. This makes electricity a commodity that is different than most others. The major features of the market are the following:

Non-storability: Electricity cannot be stored in a conventional manner. Currently, the only feasible method of "storing" high voltage electricity is in bulk by means of a pumped storage mechanism[1]. Any electricity that is generated must be either consumed instantaneously within the grid or transmitted to another location. It can be argued that electricity can be stored in terms of hydropower, however this does not mean that electricity can be produced now and used in the future. Moreover and especially in these forms, no substantial amount can be stored since reservoir and turbine capacities are not infinite hence the possibility of electricity storing becomes only partial. Moreover the non-storability factor of electricity is the main driver of the differences between power producers, such as power generators, and power suppliers. Power generators may store electricity in terms of its generating fuel, mainly for very short-term adjustments, and thus they have the flexibility of providing electricity during periods of high demand and shortages. On the other hand, suppliers of electricity cannot store the commodity, but they rather have to transport it to the consumers for immediate use. Finally, consumers do not need storing facilities, as long as they are continuously served with a sufficient quantity of power.

Transmission constraints: To compound the storability problem, electricity grids tend to be highly segmented. The transportation constraints for electricity come in the form of capacity limits in the transmission lines and losses from transportation, which can make the transmission of electricity between regions, impossible or uneconomical. These limitations make electricity contracts and prices highly local, i.e.

[1] Method of storing and producing electricity to supply high peak demands. At times of low demand, excess electrical capacity is used to pump water into an elevated reservoir. When there is higher demand, water is released back into the lower reservoir through a turbine generating hydroelectricity. Reversible turbine/generator assemblies act as pump and turbine. 70% to 85% of the electrical energy used to pump the water into the elevated reservoir can be regained in this process.

strongly dependent on the local determinants of supply and demand (such as characteristics of the local generation plants, local climate and weather conditions). Additionally, regulatory issues such as market rules and market structure may also have an impact on the behaviour of prices in competitive electricity markets and, consequently, on their differences across countries.

Seasonality and weather dependence: There are also seasonal and weather-related variations in the demand and supply of electricity, giving strong seasonal and intra-day price patterns. For instance in the Nordic countries, electricity prices have a tendency to increase over winter, when a substantial amount of electricity is used for heating and lighting purposes (as shown in Chapter 3). Generally electricity prices also peak during summer due to increased demand for air-conditioning. Moreover electricity consumption by industrial and residential users cause electricity demand peaks during the day, although they drop off at night, so a distinction is usually made between *off-peak* and *on-peak* prices.

Swing risk and delivery: Unlike other commodity markets, the supplier of power (utility) faces "swing" demand risk. At a short notice, such as one day, the buyer can demand more or less power within certain limits. From a delivery point of view, electricity poses an additional set of problems. Since electrons cannot be tagged, electricity is produced and dumped into pools much like water entering a water reservoir. A buyer merely taps into the pool supplied by many generators and marketers. The issue of who pays whom is decided on a notional basis rather than on an actual delivery basis.

Generation stack: From the supply side, the independent Transmission System Operator (TSO) will stack various available plants that have been "bid in" in an increasing order of their marginal costs[2]. That is the first plants to be turned on in any day are those with the cheapest marginal costs such as hydro and nuclear – the so-called base load[3] units. Then the more expensive coal- and oil-fired plants are used, and, finally, the most expensive gas-fired plants are turned on. The supply curve is the

[2] The TSOs own the transmission network and are responsible for ensuring a well-functioning physical system, including the system balance, that is, the equilibrium required in every instant between the energy produced/imported in any area and the sum of consumption /export and network losses in the same area.
[3] Base load covers power supplied at a constant rate through out all demand periods of the year and usually produced by plants with little operational flexibility running at very low marginal cost (e.g. nuclear, hydro).

locus of all these marginal costs arranged in increasing order. As shown in Figure 1.1, the market-clearing price is given by the intersection of the supply and demand curves. Moreover since the market-clearing price cannot be below the marginal cost of production, the marginal cost of a particular unit in operation at a particular point in time may be viewed as the lower price boundary at that particular time. Note also that the actual marginal costs of different generators can be influenced by a number of physical constraints such as start-up and rump-up costs, and in some situations, this might make fuel cost an insufficient proxy for the actual marginal cost of generation.

Figure 1.1: Supply Stack example and Equilibrium Price determination

The figure below shows a schematic example of a supply stack, with two potential demand curves superimposed on it. The spot price is determined by the intersection of demand and supply.

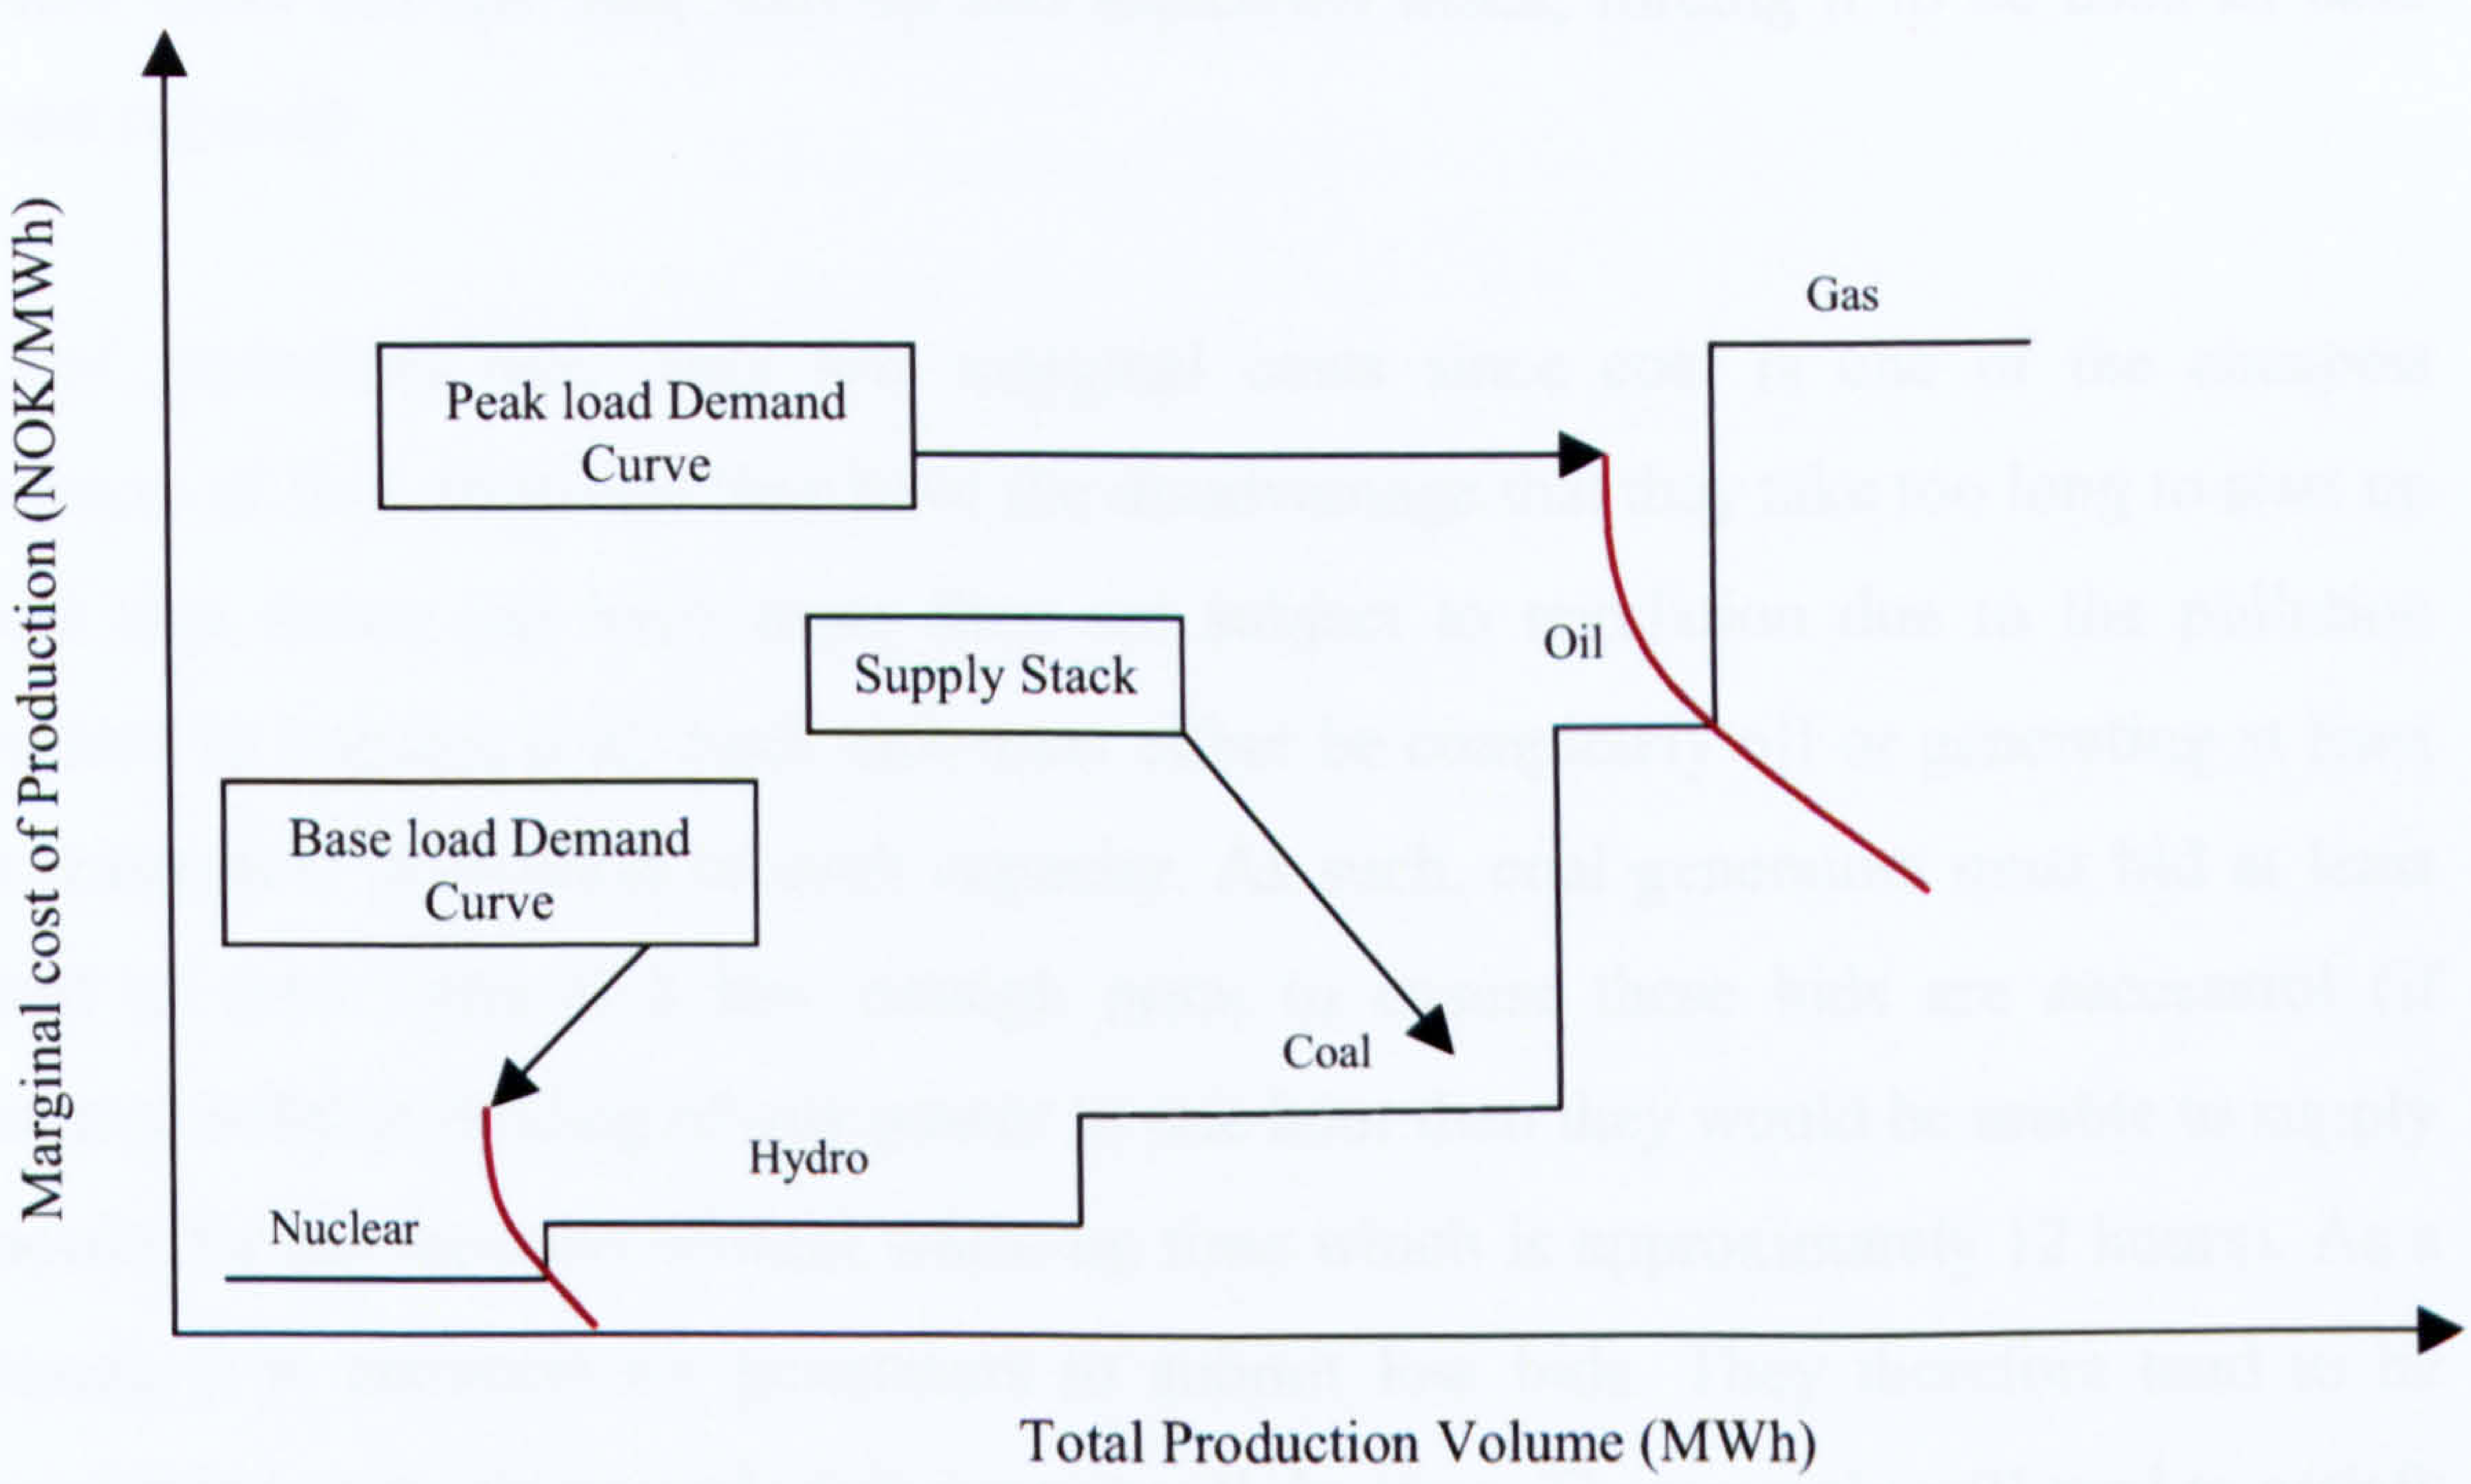

1.3.2 Different ways of Generating Power

The previous section discussed the main price determinants of electricity prices. One of the key determinants is the marginal cost of production, which mainly depends on the type of generating fuel. In this subsection, we discuss the main different types of power generators:

Hydroelectric power stations generate electricity by passing flowing water through a turbine. The water is stored in a dam, and released into a stream of river. These plants have no marginal costs of electricity production but operate under a number of constraints: they cannot produce more than the capacity of the turbine, they cannot use more water than what is in the dam (which limits their ability to operate during dry periods) and they cannot emit too much or too little water into the stream /river during a given period of time. A key advantage is their ability to start up and shut down very quickly. As a result of these factors, hydro-storage plants tend to bid their load at high prices to maximise the revenue from the water that they can send through. This method of generation can cause power prices to go to very low levels during wet periods and reduce the incidence of price spikes.

Nuclear power operates like coal in the way that it has low running costs yet high fixed costs and has long start-up and shutdown times, forcing it to be used as base load capacity.

Coal generators have very low marginal costs since coal is one of the cheapest sources of fuel. However they have the disadvantage that they take too long to start up and shut down and even more they are subject to regulation due to the pollution caused by burning coal. Each unit must either be completely off or generating at least a reasonable proportion of each capacity. As such, coal generators must bid at least half of their units at a low enough price to ensure these bids are successful (if unsuccessful in bidding of any power in one hour then they would be unable to supply power for the duration of their warm-up time which is approximately 12 hours). As a result, it is common for generators to submit low bids. They therefore tend to be generating at levels towards full capacity all the time. Their supply will tend to satisfy the demand mainly during off-peak times.

Oil and *gas* generators, in contrast, are much faster to start up yet have the highest marginal costs. They therefore tend to be used as peak capacity, rather than base capacity. They set the price during peak hours, and hence place bids in an attempt to try and maximise the prices yet supply as much as possible. Like coal generators, they do not have total flexibility in how much power to generate, and must produce more than half of their capacity for each unit operating.

Finally there are also renewable sources for electricity generation such as wind and solar energy. For instance, *wind* powered generators operate only on a very small scale. They are different from many other generator types in that although they have no marginal costs, they have high fixed costs and also operate under constraints due to their reliance on the availability of wind. They have short start-up and shutdown times, but cannot easily bid due to the uncertainty as to when wind will be available. As such they generally operate in such a way that whenever electricity can be generated, they will be used.

1.3.3 Electricity Spot price behaviour

The discussion above gives an indication of the features of electricity prices. Electricity spot prices move in a totally independent fashion than those of most other commodities or financial assets. Power prices tend to fluctuate around values determined by the cost of production and the level of demand, in other words they have a tendency of *Mean Reversion* to an equilibrium price, which is a common feature of most commodities. However, power prices tend to change by the time of day, week, month and in response to cyclical fluctuations in demand, thus introducing a *seasonality* component. In addition, *non-storability* does not allow electricity prices to follow a "smooth process" as prices of other commodities do so. Also supply shocks, such as generating or transmission constraints and unexpected outages, cause temporary price "*spikes*" due to less availability of capacity; in other words, electricity prices may jump to a new level of ten times their mean, but they do not stay there as they quickly revert back to their mean level. On the other hand, spikes may also occur due to electricity's inelasticity of demand. As it can be noticed in Figure 1.1 the supply stack has a convex shape, so temporary extreme movements in demand shift the demand curve to the right, hence the equilibrium price increases. At this high level then more generators will enter in the market to take advantage of the higher price for profit making thus forcing the price to revert back to its mean. Moreover there is another implication of the convex shape of the supply stack. At higher price levels the supply stack curve becomes steeper and steeper, hence price changes are bigger for a given change in demand, causing an asymmetry in the volatility. This is exactly the opposite from what is noticed in the stock markets, and is known as the more specific description of the stylised facts for the Scandinavian market will be given in the empirical part of this thesis, in Chapter 3.

1.4 FUTURES PRICING THEORY AND ELECTRICITY FUTURES

The previous sections explained the theory behind the spot price movements in the electricity market, by mainly examining its fundamentals. These features provide the building block on derivatives pricing when the underlying is the electricity price.

When pricing a derivative contract, one usually uses the arbitrage pricing approach. That is, if it is possible to construct a dynamically adapted portfolio that will perfectly replicate the payoff of the derivative contract, then the absence of arbitrage forces the price of this derivative to be equal to the price of the replicating portfolio. If it is always possible to build a dynamically adapted portfolio that will perfectly replicate any payoff, then the market is said to be *complete*. In this case, there will only be a single no-arbitrage price for any contingent claim, and there exists a unique probability measure called the risk-neutral probability measure, equivalent to the physical probability measure, under which the no-arbitrage price of any contingent claim is equal to the expectation of its payoff discounted at the risk-free rate.

This approach however is not applicable when it comes to pricing electricity derivatives as discussed earlier, since storage possibilities are very limited and quite expensive. Only large hydro systems do have this possibility and even then, because of constraints in reservoir and turbine capacities, this possibility is only partial. Lack of a storage relationship may prevent us from using no-arbitrage arguments for derivatives pricing because one cannot create a replicating trading strategy involving the spot price, and thus there may not exist a unique risk-neutral equivalent probability measure. Furthermore, the dynamics of electricity spot price are very complex and quite far from what one usually assumes for financial assets. In this section we will describe the various theories underpinning the pricing of the commodities futures and asserts whether they can be applied to pricing electricity derivatives.

1.4.1 Commodity Futures Pricing

The pricing theory of futures for financial assets such as bonds and stocks is different than that of commodities. The main difference is that pricing financial derivatives relies on pure arbitrage arguments, whereas commodities are more complicated due to the fact that storage is costly and that spot markets may be non existent or too thin for arbitrage. Before deriving formulas however we should state their required assumptions (*Hull 1999*):

1) *There are no transaction costs.*
2) *All net trading profits are subject to the same tax rate.*
3) *Market participants can borrow and lend money at the same risk-free rate of interest.*
4) *Market participants take advantage of arbitrage opportunities when they occur.*

Note that it is not required that these assumptions are true for all market participants. All that is needed is that they hold for a subset of all market participants. Arbitrage opportunities disappear as soon as they occur given assumption 4. An implication of this assumption is therefore that market prices are such that there are no-arbitrage opportunities. The first three assumptions are obviously not perfectly valid in commodities, however the degree of validity in each market is almost the same. Adjustments can also be made to adjust the model to "the real world". The main requirement is the arbitrage assumption, and it is highly discussible if assumption 4 is valid in the electricity market. At the same time the volatility is extremely high which makes it difficult to forecast future prices. The Pricing theories are the theory of storage and the risk premium theory which are described next:

1.4.2 Theory of Storage

Inventories play a crucial role in the price formation in markets for storable commodities, which are also referred as "cash and carry markets". The theory of storage explains the difference between current spot prices and futures prices in terms of interest foregone in storing a commodity, warehousing costs and a *convenience yield* on inventory, which was introduced by Kaldor (1939) and Working (1933). The convenience yield can be regarded as a liquidity premium and represents the benefits of ownership of the physical commodity that are not obtained by the holder of a futures contract. These benefits may include the ability to profit from temporary local shortages or the ability to keep a production process running. According to the storage theory (or stockpiling), which again is based on traditional arbitrage pricing, the futures price, $F(t,T)$, of a contract expiring at period T observed at time t is given by:

$$F(t,T) = P_t e^{r(T-t)} + U - Y \qquad (1.1)$$

Where the first term on the RHS of (1.1) is the current spot price compounded by the interest rate, r, for the period until expiration of the contract $(T-t)$. U and Y represent the storage cost and convenience yield in money terms, between now (t) and delivery (T). Alternatively the above formula can be represented in terms of the convenience yield, y, as a compound factor and the storage cost as a proportion of the spot price, u:

$$F(t,T) = P_t e^{(r+u-y)(T-t)} \qquad (1.2)$$

The above formulas were introduced by Brennan and Schwartz (1985) in their pioneering research for the valuation of commodity derivatives. The convenience yield measures the extent to which the spot price compounded with the interest rate plus the cost of storing the commodity, exceeds the futures price. It mainly holds for consumption assets, where its owners keep such a commodity in inventory for consumption purposes and not investment purposes, hence they are reluctant to sell the commodity and buy futures contracts, because futures contracts cannot be consumed. Hence the convenience yield reflects the market's expectations concerning the future availability of the commodity. The greater the possibility of shortages that may occur during the life of the futures contract, the higher the convenience yield.

1.4.3 CAPM and the Theory of Risk Premium

According to the Capital Asset Pricing Model (CAPM by Sharpe 1964 and Lintner 1965), the higher the risk an investor bears for an investment, the higher the Return demanded. CAPM suggests that risk can be split in two categories: *Specific Risk* is the risk that is common to a class of assets and hence can be eliminated in a well-diversified portfolio. *Systematic Risk* on the other hand, is the risk that cannot be diversified away and arises from the correlation of that asset's returns and the returns of the market as a whole. That is why stocks have a tendency to move together and investors are exposed to uncertainties no matter how many stocks they hold. Hence an investor demands a higher Expected Return than that of the risk-free rate, in order to bear the extra Systematic Risk.

Therefore according to the CAPM theory the expected return on the asset is the risk-free rate plus an expected premium for bearing that extra risk, which can be presented in a one-period setting by the following formula:

$$E(R_i) = R_f + \left(E(R_m) - R_f\right)\frac{Cov(R_i, R_m)}{\sigma^2(R_m)} = R_f + \left(E(R_m) - R_f\right)\beta_i \qquad (1.3)$$

$E(R_i)$: Expected Return on the *i*th asset

$E(R_m)$: Expected Return on the market

R_f: Risk-free rate of return

$Cov(R_i, R_m)$: The covariance of the returns of the *i*th asset and the market

$\sigma^2(R_m)$: Variance of market returns

β_i: Systematic Risk of the *i*th asset

Moreover the equation for the one-period expected return on an asset is:

$$E(R_i) = \frac{E(P_{iT}) - P_{it}}{P_{it}} \qquad (1.4)$$

P_{it} : The price of the *i*th asset at time *t*

$E(P_{iT})$: The expected price of the *i*th asset at time *T*, where $t<T$

Therefore combining equations (1.3) and (1.4) and solving for the price of the asset now (*t*) we have:

$$P_{it} = \frac{E(P_{iT}) - \left(E(R_m) - R_f\right) P_{it} \beta_i}{\left(1 + R_f\right)} \qquad (1.5)$$

Equation (1.5) is used to give us a formula for the futures price, which allows an investor to buy an asset now and differ the payment for one-period. Hence the current price of a future, $F_i(t,T)$, is the spot price of the asset multiplied by its future value factor:

$$F_i(t,T) = P_{it}\left(1 + R_f\right) = E(P_{iT}) - \left(E(R_m) - R_f\right) P_{it} \beta_i \qquad (1.6)$$

Following Hull (1999), the CAPM theory leads to an alternative way of estimating the futures prices on commodities than the classical theory of storage. Consider a speculator who takes a long futures position in the hope that the price of the asset will be above the futures price at maturity. Assume that the speculator puts the present value of the futures price into a risk-free investment at time *t* while simultaneously taking a long futures position. The proceeds of the risk-free investment are then used to buy the asset on the delivery date, at time *T*. The asset is then immediately sold for its market price. This means that the cash flows to the speculator are:

Time *t*: $-F(t,T)e^{-r(T-t)}$

Time *T*: $+P_T$, which is the price of the commodity at time *T*.

Hence the present value of the investment at time *t* is: $-F(t,T)e^{-r(T-t)} + E(P_T)\, e^{-\lambda(T-t)}$

That is to say that the present value of the investment at time *t*, is the present value of the money that will be given to settle the futures position at *T*, plus the expected price of the commodity at time T, discounted by an appropriate rate λ for the investment. That means that λ represents the expected return required by the speculator on the investment. Assuming that all investment opportunities have a net present value of zero (otherwise arbitrage opportunities arise), the fair price of the futures in the risk neutral world is:

$$F(t,T) = E(P_T)e^{(r-\lambda)(T-t)} = E(P_T)e^{-pr(T-t)} \tag{1.7}$$

The value of λ depends on the systematic risk of the investment that was discussed in the CAPM and hence the term pr represents the risk premium. One way of explaining the risk premium would be to look at the conditions within the specific commodity market. A majority of risk averse producers wanting to hedge their products in the futures market would probably result in futures prices lower than the expected future spot price, hence $pr>0$. The opposite relation would occur when the demand side is the most risk averse. A second way of explaining the risk premium is to compare the spot price at delivery, P_T, to other assets in the stock market. Hence, if the return on the spot price is positively correlated to the level of the stock market, the investment, using a forward contract, involves positive systematic risk and an expected return above the risk-free rate is required $(pr>0)$.

This pricing theory can also be applied in markets where the commodity is perishable (also referred as *"price discovery markets")*[4] or has limited storability. The no-arbitrage argument underlying the theory of storage, cannot be applied to non-storable commodities, as there is no possibility of obtaining a risk-free position by buying the commodity in the spot market and selling it in the futures market. The market hence is said to be incomplete, as the number of assets traded is not equal to the sources of risk, hence no risk-neutral strategies are identified.

1.4.4 Which theory is more appropriate for pricing electricity derivatives?

Comparing the two theories for futures pricing in the commodities markets, the CAPM approach argues that systematic risk should be important in the pricing of

[4] *"Price Discovery"* refers to "the market's ability to "discover" true equilibrium prices. Futures markets provide centralised trading where information about fundamental supply and demand conditions for a commodity is efficiently assimilated and acted on and, as a consequence, equilibrium prices are determined." (Edwards and Ma (1992)).

futures contracts but leaves out storage costs and convenience yields. On the other hand, the first approach ignores the possibility that systematic risk may affect the equilibrium prices of commodity futures contracts. The issue is which theory is the most appropriate for pricing derivatives in the electricity market.

As it was mentioned earlier, electricity is different than most commodities due to its non-storability and the physical requirements for the instantaneous equilibrium between local demand and supply. One can argue that especially in Scandinavia, power generators can store the commodity in terms of hydropower using water reservoirs. It can also be said that consumers have the possibility to store electricity in batteries, but this option is not available in large scale. In the future energy systems could possibly include large-scale storage capacity such as in hydrogen reservoirs. On the supply side there has been a limited amount of pumped hydro storage in the system up to-date. However, all these storage options involve substantial investment costs, and as market participants have stated (*Nord Pool Reports and Discussion*) they do not see them as possible tools for making arbitrage from the difference between spot and futures prices. Another way of exploring arbitrage in the electricity market is by relating the electricity price back to the cost of fuel, which is storable. With this approach we can determine a forward arbitrage price by introducing an additional variable, which is the relationship between the fuel and electricity price. However, this relationship is extremely volatile (as it will be discussed in more detail in section 2.5), and thus the approach is not very useful. Another argument why this approach is not efficient is due to the fact that the cost of storing the fuel is not economically efficient for the generators, especially in the long run. Hence based on this electricity cannot be stored today (at least not in substantial amounts) for future sales, and arbitrage across time and space, based on the theory of storage, is seriously limited if not completely impossible. This makes electricity more like a price discovery market.

The above argument turns to favour the theory of risk premium in the electricity futures market. A risk premium could arise if either the number of participants on the supply side differs substantially by the number on the demand side, or if the degree of risk aversion varies considerably between the two sides. Most of the companies participating in the market are both generators and load serving entities. However in terms of flexibility of adjusting the quantity on the supply and demand side, there is a significant difference. Generators can control parts of their generation on a short notice, especially in Norway, due to the large share of hydropower in the system. This allows them to take advantage from short-term price fluctuations occurring in the day ahead market or even closer to real time, by adjusting their generation. Therefore, for them it does not make sense to fix the futures market for all of the planned future generation. On the demand side however the situation is different, as the entities cannot adjust the load on the price, due to the price inelasticity of demand. Hence, assuming that market participants on the demand side are risk averse, it makes sense for them to lock in as much as possible of their expected future demand in the futures market. Moreover, the probability of spikes gives another incentive for the demand side to hedge their costs coming from the spot price. If this difference in terms of hedging needs between the demand and supply side occurs, there is excess demand in the futures contracts, translating into a negative risk premium. This is something that will be proven empirically in chapter 4.

1.5 The importance of Model Building For Uncertainty

The above discussion leads to the conclusion that electricity is a consumption good with very limited storability, in the sense that one cannot buy electricity and store it for future use or sell it in the future but rather has to consume it immediately. This implies that the market is incomplete and hence one must assume a market price of risk to price contingent claims; thus the drift in the risk neutral probability measure is the drift in the physical measure minus the market price of risk. Contingent claim prices are however internally consistent in the sense that they will not provide any arbitrage opportunities and they will be valued based on the same price of risk.

Although spot electricity is not a storable asset, the futures and forward contracts are regular financial contracts that are traded and hence can be used in a replicating strategy. Therefore the information contained in the futures price should be used when pricing derivatives on the spot prices. Similarly, spot price dynamics should also be considered when constructing a futures model or pricing derivatives on the futures contracts. Hence this implies the need for a consistent underlying framework, which allows to price derivatives on both the spot and the futures contracts. So modelling spot and futures prices dynamics, estimating the parameters of their stochastic processes and pricing derivatives on both of them are closely related issues. There are strong links tying the two processes together. Completeness in the market is in fact the key issue relating these activities. This in turn depends on the choice of the stochastic process to model the securities. For example if we consider a market composed of a single security and a deterministic risk-free bond, this market will be complete if the security's dynamics are modelled with a single factor diffusive stochastic process, but it will be incomplete if its dynamics are modelled with a multi-factor diffusive stochastic process. Hence, if one considers different stochastic processes that are assumed to describe both the spot dynamics and the futures prices of different maturities (which means that both spot and futures are considered as primary traded securities), then these stochastic processes should be estimated on historical time-series of both spot and futures prices. In such case, since futures contracts are regular financial contracts that can be traded, the market will be complete as long as the model remains diffusive and the values of the market price of the risk factors can be uniquely determined by the observations of different futures forming the term structure.

The above conclusion stems from the fact that futures/forward prices are assumed to be such as so no-arbitrage opportunities arise. This implies that there exists at least an Equivalent Martingale Probability Measure[5] (EMM) (see Harrison and Kreps, 1979, Cox et al, 1981, and Musiela and Rutowski, 1998). The no-arbitrage condition further implies that derivative prices must converge to the value of their payoff at maturity, e.g. futures prices must converge to spot prices at maturity. Hence from these properties it can be shown that futures prices must at any point in time be the expected value of the spot price at their maturity, under the EMM. However when using stochastic processes to model the spot price under the physical probability measure, there may not exist an equivalent probability measure under which all futures prices are the expectation of the spot price at their maturity, hence we are not able to write any stochastic process describing the dynamics of all the futures prices under the

[5] Martingale means that the process becomes unpredictable, i.e. $E^*[P_{t+dt}]=P_t$ and it is used in derivatives valuation in order to find their fair value, as in the derivatives market their main purpose is risk management, and if there is any observed drift component, arbitrage opportunities arise. Hence futures prices are Martingales, that is to say their price at t is equal to their expected price at $t+dt$. In other words the expected return in the position is zero.

physical probability measure that will not allow for some arbitrage opportunities. But we can find one that will fit most of them.

Following the above discussion, the main aim of this research is to introduce a class of stochastic processes that will be able to accurately describe the behaviour of electricity price movements. In addition, this model needs to be consistent with the "no-arbitrage opportunities" assumption and hence it should fit best the observed term structure and describe the spot market's dynamics. To do that, we need to specify the stochastic process that best explains the movements in the prices and define a class of EMM that is able to price all sources of uncertainty, and hence find within this class the particular martingale probability measure that is able to fit the historical term structure. This will then enable us to price other contingent claims that are widely used in electricity markets.

1.6 Outline of the Thesis

The outline of this thesis is as follows. Chapter 2 gives a detailed summary of the literature up to date in electricity modelling and pricing, by dividing the literature into four types, namely the reduced form, forward, equilibrium and hybrid models, and then discusses the main advantages of our modelling procedure, and our contribution to the literature. Chapter 3, starts by explaining the different market structures in the deregulated power markets and then gives a description of the Scandinavian electricity physical and financial markets. It then conducts an analysis on the descriptive statistics on the spot prices in Nord Pool, in order to ascertain the statistical properties of electricity prices.

Following the theoretical foundations given in the present chapter, and supported by the empirical findings of the descriptive statistics in Chapter 3, Chapter 4 then introduces three new models for derivatives pricing in Nord Pool: the spike model which is designed to capture short-term variations in spot prices, the regime switching spike model which captures the short-term variations, as in the spike model, as well as mid-term variations in the equilibrium level due to changes in the level of water in the reservoirs, and a three-factor spike model which captures short-term variations as well as general long-term variations in the equilibrium level that are expected to persist as implied by the forward term structure. The models are tested in terms of fitting the forward term structure, by providing closed-form solutions for forward prices, as well as by capturing the distributional and path characteristics of spot prices. We would also like to mention that Chapter 4 has been accepted for publication and is forthcoming at the *Journal of Applied Mathematical Finance* under the title, "Using Affine Jump Diffusion Models for Modelling and Pricing Electricity Derivatives".

Chapter 5, examines the issue of option pricing in the electricity market. It starts by discussing the shortfalls in the Black & Scholes world in terms of pricing options in the electricity market. Semi-closed form solutions are then introduced for the pricing of European options using the jump diffusion models introduced in Chapter 4. It then produces an option pricing analysis in terms of producing volatility smiles and smirks, analyses the option values given by the different proposed models and conducts a sensitivity analysis for Asian structures using Monte Carlo simulation, since their pricing in closed-from does not exist.

Chapter 6, focuses on the Swing options which are contracts that are used extensively in the OTC electricity market, for the hedging of both volumetric and price risk. The chapter first looks on the pricing theory of Swing options, discusses their advantages in hedging and then uses a modified Monte Carlo Least Squares simulation approach (Longstaff and Schwartz, 2001) to find the fair price of these contracts using the proposed models. Then, it conducts a sensitivity analysis to identify the drivers of the incremental swing option premium relatively to European strips and also explores the effect of penalties imposed if the minimum number of rights is not satisfied. Finally, Chapter 7 concludes the thesis, and also presents a few suggestions for further research.

1.7 Conclusions

This chapter is an introduction to the thesis and as such it outlines the aim and the structure of the thesis. The main aim is to model electricity spot prices for the purpose of derivatives pricing in the Scandinavian power market. The motivation for this stems from the fact that in the new deregulated power markets, derivatives are important for hedging price risk hence there is a need for an adequate model to price power derivatives.

Furthermore we provided a theoretical framework for the process of electricity prices. More specifically we explained that the spot price should revert to an equilibrium level which displays a seasonal pattern due to the inelasticity of power demand and its dependence on weather as well as human activity. Moreover, the inelasticity of demand and the non-storable nature of electricity coupled with the convex shape of the supply curve causes prices to spike. Furthermore, we analysed the theory of storage and the theory of risk premium used for the pricing of commodity derivatives, and explained that due to the nature of electricity, the theory of risk premium is the most appropriate for derivatives pricing in the electricity markets. Finally, the last section of this chapter provided the outline of the thesis.

The task of modelling prices for the valuation of contingent claims is very difficult in such a "turbulent energy industry"[6] as electricity. The next chapter will discuss the literature review on the subject, which will give an outline on what has been done and what needs to be done in order to identify the most appropriate model for Pricing and Risk management purposes.

[6] As the energy markets were characterised by the Professor of Finance at the University Paris Dauphine and ESSEC, Helyette Géman.

2. LITERATURE REVIEW

2.1 Introduction

This chapter focuses on the literature review on modelling and pricing electricity derivatives. This is presented in four sections, representing four different modelling approaches that have been proposed in the literature, using the classification by Eydeland and Wolyniec (2003). The first approach is the *Reduced-Form Modelling*, which consists of specifying a stochastic process of the price evolution by selecting a parameterised family of processes (e.g., Geometric Brownian Motion, jump-diffusion process, and so-on). Second, is the *Forward Modelling* that aims to model the evolution of Forward Prices without the need to take into account the market price of risk since this is already incorporated in the prices. Then, we have the *Fundamental Equilibrium Models* which aim to model supply/demand relations and to obtain the power process as a solution of certain optimisation, or more precisely equilibrium problems. Finally, the *Hybrid Models* use the fundamental approach to represent supply and demand relations and the stochastic techniques to represent the evolution of the underlying drivers. The main research in electricity that has been conducted for every one of the above categories will be discussed separately in the following subsections. Finally, the last section of the chapter explains our modelling and pricing approach and our contribution to the literature.

2.2 Reduced form Models

Starting with the Reduced Form Models, as shown by Eydeland and Geman (1999), in the classical Black-Scholes-Merton world, the main driver in option pricing is the process that describes the dynamics of the underlying asset. However in the modelling of price dynamics in electricity one encounters the problem of matching the fat tails of marginal and conditional distributions and the spikes in the spot price. The Geometric Brownian Motion tells us that the spot is driven by an exponential growth rate μ, plus a noise characterised by a Brownian Motion W_t scaled up by the constant standard deviation term σ, and the underlying price P_t:

$$dP_t = \mu P_t dt + \sigma P_t dW_t \tag{2.1}$$

However as explained in Chapter 1, electricity prices should mean revert to an equilibrium level. As shown by Schwartz (1997), in equilibrium setting one should expect that when prices are relatively high, supply would increase since higher cost producers of the commodity will enter the market, putting a downward pressure on prices. Conversely, when prices are low, higher cost producers will exit the market, thus putting an upward pressure on prices. On the other hand when specific events occur that affect directly the demand side (e.g. a heat-wave), then mean reversion depends on how quickly these events dissipate. Therefore in theory, equation (2.1), does not make sense as electricity prices are expected to mean revert to a long-run equilibrium level which is tied with the long-run marginal production cost.

Using that fact, Lucia and Schwartz (2002) propose a one and a two-factor model to describe the stochastic behaviour of the 24-hour average spot price in Nord Pool, and test the model fit in the forward term structure. In their general extended two-factor model, the electricity spot price is decomposed into three components; one that is totally predictable and the other two, which are stochastic. The deterministic variable $f(t)$, is supposed to capture the regularities that occur in the electricity market, such as higher prices during winter due to higher energy use. The two stochastic components are assumed to capture the short-term (X) and long-term (ε) variations that occur in the market, as in Schwartz and Smith (2000). Thus their model is as follows:

$$P_t = f(t) + X_t + \varepsilon_t$$

Where

$$\begin{aligned} dX_t &= -\kappa X_t dt + \sigma_X dW_X \\ d\varepsilon_t &= \mu_\varepsilon dt + \sigma_\varepsilon dW_\varepsilon \\ dW_X dW_\varepsilon &= \rho dt \end{aligned} \tag{2.2}$$

The variable X is assumed to capture the short-run variations in the spot price process and thus it follows an Ornstein–Uhlenbeck with a mean reversion to zero, and speed of mean reversion $\kappa > 0$. The long-term variable ε, is assumed to capture the uncertainty about the equilibrium price to which prices revert. The equilibrium price level is assumed to follow an Arithmetic Brownian Motion (ABM) with the drift reflecting expectations of improving technology for the production of the commodity, inflation, as well as political and regulatory effects, which are very crucial in electricity. Hence uncertainty in prices evolves from short–term deviations from the long-run equilibrium level as well as from changes in the long-run equilibrium level to which prices revert. Long–maturity futures prices give information about changes in the equilibrium price, and changes in the difference between near- and long-term futures prices give information about the short-term deviations.

They price futures and forwards in the Eltermin market, and the performance of each model was tested in terms of the Root Mean Square Errors between the actual forward prices and those implied by the models. The analysis was carried using both spot prices and log-spot prices, as well as by letting the equilibrium level ε to be either constant or stochastic. Their results showed that seasonality is better modelled using sinusoidal functions, rather than monthly dummies, and they give evidence that a risk-premium is of paramount importance when pricing derivatives in the electricity market. Comparing the one-factor model (where ε is assumed to be constant) against the two-factor model, the two-factor model performed better, since it captures the trend that is evident in the term-structure, but not in the spot market. Furthermore, another advantage of the two-factor model is that it implies a non-perfect correlation among forward prices of different maturities, something that is not implied by the one-factor model. However one of the major problems with the two-factor model is that it deals with non-observable state variables, and since the data they use is limited and the market illiquid, instead of the Kalman filter proposed by Schwartz (1997), they implement a filtering procedure by Cortazar and Schwarz (2003), where the state variables are estimated by minimising the Sum of Squared Errors at each observation date, and the parameters are estimated by minimising the Sum of Squared Errors for the whole sample. The main drawback of this procedure is that it does not provide standard errors from which to infer the significance of each parameter.

As shown by Lucia and Schwartz (2002), any diffusive model at this setting is not compatible with the spikes as well as the fat tails realised by the distribution of the power prices. Therefore, Lucia and Schwartz (2002) as well as Kaminski (1997), point out the need of introducing stochastic volatility and jumps. The former model is first tested by Kellerhals (2001) who uses Heston's (1993) stochastic volatility model to price short-term electricity derivatives. Heston (1993) assumes that the asset price follows a GBM but allows the volatility to be stochastic and following a mean-reverting process.

However, since the electricity market is incomplete as the spot is not storable and the volatility is not a tradable asset, one faces two market prices of risk that arise from both processes. For the electricity spot the Girsanov Transformation $dW_{1,t}^{*} = dW_{1,t} + \lambda\sqrt{v_t}dt$ (where v_t is the variance) is used, with a time invariant market price of risk λ, in order to shift the distribution from the Real to the Equivalent Martingale Probability Measure. As in Heston (1993), he assumes that the market price of risk of the volatility process is independent from the spot, but it is proportional to the relative risk aversion and the covariance between the change in consumption and change in variance. Kellerhals (2001) proposes that the volatility risk premium is proportional to the volatility and has the form $\lambda_v v_t$, under the assumption that consumption growth and spot asset returns have a constant correlation. Therefore, the processes of the log-price, X_t, and that of the variance, under the Equivalent Martingale Probability Measure are as follows:

$$
\begin{aligned}
dX_t &= \left(\mu - \lambda v_t - \frac{1}{2}v_t\right)dt + \sqrt{v_t}dW_{1,t}^{*} \\
dv_t &= \left(\kappa(\varepsilon_v - v_t) - \lambda_v v_t\right)dt + \sigma\sqrt{v_t}dW_{2,t}^{*}
\end{aligned}
\tag{2.3}
$$

However, since stochastic volatility affects the kurtosis of the distribution, he further assumes a correlation ρ between the Brownian Motions of the spot and the variance, which helps to capture the skewness effects, arising from the fact that electricity prices are more volatile at high levels, hence "spreading" the right tail of the probability density. Using the martingale methods developed by Geman et al (1995) and Scott (1997) he provides closed-form solutions for the forward prices. His model is tested empirically on the Californian electricity market, and estimated using the Kalman filter. The results show that the stochastic volatility is very important for modelling sport prices however it does not give any significant better fit in forward prices, when compared against the mean-reverting model of Schwartz (1997). This may be due to the fact that his stochastic volatility model does not take into account the mean reversion in the spot prices.

Another approach to capture the spikes and the fat tails of the distribution of power prices, are the jump diffusion models, which where first introduced by Merton (1976) for the stock market. As Eydeland and Geman (1999) show, this set of models explain intuitively the effect of extreme temperatures leading to extreme power demand, which happen to coincide with outages in power generation and/or transmission and therefore lead to large sudden increases in power prices. The general model is as follows:

$$dP_t = \mu P_t dt + \sigma P_t dW_t + J P_t dq_t \qquad (2.4)$$

Where the jump-diffusion component is represented by a Poisson process q_t, with an intensity *l*, characterising the frequency of jump occurrence and a real-valued random variable *J* representing the direction and magnitude of the jump whenever it occurs. The jump sizes, *J*, are determined by the natural logarithm of proportional jumps being normally distributed, $\ln(1+J) \sim N\left(\ln(1+\mu_J) - \frac{1}{2}\sigma_J^2, \sigma_J^2\right)$, where μ_J is the mean jump size and σ_J is the standard deviation of the proportional jump size called the jump volatility. The dq_t is a discrete time process characterised by a Poisson distribution, since it occurs at specific points in time. Therefore $dq_t=0$ most of the time and takes the value of *1*, when a jump occurs. Hence the process behaves like a GBM when no jump occurs or when the jump frequency in near to zero. The same also stands when the frequency is high but the jump volatility is low.

As stated earlier, electricity prices tend to follow a mean-reverting process, therefore equation (2.4) needs to take account of the fact. Barz and Johnson (1998) as well as Clewlow, Strickland and Kaminski (2000 and 2001), were the first papers, which combine a mean-reverting process with jumps. The latter use the Mean-Reverting Jump Diffusion (MRJD) model to capture the main features of electricity prices in the US:

$$dP_t = \kappa(\varepsilon - \ln P_t)P_t dt + \sigma P_t dW_t + P_t J dq_t \tag{2.5}$$

Hence the price process follows a standard mean-reverting process with probability equal to *one minus the jump intensity*, and it becomes a mean-reverting process plus a random jump size with probability equal to the intensity. However, since empirically the distribution of jump sizes in energy is positively skewed, they propose to have proportional jumps drawn from a normal distribution, but with different jump volatilities for the positive and negative jumps. They also introduce the *recursive filter* estimation, by defining as a jump every observation of returns whose absolute value is greater than three standard deviations of all returns. This way, they can capture the low frequency and high volatility component of jumps more accurately, compared to the Maximum Likelihood estimation of Ball and Torous (1983) which identifies the smallest and most frequent jump component of the actual data. However one major drawback of their model is that seasonality in the equilibrium level is not included.

Escribano et al (2002) propose a general benchmark model that encompasses the main features in the main deregulated electricity markets. Using Lucia and Schwartz (2002) one-factor model, they specify the process by adding a jump diffusion term and stochastic volatility. However since continuous time models are hard to estimate and computationally intensive, they use a more flexible in discrete time GARCH (1,1) Jump Diffusion model, with seasonal jump intensity, to capture the more frequent jumps that occur when the demand is high. Thus their general model is as follows:

$$P_t = f(t) + X_t$$

$$X_t = \begin{cases} \phi X_{t-1} + h_t^{1/2} u_{1t} & \text{with probability } (1 - l_t) \\ \phi X_{t-1} + h_t^{1/2} u_{1t} + \mu_J + \sigma_J u_{2t} & \text{with probability } l_t \end{cases}$$

$$\begin{aligned} \text{with} \quad & h_t = \omega + \alpha u_{t-1}^2 + \beta h_{t-1} \\ & l_t = l_1 winter_t + l_2 fall_t + l_3 Spring_t + l_4 Summer_t \\ & \phi = 1 - \kappa \\ & u_{1t} \text{ and } u_{2t} \sim \text{i.i.d N}(0,1) \end{aligned} \tag{2.6}$$

where l_i for $1 \leq i \leq 4$, are seasonal jump intensity parameters, and $winter_t$, $fall_t$, $spring_t$ and $summer_t$ are seasonal dummies.

The model is estimated using Maximum likelihood with a Poisson-Gaussian density and Likelihood Ratio Tests gave overwhelming evidence in favour of the GARCH-Poisson-Gaussian model for all markets against alternative model specifications. One main finding is that $\alpha+\beta>1$ when no jumps are included, hence there is explosion in volatility process implying that the inclusion of the Jump diffusion process is a very critical part. In Nord Pool they found that without the inclusion of jumps and time-varying volatility, mean reversion is very slow, but after incorporating these facts the parameter is significantly increased. They conclude that jumps and time varying volatility are *complementary* rather than *substitute* factors of risk.

A drawback of the research presented until this point on the jump diffusion models, is that the authors do not provide any analytic results regarding derivatives valuation. One of the first papers that deals with the issue of derivatives pricing under a pure jump diffusion framework is Villaplana (2003) who extends the research by Lucia and Schwartz (2002). The paper is based on the argument that the price of risk is related to economic risks and willingness of market participants to bear it. Therefore in order to capture most of the electricity market characteristics, he extends the two-factor model by Lucia and Schwartz by introducing jump diffusion to the short-term factor. Moreover the author proposes two different specifications for the long-term factor: Arithmetic Brownian Motion or Ornstein–Uhlenbeck process. Using the Transform Analysis by Duffie et al (2000), closed-form solutions are derived for the prices of forward contracts for both spot and log-spot price models. Under this two-factor jump diffusion approach there are four sources of uncertainty: The two diffusive factors imply a short-term and a long-term uncertainty, where as the jump component introduces risks arising from the intensity of the jumps but also the size of the jumps. Hence to price derivatives he assumes that a risk premium is demanded for all three types of risk. He also assumes seasonal jump intensity, along the lines of Escribano et al (2002), and in this way the models are able to capture the "seasonal skewness" under the objective probability measure, that translates into a "seasonal forward premium". His empirical tests are based only on one month forward contract, and the long-run stochastic variable is not estimated at all. One of the important findings of this research is that during periods of low jump probability, the risk premium is driven by the premium implied from the short-term diffusive term. That is to say that producers are willing to sell forward for these periods, since there is no significant skewness and hence, unexpected large price increases are unlikely to occur. A significant result was that the jump premium over summer (in the USA) was 40% of the total risk premium, which shows how important the jump diffusion parameter is in a model for derivatives pricing. However the paper uses just one contract rather than the whole term structure to estimate the risk premium. Furthermore, for the accurate estimation of the jump risk premium one needs to use options data as shown by Koekkebaker (2004).

Along the same lines, Cartea and Figueroa (2005) propose a similar model to Clewlow et al (2000 and 2001) and Lucia and Schwartz also (2002), where electricity prices are modelled using a one-factor mean reversion jump diffusion, by allowing the equilibrium level to follow a deterministic seasonal process. The model is applied to the England and Wales electricity market (NETA). The speed of mean reversion and the risk premium are calibrated from the forward curve of one particular date.

One main drawback of all the models up until now is that they assume that the speed of mean reversion of the spikes, when jumps are included in the model, is the same as for the diffusive shocks. However, since spikes occur mainly due to outages or major generators being off-line for no longer than a few days, hence, as it will be shown later in this thesis, their speed of mean reversion is much higher than that of normal shocks. Thus imposing the same mean reversion for both normal and jump shocks, might have serious consequences for derivatives pricing.

Among the first papers to deal with the fast mean reversion of the spikes in a jump diffusion setting is Weron (2005) who uses the one-factor model of Lucia and Schwartz (2002) with the addition of Jump diffusion, to model prices on the spot prices in Nord Pool. In order to estimate the spike parameters he uses the Recursive filter approach by Clewlow and Strickland (2000), however he defines as a spike any return observation that is above three standard deviations followed by a return in the opposite direction of similar size. The simulated paths of the model closely resembled the original spot price. However as it will be shown in later chapters, this definition might be too restrictive as a spike may last a few days, and sometimes might be negative. The main purpose of the model was to price and infer the market price of risk from Asian options using Monte Carlo simulations. No closed-form solutions were derived, since the model is not tractable enough, but captures the spiky behaviour of the market.

Geman and Roncorony (2006) introduce another one-factor model specifically to describe the distribution of electricity price in the US electricity markets. The stochastic process of the natural logarithm of electricity prices, X_t is presented as follows:

$$dX_t = D\varepsilon(t) + k_1\left(\varepsilon(t) - X(t\text{-})\right)dt + \sigma dW_t + h(t\text{-})dJ(t) \tag{2.7}$$

The first term on the right hand side of (2.7) is the first derivative of the mean, $\varepsilon(t)$ to which prices revert, and consists a predictable seasonal trend of the price modelled by the sum of a constant reflecting the fixed cost of production, a trend and two sinusoidal functions for the two peaks per year (seasonality) of possibly different magnitudes. The last term of the equation capture the spikes, which are defined as a cluster of upward shocks, of relative large size with respect to normal shocks of the deseasonalised series, followed by sharp downward shocks to the equilibrium level. They filter the jumps from the normal process using a threshold level. If price exceeds that threshold, one can yield the direction of the jump, which is identified by a switching function $h(X(t))$ that takes the value of +1, if the price is below the threshold, and –1, if the price is above the threshold. Even though the model is able to capture most of the trajectoral and statistical features of the prices, no implementation in derivatives pricing has been given yet. Furthermore, we have to comment on the way they define spikes, i.e. the use of a threshold level. As we will see in latter chapters, electricity prices in Nord Pool reach smoothly very high levels for a long period of time due to the low level of water in the reservoirs, which increases the long-run equilibrium level. However by using this model, this type of price pattern would be classified as spike, which is inaccurate.

Another area of research focused on capturing the spiky behaviour of prices in the electricity market is through the use of Regime Switching models. The first paper that introduces Regime Switching in the electricity market is by Huisman and Mahieu (2003), in which electricity prices are assumed to follow different processes during normal and spike periods. During normal periods, their model is along the lines of Lucia and Schwartz (2002) one-factor model of the log-prices. Following on, the stochastic factor X_t is modelled as a three-regime process[7]:

Regime 0: Normal- normal price dynamics

$$dx(t) = -\kappa_o x_{t-1} + \sigma_0 u(t) \tag{2.8}$$

Regime +1: Initial jump- price process at sudden increases or decreases

$$dx(t) = \mu_1 + \sigma_1 u(t) \tag{2.9}$$

Regime −1: How electricity reverts back to "normal" after a spike

$$dx(t) = -\kappa_{-1} x_{t-1} + \sigma_{-1} u(t) \tag{2.10}$$

Where $u(t) \sim N(0,1)$

The estimation of such a model requires a few assumptions especially about the Markov Transition Probability Matrix, which displays the probabilities of moving from one regime to another at each time step. Huisman and Mahieu (2003) assumed that the price cannot proceed from a normal regime to a spike reverting regime since a spike has to occur first; being in the reverting regime it is assumed that the process will be in the normal regime in one day; and also given the probability of being in the normal regime one can estimate the probability of a jump. The estimation becomes even trickier as one has to use the Kalman filter since the regimes are latent. Comparing the mean-reverting jump diffusion model as in Clewlow et al (2000 and 2001) against their Regime jump model, the latter model provided a better fit. Their results show that their Regime-Switching model captures the jumps more accurately when compared against a mean-reverting jump diffusion model. However, the three-

[7] Note that the subscripts in the parameters in each model specify to which regime model that parameter belongs.

regime model is too restrictive in the assumptions especially due to the inclusion of the third regime, which restricts the duration of a jump.

Following the above discussion, Jong and Huisman (2002) propose Regime-Switching models with only two regimes, rather than three, in order to allow spikes to last longer than only one day. That is to say that if there is a generator outage for example, we might have very high spikes for some time period until the generator is repaired; then prices will revert back to normal. Hence this time the stochastic factor X_t, is specified as follows:

Stable Regime: $X_{M,t} = X_{M,t-1} + k(\varepsilon - X_{M,t-1}) + u_{M,t}$, where $u_{M,t} \sim N(0,\sigma_M)$ (2.11)

Spike Regime: $X_{S,t} = \mu_S + u_{S,t}$, where $u_{S,t} \sim N(0,\sigma_s)$ (2.12)

Their estimation is along the lines of Huisman and Mahieu (2003). Their results show that the probability of being in the spike regime is overestimated. Their model is then fitted to the forward term structure by first decomposing the forward price into two components; one for the expected price being at the stable regime and another for the expected price being at the spike regime. They then adjust ε and make it a function of time to maturity, in order for the model to fit perfectly the term structure at a given date. Closed-form solutions are also provided for European options on the spot price. One criticism made on their model is the fact that it is mainly adequate to price short-term derivatives, as in the long-run the spikes are averaged out and the derivatives depend on longer term variations, such as those on the equilibrium level as in the two-factor model of Lucia and Schwartz (2002).

Deng (2000) provides three different models for the pricing of electricity derivatives, using the properties of all models discussed in the literature until this point, such as stochastic volatility, jump diffusion and regime-switching jumps. The regime-switching jumps capture the systematic alterations between "abnormal" (spikes) and "normal" (no-spikes) equilibrium states of supply and demand. He states that, these regime-switching jumps may be caused by weather patterns and varying precipitation[8], in markets where the majority of installed capacity is hydropower. More specifically in his main model, he assumes a joint stochastic process of the logarithmic prices of electricity, $lnP_t=X_t$ and a generating fuel, Y_t, directly under the risk neutral measure:

$$d\begin{pmatrix} X_t \\ Y_t \end{pmatrix} = \begin{pmatrix} \kappa_1(t)(\varepsilon_1(t)-X_t) \\ \kappa_2(t)(\varepsilon_2(t)-Y_t) \end{pmatrix} dt + \begin{pmatrix} \sigma_1(t) & 0 \\ \rho(t)\sigma_2(t) & \sqrt{1-\rho(t)^2}\sigma_2(t) \end{pmatrix} dW_t \qquad (2.13)$$
$$+\sum_{i=1}^{2}\Delta Z_t^i$$

where ΔZ_t^i is a compound Poisson process capturing up and down jumps and $\rho(t)$ denotes the correlation coefficient between the Brownian Motion of the electricity price and that of the generating fuel price. He further extends the above model, by proposing two additional models: In the first he includes regime switching to distinguish prices between "normal" and "abnormal" states, and thus capture the quick speed of mean reversion for spikes. In the second extension he includes a stochastic volatility process along the lines of Heston (1993). In order to derive the formulas for derivatives valuation the author uses the transform analysis in derivatives pricing for affine-jump diffusion models proposed by Duffie et al (2000). Using hypothetical values for the parameters, the paper shows that futures prices are higher when jumps occur even if the upward and downward jumps are the same. He then derives Call options prices, and shows that their values increase with time to expiration, but they converge to a long-run mean showing an increasing dependence to the fundamental characteristics of supply and demand as maturity increases. However the paper does not give any empirical results on any market, as the models are difficult to calibrate.

Elliot and Sick (2003) use one-factor diffusion models with jumps in the price explained by a finite state Markov chain describing the process of the number of generators on-line at every time. The paper focuses on the Power Pool of Alberta in the US. Suppose you have N number of generators in a pool, and that $\bar{Z}_t$ represents

[8] The quantity of water falling to earth at a specific place within a specified period of time and the quantity of water deposited (e.g. by melting the snow).

the number of generators online at time t. Hence the dynamics of the number of generators on-line are modelled using a Markov Chain in continuous time. Then the equilibrium prices depend on the number of large plants online, which determine the state of the market, $\varepsilon = (\varepsilon_1, \varepsilon_2 ... \varepsilon_N) \in R^N$. Hence the natural logarithm of the deseasonalised prices, X_t, is given by:

$$dX_t = -k(X_t - \langle \varepsilon, Z_t \rangle)dt + \sigma dW_t \tag{2.14}$$

The results show that the long-run average level to which prices revert is actually significantly different at each state, however their residual analysis still shows high right skewness and lepto-kurtosis. This is a drawback of their approach, since jumps do not only occur by the availability of generators, but also by variations in demand and the location of the equilibrium level at that instant (if it is for example at the steep part of the stack function) as well as transmission constraints. The paper derives solutions for forwards and options however no empirical evidence is given to further support the validity of the model in terms of contingent claims pricing.

Finally Benth et al (2005), model electricity prices as the sum of different stochastic variables, where each one follows a mean-reverting process and their shocks are generated by different types of jumps, each one following a Gamma distribution. They then provide closed-form solutions for forward prices and European options on forwards and swaps, which is one of the advantages of the proposed model against alternative model specifications and in particular log-models. Although the simulated paths using hypothetical values for the parameters seem to capture the behaviour of the electricity spot prices, the paper does not provide any empirical evidence. Furthermore, they do not provide any evidence on how the model captures the distributional characteristics of the spot, such as variance, skewness and kurtosis. This follows from the fact that this kind of models are very hard to calibrate and the literature on the econometric methods that can be applied to estimate them is very limited.

2.3 Forward Modelling

Up to now we have focused on spot price models of electricity, however these have some disadvantages: First Spot price models rely on unobservable state variables, such as the long- and short-term variables in Lucia and Schwartz (2002), or the convenience yield in other commodities. Secondly, in the spot models the forward price curve is an endogenous function of the model parameters and therefore will not necessarily be consistent with market observable forward prices, as most of the time the fit is not perfect. Moreover, we have to restate the fact that forwards and futures are often used by risk managers to hedge their risk therefore liquid forward prices will help the price discovery mechanism to determine the fair value for future delivery.

Hence instead of estimating the risk-neutral parameters, one can model the dynamics of the forward curve, which in the risk-neutral world only consists of the volatility term structure. In other words the drift term should be zero because futures and forward contracts have zero initial investment, hence their expected return in the risk neutral world must be zero, therefore:

$$\frac{dF(t,T)}{F(t,T)} = \sigma(t,T)dW_t, \text{ where } (t \leq T) \tag{2.15}$$

The equation implies that the change in a Forward price expiring at T, is explained by a function of the volatility of that specific forward contract scaling the standard Brownian Motion. However the only problem here is to specify the functional form of the volatility term structure. Using the one-factor model for the spot price with an Ornstein-Uhlenbeck process as in Schwartz (1997) one can derive that functional form of the volatility implied by the model for each forward price as:

$$\frac{dF(t,T)}{F(t,T)} = \sigma e^{-k(T-t)} dW_t \tag{2.16}$$

The model has two volatility parameters; σ determines the level of spot and forward price return volatility, whereas κ (mean reversion speed) determines the rate at which the volatility declines, as the forward's maturity increases. Equation (2.16) is consistent with the empirical evidence as well as with *The Samuelson hypothesis* which predicts that futures price volatility increases as the contract expiration date nears. Two main drawbacks of the one-factor model are that the volatility of longer dated to maturity contracts, $T \to \infty$, is zero, and that the correlation between the forward prices of different maturities is one, which is not the case, since it has been shown that they do not always move "identically". In the same manner one can increase the number of factors, such as in Lucia and Schwartz, and derive again the functional form of the volatilities, scaling the appropriate source of risk (Brownian Motion). Very good references on forward curve dynamic modelling can be found in Cortazar and Schwartz (1994), Amin et al (1994), Miltersen and Schwartz (1998) who develop a general framework for commodity derivatives valuation and risk management with stochastic interest rates as well as stochastic convenience yield, and Clewlow and Strickland (1999) who apply the different approaches in the energy market (oil and gas), but also adapt it to a Heath-Jarrow-Morton (1992) framework. Similarly all the above models fall in the category of the Heath-Jarrow-Morton (1992), which focuses on the dynamics of the forward curve as a whole rather than on the evolution of single contracts, hence in terms of the commodities where the risk neutral drift is zero:

$$\underbrace{\frac{dF(t,T)}{F(t,T)} = \sum_{i=1}^{K} \sigma_i(t,T) dW_{i,t}}_{\text{Class A}} \quad \textit{or alternatively} \quad \underbrace{dF(t,T) = \sum_{i=1}^{K} \sigma_i(t,T) dW_{i,t}}_{\text{Class B}} \qquad (2.17)$$

Where $0 \le t \le T$

The above model implies that the financial market in the risk-neutral world is described by K number of Risks, hence the K-dimensional Brownian Motion scaled by their respective volatilities. It has to be stated that the interest rates are assumed to be constant in order for common maturity forward and futures prices to be indistinguishable (see Cox et al (1981)). The above are multifactor models aiming to identify the sources of risk that cause movement in the whole forward curve. The first equation implies log-normality (class A) and the second normality (class B) in the time series of the forward prices. The spot price P_t is equivalent to a forward maturing at the same instant, $P_t = F(t,t) = \lim_{T \to t} F(t,T)$, hence the spot price process from the forward curve dynamics is implied to have the following form given a class A model:

$$\frac{dP_t}{P_t} = \left[\frac{\partial \ln F(0,t)}{\partial t} - \sum_{i=1}^{K} \left\{ \int_0^t \sigma_i(u,t) \frac{\partial \sigma_i(u,t)}{\partial t} du + \int_0^t \frac{\partial \sigma_i(u,t)}{\partial t} dW_i(u) \right\} \right] dt + \sum_{i=1}^{K} \sigma_i(t,t) dW_i(t) \tag{2.18}$$

The terms in the square parenthesis in the drift can be interpreted as the sum of the deterministic risk-free rate of interest and a convenience yield, or the risk premium, which in general will be stochastic as shown by Clewlow and Strickland (2000). However the drift term also involves the integration over the Brownian Motions, which means that spot prices depend on their history, hence in mathematical terms the process is said to be non-Markovian so the model is very difficult to use e.g. when constructing a tree for option valuation, the tree becomes non-recombining. On the other hand, when using the above equation and applying it to the one-factor model one ends up with the same model as Schwartz (1997), where the spot process is mean-reverting, with a time dependent drift which allows the model to fit the observed forward prices.

Bjerksund et al (2000) propose two different kinds of models belonging to class B, implemented in Nord Pool. The first model is a one-factor model with the following volatility function:

$$\sigma_1(t,T) = \frac{a}{T - t + b} + c \tag{2.19}$$

Where a, b and c are positive constants. Although the model does not imply that the volatility curve goes to zero with $T \to \infty$, but to a constant c, the decrease in the volatility curve implied by the model may be too sharp. As a remedy to that, they also propose a three-factor model with all parameters positive, as follows:

$$\sigma_1(t,T) = \frac{\alpha}{T-t+b} + c$$
$$\sigma_2(t,T) = \left(\frac{2ac}{T-t}\right)^{1/2} \qquad (2.20)$$
$$\sigma_3(t,T) = c$$

This three factor model allows a richer structure of the forward price dynamics, however it is not well behaved at all in the short end, because $\lim_{T\to t}(\sigma_2(t,T)) = \infty$, in other words the volatility of a forward maturing now (or the volatility of the spot) goes to infinity, which is not realistic. The authors state that the one-factor model is more adequate for pricing derivatives, where as the three-factor model is better for risk management. However in these models again the correlation between the forward contracts is one, hence forward prices of all maturities will move in the same direction. This conjecture contradicts with the empirical findings of Koekebakker and Ollmar (2005). They employ a Principal Components Analysis with the aim of determining the factors that explain as much of the total variation in the data as possible. The main advantage of the principal components analysis is that it does not restrict the functional form of the volatility and gives an indication of the number of risks in the market. Their results show that the method does not fit the data well at all, since 10 factors are needed to explain 94% of the variation in forward contracts, whereas in other markets 3 factors can explain more than 95%; for instance in the NYMEX Crude Oil Futures market, investigated by Clewlow and Strickland (2000), the first 3 factors explain 98.4% of the total variation. Factor 1, called the *shifting* factor, causes changes in the same direction for both short- and long-term ends of the forward curve, however the effect is bigger for the short end. Factor 2 causes movements between the two ends in the opposite direction and thus it is called the *tilting* factor. The third factor is called the *bending* factor, moving the short and long end in opposite directions of the mid-range of the term structure. Very little is gained by adding more than five factors. The fact that there are different factors affecting the two ends of the forward curve, is due to the non-storability and other features of electricity markets. For instance, if the government plans to reduce nuclear power generation in two years time, futures prices with maturity of two years will rise, but this is not expected to affect short maturity futures. Moreover the above model cannot capture the presence of time varying volatility and excess kurtosis. The results from this research indicate that contracts sold in the OTC market with maturities further into the future are even less correlated with short-term contracts, which implies that hedging long-term commitments using the short-term contracts may prove inefficient.

Audet et al (2003) extend the above research by modelling the forward curve dynamics, allowing a deterministic (not stochastic) spot volatility curve, $\sigma(.):[0,\tau]\rightarrow R_+$. They also assume that forward prices are the expected future spot prices hence forward prices are martingales under the objective probability measure $F[t,T]=E[P(T)|\mathcal{F}_t]$. They further assume that the forwards' volatility is decreasing exponentially with time to maturity; hence spot prices have the highest volatility. Moreover the correlation term structure is also decreasing exponentially with time to maturity, hence the model is as follows:

$$\frac{dF(t,T)}{F(t,T)} = e^{-\kappa(T-t)}\sigma(T)dW_T(t) \quad \text{for all} \quad t\in[0,T],\ T\in[0,\tau]$$
$$dW_{T^*}(t)dW_T(t) = e^{-\rho|T^*-T|}dt \quad \text{for all} \quad T,T^*\in[0,\tau] \tag{2.21}$$

However their analysis is more like an illustrative one as the model is calibrated several times for different seasons and different years (using 52 weekly contracts for each year) and indicates that volatility varies inside the year and also between the years. Volatility is higher in summer time (as also shown by Lucia and Schwartz, 2002) when the water reservoirs are almost empty, hence small changes in demand cause changes in the used production technologies and production marginal costs. During winter however mainly condensing power is used, hence there are no major changes in the production marginal costs and thus volatility is lower. Moreover, yearly volatility depends on the accumulation of snow; if this is lower than the normal levels it causes uncertainty on the spring flood, and hence volatility increases. However, as previous research has shown, jumps occur over the winter in Nord Pool and sometimes the high load demand cannot be met by the available capacity of hydro-generators.

Finally, Benth and Koekebakker (2005), model directly the dynamics of the forwards and futures traded in Nord Pool, by applying different volatility functions for the swap dynamics, as the forward and futures prices in Nord Pool are more similar to swaps since they contain delivery periods over a period of time. Their approach has the main benefit of modelling actual traded contracts, rather than theoretical continuous forward prices as is the case in the Heath-Jarrow-Merton (1992) framework. Their modelling framework works very well in terms of pricing options on forwards traded in Nord Pool, however, it has the major disadvantage that one cannot associate the spot dynamics from the defined swap dynamics and thus the pricing of structured products, such as Caps, Floors and Swing Options (discussed in later chapters), becomes more complicated.

As it can be seen forward curves give a better understanding about the dynamics of different contracts in the term structure, and there is no need to model the forward curve (cross sectionally), since we take it as given. However the basic models seem to fail to capture the main dynamics of electricity prices. One can argue however that Lucia and Schwartz model (2002) can also be classified as a forward curve model, even though they do not postulate directly the risk-neutral processes followed by futures prices, but define futures prices as the expected spot prices under the Equivalent Martingale Measure. One of the main benefits of the latter is that it has a high degree of intuition and furthermore, information from both the spot and the forward prices is used, which is very critical in capturing the market characteristics and completeness. Furthermore, the forward curve models rely on the complete term structure for each date. However in electricity not all contracts are traded every day, particularly long dated ones, which violates the required assumption of a liquid forward market and thus the estimated results may not be reliable. In addition, it has to be strongly stated that modelling either the forward dynamics or the swap dynamics would fail to capture the spikes shown in the spot market. This is caused first of all by the strong mean reversion of the spikes, which is not transmitted to the forward prices (e.g. if a spike occurs now and has a half life of one day, a forward maturing in 5 days will not be affected), and second the spikes are averaged out during the delivery period, something which also occurs for Asian options, as it will be examined later. For those reasons, the use of an adequate spot model that is also arbitrage free is more appropriate.

2.4 Fundamental Equilibrium Models

Now a brief discussion will be given on Fundamental Equilibrium models. Their aim is to model supply/demand relations and to obtain the power process as a solution of certain optimisation, or more precisely equilibrium problems. These models are useful for a wide range of applications, especially qualitative analysis of price behaviour. However the challenge is that these models are not designed to capture the price dynamics of power price formation in a robust quantitative manner, which is ultimately what one needs in order to develop effective hedging and derivatives pricing tools.

The most popular application of these models in power market is by Bessembinder and Lemmon (2001). The main assumption here is that market prices are only determined by industry participants and that companies are concerned about maximising their expected profits from power transactions in order to add value to the firm and minimise volatility which increases the likelihood of financial distress and affects future investment incentives. Forward prices are then derived using an equilibrium relation that they provide for forward contracts. As discussed earlier, the forward premium represents the compensation required in equilibrium by those who support the price of risk of the underlying commodity. Their findings support the above theory and prove that the sign and size of the forward premium is related to economic risks and the willingness of different market participants to bear these risks. Their model indicates that forward power prices are a downward biased predictor of the future spot price if expected power demand is low and demand risk is moderate. However, the equilibrium premium increases when either expected demand or demand variance is high, because of the positive skewness in the distribution of the spot price. Their empirical evidence for the PJM market indicates that the risk premium is positive and greater during summer months, which was also shown by Villaplana (2003) with the inclusion of a jump premium for those seasons.

Continuing the discussion on the existence of a risk premium in the market and the existence of economic risks, Longstaff and Wang (2004) also find that forward prices differ from the expected spot prices. They further provide empirical evidence that risk premia increase if demand forecast is higher, and lower forward premia result from higher volatilities in unexpected spot price and demand changes. Their results hence are consistent with Bessembinder and Lemmon. Botterud et al (2002) used weekly futures contracts and provide evidence that the electricity market in Scandinavia is in Contango, that is to say that futures prices on average have been above the spot prices in the actual week of delivery, and they find a significant evidence of a negative risk premium. This phenomenon exists in Nord Pool because of the greater flexibility of the supply side using water reservoir, so there is a greater hedging pressure from the demand side. Alizadeh and Nomikos (2003) give further support on the evidence of risk premium in Nord Pool, where they test if future contracts are unbiased estimators of the spot. In their tests, the hypothesis is rejected by simple univariate tests, however when using multivariate tests based on parameter restrictions on the cointegration between spot and futures prices, they find that near maturity contracts are unbiased predictors but, as the time to maturity increases the future contracts become biased, hence there is a risk premium in the market.

Following the discussion on electricity modelling, Barlow[9] (2002) introduces an interesting diffusion model, in which price is determined by the equilibrium level of supply and demand. Hence supply and demand are functions of price and they must be in equilibrium, $S_t(P_t)=d_t(P_t)$, at every point in time especially in electricity. Furthermore, since electricity demand is inelastic $d_t(P_t)=D_t$, for some stochastic process of demand. Supply is non-random and independent of time, and therefore becomes only a function of price, $S_t(P_t)= g_t(P_t)$. The above theory then leads to the conclusion that price is determined, by the equilibrium level by inversing the supply function: $P_t=g^{-1}D_t$. Now from the supply stack function, the supply curve is increasing, however the total supply is limited, hence one can use the following function: $g_t(P_t)=a_0+b_oP_t^a$ where $a<0$. Given that supply and demand are in equilibrium, and assuming that that demand follows an Ornstein-Uhlenbeck process and can reach a finite maximum level, electricity prices are then proved to follow a *non-linear Ornstein-Uhlenbeck* process. The results showed that the fit of the model to the spot prices was good, especially in picking up the spikes when the model was tested for the markets of Alberta and California. However with this model there are no closed-form solutions for derivatives prices, and the assumption of a constant supply is too restrictive, as spikes are mainly generated by outages.

[9] I classify this as an equilibrium model using stochastic modelling, since it relies on the dynamics of supply and demand

2.5 Hybrid Models

A recent development in the power modelling literature has focused on another class of models called *Hybrid Models*. This method combines the distinctive features of both fundamental as well as pure stochastic models. The fundamental part is used to represent supply and demand relations, while stochastic techniques are used to represent the evolution of the underlying drivers. Its goal is to find out what causes prices to move and identify primary random variables, model the evolution of these variables and hence use them to identify the movements in power prices. One of the first models introduced is by Geman and Eydeland (1999), who suggest to approximate the future prices at time t for delivery at T, $F(t,T)$, as the sum of the base load price, p_0, and the supply stack function which is a function of the expected load $L(t,T)$, and the forward prices of the marginal fuel[10] $w(t,T)$:

$$F(t,T) = p_0 + \varphi\big(w(t,T), L(t,T)\big) \tag{2.22}$$

Assuming that Load has a normal distribution and the forward fuel prices are driven by a GBM, option pricing formulas are easy to derive. A more in depth model with empirical evidence is build by Pirrong and Jemarkyan (2001). Arguing that a model for power should capture seasonality and the higher jumps occurring at hot seasons (in the US), the jump probability and magnitude is explained by changes in demand growth and capacity. Hence the equilibrium approach can give the relationship using two variables that reflect *demand* and *marginal costs* in order to identify how prices change. They propose a mean-reverting process for the demand variable (Load), whose equilibrium level follows a seasonal pattern. However in the stochastic process of Load they impose a penalty function that accounts for the fact that when load exceeds the physical capacity of the generating and transmission system, the system may fail, imposing substantial costs on power users and the ISO intervenes to reduce the use of power, as soon as this event happens. The second variable that needs to be modelled is the marginal cost, for which they use a forward model, as in section 2.3 to model the forward price of the marginal fuel. Then the cost of 1MW of electricity produced is inferred using the "heat rate" which measures the amount of fuel required

[10] The fuel that is used to produce electricity and hence sets electricity prices most often

to generate 1MW of power, and thus using a 2-step semi-parametric method they infer the equilibrium relationship between Load and marginal cost, in order to find the price of electricity.

The model is tested empirically for the PJM (Pensylvania, New Jersey and Maryland) market. They find that forward prices were above expected spot prices for the summer months, indicating the long hedging pressure for that season due to the right skewness in spot prices. Hence the probability of spikes in high demand periods is high, so the premium demanded by the seller of forwards to bear this skewness risk is quite large. For the rest of the months the risk premium is negative but not as large in magnitude as the ones estimated for high demand periods. A very interesting conclusion of the paper is that in the absence of risk takers, such as investment banks and hedge funds, there is limited availability on risk taking, hence the market price of risk becomes more volatile.

One of the main problems faced by Pirrong and Jemarkyan (2001) is that capacity data are not available, which is of paramount importance in order for the model to capture the spikes, thus in the empirical part of the paper the penalty function is not included. Furthermore they use one short-term forward contract to estimate the risk premium, rather than the whole forward term structure, thus the information about the expectations in the market is very limited. Also as pointed out by Geman and Eydeland (1999), in such models there is a problem in specifying the functional form of the transformation stack (the relationship given by load and marginal cost to calculate the price).

Eydeland and Wolyniec (2003) overcome these problems by proposing a model where power prices are a function of the drivers forming its price. First, they build forward curve models of the Heath-Jarrow-Morton form for the evolution of the forward curves for each generating fuel (e.g. natural gas and heating oil) and weather derivatives, and simulate the spot values implied by these models[11]. Next step is to find the dynamics of the load which, due to limited data availability, is approximated using a polynomial relation with temperature. Furthermore, a Poisson distribution is

[11] In the literature it has been shown empirically that for Natural gas and Oil the HJM is very successful in modelling the dynamics of the forward cure.

used to model the process of the Rate of Outages for each generator in the system. Next, the generation stack is build by sorting in an ascending order all the generating units according to their marginal cost of generation and find the marginal cost in the market for each level of demand. The relation between the generation stack and prices gives the minimum prices that generators would be expected to bid in the market, based on their marginal costs of production. To estimate what actually is bid in the market three parameters are introduced in the model that capture the premia from numerous optionalities and market constraints, the uncertainty in the reserve levels and operational constraints as well as the uncertainty about the operational characteristics of the plants in the stack. These parameters can be estimated in order to match the market price of liquid tradable derivatives.

The fundamental Hybrid Model of Eydeland and Wolyniec (2003), is perhaps one of the most successful models in terms of simulating price paths that resemble the characteristics of the physical market. As they argue, it can be used to price and hedge every type of uncertainty in the power market. To implement this model one needs detailed information such as start-up costs, rump-up and rump-down rates. However, the characteristics of each unit operating in a market are very difficult to find, especially in terms of outages but also in terms of their heat rates, capacity and emission costs. Moreover there has been evidence that the heat rate changes due to technological advances and is thus hard to estimate for every unit. Also the functional form between load and weather maybe accurate for explaining residential demand however, it may not be adequate for representing industrial demand which creates load growth over time. That is why it has also been argued in the RISK and EPRM magazines that weather derivatives may not be very efficient for hedging electricity prices and more specifically electricity demand, since the correlation between them has a high probability of breaking down. Finally, electricity in Nord Pool is mainly Hydroelectric; even though the storage cost of water is negligible, the investment costs to install water reservoirs are high, and hard to find data for. Consequently, one cannot identify the prices at which hydro units bid. Hence, although this model may look very appealing especially for market participants such as generators or retailers, it is very computational and data collection intensive and sensitive, and may provide misleading results in terms of pricing and hedging, if too many unrealistic assumptions are made.

2.6 Our contribution to the literature

From the previous discussion on the literature review, it seems that most of the papers have at least one of the following four disadvantages: First, they fail to capture the dynamics and distributional characteristics of the market correctly (e.g. Lucia and Schwartz, 2002). Second, they do not provide any empirical evidence on the fit of their proposed models, which in most cases is due to the complicated structure of the models for which econometric techniques may not be yet available (e.g. Benth et al, 2005). Third, they do not use all information given from the whole term structure and as a consequence they are not arbitrage free (e.g. Geman and Roncoroni, 2006). Fourth, if the models are too complicated, analytic solutions for derivatives pricing cannot be derived, thus their application is limited (e.g. Barlow, 2004).

The aim of this thesis is to fill these gaps in the literature. As discussed in the previous sections, our main task will be to build an arbitrage-free model that best describes the dynamics of the spot prices and fits the forward/futures prices, but at the same time it can provide analytic solutions for derivatives pricing. Each model will be assessed in terms of fitting the forward prices cross-sectionally and across time. Furthermore, we assess each model's ability to capture both the characteristic path followed in the real world by the electricity spot prices, as well as capturing the four moments (mean, variance, skewness, kurtosis) of their distribution.

More specifically, we introduce for the first time in the literature three Jump Diffusion spot models. The first is called the Spike model, which captures the seasonality and regularities in the market, but most importantly it incorporates two different speeds of mean reversion in order to capture the spiky behaviour of jumps and at the same time to distinguish the price behaviour between periods when there are jumps and when the process follows a smoother path. The second model is the Regime Switching Spike model which is an extension of the Spike model, but takes into account the large amount of hydropower generation in Nord Pool and thus captures short- and mid-run risks, by distinguishing between periods when there is enough water to produce cheap hydroelectricity and periods when users have to rely on more expensive sources of generation such as natural gas and oil. The third model is the Three-factor Spike model which is a more general extension of the Spike model but captures both short-and long-run variations, however this time the long-run dynamics are not dependent only on the water availability in the system as in the Regime Switching Spike model, but depend on general long-run risks that are expected to persist and arise mainly from regulation risk, possibility of horizontal or vertical disintegration, threat of increased competition, load growth and generation plants closing down.

In addition, we provide Closed-form solutions for the term structure for each model using the Transform analysis by Duffie et al (2000), since these models belong to the family of Affine Jump Diffusion. The models will be tested both in terms of their goodness of fit to the forward term structure, as well as in terms of capturing the distributional characteristics and price paths of the spot prices. In this way the proposed models are tested from all aspects that a model needs to be assessed in order to ascertain whether it is suitable for derivatives pricing in the power market.

To our knowledge very little evidence has been given in the literature on how well a model fits the term structure in the cross-sectional and time series level in electricity markets using all information that is available about the term structure, from futures and forward prices. Furthermore, in the power market literature no model has been proposed that is able to capture the spiky behaviour in a way that is easy to calibrate and at the same time can provide closed-form solutions for derivatives pricing, which is of paramount importance for market practitioners that need quick answers to make their decisions. In addition, the effect that water levels in the reservoirs may have on the equilibrium level of electricity prices, is an issue that has not been investigated in the literature. This is very important since the equilibrium level adjusts the level of mean reversion of prices, hence if there is a high probability of the water reservoir levels being low in the future, this will increase the expected equilibrium price level and consequently, the price of derivatives with the right to buy the underlying commodity for these maturities will increase. The opposite of course stands when water levels in the reservoirs are high. Furthermore, our spot models use a limited number of unobservable factors to summarise the shape of the whole term structure of electricity in a way that is sufficiently accurate and tractable, which is along the lines of Schwartz and Smith (2000), Manoliu and Tompaidis (2002), Sørensen (2002) and Lucia and Schwartz (2002), and thus imply non-perfect correlation across forward contracts with different times to maturity, something that is captured by the three-factor spike model. However our study not only extends the application of these models in the electricity markets, but also expands these models to incorporate Affine Jump Diffusion processes, which will let us go beyond the Gaussian framework and hence capture more accurately the distributional characteristics of power prices.

Furthermore, given that the proposed models provide a very good fit to electricity spot and forward prices, the next logical step is to examine their implications in terms of option pricing. Therefore in the thesis we also investigate the proposed models' implied volatilities generated using the Black and Scholes formula, and interpret their intuition in terms of reflecting the kurtosis and skewness of the underlying distribution. With the exception of Branger (2004) who applies general extensions of the GBM model in the equity market, such analysis lucks in the option pricing literature for electricity derivatives even for the simplest mean reversion model which has not yet been examined for any market. As it will also be shown the implied volatility for mean-reverting models produces a skewness by itself as mean reversion does not allow the underlying prices to reach extremely high or low price levels, thus reducing the probability of OTM call options ending ITM. Furthermore, based on our proposed spike model, we find that the mean jump size explains the volatility skew whereas it is the jump size volatility, rather than the jump intensity, that plays a more significant role in making the volatility smile more pronounced and thus increasing the kurtosis of the price distribution.

Next, we also explore the option prices generated by the three proposed models and provide the intuition behind them. More specifically we find two very interesting results which are pointed out for the first time in the literature; one is that mean reversion causes European Option prices to reach a constant value as time to maturity increases and then the discounting starts affecting the option values more thus decreasing the value of European options as time to maturity increases, something which is totally different to what one would get from using the Black and Scholes model where prices increase with time to maturity. On the other hand we find that using the regime switching spike model gives more intuitive results than using the three-factor spike model. Given that we are currently in a regime with no water the regime switching spike model produces a humped shape for European option prices as time to maturity increases, whereas an exponentially increasing shape is produced if we are currently in a regime with water. The three-factor spike model on the other hand may produce a similar shape, but translates the water reservoir risk as a short-term effect rather than a mid-term (or more precisely yearly) effect, which is not intuitive.

In addition, using our spike model we also perform a sensitivity analysis on the prices of Asian options. We show how Asian options prices are affected by changes in the parameters and more importantly how option prices change due to a change in the mean reversion of spikes. We also examine how the importance of spikes changes as the jump intensity increases, since the averaging effect may not diminish their contribution to the option's price. These findings are important to practitioners who are interested on the factors that will significantly change the moneyness of Asian options, when engaged in trading these products, and are presented for the first time in the literature on Asian options.

Finally, the last chapter is devoted on the so-called Swing contracts that have been developed in order to give their holder flexibility with respect to the amount purchased in the future. These contracts are mainly found in energy markets and allow for flexibility in the amount delivered depending on the customer's needs. This refers especially to electricity, but Swing contracts appear also in coal (as in Joskow 1985 and 1987) and gas markets (Clewlow and Strickland 2000). We explain the theoretical complications in pricing these exotic options, using a stochastic optimal control technique, since no closed-form solution is available due to their early exercise features. However, since our spot price processes are based on complicated reduced form models that have jumps, lattice methods and trees become impractical since the dimension of the problem is too big. We thus use the extended Least Squares Monte Carlo algorithm for swing options, by Dörr (2003). We perform for the first time in the literature, a sensitivity analysis on the main drivers of the incremental swing premium with respect to strips of European options whose maturities are within the period of the swing rights. Moreover, we provide the lower and upper bounds of the swing options, which are strips of corresponding European options and Bermudan options, and we see how the swing options converge to their lower bounds as the number of swing rights approaches the number of swing dates. We also carry a comparative analysis on the pricing of swing options, between the different proposed short-term models. Finally, we explore how the swing premium is affected by penalties imposed, when the total number of swing rights exercised during the delivery period does not satisfy special contracted restrictions.

2.6.1 Comparative advantages of Proposed Methodology of Modelling

2.6.1.1 Why is our approach better than the rest of the Reduced Form Models?

As discussed earlier, only Lucia and Schwartz (2002) consider a model for electricity prices based on both spot and forwards. Most of the reduced form literature has focused on modelling the spot only hence the models are not arbitrage free. Furthermore, most models are one-factor and do not capture the most significant characteristics of the market, while at the same time they impose a perfect correlation between forward prices of different maturities. Some of the papers, like Huisman and Mahieu (2003), specify that for long-term obligations different models have to be developed, or in Elliot and Sick (2003) it is mentioned that even though there is a trend in spot prices there exists some uncertainty about this trend. However none of the papers has specifically addressed these issues.

In our research, we want to be able to capture as many as possible sources of uncertainty, irrespective of whether they originate from short-term variations, macroeconomic uncertainty or technological advances. Our approach estimates the models in the real and the risk-neutral world, by estimating their parameters using both the spot and forward term structure (since the forward market is the most liquid market for derivatives trading in electricity). Even though the spike model captures mainly the short-term variations of electricity prices, we extend it by introducing the regime switching spike model, which can capture longer term variations mainly caused by the amount of water in the reservoirs which affects the marginal cost of production and thus the equilibrium level to which prices revert. On the other hand, the three-factor spike model is a more general extension of the spike model, and is able to capture both short- and long-term variations, where the former is inferred from the spot market and near maturity forward contracts and the latter is inferred using long-maturity contracts.

Finally, we also capture the spiky behaviour of the spot prices in a practical, easy and at the same time accurate way, using our spike models where we impose two speeds of mean reversion, to distinguish between periods when spikes occur and when the price follows a normal mean-reverting process. This in turn improves the spot-forward relationship, which is very helpful both in terms of hedging but also in terms of capturing the distributional characteristics of prices. The above facts are very important in providing an accurate model for pricing derivatives, and it is to some extend surprising that not all of them have been examined at the same time in the literature. By investigating these issues, this is then a contribution to the literature.

2.6.1.2 Why is our approach better than the Forward Curve models?

As discussed in the literature review, previous attempts have shown that modelling the Forward curve directly has failed to capture the dynamics of both spot and forward contracts. First of all the contract specification for a delivery period averages out the jumps occurring in the spot market (since the forwards in the electricity market are more like swaps), hence if one uses only financial contracts in the electricity market, he will be ignoring information about the behaviour of spot price dynamics. Second, in the electricity market, even in Nord Pool which is considered as the most liquid market internationally, long maturity contracts are not traded very often; this creates a problem of sparse data sets in which at each date there are only a few different maturity contracts traded and thus the use of daily data on the term structure needed to calibrate these models may not yield reliable results. Furthermore, our modelling procedure is more practical and tractable in pricing derivatives on the spot, and at the same time it provides a clear relationship with the forward market which can be used for *delta* hedging strategies, as it is shown in the appendix of Chapter 5.

2.6.1.3 Why prefer this approach than the Hybrid?

The Hybrid model by Eydeland and Wolyniec is considered to be the most successful in terms of simulating the spot electricity price paths. However it relies on detailed information, such as plant availability and generating capacity and generators' operating characteristics, that is perhaps available only to regulators. It is very hard to find the exact characteristics of every generator in the market such as the number of times of forced and unforced outages, the fuel it burns, its exact Heat rate and the capacity data. Therefore empirically these models will have to rely on too many assumptions, which may be too restrictive and give misleading results. On the other hand our approach relies on more liquid data that are publicly available in the market and, as it will be shown in the subsequent chapters, our models can capture the most significant characteristics in the market and at the same time derive analytic formulas for derivatives pricing. The latter is important especially for traders and market participants who need to have quick results to take decisions.

2.7 Conclusions

In this chapter we examine the literature review on electricity modelling for the purposes of derivatives pricing. We followed the same classification as in Eydeland and Wolyniec (2003) and discussed the main advantages and disadvantages of each modelling approach. More specifically we found that in the spot price modelling, none of the papers address the issues of capturing the dynamics of electricity price accurately and at the same time implementing them in derivatives pricing and thus finding their risk-neutral measure. On the other hand, the Forward curve models have not been successful up to date since a large number of factors is required to capture the dynamics of the whole term structure. At the same time modelling the forward curve directly cannot capture the spiky behaviour of the spot prices, since the spikes are averaged out during the delivery period. We also explained that the fundamental equilibrium models are not suitable for derivatives pricing since they do not capture the dynamics of prices in a robust quantitative manner. Finally, the hybrid models, even though they are very appealing, their required dataset is difficult to obtain.

The last section of the chapter discussed our contribution to the literature, where we pointed out that our modelling framework aims to provide an arbitrage free model that is able to capture the spot dynamics accurately, and at the same time provide closed-form solutions for derivatives prices and therefore analyse the implications of each model on the prices of plain vanilla as well as exotic products. The next chapter describes the operational functions of the Scandinavian power market, which is investigated in this thesis, and examines its stylised facts in order to identify the main features a successful model needs to capture.

3. DESCRIPTION OF THE NORDIC MARKET AND PROPERTIES OF THE POWER SERIES

3.1 Introduction

Before we start building models for the pricing of derivatives and risk management in the electricity market, we first have to understand how the market under investigation functions and the properties of its spot series. Thus, this chapter is an introduction to the empirical part of the thesis. The second section of this chapter provides an in depth description of the different market structures that exist in electricity markets internationally. In section 3 we focus more specifically on the Scandinavian electricity market, by first giving a brief history of its development and then explain its two major market segments, i.e. the spot and the financial market. After explaining how the market functions, section 4 starts the empirical analysis, by first looking at the descriptive statistics and thus the distributional and path characteristics of the Scandinavian electricity spot market. This analysis will then give a clear understanding on the main features a model needs to capture for the pricing of derivatives in the Nordic market. Section 5, analyses the relationship between the level of water in the reservoirs and the spot price, which will lead us to one of the main driving forces for the changes in the equilibrium level that occur at different periods in time, and finally section 6 concludes the chapter.

3.2 Market Structure and Regulation

3.2.1 The Services and regulation in the Electricity Market

The electricity business, competitive or otherwise, generally comprises five independent services. First of all the *Generation process* is where production takes place for the wholesale quantities of power. *Transmission* then is responsible for the transportation of the generated wholesale power over large distances using high-voltage cable networks. *Ancillary Services* are provided by generation units (through resources which are brought in full capacity to maintain balance) and used by transmission operators to balance supply and demand in real time and to maintain overall system security, in order for demand and supply to be always in equilibrium. *Distribution* then takes place for the transportation of power from the transmission system to the consumer. Finally the *Wholesale/ Retail Supply* services facilitate the purchase and sale of the physical commodity.

The common thread running through every reconstruction is the realisation that providing at least some of these services does not necessarily require a monopoly market structure. Consequently each of these services must be unbundled and treated as a separate market. In particular, the essential feature of every reformed electricity market is the separation of the *Generating Function* from the *Transmission Function* and the provision for equal access to and fair use of transmission service. Separation generally means the privatisation of generating, transmission and distribution assets as three or more specialised entities. Generation and Supply of wholesale power are made into an explicit competitive market, while transmission and distribution continue to be considered as natural monopolies and regulated as such.

3.2.2 Primary Market Structures

After reconstruction the cash (physical) markets in electricity fall under one of the two contracting structures: pools and bilateral markets.

The main characteristic of the *pool market* is the formal establishment of the market clearing price, also known as the system price, at which all cash transactions clear. The generators sell and the suppliers buy electricity through the same pool, therefore physical risk is removed from individuals and concentrated to the pool. Examples of this market structure include Nord Pool in the Nordic countries, New England Power Pool (NEPOOL) and the California Independent System Operator (CAISO).

In the *Bilateral Markets* all transactions are entered by two parties and are independent of any other transactions in the market. Utilities do not own generating assets and must buy power to serve their native load through long-term "full requirement" contracts. Examples of these markets include the East Central Area

Reliability Council (ECAR) in North America (more specifically east central US) and the Electricity Reliability Council of Texas (ERCOT).

Moreover in any given market, energy cash settlement can be handled in different ways. *Day-ahead* markets, transact for generation of energy for the next day, and each hour is transacted separately. *Day-of* markets, transact for generation of energy for the rest of the day, and every hour is transacted separately. *Hour-ahead* markets, transact generation of energy for the next hour. *Real-time* markets are reconciliation markets that clear any deviations from the predicted schedules entered in the previous markets (day-ahead, day-of and hour-ahead).

3.3 THE SCANDINAVIAN ELECTRICITY MARKET

3.3.1 History

The history of the Scandinavian electricity market starts from 1971 when a new market structure was established, called Samkjøringen, to coordinate the Norwegian electricity production. Every week Samkjøringen set the daily or part-of-the-day price for electricity. This price was used to decide the Norwegian electricity production and the exchange with other countries. However the market changed in 1991, when the Norwegian Parliament approved a new Energy Law. This law introduced market-based principles for the production and consumption of electricity in Norway. It introduced competition as a tool for ensuring a more efficient and reliable energy supply. The act mandated separation of grid transmission activities from competitive activities such as generation and supply of wholesale power. The main aims of the Power market reform were to obtain a better balance between power generation capacity and power demand, increase efficiency within the power industry and reduce regional differences in electricity prices to end-users. Therefore, after England and Wales in 1989, Norway was the second country in Europe to deregulate the electricity market.

In 1993 Samkjøringen merged with Statnett SF to create a new company called Statnett Marked AS. Statnett Marked AS organised the new Norwegian market place for electricity from 1993 to 1996. In 1996 the Swedish grid company, Svenske Kraftnät, bought 50% of Statnett Marked AS and became part of the power exchange area. At the same time Statnett Marked AS was renamed to Nord Pool ASA. Finland joined the power exchange area in 1998, western Denmark in 1999 and eastern Denmark in 2000, and the first successful multinational market was formed.

As a result of the deregulation process in the Nordic electricity system, the transmission network is owned and operated by a number of independent Transmission System Operators (TSOs), whose activity is subject to regulation and control by public authorities. This guarantees a non-discriminatory access to the grid to all market participants in the new electricity market.

3.3.2 Power Generation and Storability in the Nordic Area

As we argued before storability of electricity is in general not possible for consumers of electricity unless they have some kind of pumping storage facilities or very large batteries, which very rarely happens. The consumers do not need storing facilities either, as long as they are continuously served with a sufficient quantity of power. The electricity market is in this way different from other commodity markets, where supply does not take place continuously. For suppliers with water reservoirs on the other hand, it is possible to store energy in the form of water behind dams. The volume kept in the reservoirs is usually between 30% and 100% of the total electricity generating capacity in Norway, depending on the time of the year, and is always well above the demand for the following month (for Norway only though). As it will be shown later in this chapter, 1996 was a year with high import of power because of the extremely low inflow to the reservoirs. The prices were accordingly very high. The limited storage capacity does therefore also contribute to uncertainty about prices.

More generally power generation in Scandinavia is mixed as shown in Figure 3.1, Denmark uses 85-90% fossil fuel-based generation and 10-15% wind power. Norway has nearly 100% hydropower production. Sweden and Finland rely on a mix of hydropower, nuclear power, and conventional thermal generation.

Figure 3.1: Power generation by source in Scandinavia, 2003

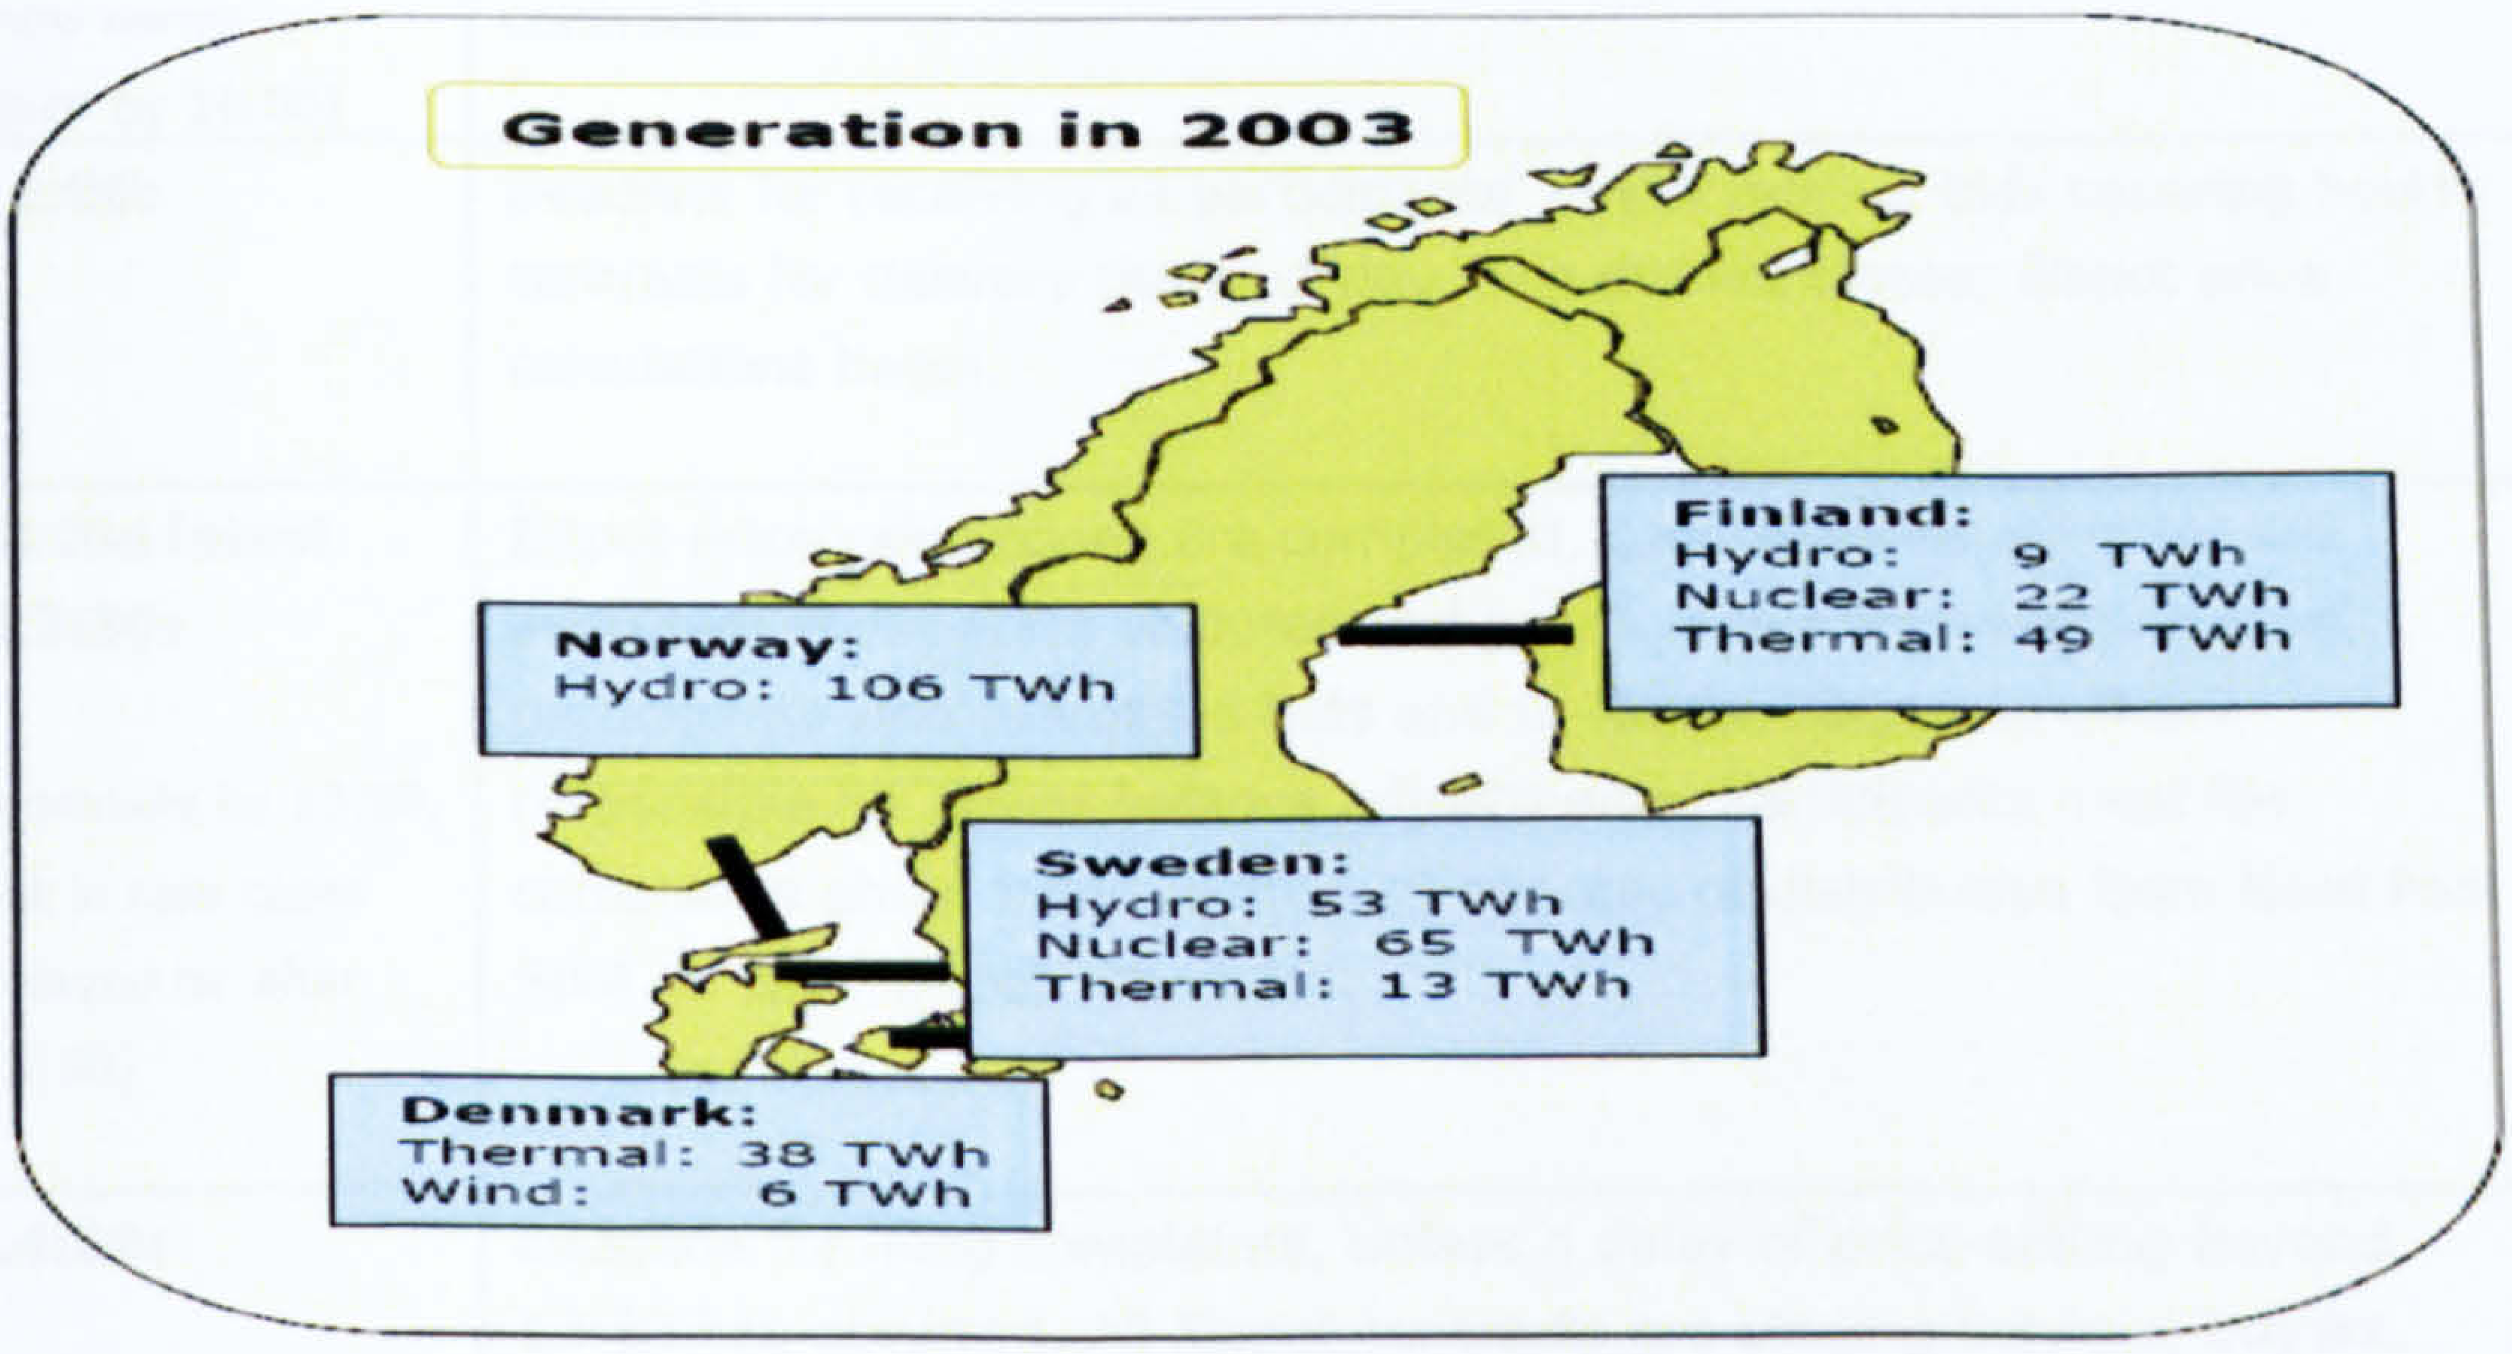

Source: Nord Pool report

3.3.3 The Physical Market

The physical market is called *Elspot*, which is a "spot" market where day-ahead electric power contracts are traded for physical delivery for each one of the 24 hours during the following day. It provides a neutral, transparent reference price for both the wholesale and retail markets. Every contract in Elspot refers to a load, in megawatt-hours (MWh, 1 MWh equals 1,000 kWh), during a given hour, and a price per MWh. A price, called the *System Price*, is fixed separately for each hour for the next day, based on the balance between aggregate supply and demand for all participants in the whole market area (the so-called Nordic Power Exchange Area), without considering capacity limits in the grid among countries. The participants in the spot market submit sealed bids and offers for the following day to Nord Pool before 12.00. These orders are demand and supply schedules that specify the price-quantity combinations at which buyers and sellers are willing to trade for each single hour the following day (see Table 3.1).

Table 3.1: Daily routines for trading in Nord Pool's spot market.

10:30 : (now normally given by 10:00)	Deadline for TSOs to submit their capacity allocations for Elspot contracts.
12:00:	Deadline for receiving all participants' Elspot market bids covering hourly contracts for delivery the next day. The market closes; Elspot price calculations begin.
At the latest 13:30: (normally by 13:00, but in rare cases delayed to after 13:30)	Elspot price calculations are completed, confirmations of trades are executed; these state volumes and prices, and are dispatched to all participants who submitted bids and to Nordic TSOs and others responsible for power balance adjustments. Participants must file complaints about trades within 30 minutes of distribution from Nord Pool Spot of their Elspot schedules.
14:00:	Deadline for filing complaints, unless a delay of price-setting beyond 13:30 has occurred. All Elspot contracts are binding between buyers, sellers, and Nord Pool Spot.

Source: Nord Pool reports

A *system price* can thus be defined as the market-clearing price at which market participants trade electricity for the entire exchange area when no transmission constraints apply (see Figure 1.1). The spot market is also the primary Nordic market place for handling potential grid congestion (called grid bottlenecks) that is, insufficient transmission capacity in a sector of the grid. In other words, Elspot is a market place where energy and capacity is combined in to one simultaneous auction. Within the Norwegian power system and at the border interconnectors between Nordic countries, the spot market price mechanism is used to alleviate grid congestion by establishing different area prices. The Nordic market is partitioned into separate bidding areas, which can become separate price areas, if the contractual flow between bidding areas exceeds the capacity allocated by TSOs for spot contracts. If there are no such capacity constraints, the spot system price is also the spot price throughout the entire Nordic Power Exchange Area. If contractual flow at System Price exceeds grid capacity limits, two or more area prices are calculated for the affected spot market delivery hour. Once spot market prices and volumes are determined, it can be said that the Nordic Market is in planned balance according to predicted generation, loads and contractual flows. Since 1993 the turnover in Elspot market has increased steadily from about 10 TWh (10^{12} Watts) in 1993 to almost 120 TWh in 2004, and the trend seems to continue. In 1999, more than one fifth of the total consumption of electricity power in the Nordic countries was traded via Nord Pool, and in 2003 this increased to one third (as reported by Nord Pool statistics).

In addition, in order to handle any unpredictable differences between the planned and the real exchange during delivery, once the Elspot market is closed, the national system operators have additionally set up *regulating* or *balance markets* from which the required upward or downward regulation is obtained on short notice[12].

3.3.4 The Financial Market

Eltermin and Eloption is Nord Pool's financial market where derivatives are traded, and it has been designed to serve as a risk management tool for generators and retailers that want to hedge their future profits. At the same time, the market also tries to attract speculators who seek to profit from the highly volatile electricity prices in order to increase the liquidity in the market. Eltermin allows trading in financial contracts such as forward and futures with delivery periods up to three years in advance. Since September 1995, none of these contracts entails physical delivery; they are all settled in cash against the system price in the spot market. They refer to a base load of 1 MW during every hour for a given delivery period of one day, one week, one month (replaced block contracts = four weeks), one season, and one year that may be available for trading depending on the type of contract. European style option contracts are also available for trading since October 29, 1999.

The *futures market* trades mainly for short- to medium-term purposes. Until the end of 1999, futures contracts with delivery periods of one season were available for trading up to three years in advance. There were three season contracts during a given year: Season 1 or V1 (with delivery period including weeks 1 to 16 of any given year), Season 2 or SO (weeks 17 to 40), and Season 3 or V2 (weeks 41 to 52/53). They were available for trading until the beginning of the second previous season. Then, the

---

[12] Nord Pool is involved in *Elbas*, a physical market for short-term trade launched in March 1999 by the Finish electricity exchange El-Ex, Electricity Exchange Ltd., that allows traders mainly in Sweden and Finland to adjust imbalances after the day-ahead spot market is closed (see Nord Pool (2004)).

season contracts were split into 3 to 6 block contracts. Each block contract had a delivery period of 4 weeks (5 weeks for the case of the last block for a year with 53 weeks), and each one was available for trading until the beginning of the delivery period of the previous block. Then, they were split into weekly contracts with delivery period of one week each, which stop trading before the beginning of their respective delivery periods. Additionally, futures contracts with a delivery period of one day are also available for trading several days in advance until the day before they are due for delivery. As of fall 2003 the block contracts have been replaced by monthly forward contracts, whose delivery period equal to the number of days in a given month.

Since 2003 the futures market has a time horizon of 8 weeks. The settlement of the futures contracts involves a daily mark-to-market settlement during the trading period, and a final settlement in the delivery period. The mark-to-market settlement covers the gains and losses from the daily changes in the market price of the futures contracts. The final price-securing settlement covers the difference between the last closing price of the futures contract and the system price during the delivery period. Figure 3.2 gives an illustrative example of how the settlement procedure works in the futures market. By taking a position in the futures market, and making the corresponding trade in the spot market during the delivery period, a participant is completely hedged for the contractual volumes. The settlement procedure therefore removes the basis risk form the electricity futures market. Still the participants cannot use the futures market to hedge against uncertainties concerning future load (volume risk).

The *forward market* facilitates hedging of positions further ahead in the future, and consists of Month contracts which have replaced Block futures contracts and Seasonal contracts which have gradually been replaced by Quarter contracts since the beginning of 2004 and Yearly contracts. Year contracts with delivery period corresponding to the entire calendar year are available for trading till two days before the beginning of the delivery period. These are then split into Quarter contracts (which replaced Season contracts since 2004) which are available for trading until the beginning of their respective delivery period. There are four Quarter contracts each one corresponding to the appropriate quarter with its three months. Quarterly contracts are then split into Monthly contracts available for up to six months in advance. As opposed to the futures market, there is no mark-to-market settlement in the forward market. Therefore, the accumulated profit and loss during the trading period is not realised until the delivery period. This is so as to increase liquidity for the long-term forward contracts, since no cash payment is required during the trading period. The additional settlement throughout the delivery period is, however, organised in the same way as the futures contracts.

Figure 3.2: Settlement procedures in the futures market.

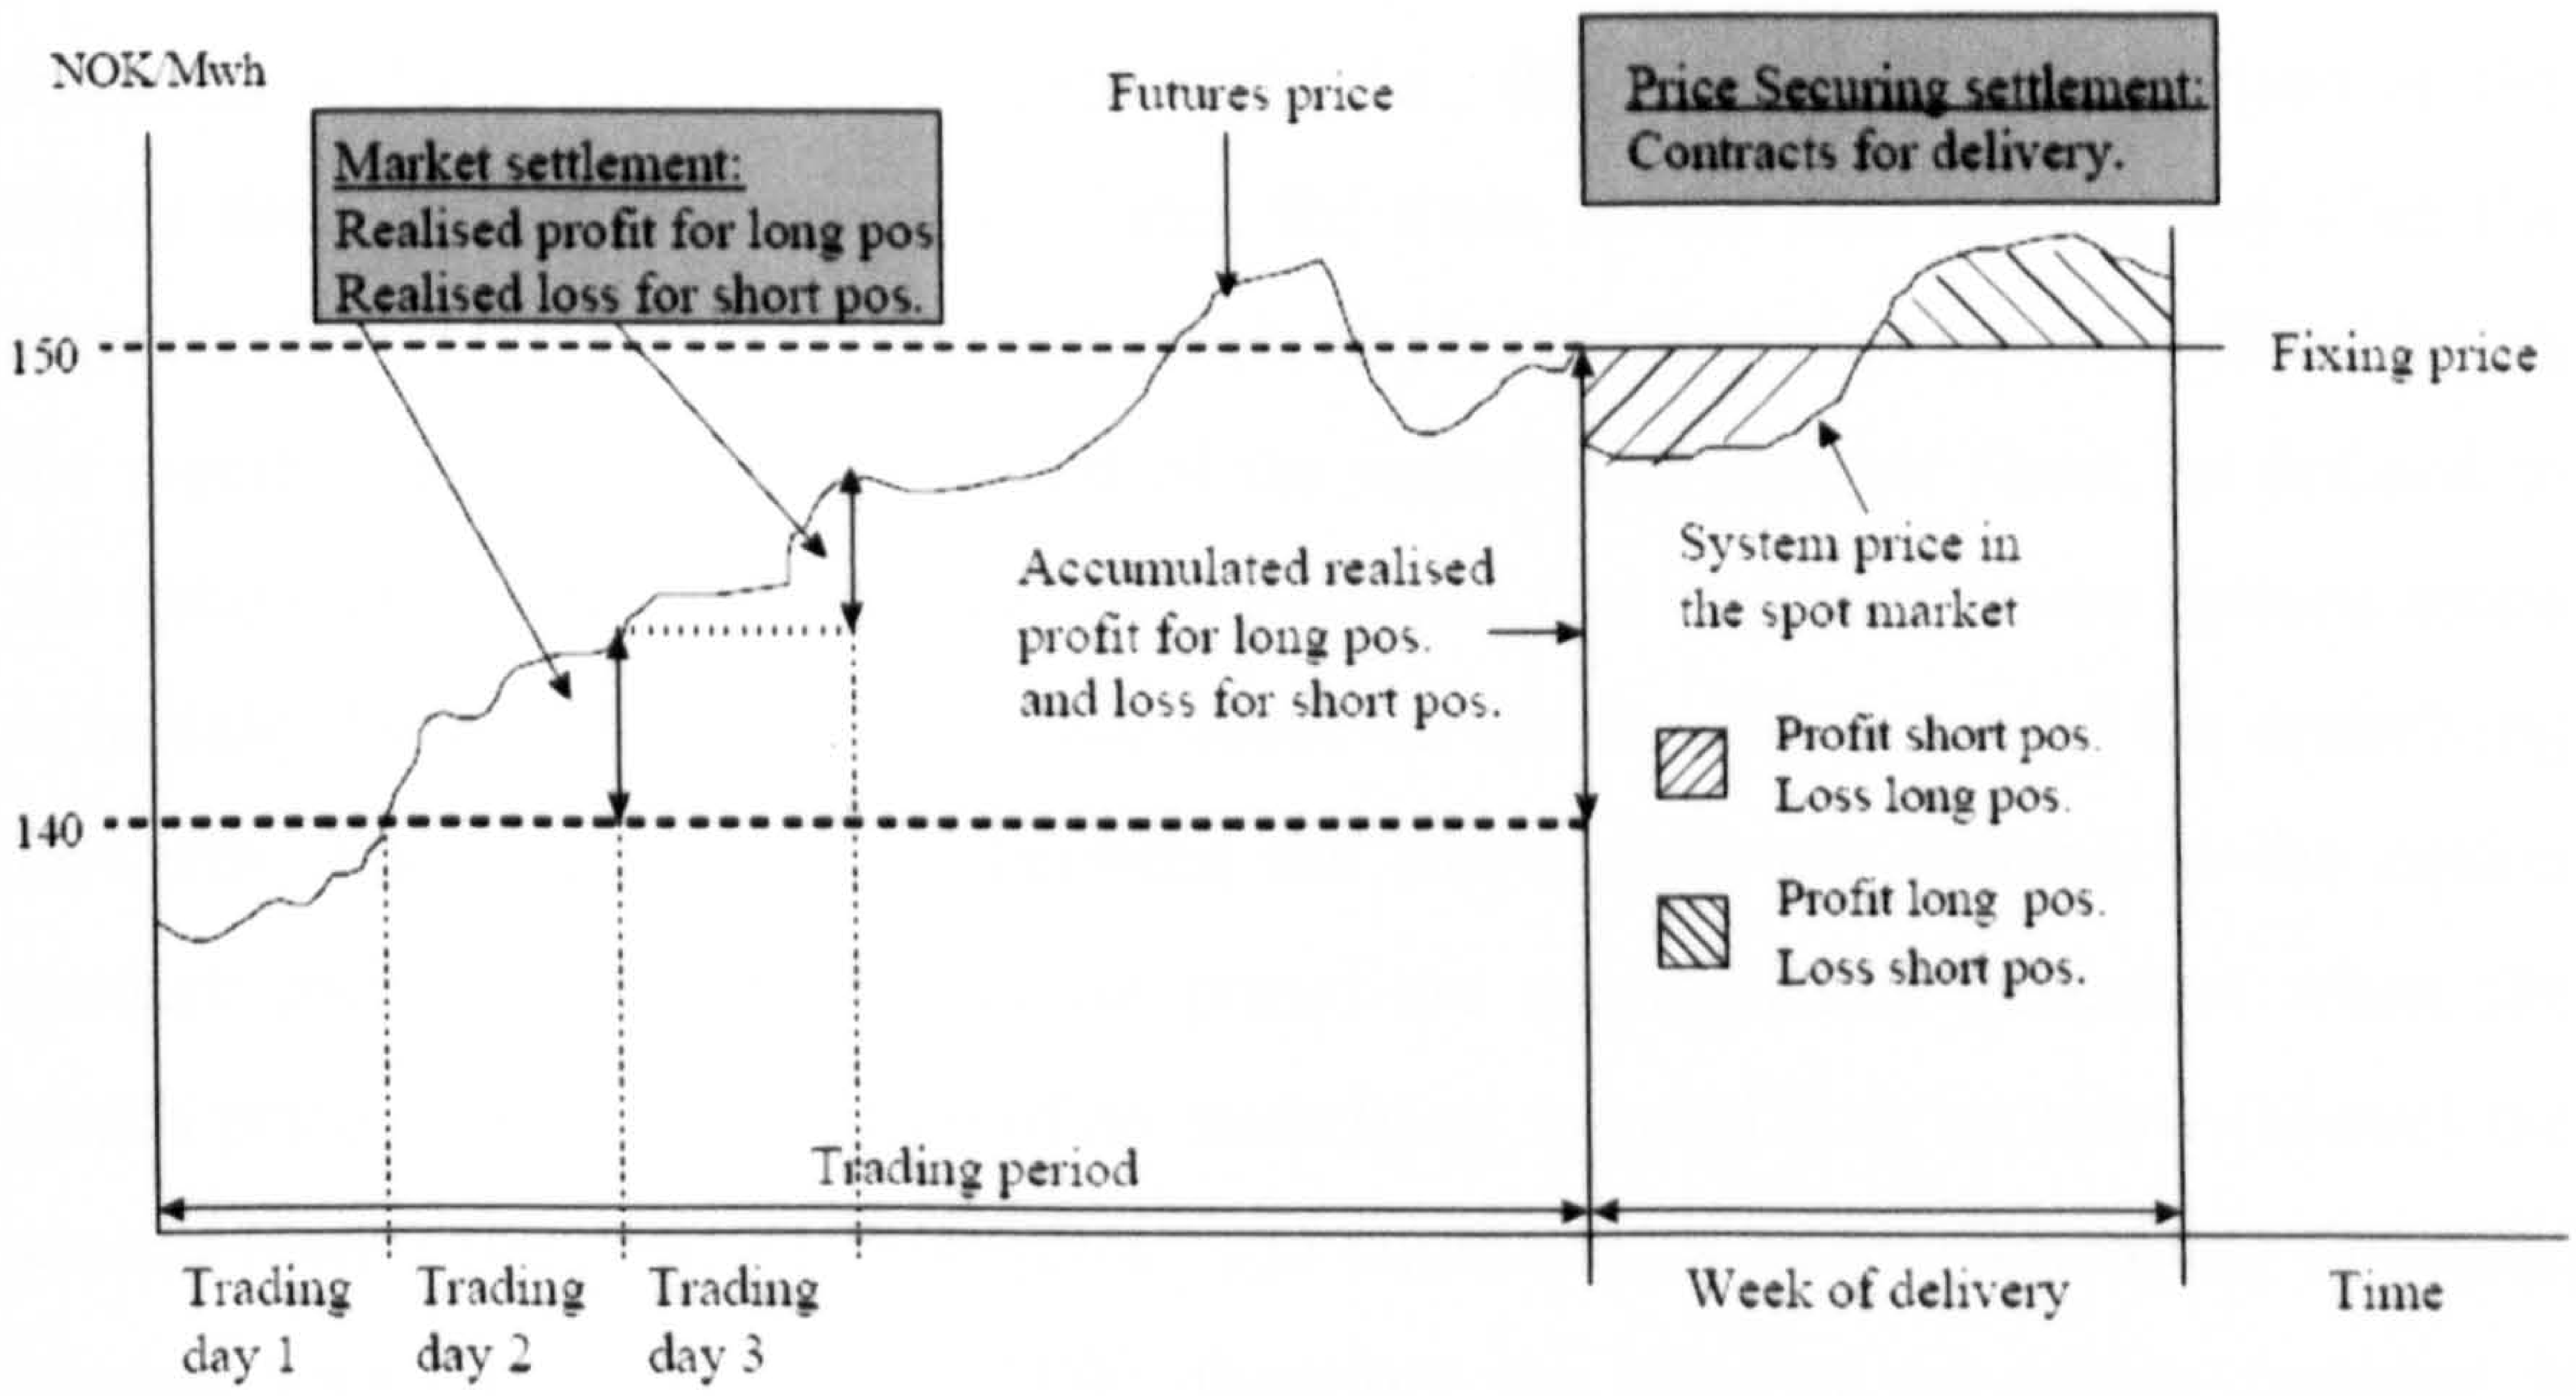

The figure above gives an illustration of the settlement procedure for a futures contract traded at Nord Pool. The purchaser of the contract receives 10 NOK/MWh in the mark-to-market settlement. Deviations from the futures closing price on the last trading day (fixed price) is taken care of the price-securing settlement, so that the contract holder ends up with a final price equal to the initial price of the futures contract, when buying the contractual amount in the spot market.
Source: Nord Pool reports

As an example on hedging using the futures or forward market, assume there is a load-serving entity, ABC, which is searching for ways to provide power to its clients during winter. ABC however has concerns that the winter power prices may be high, and it wants to shield itself from the price volatility. One solution is to lock in the winter prices as early as possibly, e.g. in June. ABC can go long on Winter Quarter contracts, to ensure a fixed price of electricity for the whole delivery period for the contractual amount. Of course that means that in the event of lower spot prices, ABC cannot take advantage of it and may end up losing money with respect to the market. However the aim of futures and forwards is not to profit, but to have a constant price at highly volatile periods.

The *Options market* in Nord Pool was launched in 1999 for standardised option products. An option in Nord Pool is a right to buy or sell an underlying Financial Market contract (Forward), at a predetermined price (*Exercise* or *Strike Price*) at a predefined date in the future (*Exercise Date*). Options are available with the two nearest Quarter (replaced Season) contracts and two nearest year contracts as the underlying forward contract. Option contracts to buy are termed *call* options, and option contracts to sell are termed *put* options. Thus, the holder of a call has the right to buy the underlying contract, and the holder of a put has the right to sell. The option contracts at Nord Pool are European-style, i.e., they can only be exercised at the expiry date. New option series are listed for trade on the first trading day after the exercise day of the previous contract series. Expiry day is set as the third Thursday in the month before the delivery period of the underlying contract starts, as defined in the option's product specifications. Nord Pool sets five strikes when an options series is initially listed for trade. Strikes are based on the closing prices of the underlying forwards. The spread intervals between the five strikes are defined in the option product specifications. New exercise prices are automatically generated when the traded price or the closing price of an underlying forward is at or below (above) the second lowest (highest) exercise price. The contract size of the underlying Financial Market forward contract is 1 MW; therefore the size of an option contract is calculated by multiplying the number of MW by the number of hours in the underlying contract. There are four different contract sizes, in MWh (with the exception of leap years):

FWV1: 1 MW * 2,879 hours = 2,879 MWh
FWSO: 1 MW * 3,672 hours = 3,672 MWh
FWV2: 1 MW * 2,209 hours = 2,209 MWh
FWYR: 1 MW * 8,760 hours = 8,760 MWh

The option premium is quoted in Norwegian Kroner, NOK/MWh and it is settled the day after the option is traded. On the exercise day, options have four basic profit and loss curves. Potential profits and losses for each option strategy depend on the price of the underlying forward contracts, exercise prices, and the premium paid. In the case of a purchased option, the risk of loss is limited to the premium paid and the potential profit is unlimited. The contrary applies to written options; i.e. the risk of loss is unlimited while the potential profit is the premium received. The pay-offs for each type of option contract depending on the position taken are illustrated in Figure 3.3:

Figure 3.3: Option Pay-offs at Expiry Date

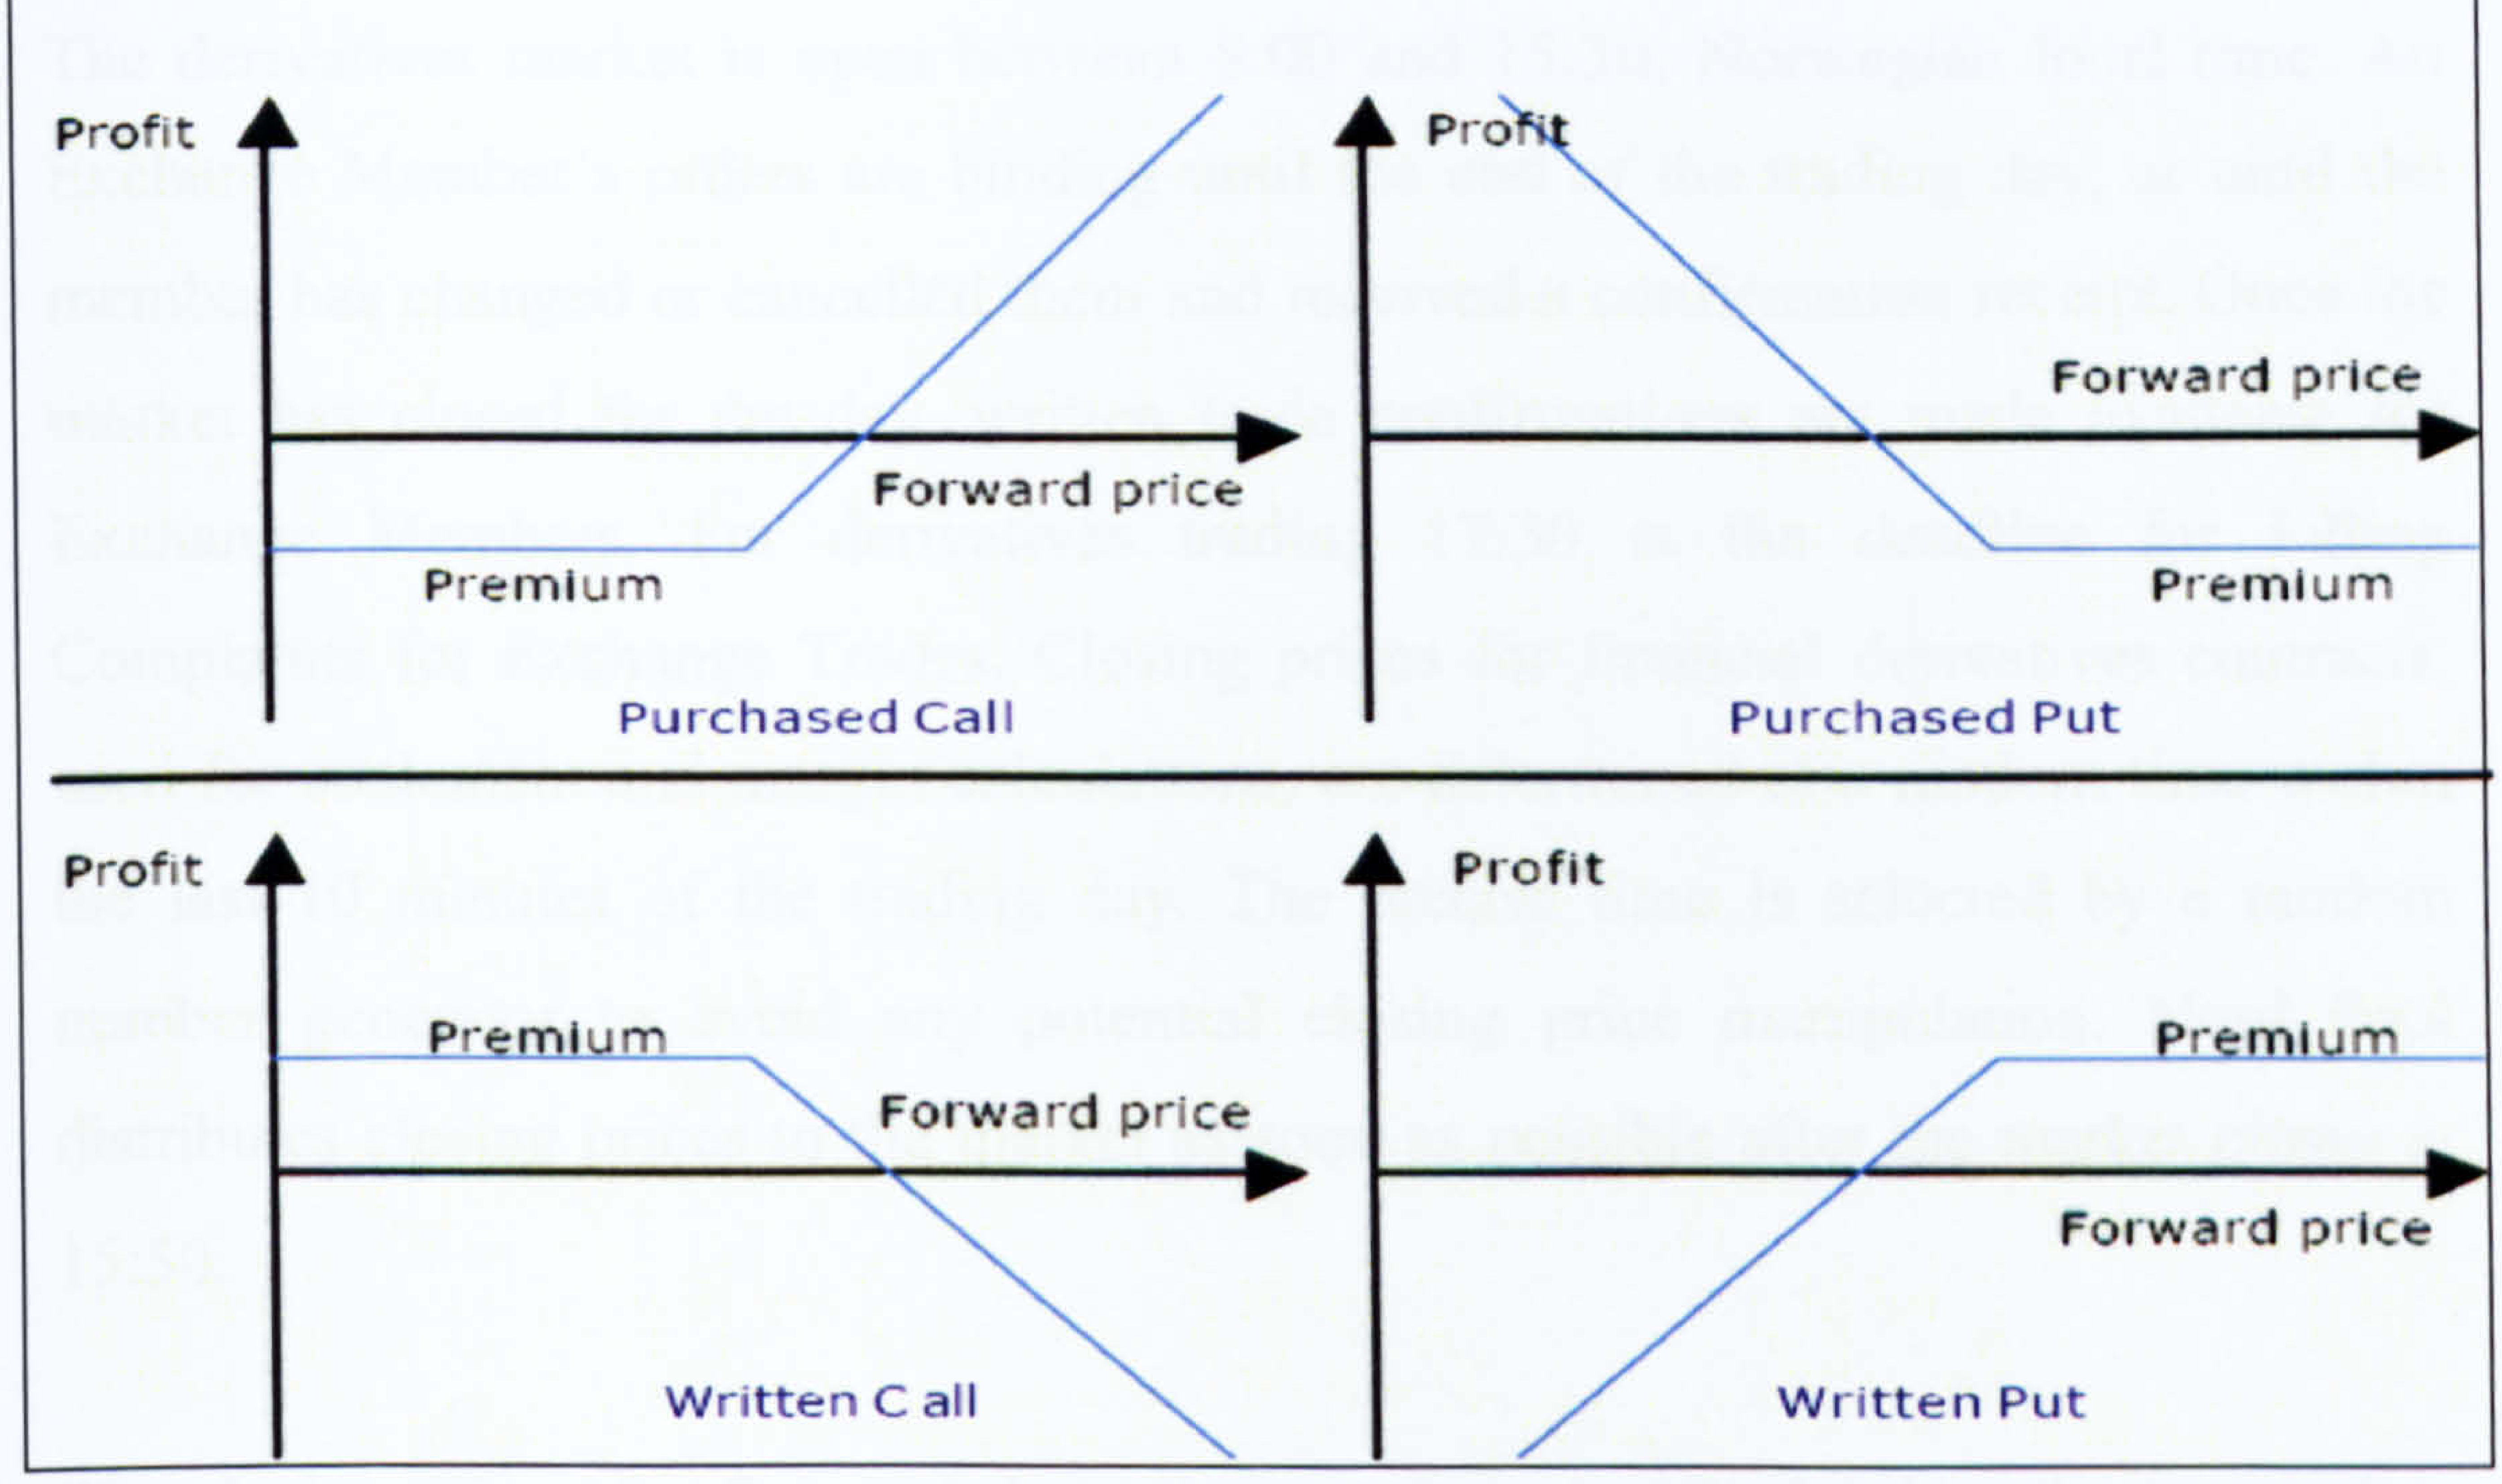

A purchased call option can be compared to an insurance against increases in the forward price of power. A buyer can use call options to insure that his fixed price of future power will not exceed certain levels, while at the same time enjoying a fall in power prices should it occur. A producer that wants to insure future sale of power can go long a put option. A variety of options strategies are used for price hedging. Because of their flexibility, options are useful to both hedgers and traders. An Exchange Member seeking trading profit, rather than price hedging, might combine a put and a call in positions known as a straddle or a strangle, if prices are expected to move sharply up or down.

An option can be described as "in-the-money" if it would result in a profit if exercised immediately (i.e., the underlying forward price is above (below) the strike price for a call (put) option). An option that would result in loss of premium is "out-of-the-money" while it is "at-the-money" if the underlying forward price is equal to the strike price. The profitability of an in-the-money purchased call option is similar to the profit potential of holding a long position in the underlying instrument. Out-of-the-money puts or calls offer potential profits, if they move as anticipated past at-the-money value into an in-the-money position.

3.3.5 Daily Trading in Nord Pool's Financial Market

The derivatives market is open between 8:00 and 15:30, Norwegian local time. An Exchange Member's orders are binding until the end of the trading day, or until the member has changed or cancelled them and received a confirmation receipt. Once the market has closed for the day, written trade confirmations are made available for Exchange Members. For derivatives trading 17:30 is the deadline for Filling Complaints for Exchange Trades. Closing prices for financial derivatives contracts, used for settlement and margin calculations, are determined at a random time within the last 10 minutes of the trading day. The precise time is selected by a random number generator to avoid any potential closing price manipulation. Nord Pool distributes closing prices to the market as soon as possible after the market closes at 15:30.

The closing price of a financial derivatives contract is calculated as the last trading price if the traded price is within the buy and sell spread, or bid and offer price, at the randomly selected time. For contracts outside the spread, or contracts that have not been traded, the closing price is defined as the average of the bid and offer, as specified by the rulebook for the financial electricity market.

Through its clearing function, currently conducted by a separate business area called the Nordic Electricity Clearing (NEC), Nord Pool guarantees settlement and delivery of all trades made at the market, by entering into the contracts as a legal counter-party for both the buyer and the seller. NEC also offers clearing services of standardized OTC bilateral financial contracts registered in the market for that purpose.

Figure 3.4 displays the time series of the yearly total volume of traded financial contracts in Nord Pool since 1993. Liquidity has increased substantially since the beginning of the market, however as Nord Pool report (2004) states, liquidity is very low for long maturity contracts, such as the second nearest quarter contracts. In addition the options market is also very illiquid. The reasons behind the latter will be explained in more detail in chapter 5.

Figure 3.4: Total volume of financial contracts traded in Nord Pool

The figure shows the total volume of financial contracts that have been traded in Nord Pool each year since 1993, in TWh.

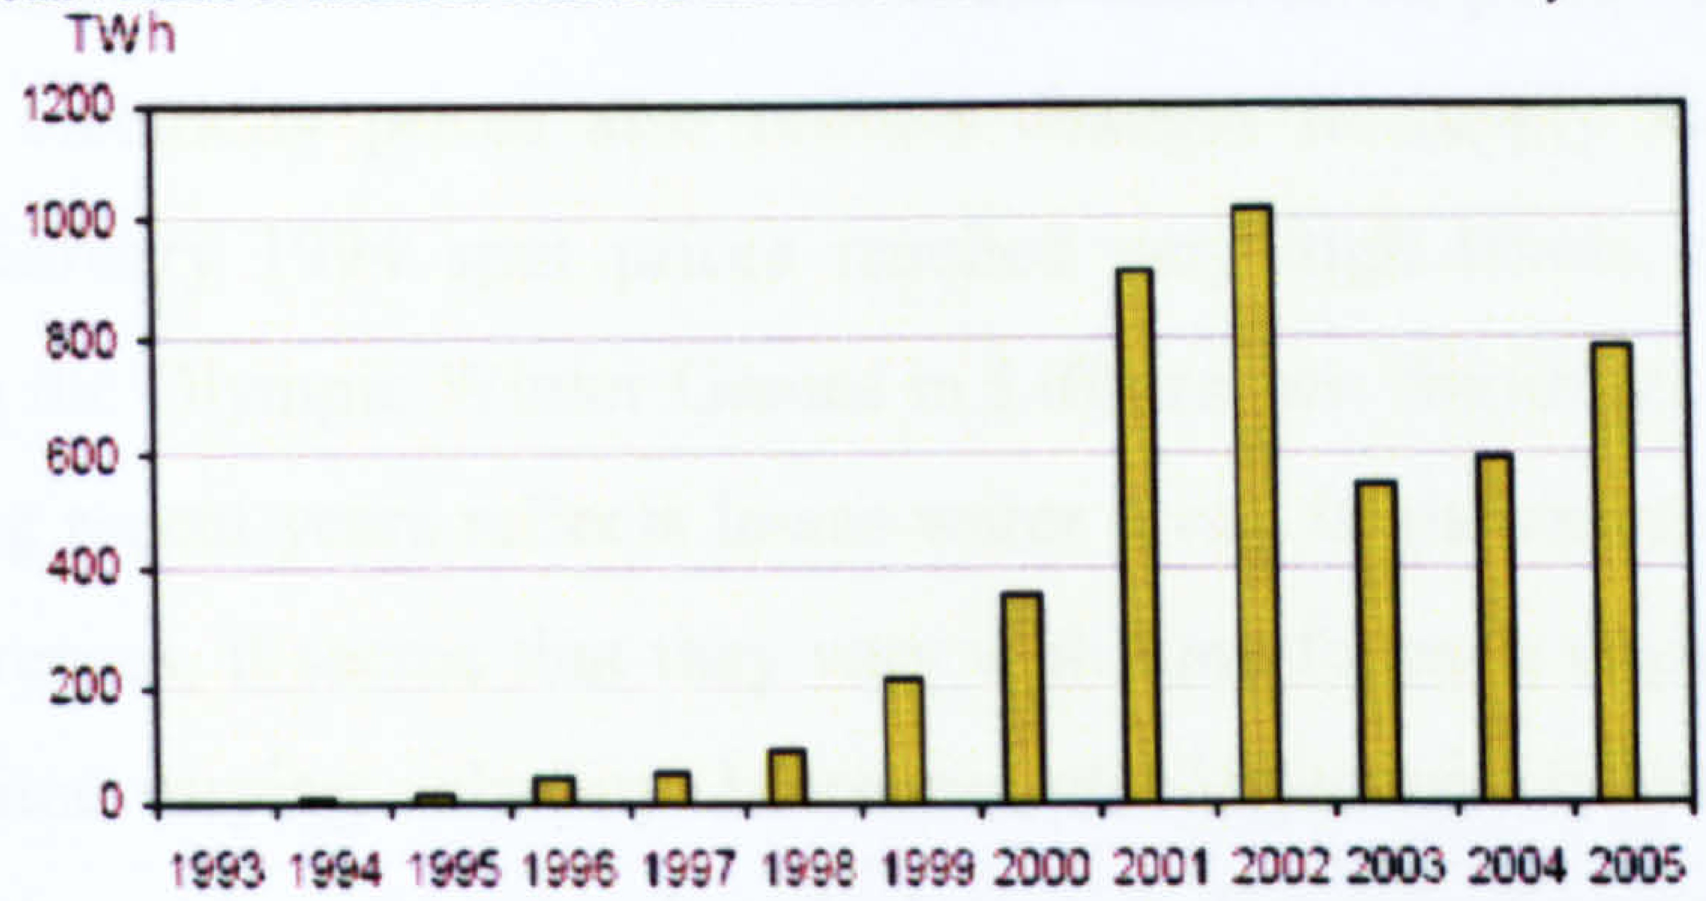

Source: Nord Pool reports

3.4 Stylised Facts of System Price

To model a series and identify its risks empirically, one has to first look at the distributional characteristics of the market to be examined. For that purpose we collect daily system prices from Nord Pool's FTP server files. The data consists of twenty-four time series (one for each hour per day), of the daily system prices (seven days a week since electricity is "traded" every day). The nominated currency is Norwegian Kroner (NOK) since the financial market is settled in that currency. The prices represent the amount of NOK per MWh in two decimal points, and the data set is collected from 1st of January 1993 to 29th of February 2004. The twenty-four hourly series are highly correlated (Table 3.2), with an average linear correlation between any two hours of 0.96, and are always above 0.96 between any consecutive hours. Eltermin uses the arithmetic average of all hourly prices for a given day, as a reference price in the cash-settlement calculations at expiration of the derivatives contracts. It makes sense then to generate a new series of the arithmetic average daily prices, and carry the analysis on that series which will be referred from now on as the "Spot" price.

Figure 3.5 shows the time series of the spot price and the logarithmic spot price, as well as their changes. It can be seen that the electricity market is very volatile. There is however a characteristic behaviour in terms of upward spikes, which seem to occur primarily during winter (December to February). This is due to the fact that in cold seasons severe weather conditions are more likely in Scandinavia, causing extreme load fluctuations, which in combination with generating outages (especially of large generating plants) or transmission failures cause short-lived price "jumps". Clearly, the pattern in electricity prices also reflects changes in supply and demand. For instance, in February 1994 spot prices reached very high levels, due to the high demand during the Olympic Winter Games in Lillhammer. Similarly, the higher level of prices during recent years reflects lower water levels in the reservoirs. Looking at the price differences, it seems that they vary with time forming clusters, which is an indication of time-varying volatility. Moreover, the log-returns in Panel B show that volatility is higher at low price levels, which is an indication of seasonality. However care has to be taken since this might be a result of the log transformation of prices as we explain in more detail later. Also log-returns show very high percentage fluctuations, being as high as 100% (in absolute terms), on some dates.

Table 3.2: Correlation Coefficient between Hourly Time Series

The table bellow shows the linear correlation coefficient between hourly electricity prices (in Levels) for any two hours out of the 24 in a particular day. The mean, maximum and minimum values are also given below the table. The estimation period is from 1st of January 1993 to 29th of February 2004.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23
1																							
2	0.997																						
3	0.991	0.993																					
4	0.992	0.996	0.995																				
5	0.989	0.993	0.992	0.999																			
6	0.990	0.992	0.989	0.995	0.997																		
7	0.986	0.985	0.981	0.986	0.987	0.994																	
8	0.938	0.935	0.931	0.936	0.937	0.946	0.961																
9	0.816	0.814	0.809	0.813	0.815	0.824	0.843	0.950															
10	0.895	0.891	0.885	0.888	0.888	0.897	0.914	0.982	0.980														
11	0.940	0.935	0.928	0.930	0.929	0.938	0.954	0.988	0.948	0.988													
12	0.967	0.962	0.954	0.956	0.954	0.963	0.976	0.976	0.896	0.957	0.988												
13	0.980	0.974	0.967	0.968	0.966	0.975	0.986	0.967	0.869	0.939	0.976	0.995											
14	0.982	0.976	0.969	0.971	0.969	0.978	0.988	0.964	0.860	0.932	0.971	0.992	0.999										
15	0.983	0.978	0.971	0.972	0.971	0.980	0.990	0.961	0.853	0.926	0.966	0.990	0.998	0.999									
16	0.981	0.976	0.969	0.971	0.970	0.978	0.987	0.960	0.853	0.924	0.963	0.986	0.995	0.996	0.998								
17	0.959	0.955	0.948	0.950	0.950	0.957	0.965	0.956	0.869	0.930	0.961	0.978	0.978	0.978	0.980	0.988							
18	0.921	0.917	0.912	0.914	0.913	0.919	0.930	0.955	0.906	0.946	0.959	0.959	0.949	0.947	0.947	0.955	0.980						
19	0.954	0.950	0.944	0.946	0.945	0.951	0.960	0.969	0.907	0.954	0.976	0.980	0.974	0.972	0.970	0.974	0.984	0.979					
20	0.979	0.975	0.969	0.971	0.969	0.975	0.983	0.965	0.872	0.935	0.970	0.988	0.989	0.989	0.988	0.988	0.982	0.960	0.987				
21	0.990	0.986	0.980	0.982	0.980	0.985	0.989	0.951	0.837	0.913	0.955	0.980	0.990	0.992	0.993	0.991	0.971	0.937	0.968	0.992			
22	0.992	0.987	0.981	0.982	0.980	0.985	0.988	0.947	0.829	0.907	0.952	0.978	0.990	0.991	0.992	0.990	0.968	0.932	0.964	0.988	0.999		
23	0.994	0.989	0.982	0.983	0.980	0.984	0.987	0.941	0.821	0.901	0.947	0.975	0.988	0.989	0.990	0.987	0.964	0.926	0.959	0.985	0.997	0.998	
24	0.996	0.993	0.987	0.988	0.985	0.987	0.986	0.938	0.817	0.896	0.942	0.970	0.983	0.985	0.986	0.984	0.961	0.923	0.956	0.982	0.994	0.996	0.998

Max. Correlation	0.999
Min. Correlation	0.81
Mean Correlation	0.96

Figure 3.5: Time Series

These figures show the time series of the spot prices in Nord Pool from 1st of January 1993 to 29th of February 2004. Panel A displays the daily average system price in Nord Pool, as well as the daily price changes, both measured in NOK/MWh. Panel B displays the logarithm of the daily average system price as well as the log returns.

(A): Daily Prices and Price Changes (NOK / MWH)

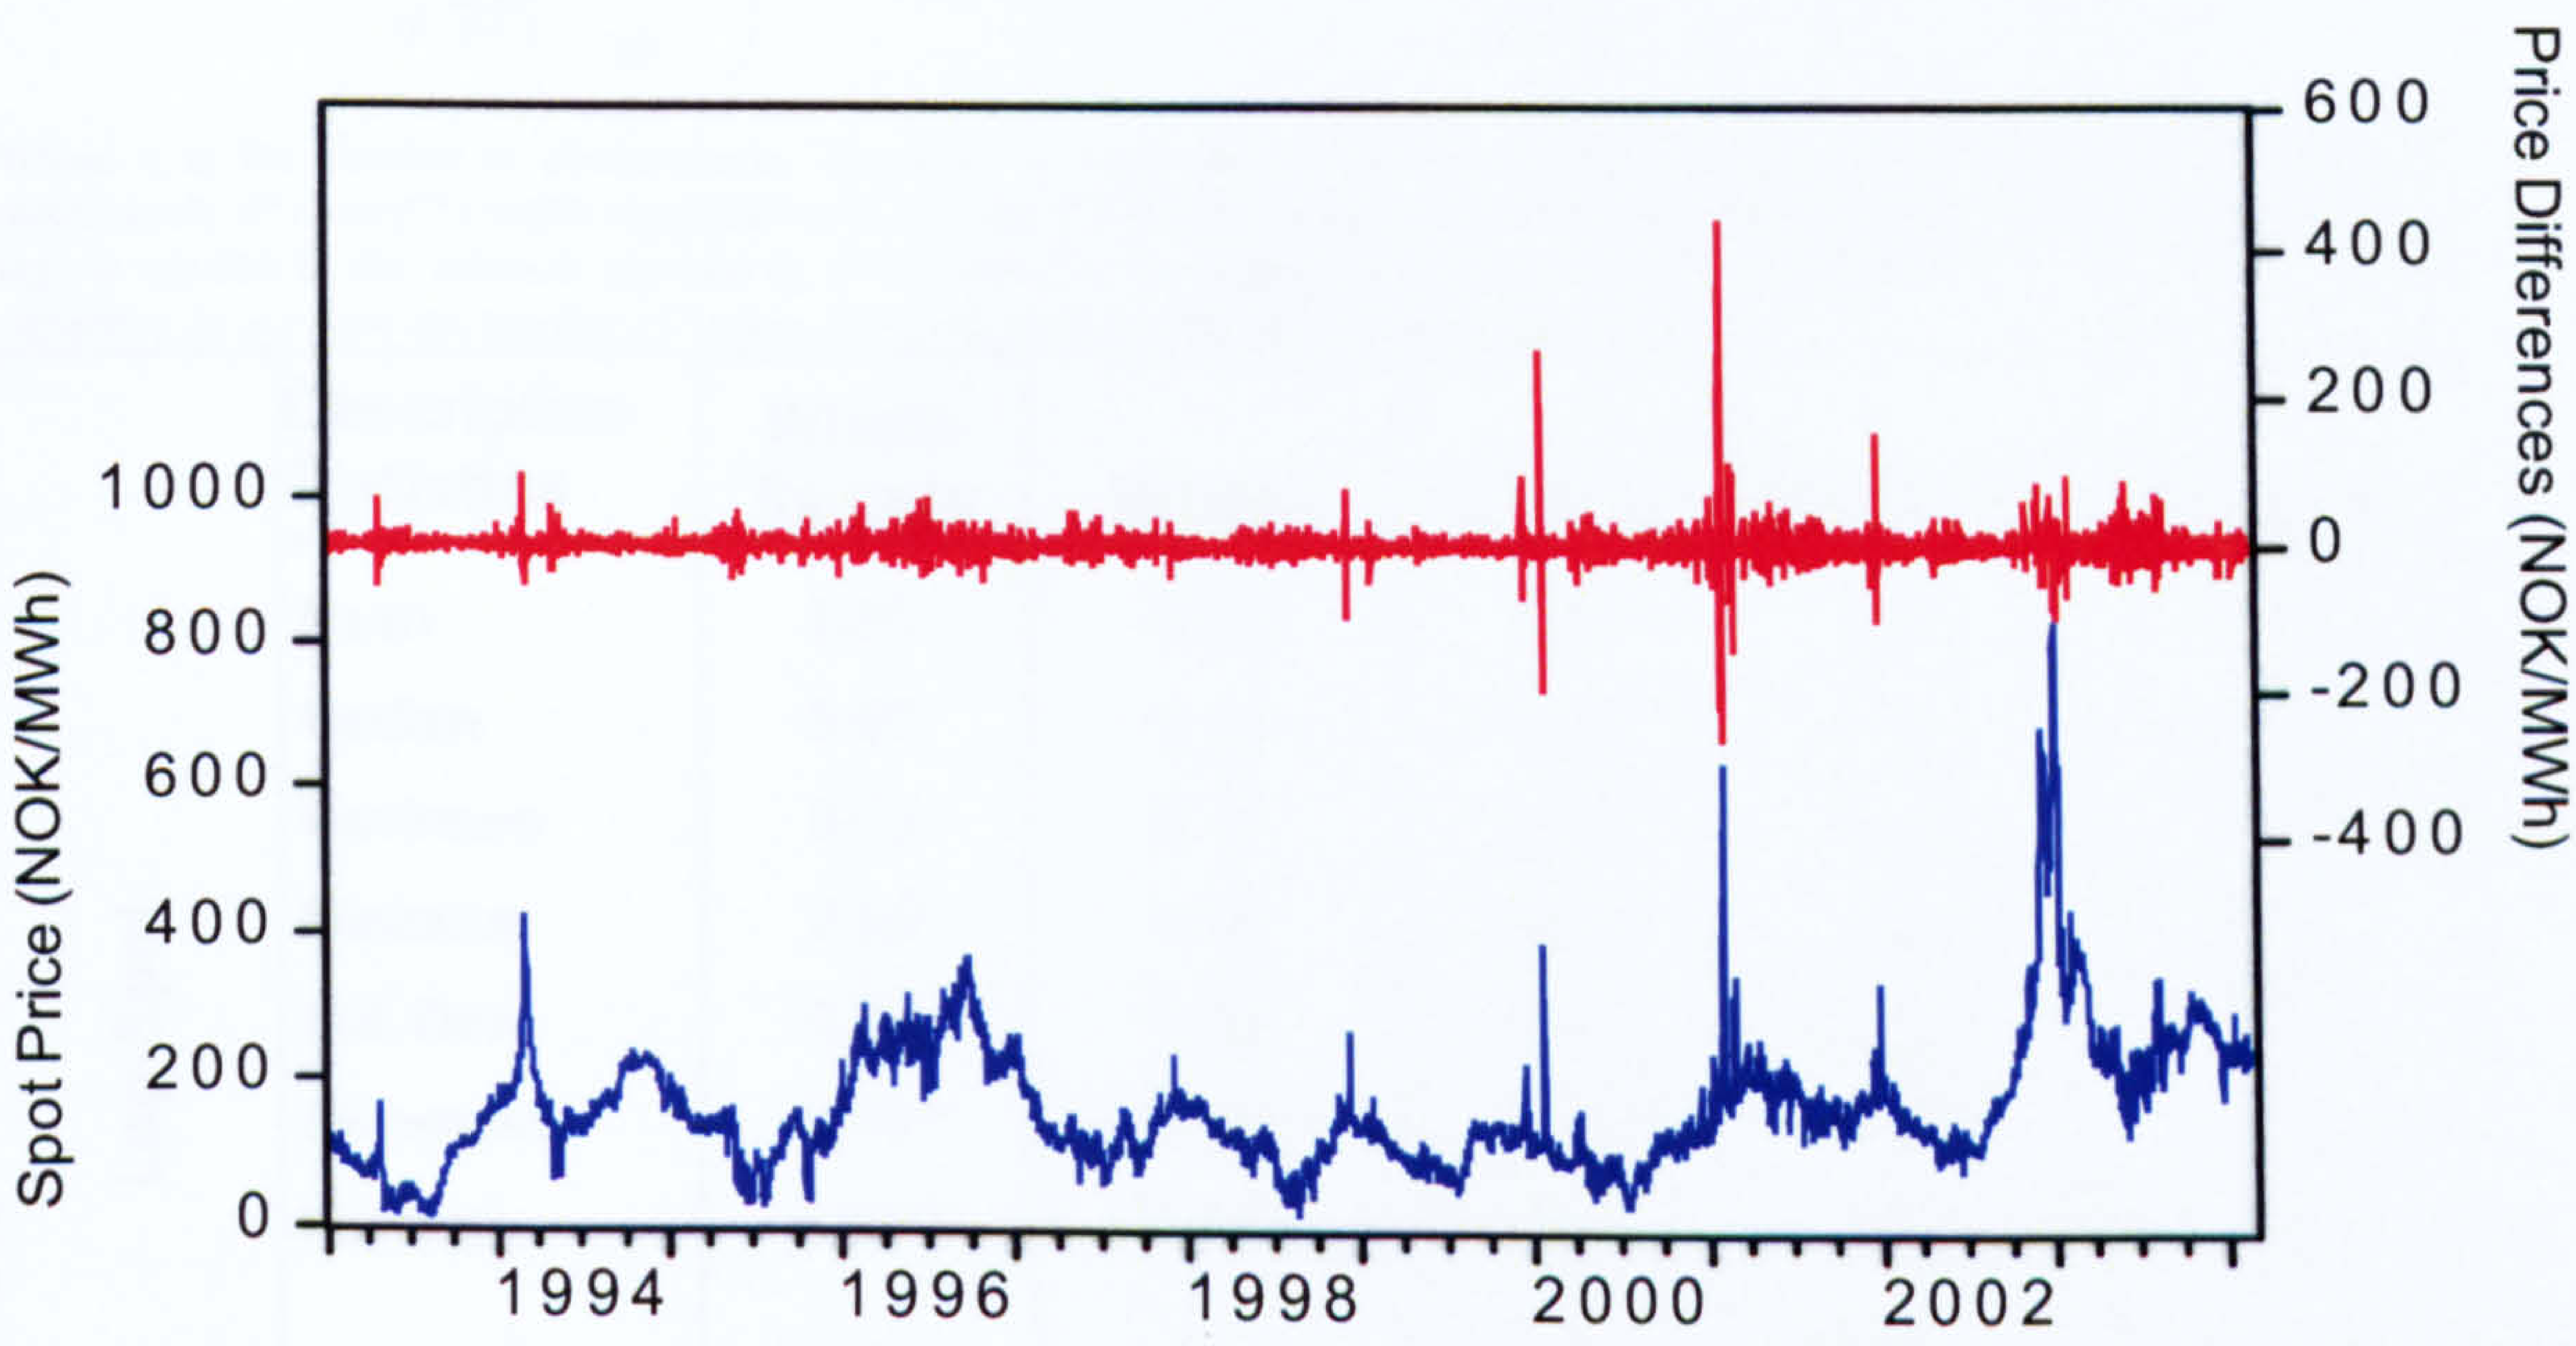

(B): Daily Log Prices and Log Returns

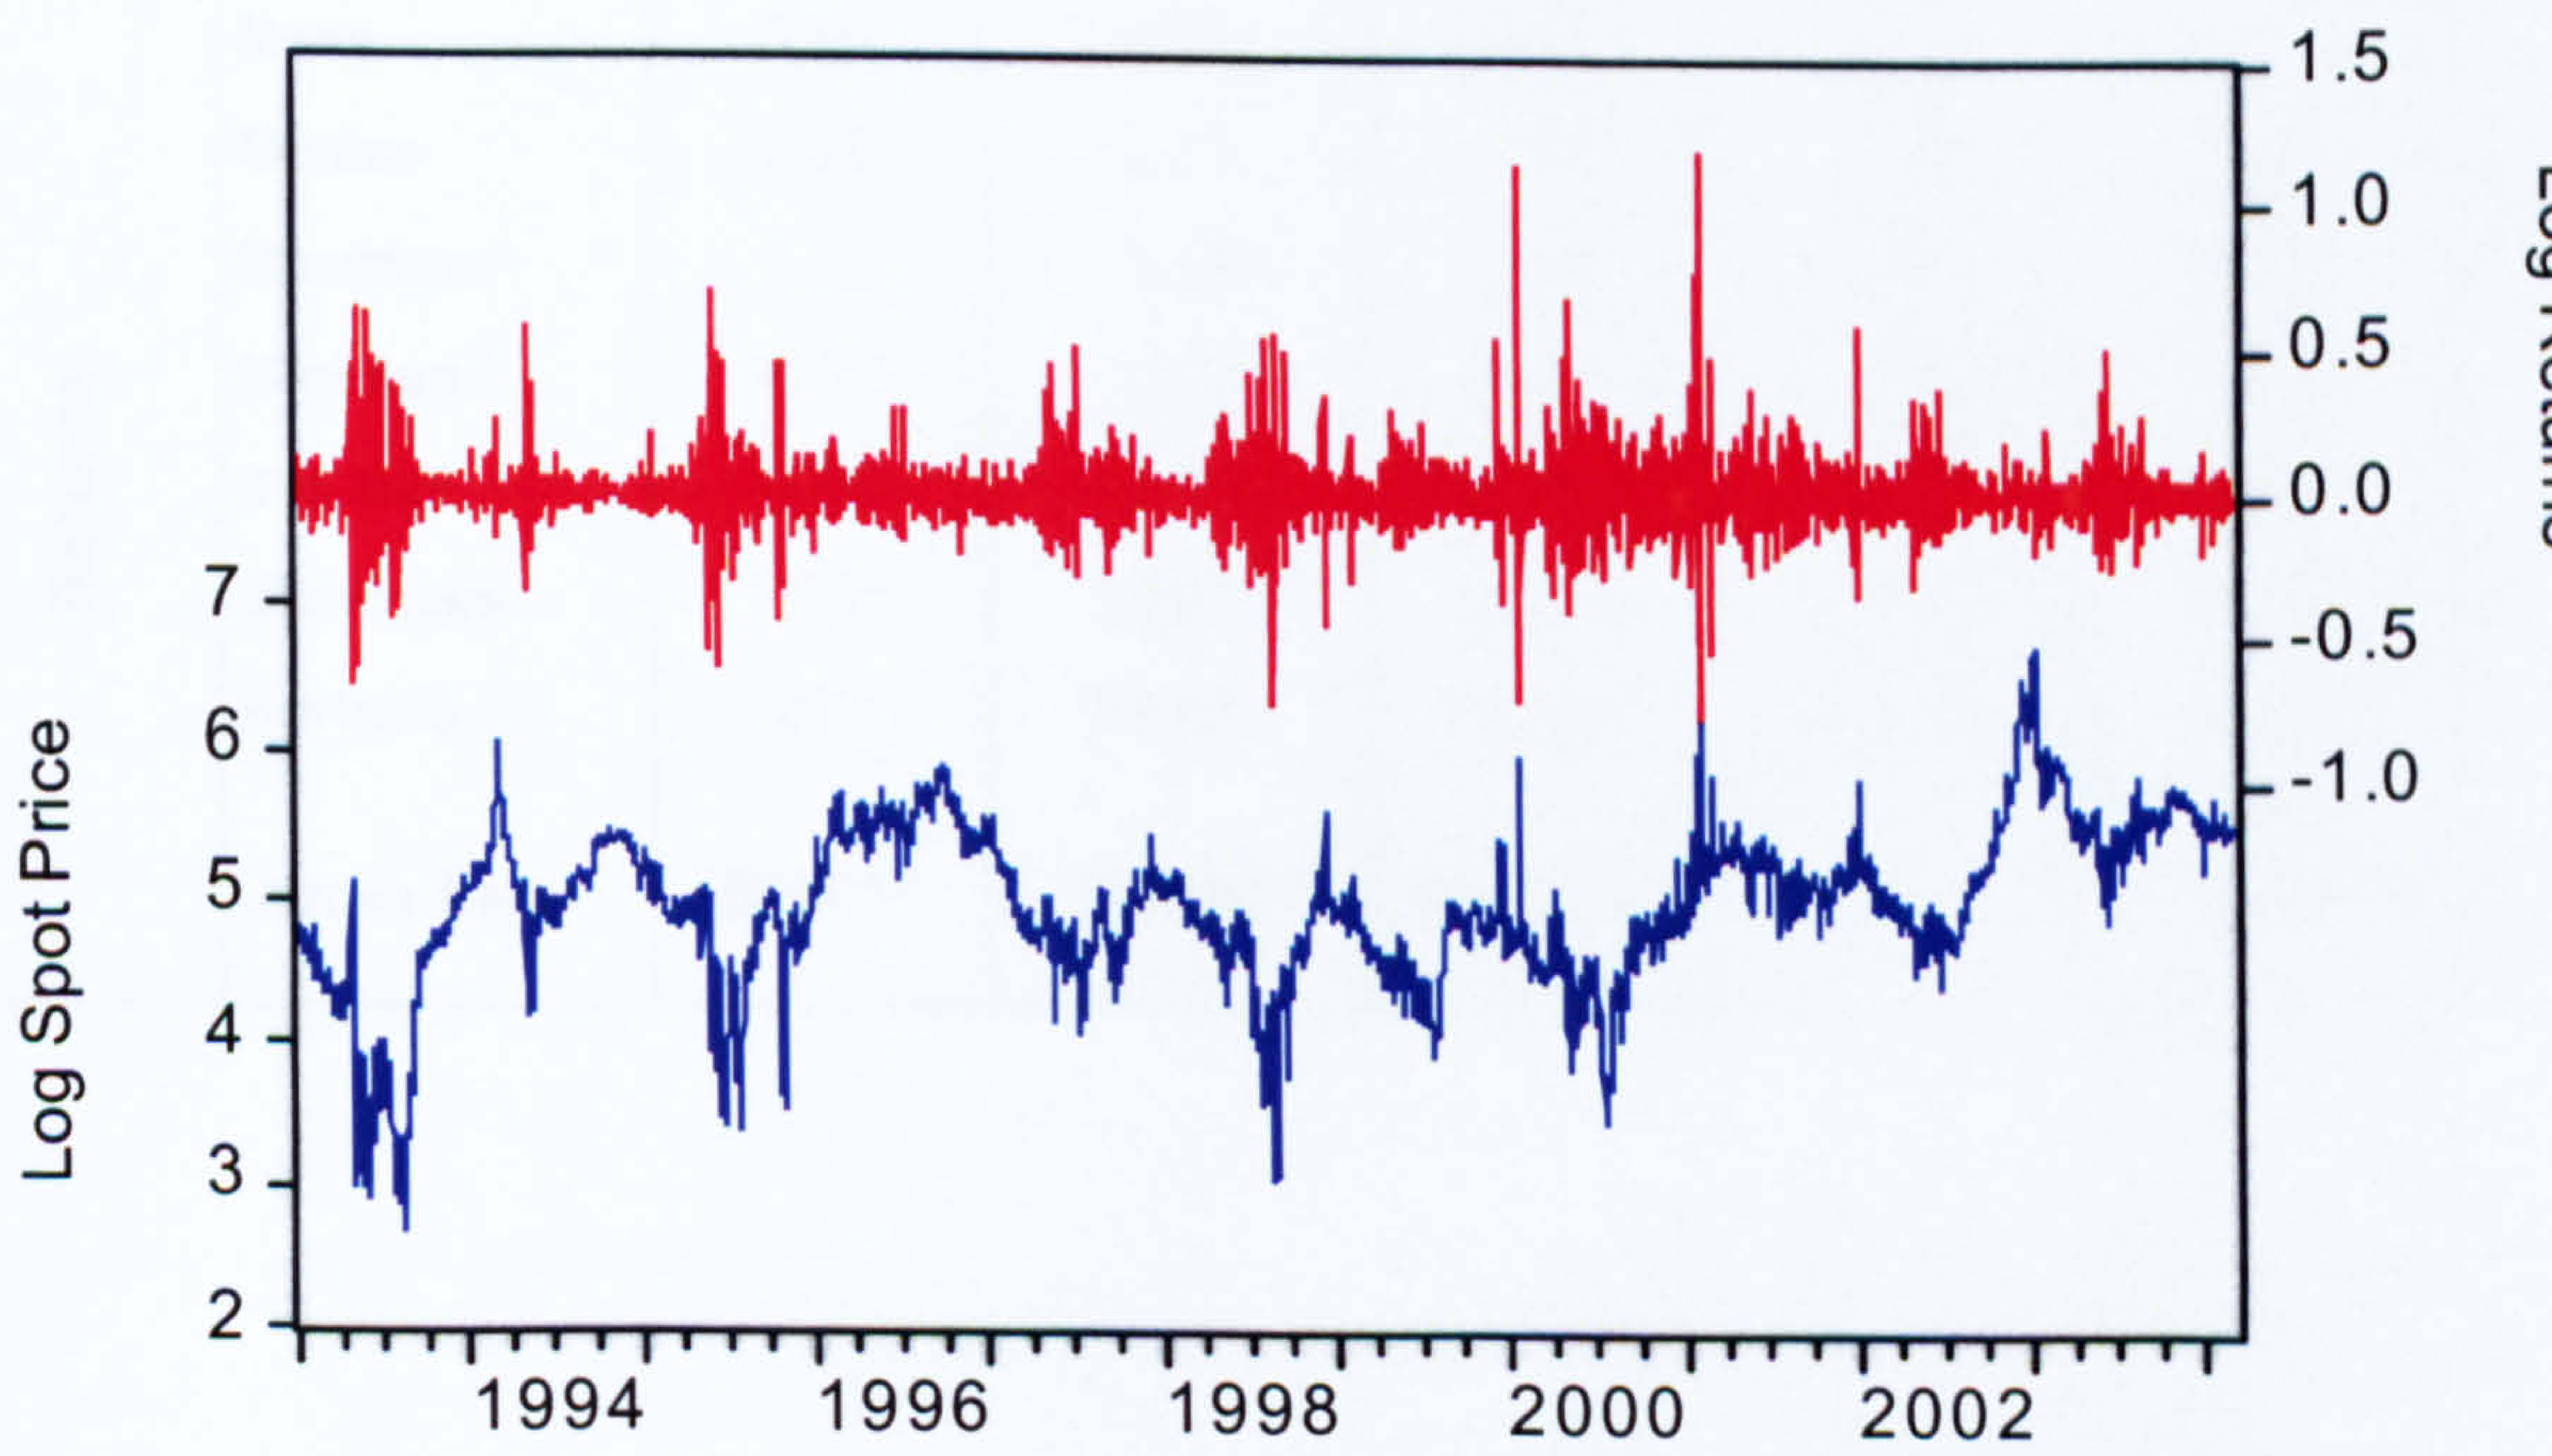

Table 3.3: Descriptive Statistics of Nord Pool Prices

The table displays the descriptive statistics and the Unit Root tests (see e.g. Green, 1993 and Hamilton, 1994) of the average daily System Price in log-levels and log-differences (returns). The data sample is from January 1, 1993 to February 29, 2004. The first column is for the whole sample, and the remaining divides the data into the four seasons. The definition for each season is as follows: Winter (December, January and February), Spring (March, April and May), Summer (June, July and August), and Autumn or Fall (September, October and November). For the estimation and hypothesis testing on the coefficients of Skewness (*S*) and Kurtosis (*K*) the following formulas are used respectively:

$$S = \frac{1}{n}\sum_{t=1}^{n}\left(\frac{r_t - \bar{r}}{\hat{\sigma}}\right)^3 \sim N(0,6/n) \qquad K = \frac{1}{n}\sum_{t=1}^{n}\left(\frac{r_t - \bar{r}}{\hat{\sigma}}\right)^4 \sim N(0,24/n)$$

Where n is the number of observations. The null for each test is that the population Skewness and Kurtosis are 0 and 3, respectively. (*) and (**) imply significance at 5% and 1% levels, respectively. Note also that for the ADF test, a large number of lags is needed in the relevant regression, to account for the serial correlation that is present in the changes of the relevant variables. In our case the number of lags is 21 based on Schwartz Information Criterion.[13]

	Descriptive Statistics	Whole Sample	Winter	Spring	Summer	Autumn
Log-Spot	Mean	5.00	5.20	4.94	4.75	5.10
	Median	5.00	5.15	4.90	4.81	5.03
	Maximum	6.72	6.72	6.04	5.84	5.93
	Minimum	2.69	4.35	2.89	2.69	3.56
	Std. Dev.	0.54	0.40	0.50	0.66	0.41
	Skewness	-0.70**	1.06**	-0.76**	-0.54**	-0.21**
	Kurtosis	4.47**	4.86**	4.43**	2.67*	2.84
	Jarque-Bera	758 **	356**	200**	58.81**	9.09*
	ADF	-3.37*				
	PP	-4.73**				
Returns	Mean	0.00	0.00	0.00	0.00	0.00
	Median	0.00	-0.01	-0.01	0.00	0.00
	Maximum	1.19	1.19	0.69	0.70	0.50
	Minimum	-0.77	-0.77	-0.67	-0.72	-0.43
	Std. Dev.	0.10	0.10	0.11	0.12	0.07
	Skewness	1.06**	2.52**	0.37**	0.71**	0.87**
	Kurtosis	21**	50.68**	12.35**	8.78**	14.05**
	Jarque-Bera	59941**	103448**	4047.86**	1627**	5689**

[13] We also performed the ADF and PP tests for the log-spot series, and the results revealed that unit root is rejected at 5% and 1% significance levels for this series too.

In order to build successful stochastic models for electricity it is essential to understand its distributional characteristics. Modelling methodologies, model testing and acceptance, and model parameters, depend on the choice of the relevant distributions. In financial pricing models, we find that most of them are based on the assumption of normality or log-normality (such as the Arithmetic and Geometric Brownian Motion) of the underlying price distribution. Even a histogram of electricity price returns can give an indication of which model can capture realistically the properties of the empirical distribution. Alternative statistical tests such as the Jarque-Berra (1980), allows us to quantify our intuition and provide a formal foundation for acceptance or rejection of a particular choice of distributions build into the pricing models.

Table 3.3 presents the descriptive statistics of the daily electricity data, in log-levels and log differences. The logarithm of the spot for the whole sample period has a mean value of 5.00, and has reached maximum and minimum values of 6.72 and 2.69, respectively. However, prices seem to be different between warm and cold seasons. The mean value during winter reaches its peak at 5.20 and it smoothly decreases to 4.75 during summer. This difference is caused by fluctuations in residential demand, which increases in winter mainly for heating (as the winter is very cold in the Nordic Countries) and lighting purposes, and then drops, as the summer is milder and the days are longer.

Using the standard volatility measure (standard deviation of log-returns), it is observed that electricity is a highly volatile market. The standard deviation of the daily log-returns is 0.10, which translates into an annualised volatility of 199%[14]. This is very high compared to the volatility in other markets such as LIBOR rates (10% - 20%), or NASDAQ (30% - 50%). Volatility is also different between cold and warm seasons; the annualised volatility of log-returns is 229% during summer and smoothly declines to 191% in winter. One can thus state that there exists some distinguishable seasonal pattern in volatility. Different theories have been suggested for that, such as the fact that during warmer periods, water reservoirs are almost empty, thus the equilibrium moves to the steep part of the supply stack. On the other

[14] Annualised volatility = Daily Standard Deviation * $\sqrt{365}$

hand, one can also argue that during warmer periods, demand volatility in Scandinavia is much lower; therefore price volatility should be lower. However, we see that volatility is negatively correlated to prices, meaning that at lower price levels volatility is higher than at high price levels. This is caused by the fact that the slope of the logarithmic function is a decreasing function of price; therefore it makes price differences to be more volatile at low-price levels than at higher price levels, a fact that has been discussed by Lucia and Schwartz (2002). As argued by Simonsen et al (2004), for the notion of log-returns to make sense as a measure of gain and loss, the spot prices have to possess no seasonality, an assumption that is clearly rejected in our data set. Thus in order to interpret returns and their volatilities consistently across seasons, one has to first pre-process the spot prices in order to remove any apparent periodicity. We thus follow an approach similar to Weron et al (2005), by firstly deseasonalising the price level data (using a simple sinusoidal function, discussed in Chapter 4) and then taking the logarithm of the residual terms. Using this method, the annualised volatility is found to be 224%, 169%, 162% and 157% during the periods of winter, spring, summer and autumn, respectively. Thus the distinguishable pattern we find in volatility is that it is significantly higher during winter. This is a consequence of the winter peak demand, bringing the equilibrium at the steep part of the supply stack. Therefore any outage or transmission failure can cause more severe price movements in winter.

The kurtosis coefficient gives an indication of the probability of extreme values in a series. In Nord Pool, the estimate of kurtosis for the logarithmic spot price series is 4.47, implying that log-prices are leptokurtic. This means that the probability of high or low prices is much higher in the market, than that dictated by the normal distribution. Leptokurtosis is mainly evident in the winter and spring but not so much in warm periods. However are these extreme values due to abnormal larger variations in prices called jumps? The parameter of kurtosis (21) of log-returns is significantly larger than that of a Normal distribution (3), which implies that daily variations are indeed relative more frequent than what a Normally distributed variable would capture. The case is far more extreme in winter than during the rest of the seasons, and this can be explained by the fact that during winter, the occurrence of jumps is high.

Moreover the coefficient of skewness measures the degree of symmetry of the distribution and shows whether high prices are indeed more probable than low values. The positive sign of the skewness coefficient indicates that extreme positive returns are more likely than negative, especially during winter. From these stylised facts it can be concluded that the empirical distribution of electricity in Nord Pool is not normal. The Jarque-Berra statistic (1980) is based on the empirical kurtosis and skewness ($JB = n\left[\frac{skewness^2}{6} + \frac{(kurtosis-3)^2}{24}\right]$) $\sim \chi^2(2)$ and is used to test the hypothesis of normality in the distribution of electricity prices. In all series normality is overwhelmingly rejected. From the histograms on log returns (Figure 3.6) and the price history (Figure 3.5) one can reveal that non-normality occurs due to the large price movements/spikes yielding exceptional fat tails. Overall, it can be argued that extreme price movements (spikes) mainly occur at peak demand periods and are predominantly positive, reflecting outages causing very short-lived supply disruptions.

Figure 3.6: Histogram of Log-returns

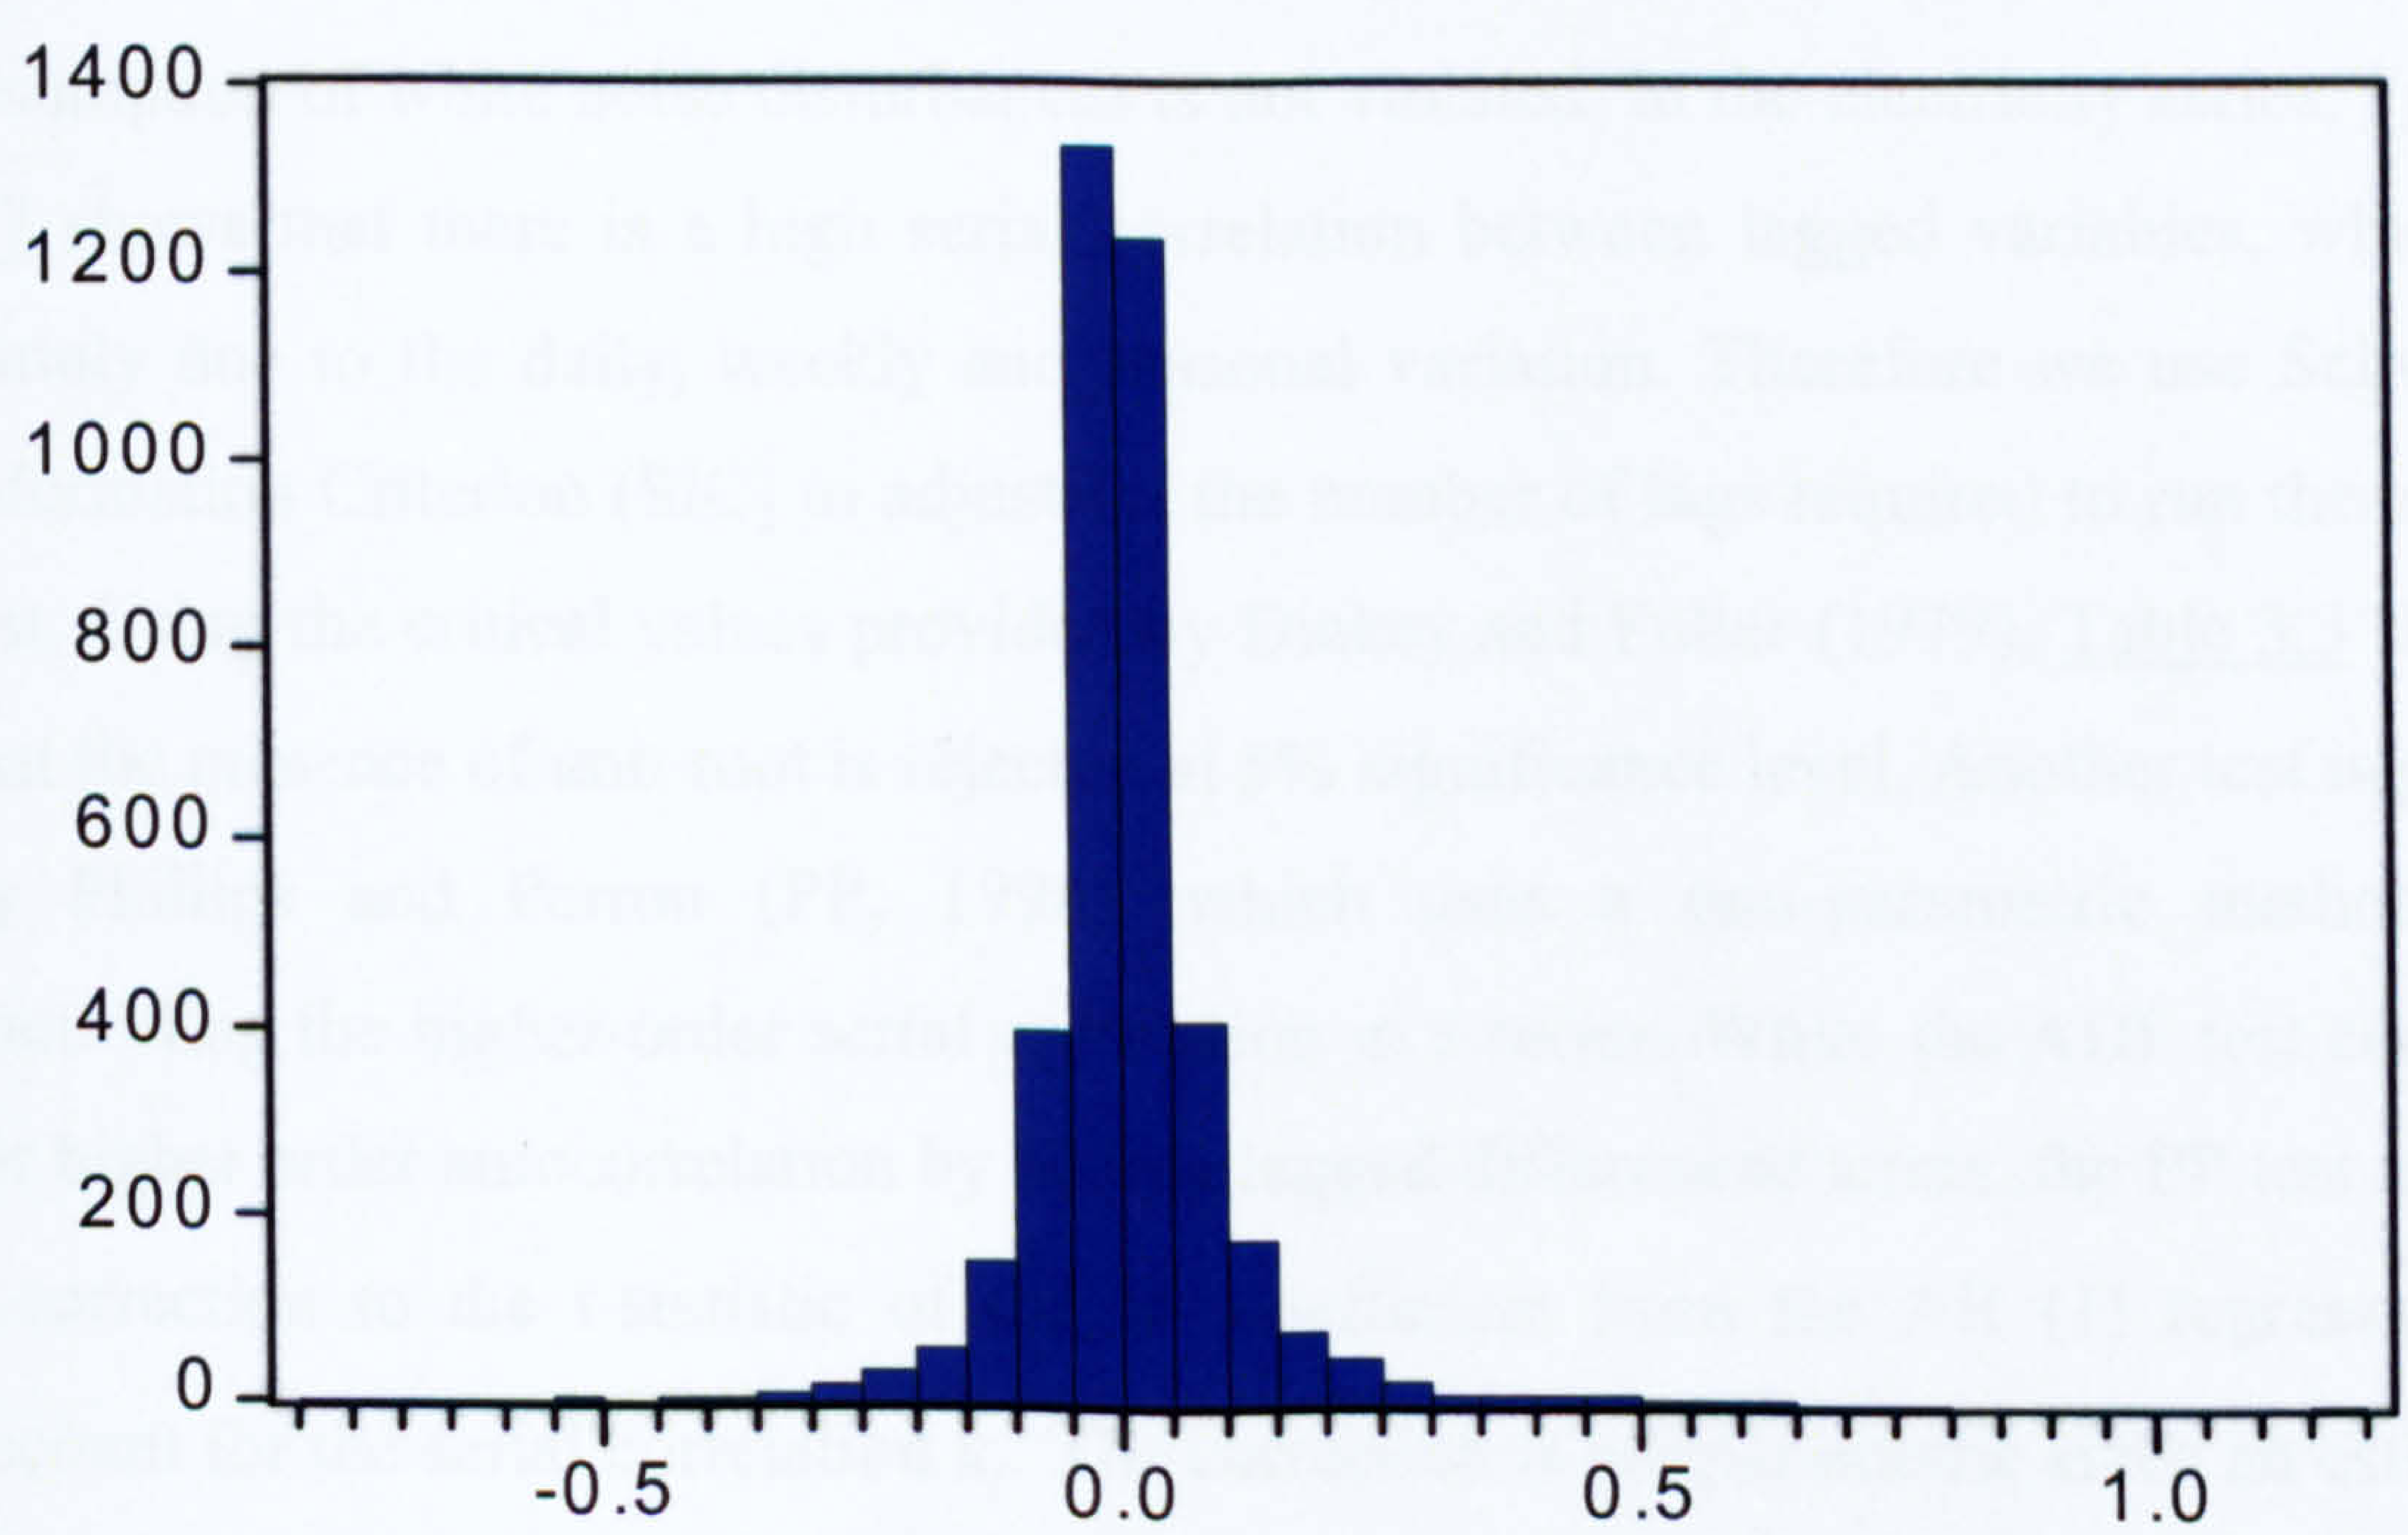

To carry on the analysis we have to test whether electricity prices are stationary or not. A process is stationary if its distributional properties are time invariant; that is to say, its joint distributions $P_{t_1} \ldots P_{t_n}$ are the same as those for $P_{t_{1+d}} \ldots P_{t_{n+d}}$ for all d.

Stationarity thus implies that the process has no trend, so that prices may either be drift-less or may stay roughly within some range. If we model a non-stationary series with a stationary model, regression results will be misleading and spurious as shown by Granger and Newbold (1974). Non-stationarity implies that an innovation may have a permanent impact in the future. An AR(1) model, $P_t = \theta + \phi P_{t-1} + \varepsilon_t$, gives an indication of whether the series is stationary or not. If the coefficient is not $-1 < \phi < 1$, then the series is non-stationary; that is to say that if the process is started at some point, the variance of P_t increases with time to infinity and explodes. A formal test for unit root in a series is the Augmented Dickey Fuller (ADF, 1979) Test:

$$\Delta P_t = \theta + \phi^* S_{t-1} + \sum_{j=1}^{p-1} \phi_j^* \Delta P_{t-j} + u_t \qquad (3.1)$$
$$\text{where } u_t \sim iid(0, \sigma^2)$$

In equation (3.1) the null hypothesis is $\phi^* = 0$, which is equivalent to say, that the series has a unit root. Moreover, by adding lagged difference terms of the dependent variables, we are able to control for the higher-order correlation and hence the assumption of white noise disturbances is not violated. In the electricity series, Figure 3.7 shows that there is a high serial correlation between lagged variables, which is mainly due to the daily, weekly and seasonal variation. Therefore we use Schwartz Information Criterion (SIC) to adjust for the number of lags required to run the above test. Using the critical values provided by Dickey and Fuller (1979), Table 3.3 shows that the presence of unit root is rejected at 5% significance level. Another test is given by Phillips and Perron (PP, 1998), which uses a non-parametric method for controlling the higher-order serial correlation in a series. While the ADF test corrects for higher order autocorrelation by adding lagged differenced terms, the PP test makes a correction to the t-statistic of the ϕ^* coefficient from the AR (1) regression to account for the serial correlation u_t. The correction is nonparametric since an estimate of the spectrum of u_t at frequency zero is used that is robust to heteroskedasticity and autocorrelation of unknown form. The PP test in Table 3.3 gives overwhelming evidence that the electricity price series does not contain a unit root. Therefore these tests also give statistical evidence to support the theory that electricity prices in Nord Pool follow a mean-reverting process. This result is surprising, because we have jumps in the series. However these jumps do not seem to affect so much the stationarity of the series because they are mean-reverting, and hence the shocks generated by these large price movements do not have a persistent effect on the price process.

Figure 3.7: Autocorrelation between time lags of the series

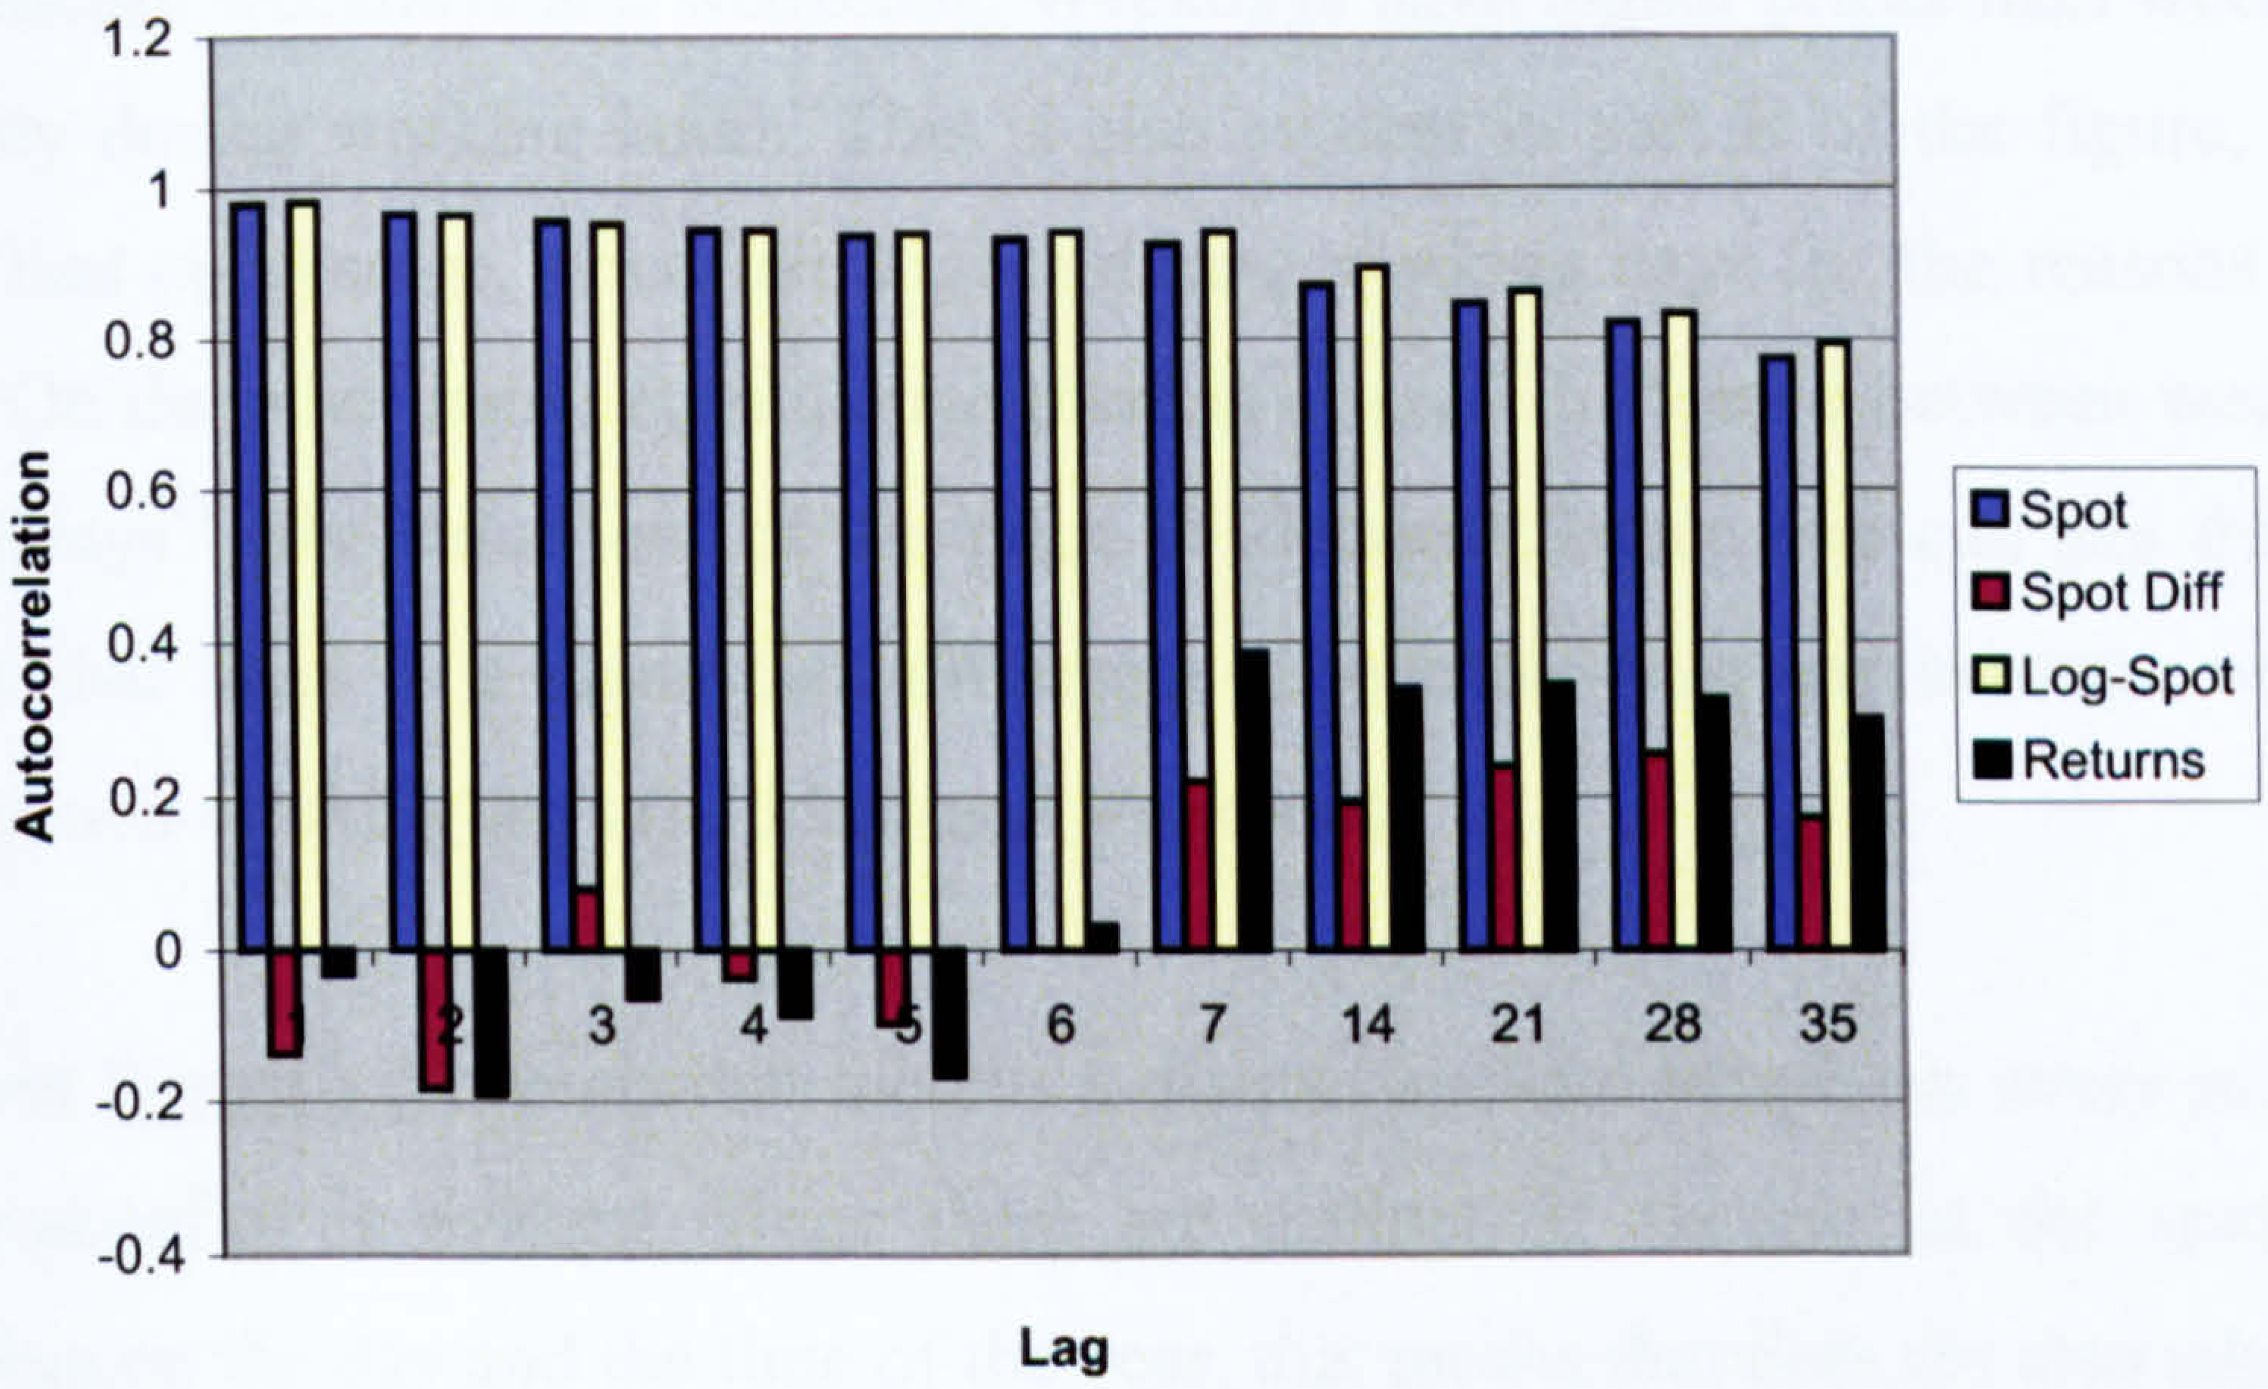

The above analysis on the descriptive statistics for the different seasons, and the autocorrelation between time lags, indicates that prices follow some regular distinct pattern. The correlogram above shows that the first differences in terms of prices and log-prices, are significantly correlated at several time lags, which are multiples of seven. For example, the autocorrelation at lag seven of the daily change of price (log-price) is 22.10% (38.6%) which implies that if we regress the first-differences of the prices (log-prices) on a constant and its seventh lag the R^2 will be almost 5%(15%)(i.e. the square of the autocorrelation coefficient). Hence 5% (15%) of the variation in the daily price (log-price) increments is predictable using the daily increments seven days apart. The same can be done for several consecutive weeks in the future. This is due to the fact that demand for electricity varies following a noticeable regular pattern within the week.

Given the above analysis, it is also interesting to examine how human activity and hence regularities in demand, affect the spot price levels. Figure 3.8(A) shows the average price for the whole sample for every one of the 24 hours in every day of the week. It is clearly observed that prices are higher during working hours and reach their low levels during non-working hours. In particular there are two peaks during a day; one in the morning when people go to work (at 9 a.m.), and one in the afternoon (between 5 and 6 p.m.) when they come back from work and start preparing dinner, watch TV etc. Moreover there is a clear difference in terms of the mean and shape level between weekdays and weekends. Weekdays have higher prices than weekends, especially during working hours. This is also evident in part B of the figure, which reveals that on average, prices are higher during working days for the reasons stated before. On the other hand, even though there is a small difference between weekends and holidays[15], the behaviour of the price is identical; hence one can say from the analysis that there is a significant difference of electricity prices between working days and non-working days (holidays and weekends).

Moreover Figure 3.5 reveals that there is a distinct cyclical behaviour every year, and thus seasonality is evident. Since there are distinctive features in the spot price depending on the day and the time of the year, this means that they are also taken into account for the valuation of derivatives. One would thus expect futures prices for warm seasons, to be lower than futures prices for cold seasons. To provide evidence for the above fact, Figure 3.9 plots a sample of the Forward Curve for March 27, 2000. It is clear from the shape of the curve that the term structure displays a cyclical pattern during a given year. This indicates, that market participants in the valuation of derivatives take into account the seasonality observed in the spot market. Hence there is always a seasonal deterministic component when valuing derivatives in electricity.

[15] Official Public Holidays available from the ministries of Foreign Affairs of Norway and Sweden. Most holidays are the same for all Nordic counties. Also more than 75% of spot trading volume is represented by Norway and Sweden (see Nord Pool reports).

Figure 3.8: Hourly Average Prices

The graphs present the average behaviour for each hour of the day during the period January 1993 to February 2004, across the different days of the Week (A) and between Working and Non-Working days (B).

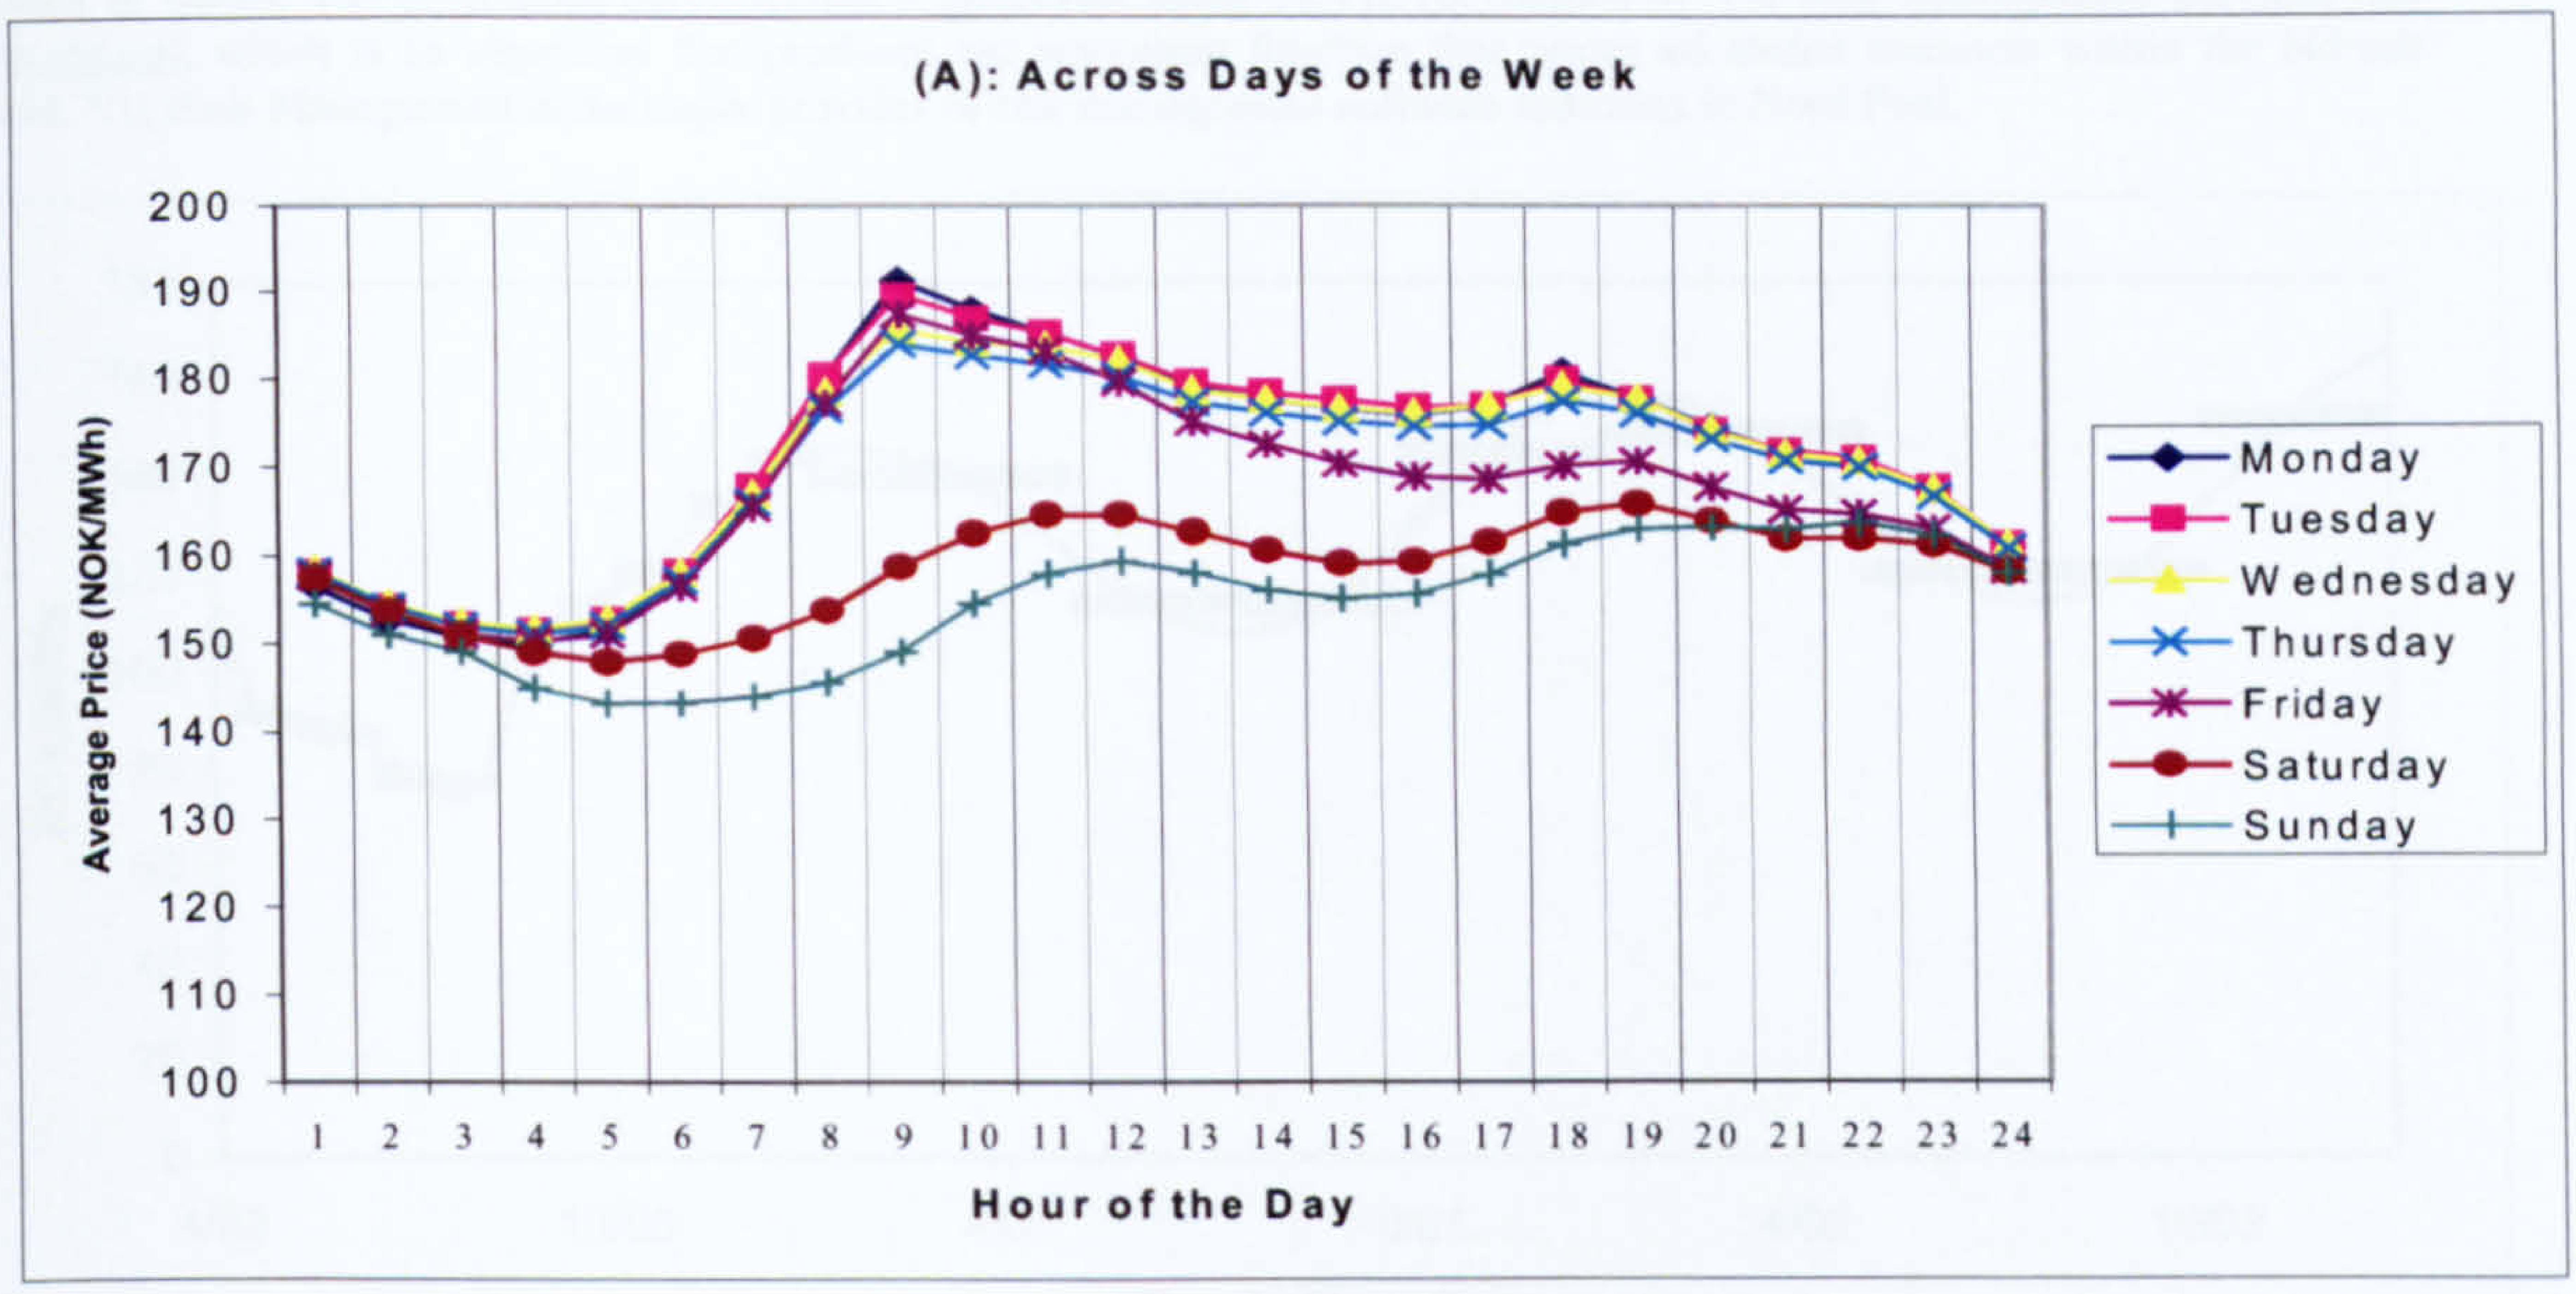

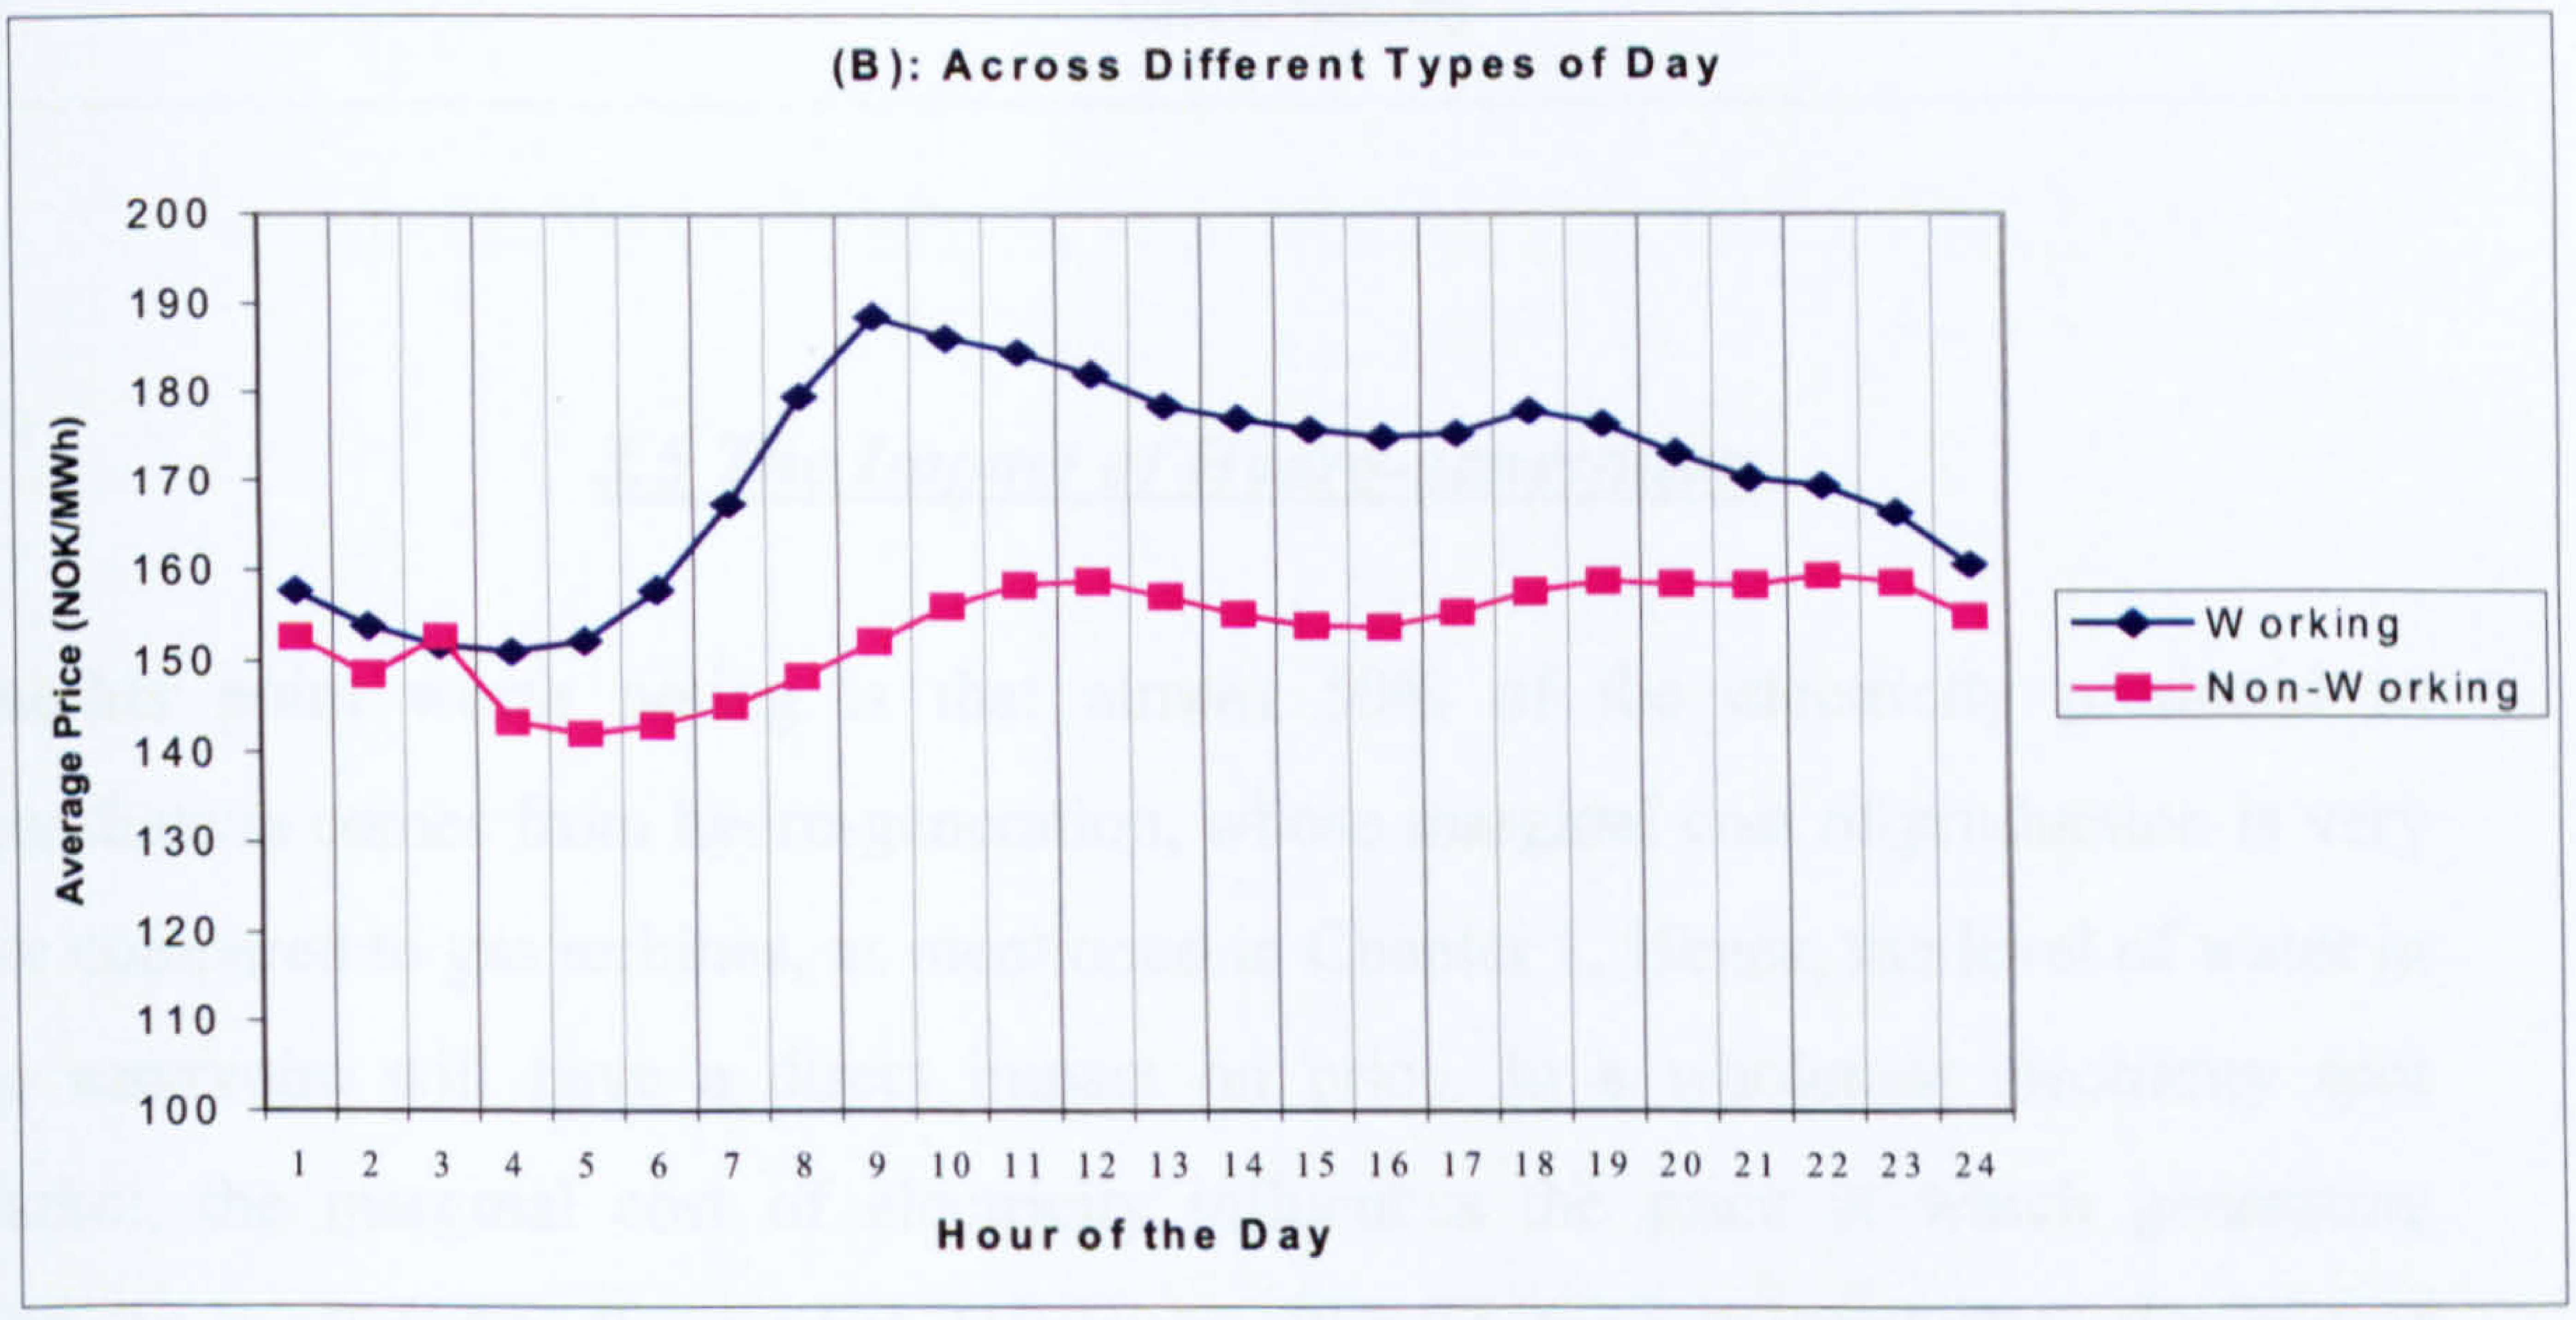

Figure 3.9: Term Structure in Nord Pool on 27/3/2000

The graph shows the seasonality pattern in the forward market. The contracts used are Week, Block Futures and Forward Season contracts listed on that particular date. The flat section (in dots) represents the actual price, during the delivery period of a contract in weeks. The continuous curve section is generated using ELVIZ (developed by Viz Risk Management Services AS, www.viz.no), which is an algorithm that produces the smoothest function that prices all traded contracts within the bid/ask spread. Viz Risk Management is the major provider of risk management software solutions in Nord Pool.

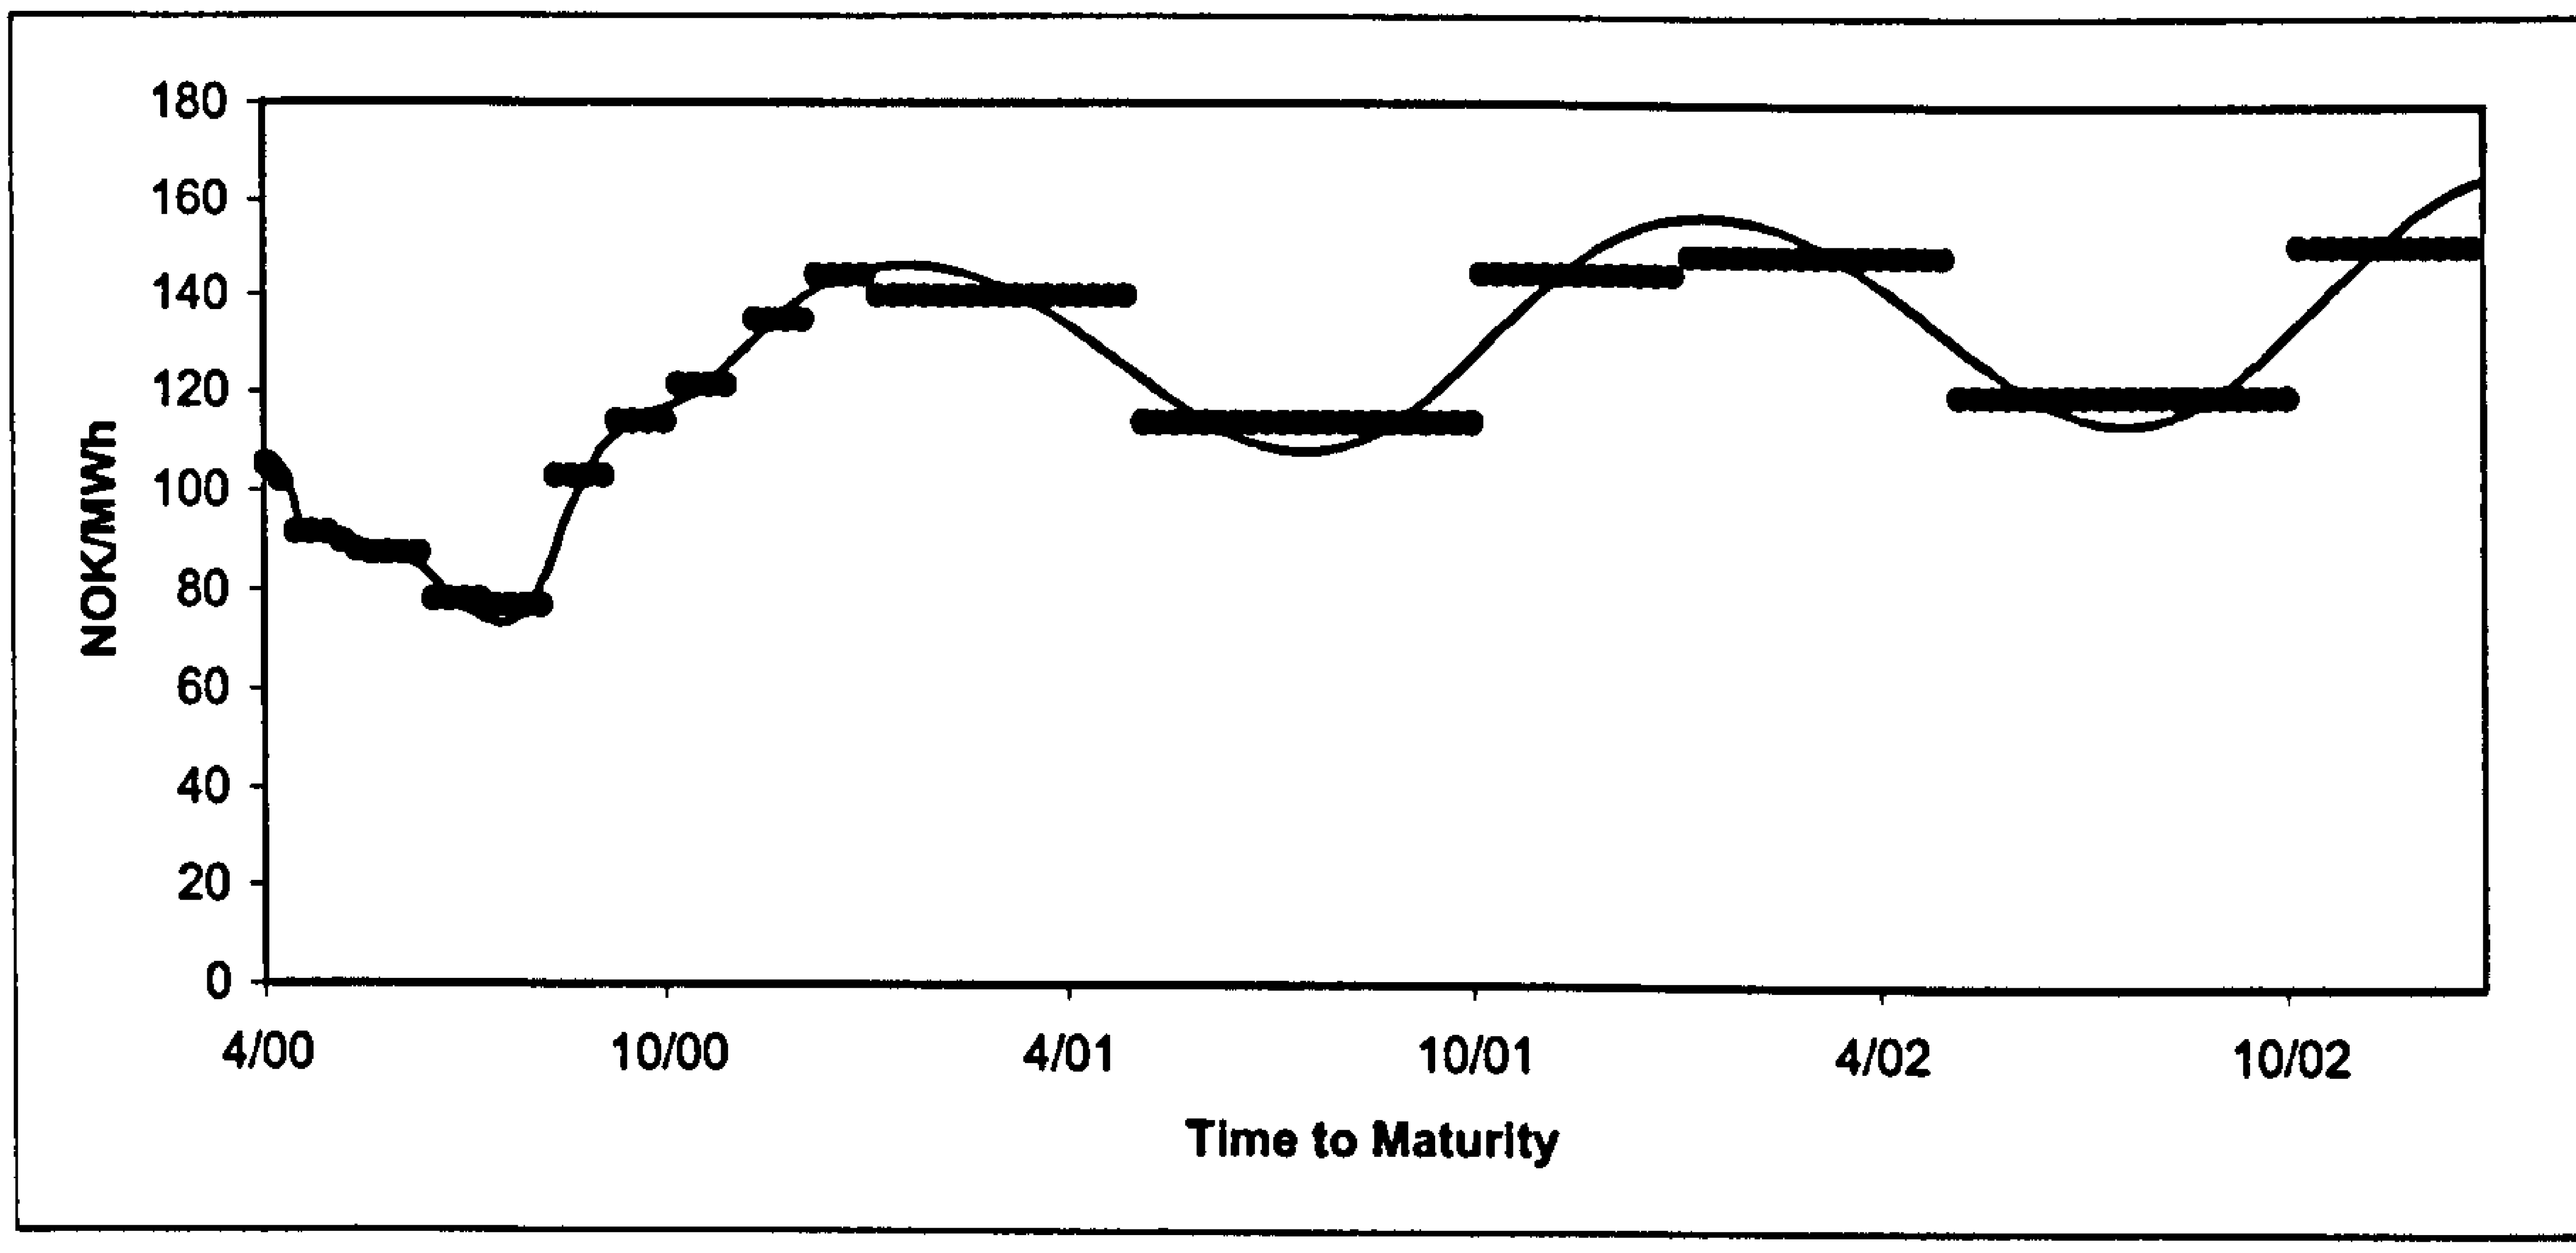

3.5 The Impact of Hydro-generation

Another point worth noting is that almost 50% of the electricity produced in Scandinavia comes from hydro-generation, whose marginal cost of production is very low compared to gas turbines, as mentioned in Chapter 1. Hence, the level of water in the reservoirs will have a direct impact on price. In a wholesale electricity spot market, the marginal cost of electricity influences the price at which generating capacity is offered to the market. However, although for a thermal plant the cost of generation is relatively clear-cut – the input fuel is purchased and the cost of combustion can be estimated – for a hydro plant, the input "fuel" is water, hence for the efficient dispatch of generation, some assessment is necessary about the (marginal) value of water in storage that the generator can "buy" from the reservoirs as fuel for generation.

Figure 3.10: Actual Reservoir Levels Vs Seasonal Average

The graph shows the time series of weekly observations for the percentage level of water in the reservoirs (weighted average from Sweden and Norway), the water reservoir seasonal average, and the spot prices.

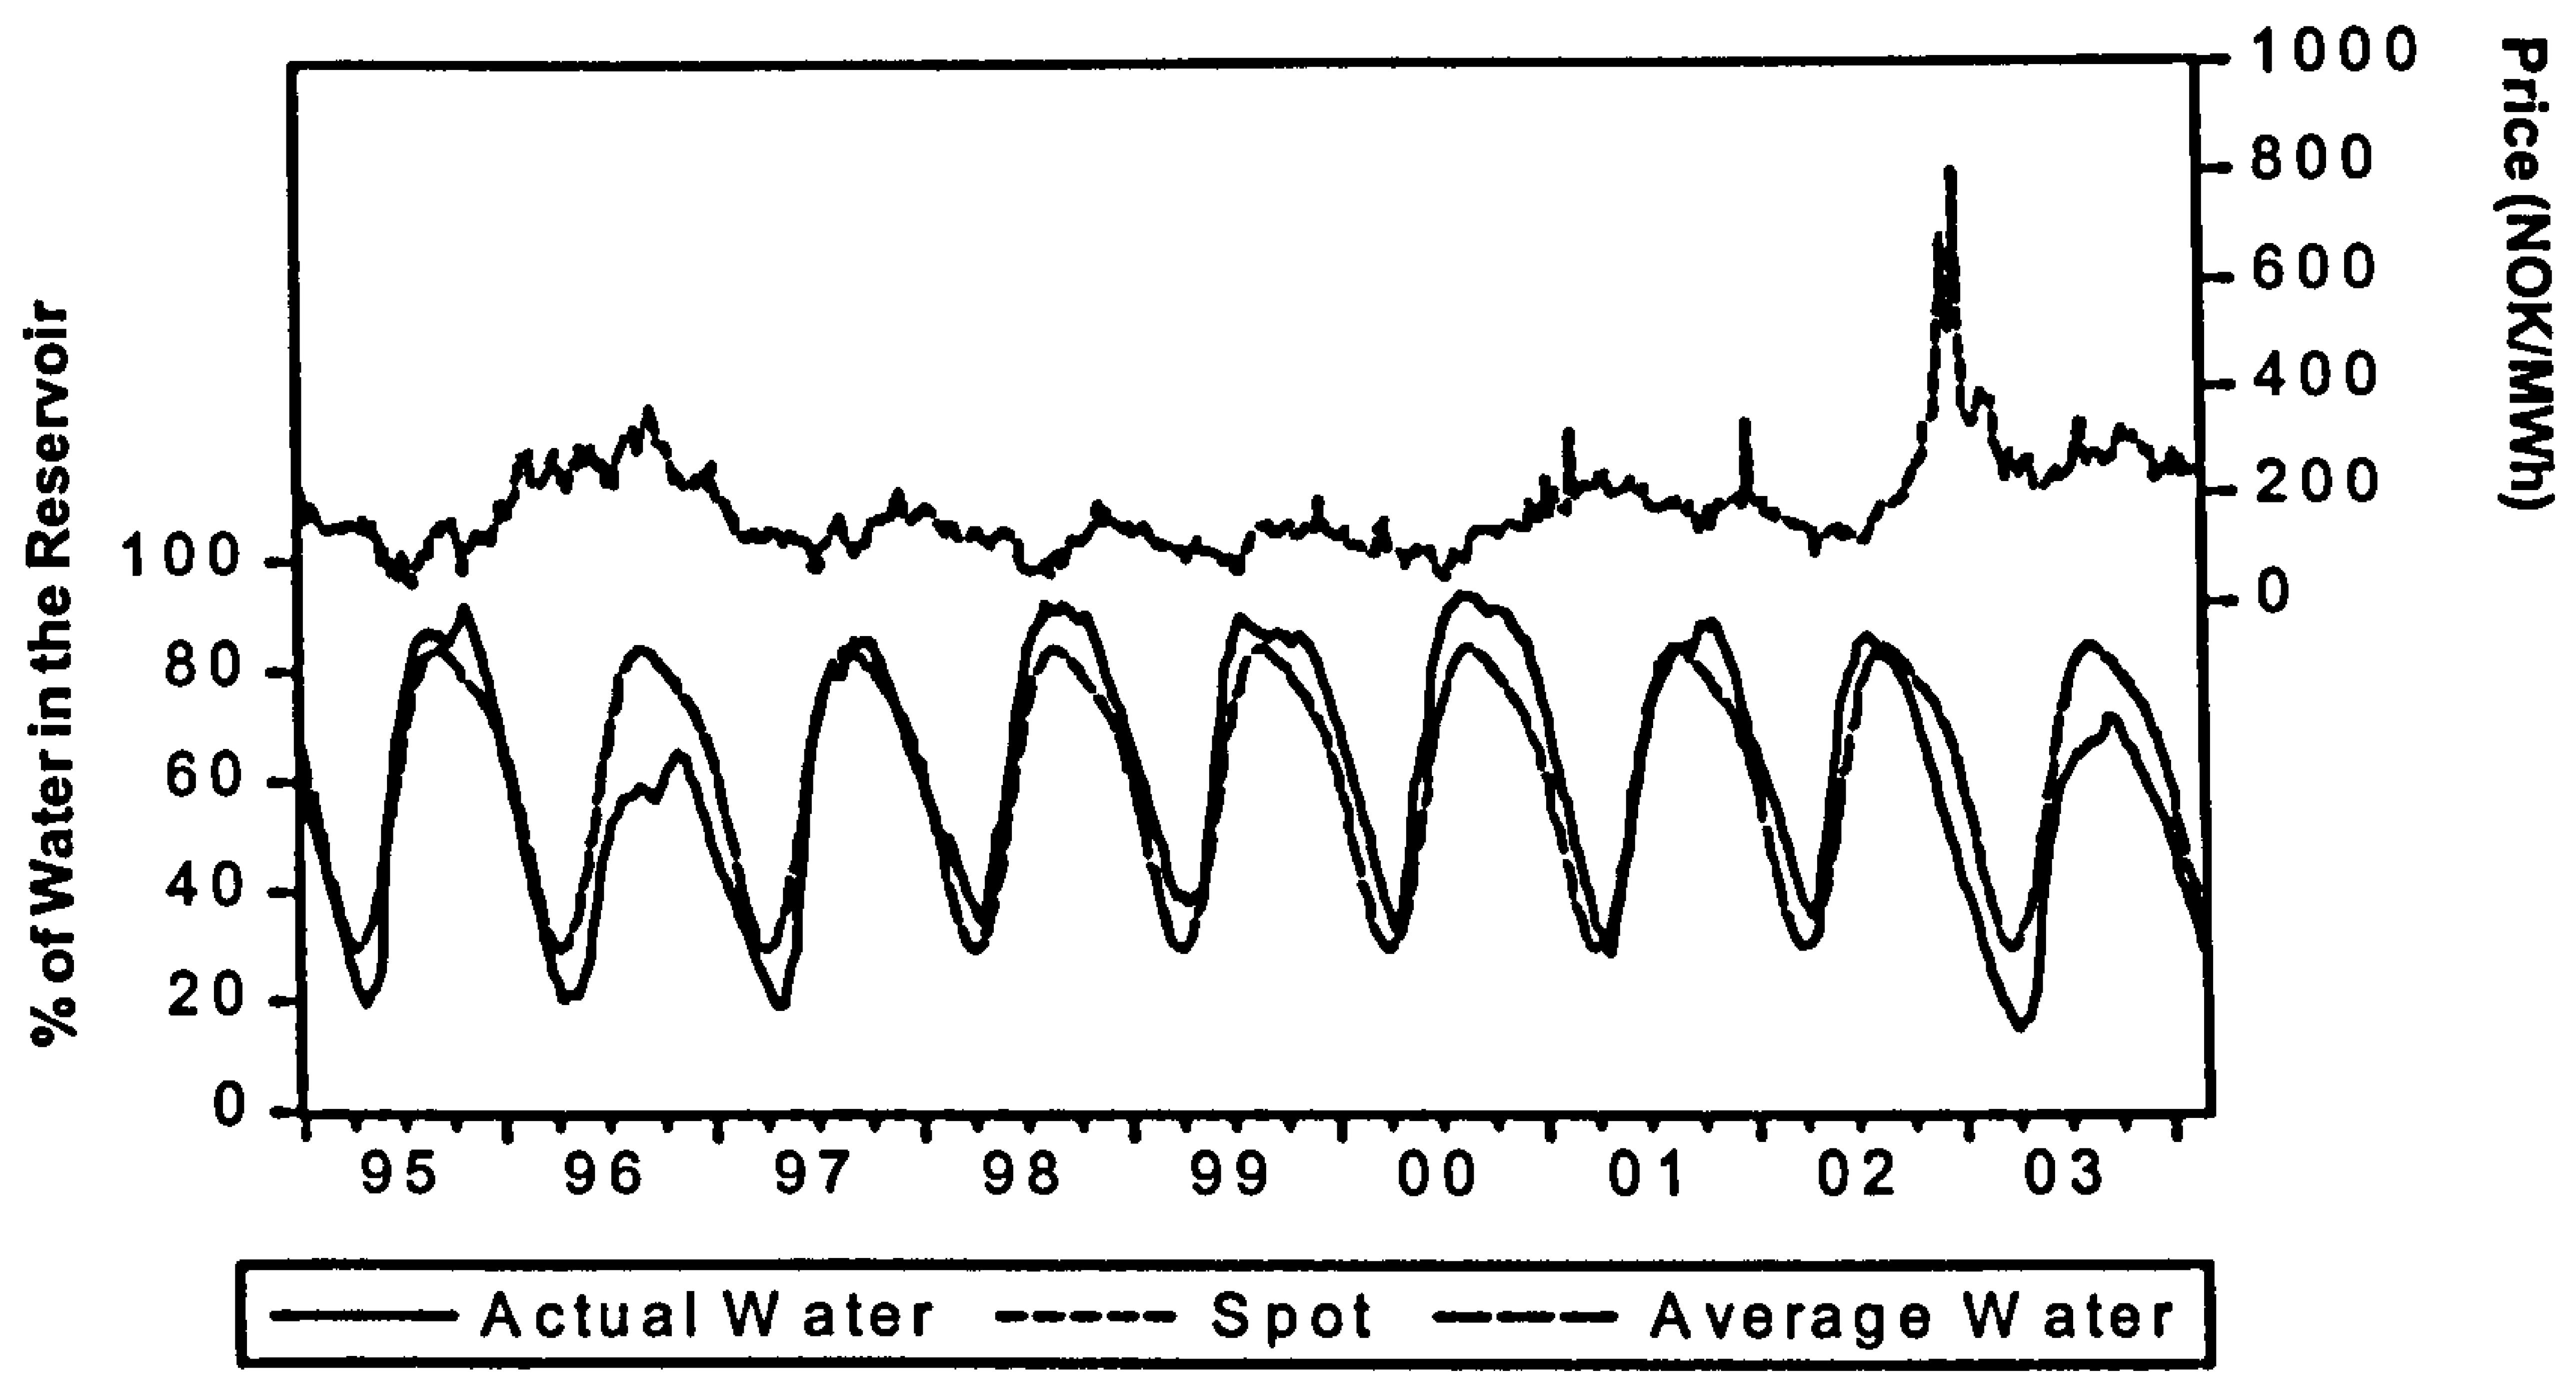

Table 3.4: Parameter Estimates of Seasonality in the Reservoir Levels

Deterministic Seasonality Parameters	
β_0	61.07** (128.36)
β_1	-25.89** (-39.55)
τ_1	-0.92** (-4.18)
β_2	6.10** (9.23)
τ_2	3.03** (6.55)

The table shows the parameter estimates for the seasonality in the percentage of water in the reservoirs as shown in equation (3.2). Estimates followed by (*) or (**), indicate that they are significant at the 5% or 1% level, respectively. Standard errors are (White, 1980) Heteroskedasticity consistent. The regression model is as follows:

$$Seasonality = \left(\beta_0 + \beta_1 \sin\left((t+\tau_1)\frac{2\pi}{365}\right) + \beta_2 \sin\left((t+\tau_2)\frac{4\pi}{365}\right)\right)$$

Figure 3.10 shows the time series of the weekly percentage level of water in the reservoirs in Norway and Sweden (since 96% of the hydro-production comes from these countries), against their seasonal average and the spot price. The seasonal average is estimated using a sinusoidal function with two terms to capture the annual and semi-annual peaks that occur during the winter as follows:

$$\left(\beta_0 + \beta_1 \sin\left((t+\tau_1)\frac{2\pi}{365}\right) + \beta_2 \sin\left((t+\tau_2)\frac{4\pi}{365}\right)\right) \quad (3.2)$$

The estimates of the regression are shown in Table 3.4, where all parameters are found to be highly significant at the 1% level. Focusing on this clearly distinguishable pattern shown in Figure 3.10, water levels start to increase in the mid-spring time mainly from snow melting and precipitation and reach their peak between the months of August and September. Then, water supplies are gradually used during the high demand cold seasons, and reach their lowest level in April. Water levels seem also to vary between different years. For instance, water reservoir levels were high from 1998 to 2001. In the second quarter of 2001 they dropped below their long-run seasonal average, which translated into higher spot prices. In 2003, reservoir levels reached their lowest level (during our sample) resulting in high spot prices, especially during winter. The percentage water level was also very low between the periods of 1995 to mid-1997, which again resulted in very high spot prices. From this graphical analysis, it seems that system prices are mainly affected by the deviation of water levels in the reservoirs from their seasonal average. It seems that this is the critical factor affecting the equilibrium level of system prices, rather than the water level in the reservoirs. For example during early winter when the reservoirs are almost full, the spot price is still higher than the rest of the seasons and the opposite occurs during summer. However when comparing periods where the actual water levels is lower than its seasonal average (e.g. in 1997 and 2003), we can clearly see that the spot is higher, compared to other periods when water levels were higher (e.g. in 1998).

3.6 Conclusions

In the current chapter, we provided a summary of the different market structures that exist in electricity markets, and then went into more depth on how the Nordic power market functions. More specifically, we explained that Nord Pool is a pool market, where generators and suppliers sell and buy electricity respectively through the same pool. Then we showed how the spot price is settled and described the different financial products offered in the financial market of Nord Pool (Eltermin and Eloption). We then analysed the stylised facts of electricity spot prices and the shape of the forward curve. Our main findings are that electricity spot prices are highly volatile, and the major contributing factor to that are the short-lived spikes, caused mainly by outages or transmission failures. We also showed evidence of a distinct pattern between different seasons and days of the week (working and non-working days), which is mainly driven by variation in electricity demand due to human activity. This seasonal pattern is also taken into account when valuing derivatives on the underlying electricity price. Our analysis also showed that due to the non-storable nature of electricity, the spot price is mean-reverting, however since in Scandinavia, 50% of electricity supply comes from hydro-generation, the equilibrium level seems to be affected by the level of water in the reservoirs. Since the scope of this thesis is to value electricity derivatives, the next chapter is devoted on capturing all of the above stylised facts in reduced form models, that are tractable, easy to estimate, and can provide closed-form solutions for derivatives prices.

4. MODELS FOR ELECTRICITY PRICES

4.1 Introduction

In the previous chapters we saw that reconstruction and deregulation, let prices to be determined by the fundamental rules of Supply and Demand, which brought a substitution from *Supply Risk* in the pre-regulation period, to *Price risk*, thus increasing the need for hedging using derivatives such as futures and options. We further explained that electricity is a commodity with some unique features and under the new market framework, the issue of price risk management is becoming of paramount importance. First, electricity is a non-storable commodity and a non-tradable asset; hence arbitrage across time and space is limited. Second, electricity prices exhibit extreme movements and volatility over short periods of time to an extent, which is not evidenced in any other market.

The above features present challenges for researchers in this market. More specifically, the non-storability of the underlying affects the relationship between spot and forward prices. Although spot electricity is not a storable asset, the futures and forward contracts are regular financial contracts that are traded and hence can be used in a replicating strategy. Completeness of the market is in fact the key issue relating these activities. This in turn depends on the choice of the stochastic process that is used to model the securities; the market will be complete as long as the model remains diffusive and the price of risk can be uniquely determined by the term structure. However, this may not always be the case as there may not exist a unique price of risk under which forward prices are the expected risk neutral spot. In this case, it is not possible to develop a model that describes the entire term structure under the real world and, at the same time eliminate all arbitrage opportunities. One way round this is to estimate the spot price dynamics and use a probability measure implied by the risk premia in the market, which fits best the observed term-structure of forward prices.

As examined in the descriptive statistics section of the previous chapter, electricity prices also exhibit very high volatility, the main feature of power prices being the large jumps in prices, called "spikes". Spikes are extreme very short-lived price movements in the spot market that usually do not spill over to the forward market [16]. They are caused by extreme load fluctuations, combined with generating outages or transmission failures, which may take from a few hours up to a few days to fix. From the modelling point of view, the existence of spikes implies that the use of a simple mean reversion model will not be appropriate since it will not capture the dynamics in the spot market. On the other hand, modelling prices using an ordinary jump-diffusion model (as in Clewlow and Strickland, 2000) will result in the jump component being transferred to the forward prices as well, particularly if the speed of mean reversion in the market is not fast enough. Also, such a model will not be able to capture the spiky nature of the jump, especially in cases where shocks in the market die out at a slower rate.

The aim of this chapter is to model the distributional and trajectorial characteristics of the Scandinavian electricity spot market (Nord Pool) and identify the most significant risks for which market participants ask for insurance. Nord Pool is one of the most successful deregulated power markets; it has a very liquid financial (i.e. electricity derivatives) market and its market structure model has been replicated in other markets where deregulation has taken place. Another significant factor that has been shown to affect electricity prices at hydropower-dominated areas (such as in Scandinavia and New Zealand) is the level of water in the reservoirs. It has been shown that during periods of lower water levels, suppliers have to switch to more expensive generators, such as thermal units, with higher marginal costs of production thus increasing the equilibrium price. This empirical fact is significant and affects both the short- and the long-run dynamics of prices, as well as the speed with which prices revert back to their equilibrium level. As it was shown in Chapter 3, the major factor affecting the spot prices is whether the reservoir levels are above or below their significant seasonal pattern.

Our modelling methodology is as follows. In order to capture the fast mean reversion of spikes, we model electricity prices as a two-factor spike model along the lines proposed by Weron (2005). This model has different coefficients of mean reversion in

---

[16] See e.g. Geman (2005).

the spike and normal variables, and provides a better fit to the observed spot and forward prices over alternative model specifications, particularly at time periods when spikes occur. We provide evidence that our model is able to capture the trajectorial properties of the spot market and also, that it improves the fit to the forward prices observed in the market, since the larger coefficient of mean reversion for spikes is able to discount the impact of spikes in prices much faster than a model with a single coefficient of mean reversion for both normal and spike shocks. Since electricity is a non-storable asset, specifying a process that yields price paths resembling the paths observed in the market, and improves the fit to the observed forward prices, is very important for the pricing and hedging of derivative instruments.

Although the stylised fact of mean reversion of spikes is well documented in the literature, and relevant models have been proposed, as in Benth et al (2005), no evidence has been given up to date on such model's performance in fitting the forward term structure and also capturing its distributional characteristics. Therefore, in this chapter we aim to introduce a model that is compared empirically against simpler ones, in terms of fitting the forward prices as well as capturing the first four moments of the actual distribution and yielding similar price paths to the real world.

Furthermore, we also show that spikes are able to explain, to some extent, the seasonal risk premia observed in the market. This is in line with previous research by Bessembinder and Lemmon (2002), Longstaff and Wang (2004) and Villaplana (2003) who show that periods of high demand and volatility correspond to periods of excess skewness in the spot market, which result in higher risk premia in the American power market. It is during these periods that spikes are more likely to occur since any outages or transmission failures will translate into jumps in the prices due to the high inelasticity of demand. Hence, at these periods market participants, such as electricity suppliers, need to transfer their risk through the use of the financial market. At the same time, speculators in the market require a risk premium in order to assume the risk of the suppliers. As a consequence, at periods of high probability of spikes the demand for hedging increases and the number of un-hedged risks rises, thus resulting in higher futures prices.

In addition, we also implement, for the first time in the literature, a Regime-Switching Spike model in order to capture the differences in the equilibrium price level between periods of low and high water reservoir levels. Such a model distinguishes in an intuitive way the medium-run dynamics, mainly caused by fluctuations in the water levels, from the short-run dynamics, and results in a lower speed of mean reversion for the short-run component of the model. Finally, another contribution of the chapter is the development of closed-form solutions for forward prices. Closed-form formulas are more insightful and convenient to use than simulation-based calculations particularly for electricity traders who often need to get quick answers in their day-to-day activities on the relative pricing of different options in the market. For them, speed is often so important that it is necessary to use closed-form formulas instead of simulation-based methods. Closed-form formulas are also useful if options are being valued in a risk management application; for instance, Value-at-Risk statistics can be computed much more quickly if closed-form solutions are available for individual products in the portfolio.

Finally, we extend the spike model to a three-factor spike model, where we let the equilibrium level to follow a stochastic process. The main advantage of this model is that is able to capture both short- and long-term dynamics of the market, and is very similar to models proposed by Schwartz and Smith (2000), Søresnsen (2002) and Manoliu and Tompaidis (2002). This is particularly important for investments that depend on both short-term and long-term market uncertainties, such as in the valuation of power plants and other Real Options. In this way we need to get information mainly from forward prices that mature after one year, for the more precise estimation of the long-term state variables.

The structure of this chapter is as follows: Section 2 introduces the Spike and Regime Switching Spike models. Section 3, presents the calibration and estimation results of the models, and also compares their performance to that of an extended Vasicek model (1977) in terms of fitting the forward prices. Section 4 compares the models in terms of capturing the trajectorial and statistical properties through simulation evidence, and provides closed-form solution for the moments of each model. Section 5 introduces the three-factor model and shows its calibration via the Kalman Filter, and its fit to the forward term structure. Section 6 gives an intuitive comparison between the proposed models, and finally, Section 7 concludes this chapter.

4.2 The Spike Model and the Regime-Switching Spike Model

Electricity prices and more generally energy prices tend to fluctuate around values determined by the marginal cost of production and the level of demand, in other words they have a tendency of *mean reversion* to an equilibrium price. Empirical evidence on mean reversion in Nord Pool has been given by Lucia and Schwartz (2002) using the Augmented Dickey Fuller (1981) test, Escribano et al (2002) after accounting for time-varying volatility and jumps, Weron (2005) via the use of the Hurst exponent, and in the previous chapter of the thesis. Another feature of electricity prices are the extreme short-lived price movements called spikes. The impact of spikes is very important especially for suppliers of electricity, since their costs depend on the variable price of electricity whereas their revenues are fixed as they supply electricity to customers at fixed prices. Despite the fact that spikes occur rarely (and not continuously), they are one of the most important motives for hedging in the power market.

As discussed earlier spikes occur due to intense events, such as extreme load fluctuations, caused mainly by severe weather conditions in combination with generation outages of large generators or transmission failures. These extreme events are normally short-lived and as soon as the outage is fixed (e.g. when a significant generator is back on-line) or the weather phenomenon is over, prices fall back to their normal levels. On the other hand, we also have some downward spikes which are negative jumps. These may occur during spring when traders receive news about the inflow in the water reservoirs from the melting of the snow pack, and also, generators need to get rid of the excess water in the reservoirs to minimise their storage costs. Hence, there is excess capacity in the system and buy-side traders may bid at much lower prices than expected. This occurs primarily during non-working days or days when demand is low.

On the other hand, one should distinguish between short-term fluctuations in supply or demand, and spikes. The normal short-term fluctuations may be caused by shifts in demand, pushing prices upwards and thus increasing the economic incentives of more expensive generators to enter the supply side thus inducing mean reversion in prices. Another factor might be due to weather, which follows a cyclical mean-reverting process, thus affecting the evolution of demand for electricity and influencing the equilibrium spot price. Overall, one should expect that the duration of a spike is much shorter than the life of any other short-term shock caused by shifts in supply or demand, and may hence take from a few hours up to a few days to fix. Another important point is that the speed of mean reversion and its significance are highly affected by the equilibrium level to which power prices revert. As explained in the previous section, the equilibrium price level is not constant, but seems to follow a yearly seasonal pattern and is also affected by the availability of hydropower. These are factors that are considered in our modelling methodology, which is described next.

We model system prices (P_t) as the sum of a predictable component ($f(t)$) and the exponential sum of two stochastic components as shown in equation (4.1). The predictable component takes into account the deterministic regularities in the evolution of prices. The first stochastic component (X_t) is assumed to follow a stationary (as in Vasicek, 1977) process reverting to an equilibrium value, in order to account for the short-term deviations from the equilibrium level due to short-term changes in demand resulting from variations in weather as well as due to market behaviour. The second stochastic component (Y_t) represents the spike component, which is assumed to follow a mean-reverting process whose shocks are modelled by a compound Poisson process.

$$\begin{aligned} P_t &= f(t) + \exp(X_t + Y_t) \\ dX_t &= k_1\left(\varepsilon - X_t\right)dt + \sigma_X dW_X \\ dY_t &= -k_2 Y_t dt + J(\mu_{J_i}, \sigma_{J_i}^2)dq(l_i) \end{aligned} \tag{4.1}$$

In equation (4.1), k_1 represents the speed at which X reverts to its mean-equilibrium value under the real probability measure, ε, after a shock has occurred. dW_X represents the increments of the Brownian motion that cause the random shocks in the short-term factor, X, and it is scaled up by the volatility factor σ_X. For the spike factor Y, k_2 is the speed of its mean reversion and its shocks are modelled via a compound Poisson process. The arrival of jumps is modelled via a Poisson process with intensity l_i, where i accounts for the difference of jump arrivals between cold seasons and the rest of the year. The distribution of the jump size is assumed to be Normal with mean μ_{Ji} and standard deviation σ_{Ji}, which are again assumed to be different between winter and the rest of the year. Note that the seasonality in the jump distribution implies a higher degree of skewness and kurtosis during winter, hence capturing the higher leptokurtosis and asymmetry in the returns distribution observed in winter. Since the spike shock is expected to have a much shorter life than a normal shock, k_2 should be much higher than k_1. This stronger mean reversion for the spike shock is what actually improves the relationship with the futures price, as it will be shown later.

In addition, given the fact that the level of water in the reservoirs affects the equilibrium price level, ε, we also specify a model that takes into account this feature. This model, presented in equation (4.2), is very similar to the first, but we now allow the equilibrium level to vary between periods when the water level is higher (state w), and periods when it is not (state d), with the associated probabilities.

$$
\begin{aligned}
&P_t = f(t) + \exp(X_t + Y_t) \\
&dX_t = k_1\left(\varepsilon - X_t\right)dt + \sigma_X dW_X \\
&dY_t = -k_2 Y_t dt + J(\mu_{J_i}, \sigma_{J_i}^2)dq(l_i) \\
&\varepsilon = \begin{cases} \varepsilon_w & \textit{if water reservoir levels are above seasonal average regime, } R_t = w \\ \varepsilon_d & \textit{if water reservoir levels are below seasonal average regime, } R_t = d \end{cases}
\end{aligned}
\tag{4.2}
$$

Thus ε is a continuous time Markov chain assuming values ε_i, for $i=w,d$. The variable R_t determines the current state and is a random variable that follows a Markov chain with two possible states $R_t=\{w,d\}$, which are assumed to be observable based on whether the water levels in the reservoirs are above or below their seasonal average, at time t. Thus the transition matrix Q contains the probabilities p_{ij} of switching from regime i at the current time t, to regime j at $t+1$:

$$Q = (p_{i,j}) = \begin{pmatrix} p_{ww} & 1-p_{ww} \\ 1-p_{dd} & p_{dd} \end{pmatrix} \qquad (4.3)$$

The Transition Matrix gives the probability of staying at the same regime (e.g. p_{ww}) in the next time step, or moving from one regime to the other (e.g. 1-p_{ww}=p_{wd}). Thus the probability of being in state j at time $t+m$, starting from state i at the current time, t is given by:

$$\begin{pmatrix} p_w \\ p_d \end{pmatrix} = \begin{pmatrix} P(R_{t+m} = w \mid R_t = i) \\ P(R_{t+m} = d \mid R_t = i) \end{pmatrix} = (Q')^m e_i \qquad (4.4)$$

where Q' is the transpose of the transition matrix, and e_i denotes the *ith* column of a 2x2 identity matrix. In Equation (4.4), p_d measures the probability of the water levels being below their seasonal average in the future and p_w the opposite. According to this, we should expect ε_w to be significantly lower than ε_d.

4.2.1 Forward and Futures Contracts Valuation

The value of any derivative must be the expected value of its payoff, under the risk-neutral measure, discounted to the valuation date at the risk-free rate of return. However for futures contracts there is a daily settlement, taking place during the trading period, where as for forward contracts settlement occurs only at maturity. This difference in cash settlement may create a price difference between the two contracts. It has been shown though (see for example Duffie (1992)) that, if interest rates are deterministic and independent from the spot, *futures prices* are equal to the *forward prices* for contracts that have identical terms. Clewlow and Strickland (2000) also show that interest rates do not affect significantly forward prices of non-storable commodities. Consequently in the ensuing analysis, we assume that forward and futures prices are equal. Letting $v_0(P_T,T)$ represent the value of a forward contract on the spot maturing at T, and assuming that the initial cost of entering a forward contract is zero, the forward price $F(0,T,P_0)$, is derived as follows:

$$\begin{aligned} v_0(P_T,T) &= e^{-rT}E_0^*\left[P_T - F(0,T,P_0)\right] = 0 \\ F(0,T,P_0) &= E_0^*(P_T) \end{aligned} \tag{4.5}$$

Where E_0^* is the expectation at $t=0$, under the Equivalent Martingale Measure. Electricity is not a storable asset, and thus we are left with an incomplete market. On the other hand, since forward contracts are tradable assets and can thus be used for hedging, we may find a risk-neutral probability measure that best fits the forward term structure (Harrison and Kreps, 1979 and Cox et al, 1981). This is achieved by taking into account the risk premium or the extra return required per unit of risk. Letting λ to represent the market price per diffusive risk or the degree of risk aversion, which is assumed constant[17], then the market risk implied from the market is $\lambda\sigma_X = \lambda_X$. Therefore using Girsanov's Theorem, the drift is adjusted to the risk-neutral rate of return. By allowing $dW_X^* = dW_X + \lambda_X dt$, to represent the increments of a standard Brownian Motion under the risk-neutral measure, the theoretical valuation of

[17] The assumption of constant market price of risk is consistent with the changes in the state variables and aggregate wealth in the economy being constant and have constant covariance (Merton, 1973 and Cox et al, 1985). Another proposed method is to allow λ to be function of time, and thus calibrate perfectly the model to the forward curve. However, since our aim is to examine which model provides the best fit for derivatives pricing, the use of a constant λ is recommended. See also Lucia and Schwartz (2002).

contingent claims under the assumption of no-arbitrage opportunities can then be formulated.

Investors in the market may also require a premium for jump risk. However there are some difficulties when it comes down to estimating the jump risk premium. First of all as Pan (2002) shows, when minimising the Root Mean Square error between the actual price of the derivative and the theoretical price from the model, the procedure cannot distinguish whether the jump risk is coming from the jump timing parameter, l_t or the mean jump size, μ_J. Another problem stems from the fact that in order to estimate accurately the risk neutral parameters for the jump part, we need a very long liquid daily time series for the forward prices, as discussed by Koekebakker and Lien (2004), otherwise we need liquid option prices. The main reason for this being that jumps violate the perfect hedge assumption in the Black and Scholes framework (as it will be shown in Chapter 5). Without a riskless hedge, option prices will mainly be determined by supply and demand. For instance, if market participants believe that the possibility of upward jumps in the market is high, they would prefer to buy high strike call prices which are cheaper, and this will create excess demand for Out-Of-the-Money Call options and their price would increase. This will then create a smile when using the implied volatility from the Black and Scholes model such as for the model to perfectly fit the option prices. In this respect then, the smile in the option prices is the risk premium for jumps. However, the daily liquidity of forward contracts, and more specifically that of long dated ones is low in Nord Pool. On the other hand, options in Nord Pool are traded once a month on average, as it will be discussed in more detail in Chapter 5. Since liquidity is low, results from the processes may not be reliable and thus we proceed with the proposed risk-neutral measure for the diffusive risk, which is also used by Cartea and Figueroa (2001), in their MRJD model.

Using the results from the Appendix of the current chapter (section 4.8), we can write analytic forms for the distributions of the state variables and spot prices. Since our model is in line with an Affine Jump-Diffusion process, we apply the elementary transform function, which was originally developed by Duffie et al, (2000). Hence equation (4.6) shows the expected risk-neutral value of the spot price in the future implied by the Spike model.[18]

$$E_0^{*}(P_T)=f(T)+\exp\left(e^{-k_1T}X_0+Y_0e^{-k_2T}+\sum_{i=1}^{2}l_i\int_0^T 1_{s\in i}\left(\exp\left(\mu_{J_i}e^{-k_2s}+\frac{1}{2}\sigma_{J_i}^2e^{-2k_2s}\right)-1\right)ds+A_T\right)$$

$$A_T=\frac{1}{2}\frac{\sigma_x^2}{2k_1}\left(1-e^{-2k_1T}\right)+\left(\varepsilon-\frac{\lambda_x}{k_1}\right)(1-e^{-k_1T}) \tag{4.6}$$

The integral in equation (4.6) represents the expected number of jumps times the expected jump size, over the entire life of the forward contract, and how each jump affects the forward price. Thus the integration is performed numerically as follows: if s is an element of 1, meaning that s is in winter, then the jump parameters take the values estimated for winter, and if s passes from spring, summer or autumn, then the jump parameters take the respective values for the rest of the seasons. In this manner, at least theoretically, the model seems to improve the spot-forward relationship especially during spiky periods.

Several points can be noticed about the importance of the state variables to the forward prices, in equation (4.6). First, as the time to maturity increases, the impact of the state variables decreases. Second, the significance of the state variables in the forward price depends on the speed of mean reversion, which acts as a discounting factor. Also this implies that a jump in the price, and hence an increase/decrease in the state variable Y, will have no significant effect in the forward price, as long as k_2 is very high. This is in line with several empirical findings that a spike in the spot prices does not disseminate as a spike in the forward price, since this shock is very short lived. Moreover the higher value of k_2, compared to k_1, suggests that the value for short-term contracts is mainly governed by the jump component, where as for longer-term contracts the premium is governed by the sum of all risk factors, with the jump component playing a lesser role. In this manner, at least theoretically, the model seems to improve the spot-forward relationship especially during spiky periods.

[18] Note that equation (4.6) contains an integral term. In our case we expand the integrand in a first order Taylor series and perform the integration analytically (see Cartea and Figueroa, 2005).

Finally equation (4.6) shows that in the long run as $T \to \infty$, the spot price in the risk neutral world tends to a mean value of:

$$f(T) + \exp\left(\frac{1}{2}\frac{\sigma_x^2}{2k_1} + \left(\varepsilon - \frac{\lambda_X}{k_1}\right) + \sum_{i=1}^{2} l_i \int_0^T 1_{s\in i}\left(\exp\left(\mu_{J_i} e^{-k_2 s} + \frac{1}{2}\sigma_{J_i}^2 e^{-2k_2 s}\right) - 1\right) ds\right) \quad (4.7)$$

Turning now to equation (4.8), this shows the expected risk-neutral spot price using the Regime Switching model. Here $E_0^*(P_t)_d$ denotes the risk neutral expected spot price for time t using the equilibrium level ε_d , and $E_0^*(P_t)_w$ is the risk neutral expected spot price for time t using the equilibrium level ε_w

$$E_0^*(P_t) = p_w^* E_0^*(P_t)_w + p_d^* E_0^*(P_t)_d \quad (4.8)$$

In other words, we find the futures prices in the two states and multiply each value by the corresponding risk neutral (p_w^* and p_d^*) probability, which is calculated using equation (4.4). However due to the same reasons described before for the jump premium, the calibration of the risk adjusted state probabilities is very hard. Therefore in the ensuing analysis we assume that the state probabilities are the same under the real and the risk-adjusted world. Hence formula (4.8) decomposes the forward prices into two values; one for being in the regime of having higher water levels and hence cheaper electricity (state *w*), and another for the possibility of the water reservoir levels running low and relying on more expensive generators to supply electricity (state *d*).

The previous discussion presented the valuation for ordinary forward or futures contracts for a delivery period of one day. However, futures and forwards traded in Nord Pool's financial market are very similar to interest rate swaps, since they trade to fix the electricity price level for a specific delivery period (one day, week, block (month), season (quarter) or year) in the future. A fixed price reference is agreed before the delivery period and the difference from the floating realisation of the spot price is settled financially. Hence the forward price in Nord Pool (F_{0,T_1,T_2}), can be shown to be intuitively the average of the forward prices over the delivery period

$[T_1,T_2]$, where the weights are adjusted by the time value of money function, as in equation (4.9).

$$v_0(P_s,T_1,T_2) = E_0^*\left[\sum_{s=T_1}^{T_2} e^{-rs}\left(P_s - F_{0,T_1,T_2}\right)\right] = 0$$

$$F_{0,T_1,T_2} = E_0^*\left[\frac{\sum_{s=T_1}^{T_2} e^{-rs}P_s}{\sum_{s=T_1}^{T_2} e^{-rs}}\right] = \frac{\sum_{s=T_1}^{T_2} e^{-rs}F(0,s,P_0)}{\sum_{s=T_1}^{T_2} e^{-rs}} \qquad (4.9)$$

4.3 Model Calibration and Results

4.3.1 Description of Dataset

For the empirical estimation of the model, we use spot data for the period March 1, 1997 to February 29, 2004. This period reflects the period after the entrance of Sweden in the Nord Pool market, which together with Norway are the major players accounting for more than 60% of the market. On the other hand, the derivatives data start from January 1, 1998 when liquidity in the financial market increased more (See Nord Pool reports on derivatives). Furthermore, investigation of this period is also interesting since it covers periods of very low water levels in the reservoirs, thus resulting in higher system prices such as in 2001 and 2003-2004. Even though it is expected that this will result in a lower speed of mean reversion, we believe that the estimator is more precise since it also accounts in general how shocks die out in the market on average when such an event occurs. Previous studies in the literature, have not considered this period and hence when such an event occurs, the speed of mean reversion is too high, thus discounting these shocks too fast.

Regarding the forward and futures prices in the dataset, we use constant maturity prices, provided by Viz Risk Management AS. The use of the dataset is motivated by the following reasons; first, they make the calibration of our models easier in terms of computational complexity and intensiveness and, second, it enables us to consistently compare the differences in the pricing errors across maturities. On the other hand as it will be explained in more detail in section 4.5, equation (4.9) does not allow the

Kalman filter that will be used for the three-factor model, to estimate the state variables efficiently. Note also that by using standard maturity forward prices the models can be further adjusted to match exactly the term structure as in Clewlow et al (1999) (an example of the forward curve using ELVIZ was given in Chapter 3 in Figure 3.9).[19]

For the empirical results of the Spike and the Regime Switching Spike models the estimation under the objective probability measure is carried using daily spot prices and weekly observations of the water level in the reservoirs. However to obtain the risk-neutral parameters, we use weekly data for the futures. The weekly frequency for the futures observations is chosen in order to avoid any liquidity problems (especially since longer term contracts are not traded every day), and weekend effects caused by the fact that spot prices are available for all days where as derivatives prices are available only for working days[20]. We use forward prices for every Wednesday, as it is the only day that is less affected by holidays. If Wednesday is a holiday, we use Thursday's prices. The forwards used in this part of the analysis reflect short- to medium-term maturities (1, 2, 4, 7, 12, 16, 20 and 24 weeks to maturity). These maturities are selected to reflect actual traded contracts in the market. Figure 4.1 graphs the time series of the 1 and 7 weeks to maturity forward prices. A very distinctive pattern is that the longer maturity contract moves similarly to the shorter, but in a much smoother manner. It is also interesting to note that forward prices increased, especially in 2003, as a result of the extremely low reservoir levels. These facts, show that forward prices in Nord Pool depend to a large extent on the spot market, but the effect of spikes is not as pronounced in the forward market. Thus we should expect a lower coefficient of mean reversion, for normal diffusive shocks, than the one noticed in other markets[21].

---

[19] Similar data have also been used by other studies such as Koekebakker and Ollmar (2001). The forward curve thus is computed by a program called ELVIZ ,which uses a maximum smoothness function with a sinusoidal prior continuous forward price function, that prices all traded contracts within the bid/ask spread, using equation (4.9). For more details on the forward curve calculation see www.viz.no.

[20] Lucia and Schwartz (2002) for example use actual traded forward contracts with 28-day frequency.

[21] See section 4.B.2 for further discussion on the difference in the speed of mean reversion between different power markets.

Figure 4.1: Time series of Forward Prices

The figure shows the time series of the weekly observations for the 1 and 7 weeks to maturity contracts.

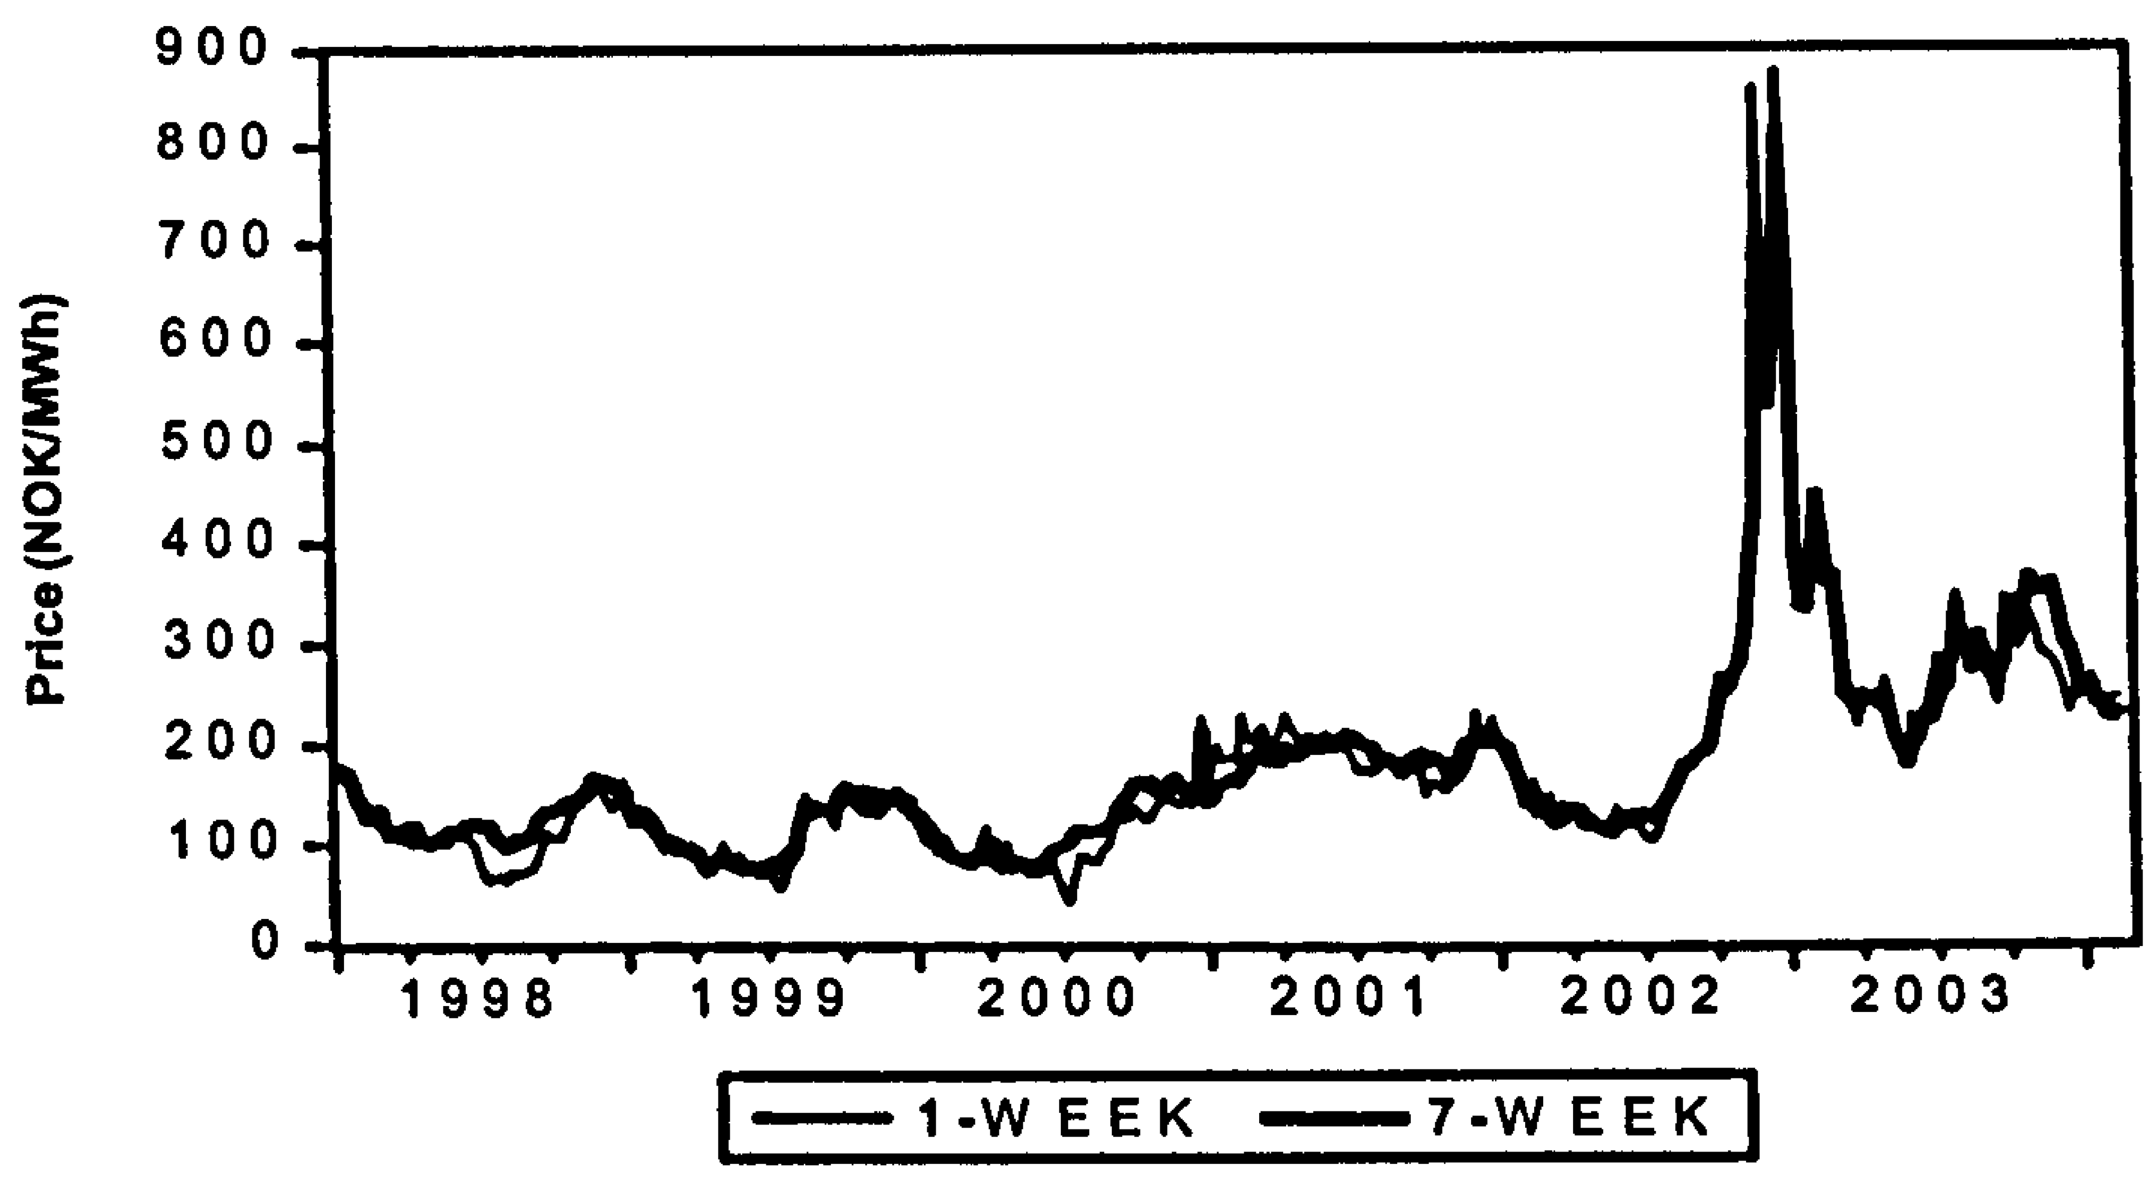

4.3.2 Model Estimation

4.3.2.1 The Deterministic Component

Before estimating the model, we have to define the functional form of the deterministic component. Figure 3.8 indicates that there is a difference between the price level in working and non-working days. Furthermore the descriptive statistics (Table 3.3) and graphical presentation of the spot, together with the shape of the term structure, provide evidence of a significant seasonal behaviour of electricity prices that seem to follow an annual cycle. Hence in order to account for this periodic behaviour we follow Pilipovic (1998) who suggests fitting a sinusoidal function to the system price.[22] Therefore, the deterministic component is modelled via a combination of sine functions, to capture seasonality, and a dummy variable, to distinguish between working and non-working days, and is estimated via Maximum Likelihood (ML). The first sine term has an annual frequency and captures the main yearly cycle, with the peak occurring during winter, and the second term captures the half yearly

[22] Alternatively we can also use a piecewise constant function of one year period, using monthly or seasonal dummies, as in Manoliu and Tompaidis (2002). Despite the fact that this is a flexible approach, it can be criticised on the grounds that dummy variables are very sensitive to anomalies in the sample such as jumps; hence the method does not provide a smooth function for the seasonal component, which may create discontinuities in the forward prices (Lucia and Schwartz, 2002).

cycle in the data. The inclusion of the second sine function is justified on the grounds that it accounts for the longer period the spot price remains at low levels (from Spring to mid-Autumn).

$$f(t) = \beta D_t + \gamma \sin\left((t+\tau)\frac{2\pi}{365} \right) + \delta \sin\left((t+\xi)\frac{4\pi}{365} \right)$$
$$D_t = \begin{cases} 1 & \text{if date } t \text{ is Non-Working day} \\ 0 & \text{otherwise} \end{cases} \qquad (4.10)$$

4.3.2.2 Estimation of Spike, Regime Probabilities and Short-term factors and Risk-Neutral Parameters

In order to identify the spiky regime, and since we want to capture the most extreme price movements we use a modified Clewlow and Strickland (2000) approach. Our definition of a spike is as follows: A spike is a return of the deseasonalised series which is greater in absolute value than 3 standard deviations of all returns, followed by a reduction or extreme returns of the opposite sign which accounts for at least half of the original movement within the most 5 days. Hence this definition is consistent with the spike having a half-life of at most 5 days. If another jump of the same sign occurs within these 5 days, the day-count starts from the occurrence of the last jump[23]. After the spikes are identified, they are removed from the spot series and the missing values are substituted by the average of the two neighbouring values. Y_t is then estimated as the difference between the actual price at that instant and the replaced price. This algorithm is then repeated until no further spikes can be identified. The choice of the 5-day window is motivated by the fact that according to market sources this is the most that an outage or transmission failure, combined with extreme load fluctuations, lasts in the market.

This approach has a number of advantages over alternative approaches for estimating spikes. First, in Clewlow and Strickland (2000) the filtering is implemented along the same lines, however it captures all general jumps in the market (spikes or long-term

[23] Note that the probability of the Brownian Motion scaled by the volatility parameter, in capturing these returns is almost zero.

shifts). Moreover, in their model the speed of mean reversion is the same for both spikes and normal shocks in the market, inducing some persistence in the jump particularly if the coefficient of mean reversion is not high. However, due to the reasons discussed earlier, spikes live much shorter than other shocks in the market. Consequently, the use of two different coefficients of mean reversion, one for normal small shocks and another for abnormal rare extreme shocks causing the spikes, is more appropriate. Second, Weron (2005) imposes the restriction of allowing a spike to last one day, thus disregarding spikes that die out at a lower rate, as it was the case during the winter of 2001. Finally, as Clewlow and Strickland (2000) state, electricity price returns have "numerous different jump components, typically ranging from very high frequency, low volatility to low frequency, high volatility jumps". The proposed method picks out the lower frequency higher volatility jumps, which is what we mainly want to capture; this offers an advantage compared to the ML method for jump estimation (as in Ball and Torous, 1983, 1985), which captures the higher frequency lower volatility jump components.

After the algorithm is performed, the intensity is then calculated as the number of spikes in the sample divided by the number of observations, and the mean jump size is calculated as the average of the spiky returns. Then, the speed of mean reversion for Y is estimated from the extracted data series by the following regression:

$$Y_t \equiv Y_{t-1}e^{-k_2\Delta t} + u_t;\; u_t \sim iid(0,\sigma_u^2) \; and \; \Delta t = 1/365 \qquad (4.11)$$

Figure 4.2 shows the time series of the Spiky State variable. It is quite distinguishable that upward spikes occur mainly during winter, and their size is much bigger than that of downward spikes. One important empirical finding is the fact that out of a total number of 39 spikes, 21 occur during December, January and February from which 20 are upward spikes; this is expected since, as stated earlier, in winter, equilibrium takes place at the steep part of the supply stack which, combined with any outage or transmission failure, may cause these extreme price movements. The downward spikes on the other hand, occur mainly during the spring and summer periods and non-working days. According to the above findings therefore we have very strong evidence of seasonality in the spikes.

Figure 4.2: Time series of *Y*

The figure shows the time series of spike factor *Y* using the proposed algorithm.

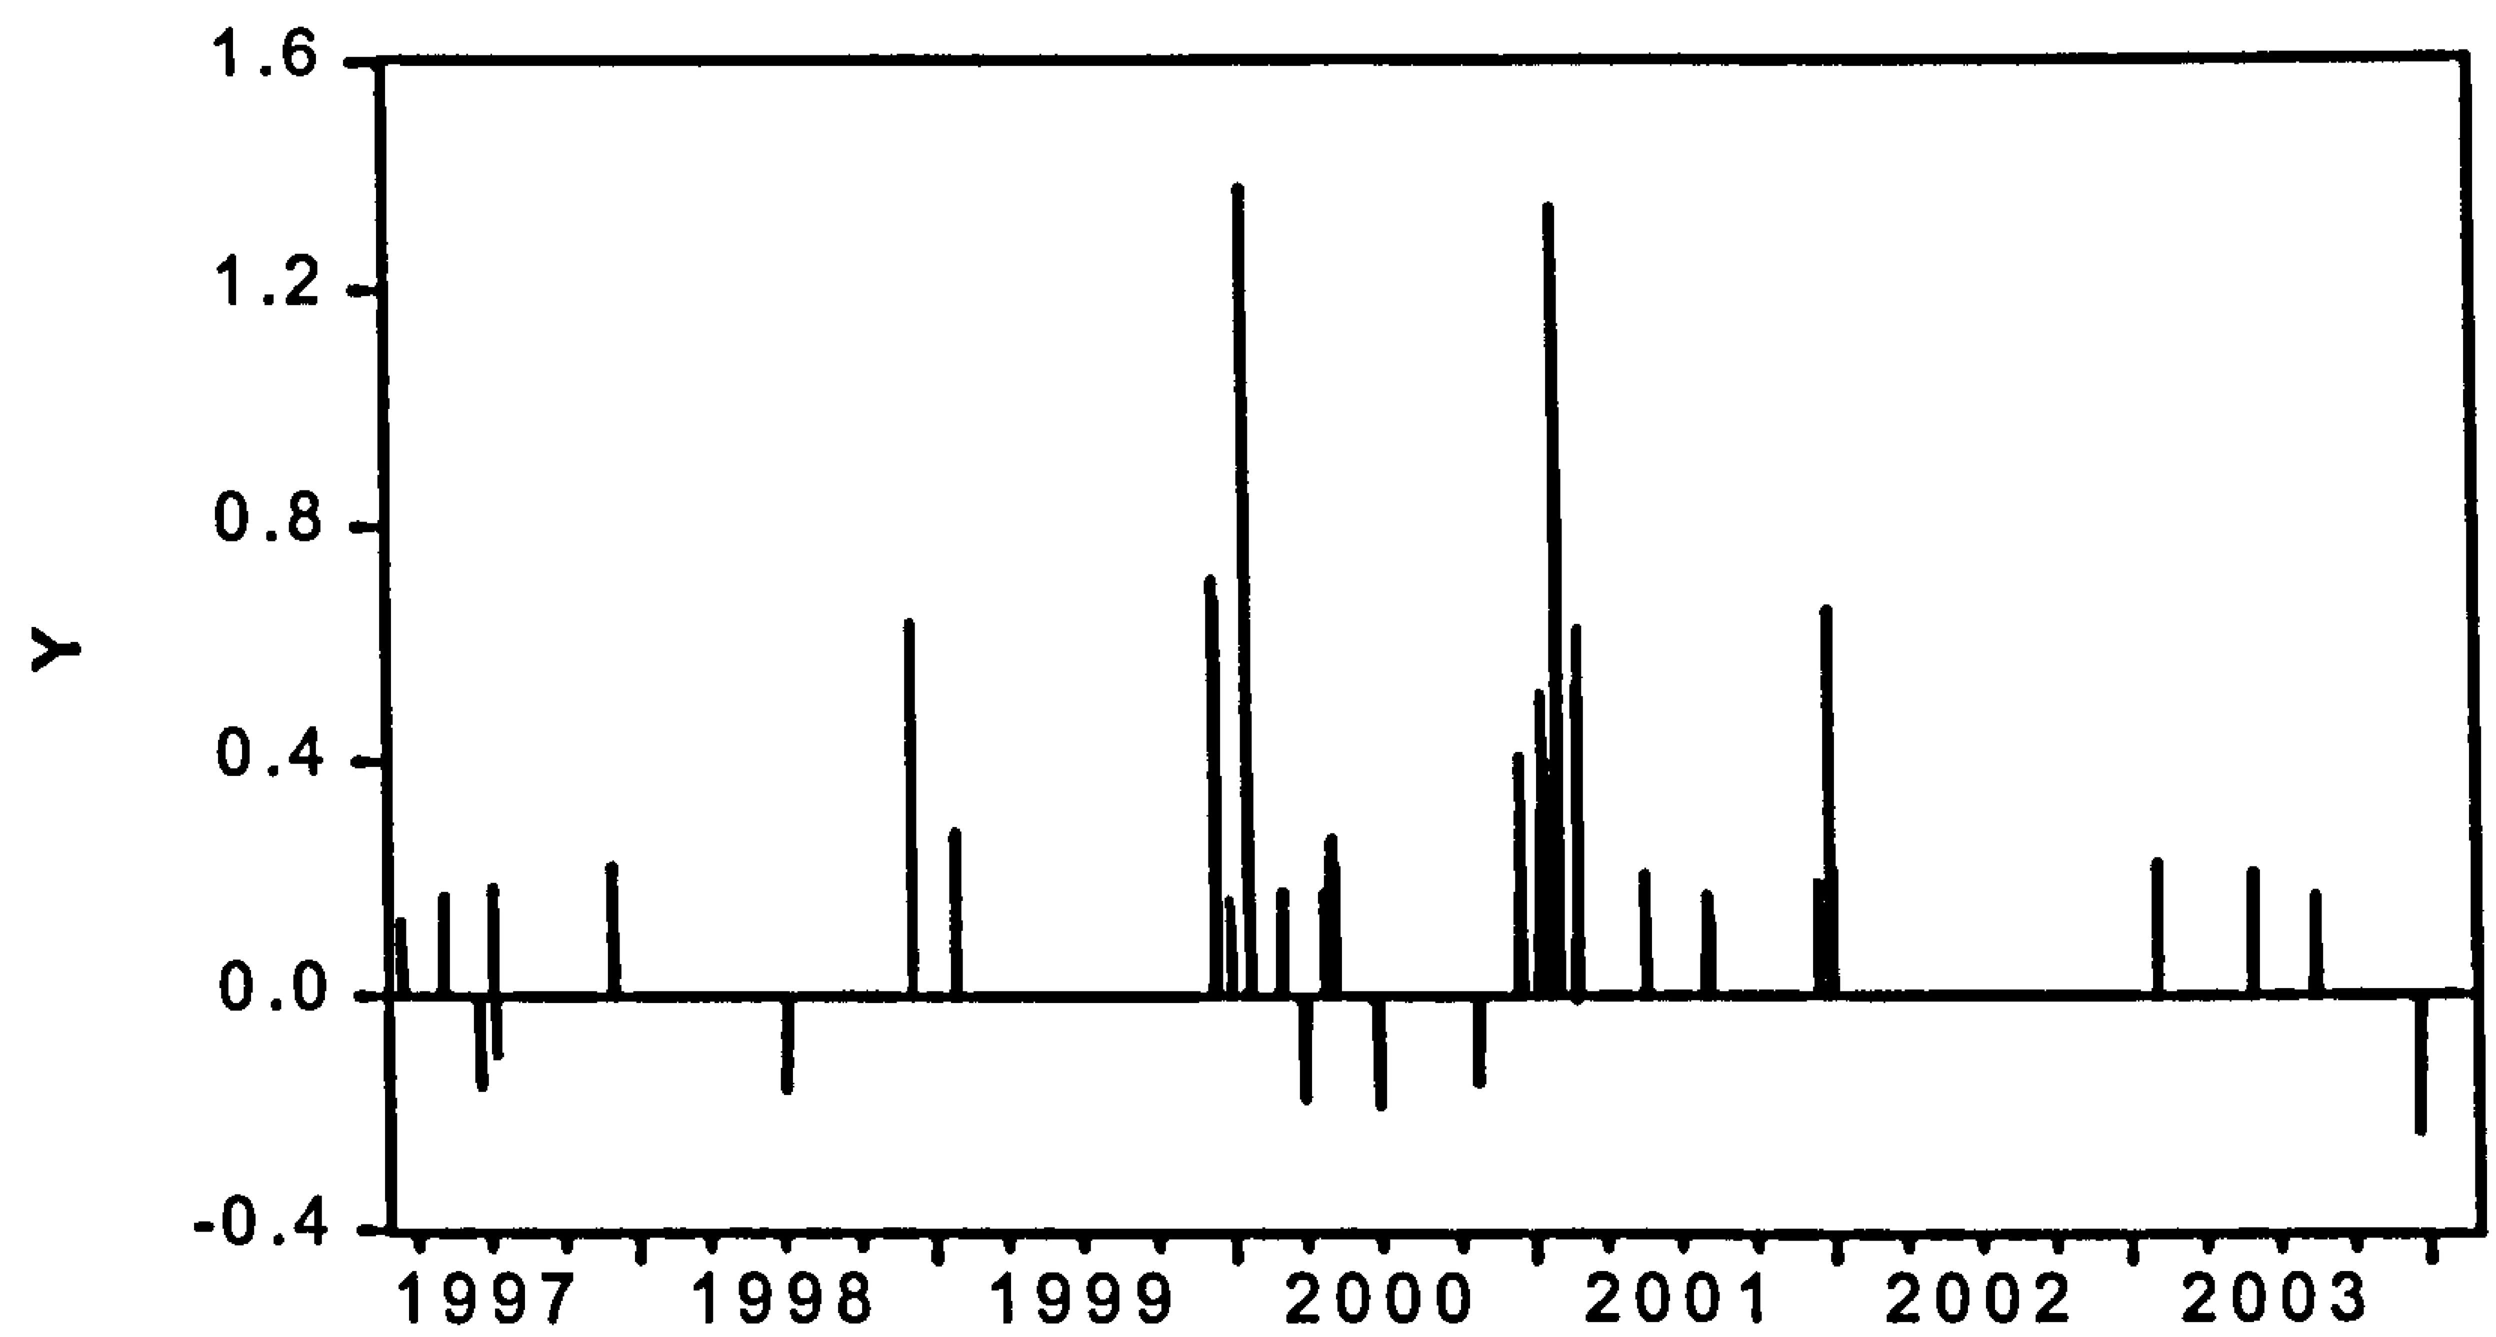

In addition, the existence of spikes implies fat tails for the returns distribution. Figure 4.3 compares the probability distribution of the original deseasonalised return series, against that of the deaseasonalised and filtered return series. In order for the returns to be Normally Distributed, the QQ-Plots have to lie on a 45° line. Looking at the original series, the existence of fat tails suggests that the probability of rare events occurring is much higher than that captured by a normal distribution. However looking at the filtered series' QQ-Plot, it is clearly distinguishable that the Gaussian distribution test of the returns series is significantly improved.

Figure 4.3: QQ plots of Original Returns and Filtered Returns Vs Normal Distribution

The first figure shows the QQ plot for the returns of the original deseasonalised series, where as the second shows the QQ plot of the deseasonalised and spike-less series.

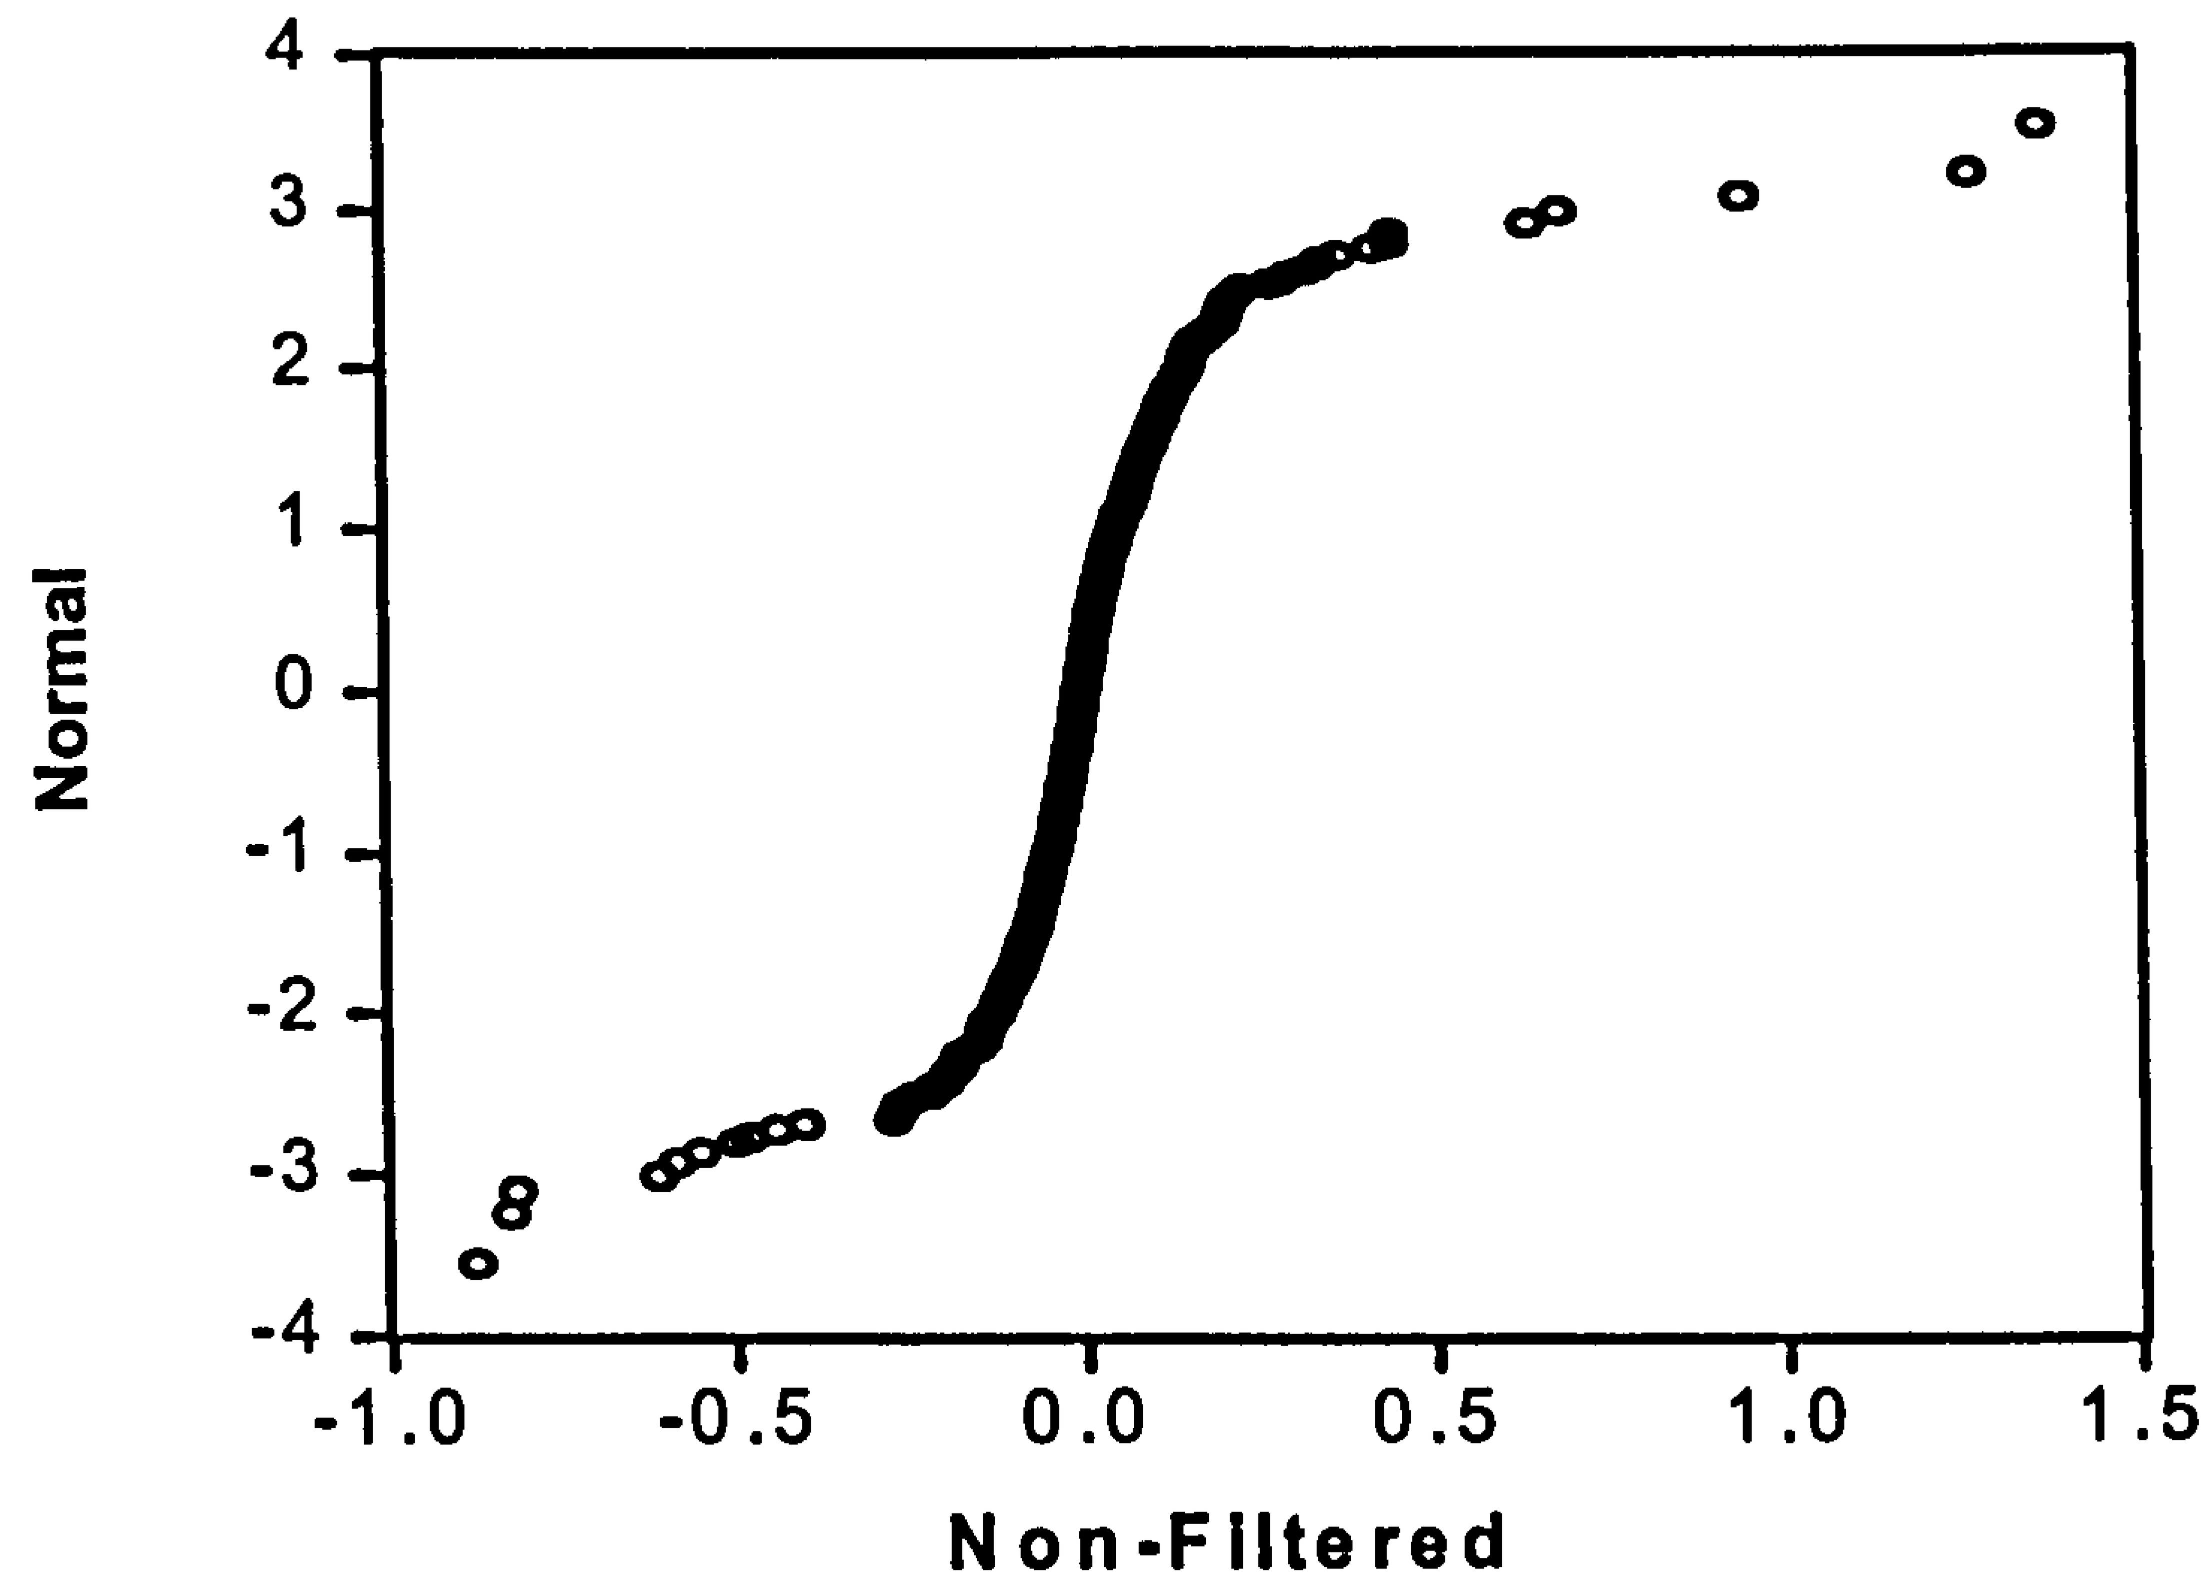

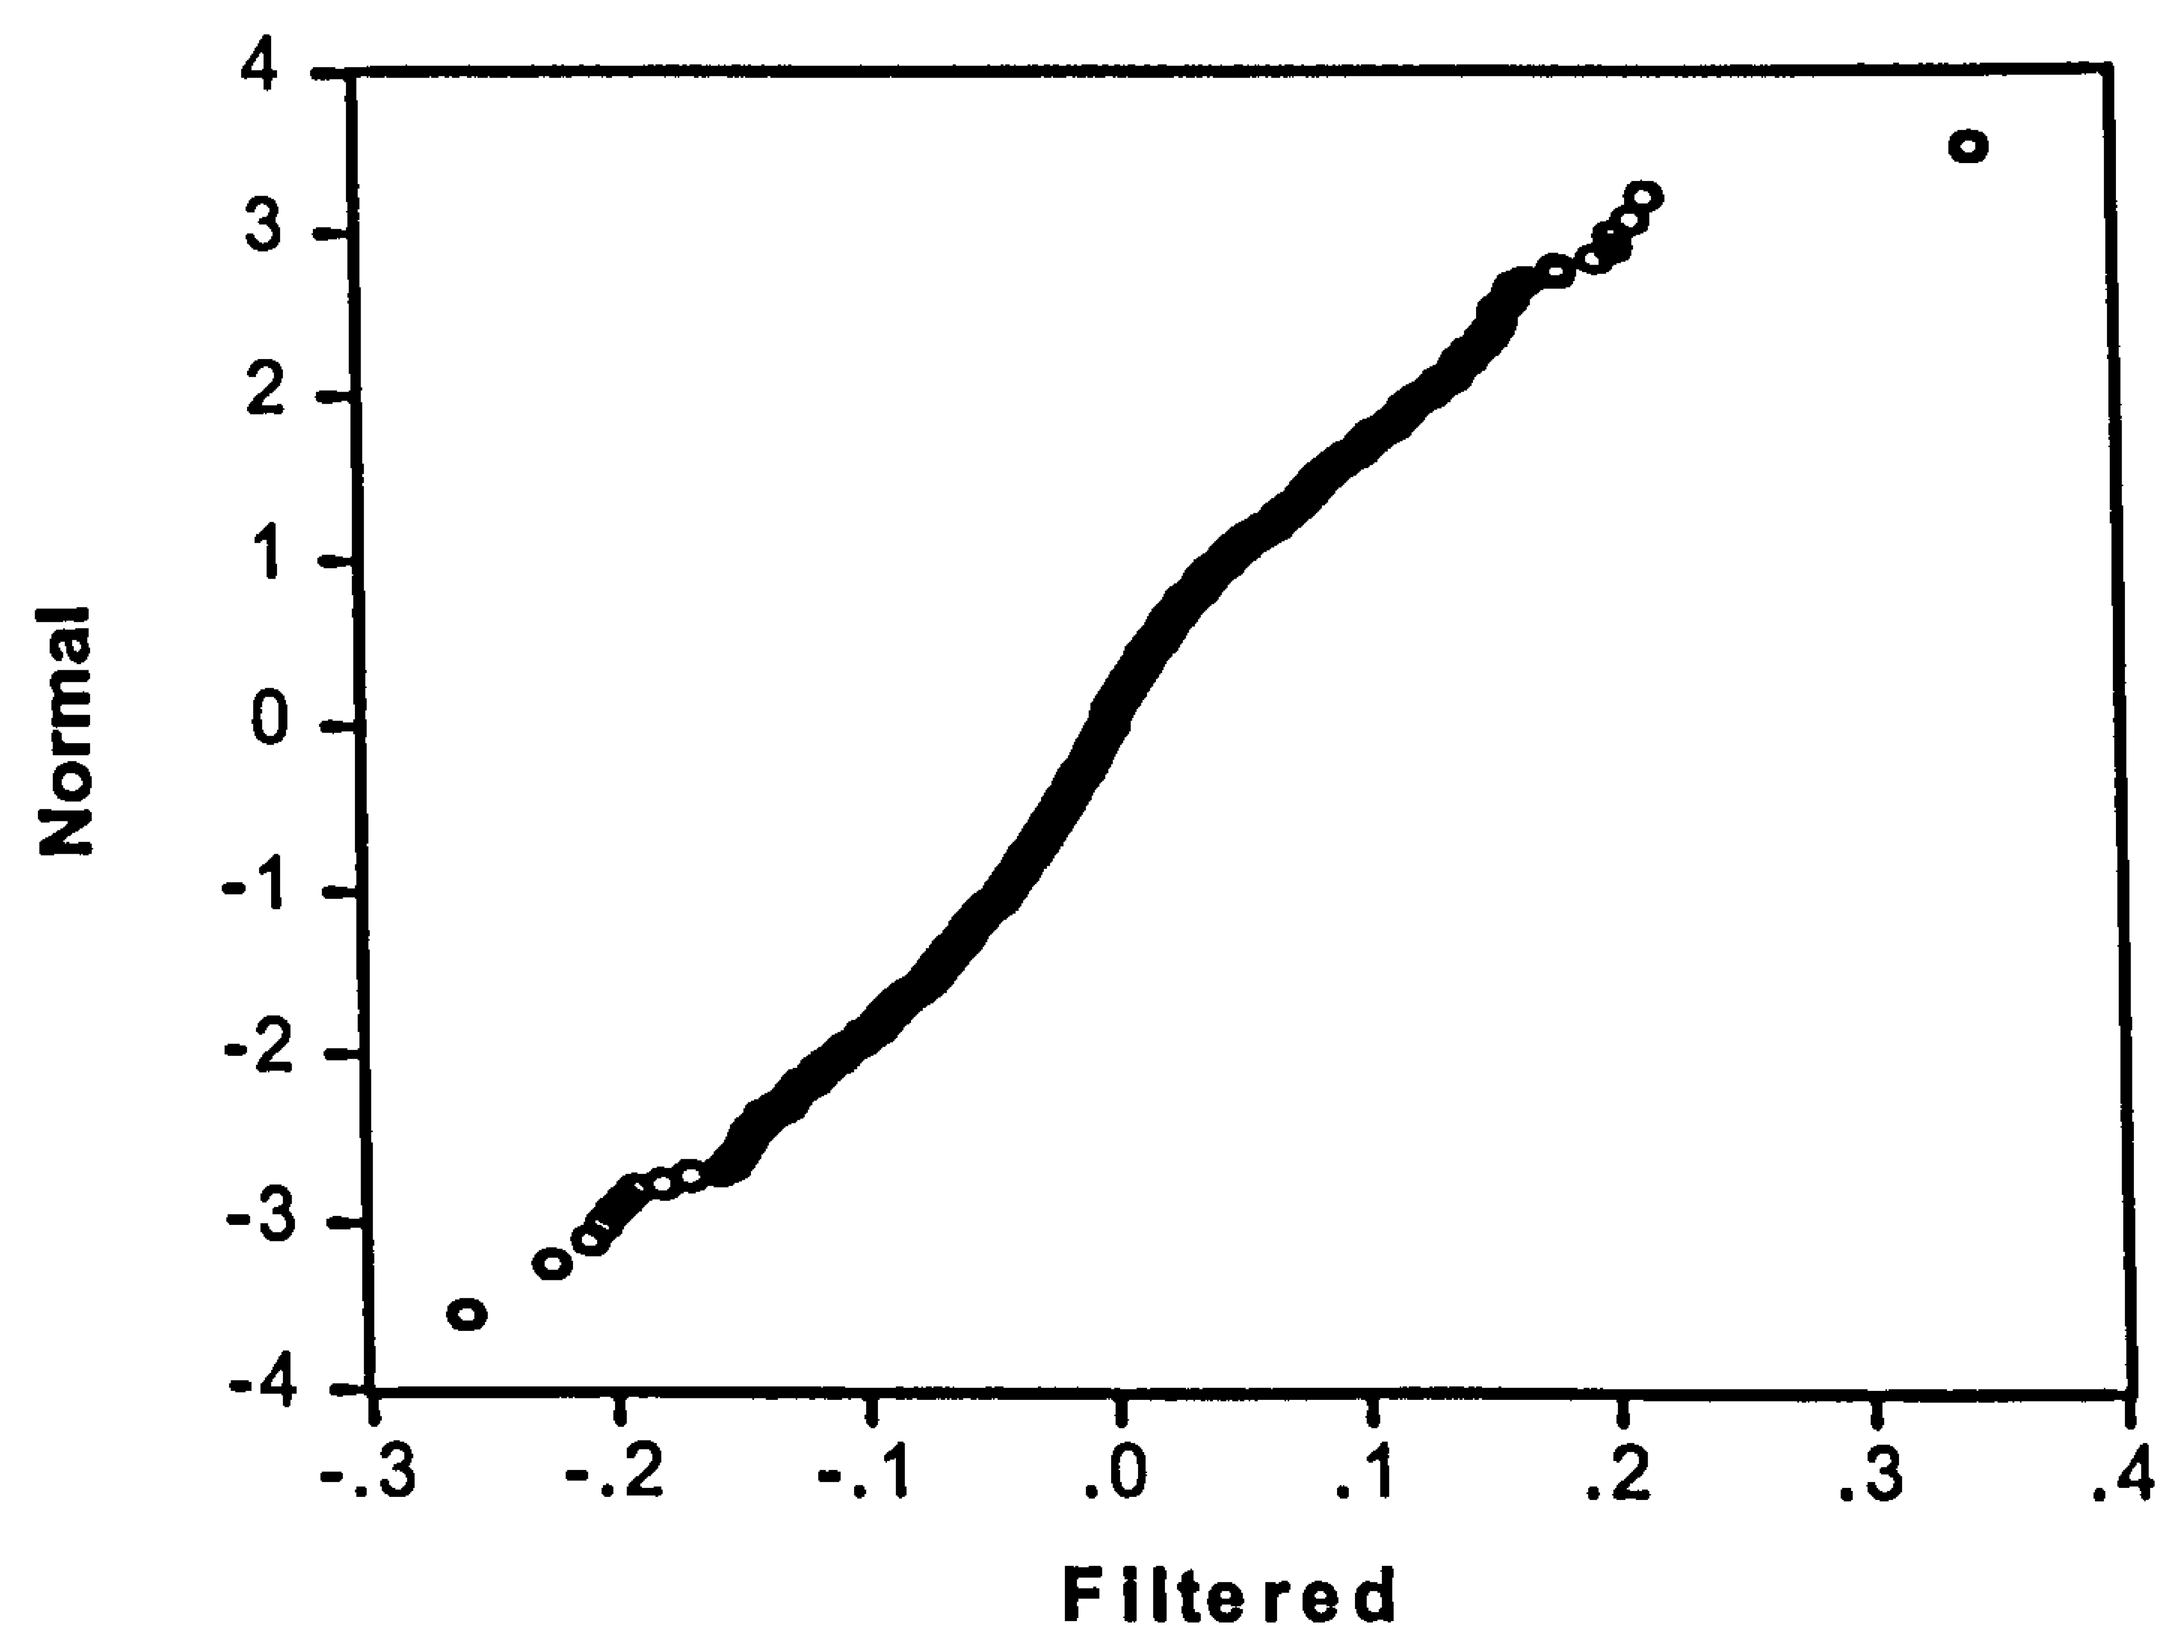

Next, we use the filtered series to estimate the parameters for variable X, in continuous time, using ML. For the estimate of σ_X, we use the standard error from the following regression:

$$X_t \equiv \varepsilon(1 - e^{-k_1 \Delta t}) + X_{t-1} e^{-k_1 \Delta t} + u_t; \quad u_t \sim iid(0, \sigma_u^2) \tag{4.12}$$

For the Regime-Switching Spike model of equation (4.2), the regimes relating to the level of water in the reservoirs are assumed to be observable. The parameters of the model are estimated as follows: first, we take the deviations of the percentage level of water in the reservoirs from their seasonal adjusted mean (modelled via a sine function as discussed in section 4.3.2). We then define a dummy variable (d_t) taking the value of 1 if this deviation is negative (i.e. if the water level is below the seasonal average at day t) and 0 otherwise as in equation (4.13). Then using these observations we estimate the transition probabilities of equation (4.2), and then regress X as in (4.13), to identify the speed of mean reversion.

$$\begin{aligned} X_t &\equiv (\varepsilon + w d_t)(1 - e^{-k_1 \Delta t}) + X_{t-1} e^{-k_1 \Delta t} + u_t; \quad u_t \sim iid(0, \sigma_u^2) \\ d_t &= \begin{cases} 1 & \text{if date } t \text{ has low water} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{4.13}$$

The transition probabilities are calculated as follows: divide the total number of observations whose current state (t) and previous state (t-1) are the same, by the total number of observations. Since we have weekly observations for the water levels, we estimate the daily levels using linear interpolation between successive weekly values. This kind of approximation does not affect our results since, if one week is in regime one, it is found that it would remain to this regime for a long period of time. Using the forward prices, we then estimate the remaining parameters, which represent the market price of risk (λ_X), by minimising the Root Mean Square Error (RMSE) between the actual forward prices and the theoretical prices implied by each model, as in equations (4.6) and (4.8). For comparison purposes, we also estimate a simple Mean-Reverting model (MR) in which the spot prices are the same as in equation (4.1), but without Y and its parameters. Therefore the forward price in this case is given by:

$$\begin{aligned} F(0,T,P_0) &= E_0^*(P_T) = f(T) + \exp\left(e^{-k_1 T} X_0 + A_T\right) \\ A_T &= \frac{1}{2}\frac{\sigma_X^2}{2k_1}\left(1 - e^{-2k_1 T}\right) + \left(\varepsilon - \frac{\lambda_X}{k_1}\right)(1 - e^{-k_1 T}) \end{aligned} \tag{4.14}$$

Table 4.1: Annualised estimated Parameters

	MR	Spike	Spike & Regime
Panel A: Deterministic Seasonality Parameters			
γ		48.61** (2.767)	
τ		432.17** (2.303)	
δ		11.92** (2.578)	
ξ		-156.98** (5.292)	
β		-15.06** (3.659)	
Panel B: Parameters for Stochastic Processes			
ε	5.05** (0.079)	5.06** (0.13)	4.83** (0.085)
w			0.60** (0.13)
k_1	8.07** (1.85)	2.98** (1.03)	5.78** (1.47)
k_2		287.61** (52.99)	287.61** (52.99)
σ_X	167.33%	103%	103%
$l_{(1)}$		12.17	12.17
$l_{(2)}$		3.41	3.41
$\mu_{(1)}$		41.86%	41.86%
$\mu_{(2)}$		8.41%	8.41%
$\sigma_{J(1)}$		38.12%	38.12%
$\sigma_{J(2)}$		20.40%	20.40%
$p_{(ww)}$			99.73%
$p_{(dd)}$			99.60%
λ_X	-0.55	-0.23	-0.45

The table shows the parameter estimates for the simple MR (equation (4.12)), the spike (equation (4.1)) and the Regime Spike (equation (4.2)) models. Estimates followed by (*) or (**), indicate that they are significant at the 5% or 1% level, respectively. Standard errors are White (1980) Heteroskedasticity consistent. For the general model such as the Regime Switching Model:

$$P_t = f(t) + \exp(X_t + Y_t)$$

$$f(t) = \beta D_t + \gamma \sin\left((t+\tau)\frac{2\pi}{365}\right) + \delta \sin\left((t+\xi)\frac{4\pi}{365}\right)$$

$$D_t = \begin{cases} 1 & \text{if date } t \text{ is Non-Working day} \\ 0 & \text{otherwise} \end{cases}$$

$$dX_t = k_1(\varepsilon - X_t)dt + \sigma_X dW_X$$

$$dY_t = -k_2 Y_t dt + J(\mu_{J_i}, \sigma^2_{J_i})\, dq(l_i)$$

$$\varepsilon = \begin{cases} \varepsilon_w & \text{if water reservoir levels are above seasonal average regime, } R_t = w \\ \varepsilon_d & \text{if water reservoir levels are below seasonal average regime, } R_t = d \end{cases}$$

$$\text{Transition Matrix} = \begin{pmatrix} p_{ww} & 1 - p_{ww} \\ 1 - p_{dd} & p_{dd} \end{pmatrix}$$

Where *i=1 for Winter or 2 for Spring, Summer, Autumn* and $\varepsilon_d = \varepsilon + w$.

Table 4.1 shows the estimated parameters of the two models. Several points merit discussion: First, the seasonal parameters are all highly significant reflecting the strong seasonal pattern in the series. Also β is significantly negative, indicating that electricity prices are lower during non-working days. Second, the degree of mean reversion in Nord Pool is lower than that evidenced in other markets (e.g. in PJM investigated by Escribano et al (2002)). This is due to the fact that in Nord Pool about 50% of the electricity is produced using hydro generators; this implies that electricity can be stored indirectly (although in limited amounts) and hence there is some inter-temporal substitution for generating electricity which can dampen very short-term variations, thus lowering the coefficient of mean reversion.

The half-life of a shock, $(ln(2)/k)$, caused by the short term Gaussian factor X_t, will be on average 1 month for the simple MR model. However when extracting the spikes from the sample we see that the coefficient reduces, and the half-life of the shocks is on average two and a half months. On the other hand, the half-life of a spike, k_2, is found to be less than a day. This reflects the effect of the strong mean reversion caused when a spike occurs. The high speed of mean reversion in the spiky regime is one of the significant features of this model, which also improves the relationship between forward prices and the spot. Also notice that the volatility parameter falls from 167% to 103%, thus, spikes seem to play a very significant role in terms of explaining the volatility in the market.

For the Regime Switching Spike model, the speed of mean reversion of the diffusive shock is almost twice that of the Spike model; hence the half-life of a shock is almost one month. This is due to the fact that econometrically the model distinguishes the difference in the equilibrium level of power, when water reservoirs are low. To that extent, we find very strong evidence that the equilibrium level is almost 100 NOK/MWh higher during periods of lower water, as the marginal costs of production increase. Furthermore, the transition probabilities show that if the spot is in one state, it will remain at this state for a long time, thus inducing some persistence. Hamilton (1989) suggested the following formula to calculate the average expected duration of being in state w and d respectively:

$$\sum_{j=1}^{\infty} j p_{ww}^{j-1}(1-p_{ww}) = (1-p_{ww})^{-1} = 370 \text{ days}$$
$$\sum_{j=1}^{\infty} j p_{dd}^{j-1}(1-p_{dd}) = (1-p_{dd})^{-1} = 250 \text{ days} \qquad (4.15)$$

Thus the expected duration when we are in a regime with high water reservoir levels, is one year and 5 days, where as when we are in a regime with low water reservoir levels, the expected duration is almost 8 months and a half. This is expected since if a year is dry or the water level in the reservoirs is low, it will remain like this for approximately one year. This also implies that for short maturities, reservoir risk has little impact on forward prices; as the time horizon increases, uncertainty regarding future water levels increases as well which has a greater impact on forward prices in the market.

In terms of the intensity parameters, upward spikes occur mainly during winter, causing severe upward price movements of approximately 42%. In the rest of the seasons, spikes occur mainly during spring and summer, but the intensity is much less than that of the cold season, with a mean size of 8.4%. This also induces seasonal volatility, which is found to be particularly high during winter.

Turning next into the estimation of the price of risk, the magnitude of the risk premium is smaller for the Spike and Regime switching spike models, compared to the MR model. Generally, the risk premium in the market is negative, which is explained by the fact that in Nord Pool there is greater flexibility on the supply side using the water reservoirs, which in turn induces greater hedging pressures from the demand side of the market that buys derivatives in order to hedge against higher prices. The seasonality on the spike factor makes the derivatives maturing during winter, more expensive than those maturing the rest of the year. This is in line with previous studies (Villaplana 2003, Bessembider et al 2002 and Pirrong and Jermakyan 2000) who argue that during periods of extreme price movements and especially spikes, distributors of electricity increase their demand to hedge their price risk and are willing to buy at even higher prices, in order to persuade "speculators" who cannot hedge their risk to trade with them. Furthermore, the results show that for the short-term forwards, their price is mainly governed by the expected value added from the spike factor, as the value of k_2 is greater than that of k_1. However, as the maturity increases the forward price is governed by the value of all risk factors as well as the risk premium.

Table 4.2: RMSE

The table reports the RMSE (NOK/MWh) for the simple Mean Reverting model (MR) the Regime Switching Spike (RSSP) and the Spike model. Errors are reported for the entire sample, when a jump occurs (at *Y*), during periods of probable upward jumps (December, January, February, which we refer here as Winter), and the rest of the year (March to November).

Weeks to maturity	Whole Sample			Winter		Rest of the Year		At Y	
	MR	Spike	RSSP	MR	Spike	MR	Spike	MR	Spike
1	36.98	30.04	30.30	66.58	53.88	16.24	13.05	53.88	20.79
2	47.52	34.74	35.16	83.83	62.11	23.83	16.34	48.95	24.45
4	63.51	42.99	44.46	106.69	72.81	38.89	25.75	46.26	26.50
7	71.92	46.81	47.48	114.30	73.97	50.53	33.21	45.49	23.40
12	76.58	52.35	53.82	84.04	64.38	74.02	47.86	43.86	26.33
16	76.90	56.21	55.85	71.99	58.90	78.42	55.31	39.40	25.11
20	68.67	52.20	51.83	71.37	58.50	67.78	50.01	42.98	28.72
24	60.82	48.08	47.12	71.24	60.84	57.07	43.20	52.31	38.54

Table 4.2 presents the Root Mean Square Error (RMSE) between the theoretical and actual forward prices, for the entire sample and also between winter and rest of the year periods and at days when a spike occurs. First, looking at the days when a spike occurs, the spike model reduces the RMSE by almost 33 NOK/MWh for the 1-week to maturity price compared to the MR model[24]. This is because when a spike occurs in the spot market, this does not translate as a spike in the forward market since the spike lasts for a very short time, where as our contract matures in a week. Thus, the large value of the coefficient of mean reversion on the spike state variable discounts that shock making its impact on the forward price very small. Compared to the MR model our Spike Model is better on average, by 20.58 NOK/MWh in Winter (which coincides with the period of more frequent and higher returns spikes), and by 15.26 NOK/MWh for the entire sample period. Thus the above analysis indicates, that our model significantly improves the fit in the forward market, especially at times when spikes occur, and is also able to explain the seasonal premium during high demand periods when the probability of a spike increases. Moreover, the RMSE seem to

[24] Note that we found that especially for the upward spikes the difference in RMSE was almost 60 NOK/MWh

increase with time to maturity, which is in line with the findings of Lucia and Schwartz (2002). This may be due to the fact that as time to maturity increases, forward prices start to depend more on other factors rather than the short-term spot price.

Finally, comparing the Regime-Switching Spike model (RSSP), with the Spike model over the entire period, we can note that the results are very similar. One reason why might this be happening is the fact that k_1 in the spike model, is almost half of that in the RSSP model, thus the shocks in the spike model are discounted at a much slower rate. Therefore, when we are in a regime where the water reservoir levels are low, spot prices are high, thus the value of X increases and the discount term $exp(-k_1T)$ is higher in the spike model than in the RSSP model, hence affecting forward prices more, and increasing their value. On the other hand, in the RSSP model this effect is captured by decreasing the effect of the X (since k_1 is higher), and at the same time increasing the weight of the higher risk neutral equilibrium level. Overall, it seems that the impact of theses factors is offsetting each other hence overall, there is no significant difference between the two models.

Table 4.3: Models Pricing Performance using Diebold-Mariano's test

The table reports the ratios of the RMSE of the model on the row to the RMSE of the model on the column. The Diebold-Mariano (1995) pair-wise test of the hypothesis that the RMSEs from two competing models are equal is estimated using the Newey-West (1987) covariance estimator with a truncation of 5 lags. * and **, indicates that the ratio is significantly different from 1 at the 5% and 1% significance level, respectively.

Weeks to maturity		MR		MR		Spike
1		0.81^{*}		0.82^{*}		1.01
2		0.73^{*}		0.74^{*}		1.01
4		0.68^{*}		0.70^{*}		1.03
7	Spike	0.65^{*}	RSSP	0.66^{*}	RSSP	1.01
12		0.68^{**}		0.70^{**}		1.03
16		0.73^{**}		0.73^{**}		0.99
20		0.76^{**}		0.75^{**}		0.99
24		0.79^{**}		0.77^{**}		0.98

In order to test the significance of the difference in the pricing performance between the models, we also carried out the Diebold-Mariano (1995) pair-wise test statistic. Table 4.3 shows the ratios of the RMSE of the model on the row to the RMSE of the model on the column; hence if the ratio is less than 1, the model on the row performs better than the model on the column. Both our models perform significantly better than the simple MR model, and, as the time to maturity increases, the significance in the RMSE increases as well. The difference between the Spike and RSSP models is not statistically significant, however, we find that as the time to maturity increases so does the test statistic between the two models in favour of the RSSP model, in nominal terms.

4.4 Simulation evidence of the models

In Figure 4.4 we show simulated paths for the spot price, which result from discretising each stochastic model and using the estimated parameters under the real probability measure. For comparison purposes, we also include simulated paths generated by a Mean reversion Jump Diffusion process (MRJD), which allows factor X to be affected directly by jumps, along the lines of Clewlow and Strickland (2000). Several points are worth noting. First, although the MR model captures the seasonality and volatility in the series, it is not able to capture the extreme jumps (Spikes) caused in the market. On the other hand the MRJD is able to capture the jumps, however it does not revert back to the normal price level fast enough, since the speed of mean reversion in the model is the same for normal and jump shocks. Consequently, once a jump has occurred the price level remains high for a considerable period of time. The Spike and the Regime Switching Spike models on the other hand, capture most of the main features of the market; extreme upward spikes that occur mainly in the cold season, small spikes and their very fast speed of mean reversion, seasonality and high volatility. What actually distinguishes the two models is that the Regime Switching Spike model is more stationary, reverting faster and more steadily to its equilibrium mean depending on the state, where as the Spike model forces shocks to last for longer time, and mean reverts to a higher price than what is actually the case in the market, when water levels are low. Hence according to the above, both of our models describe the short-term dynamics of price better, and the Regime Switching Spike seems to be more accurate in the long run.

Figure 4.4: Simulated Spot prices under different model specifications

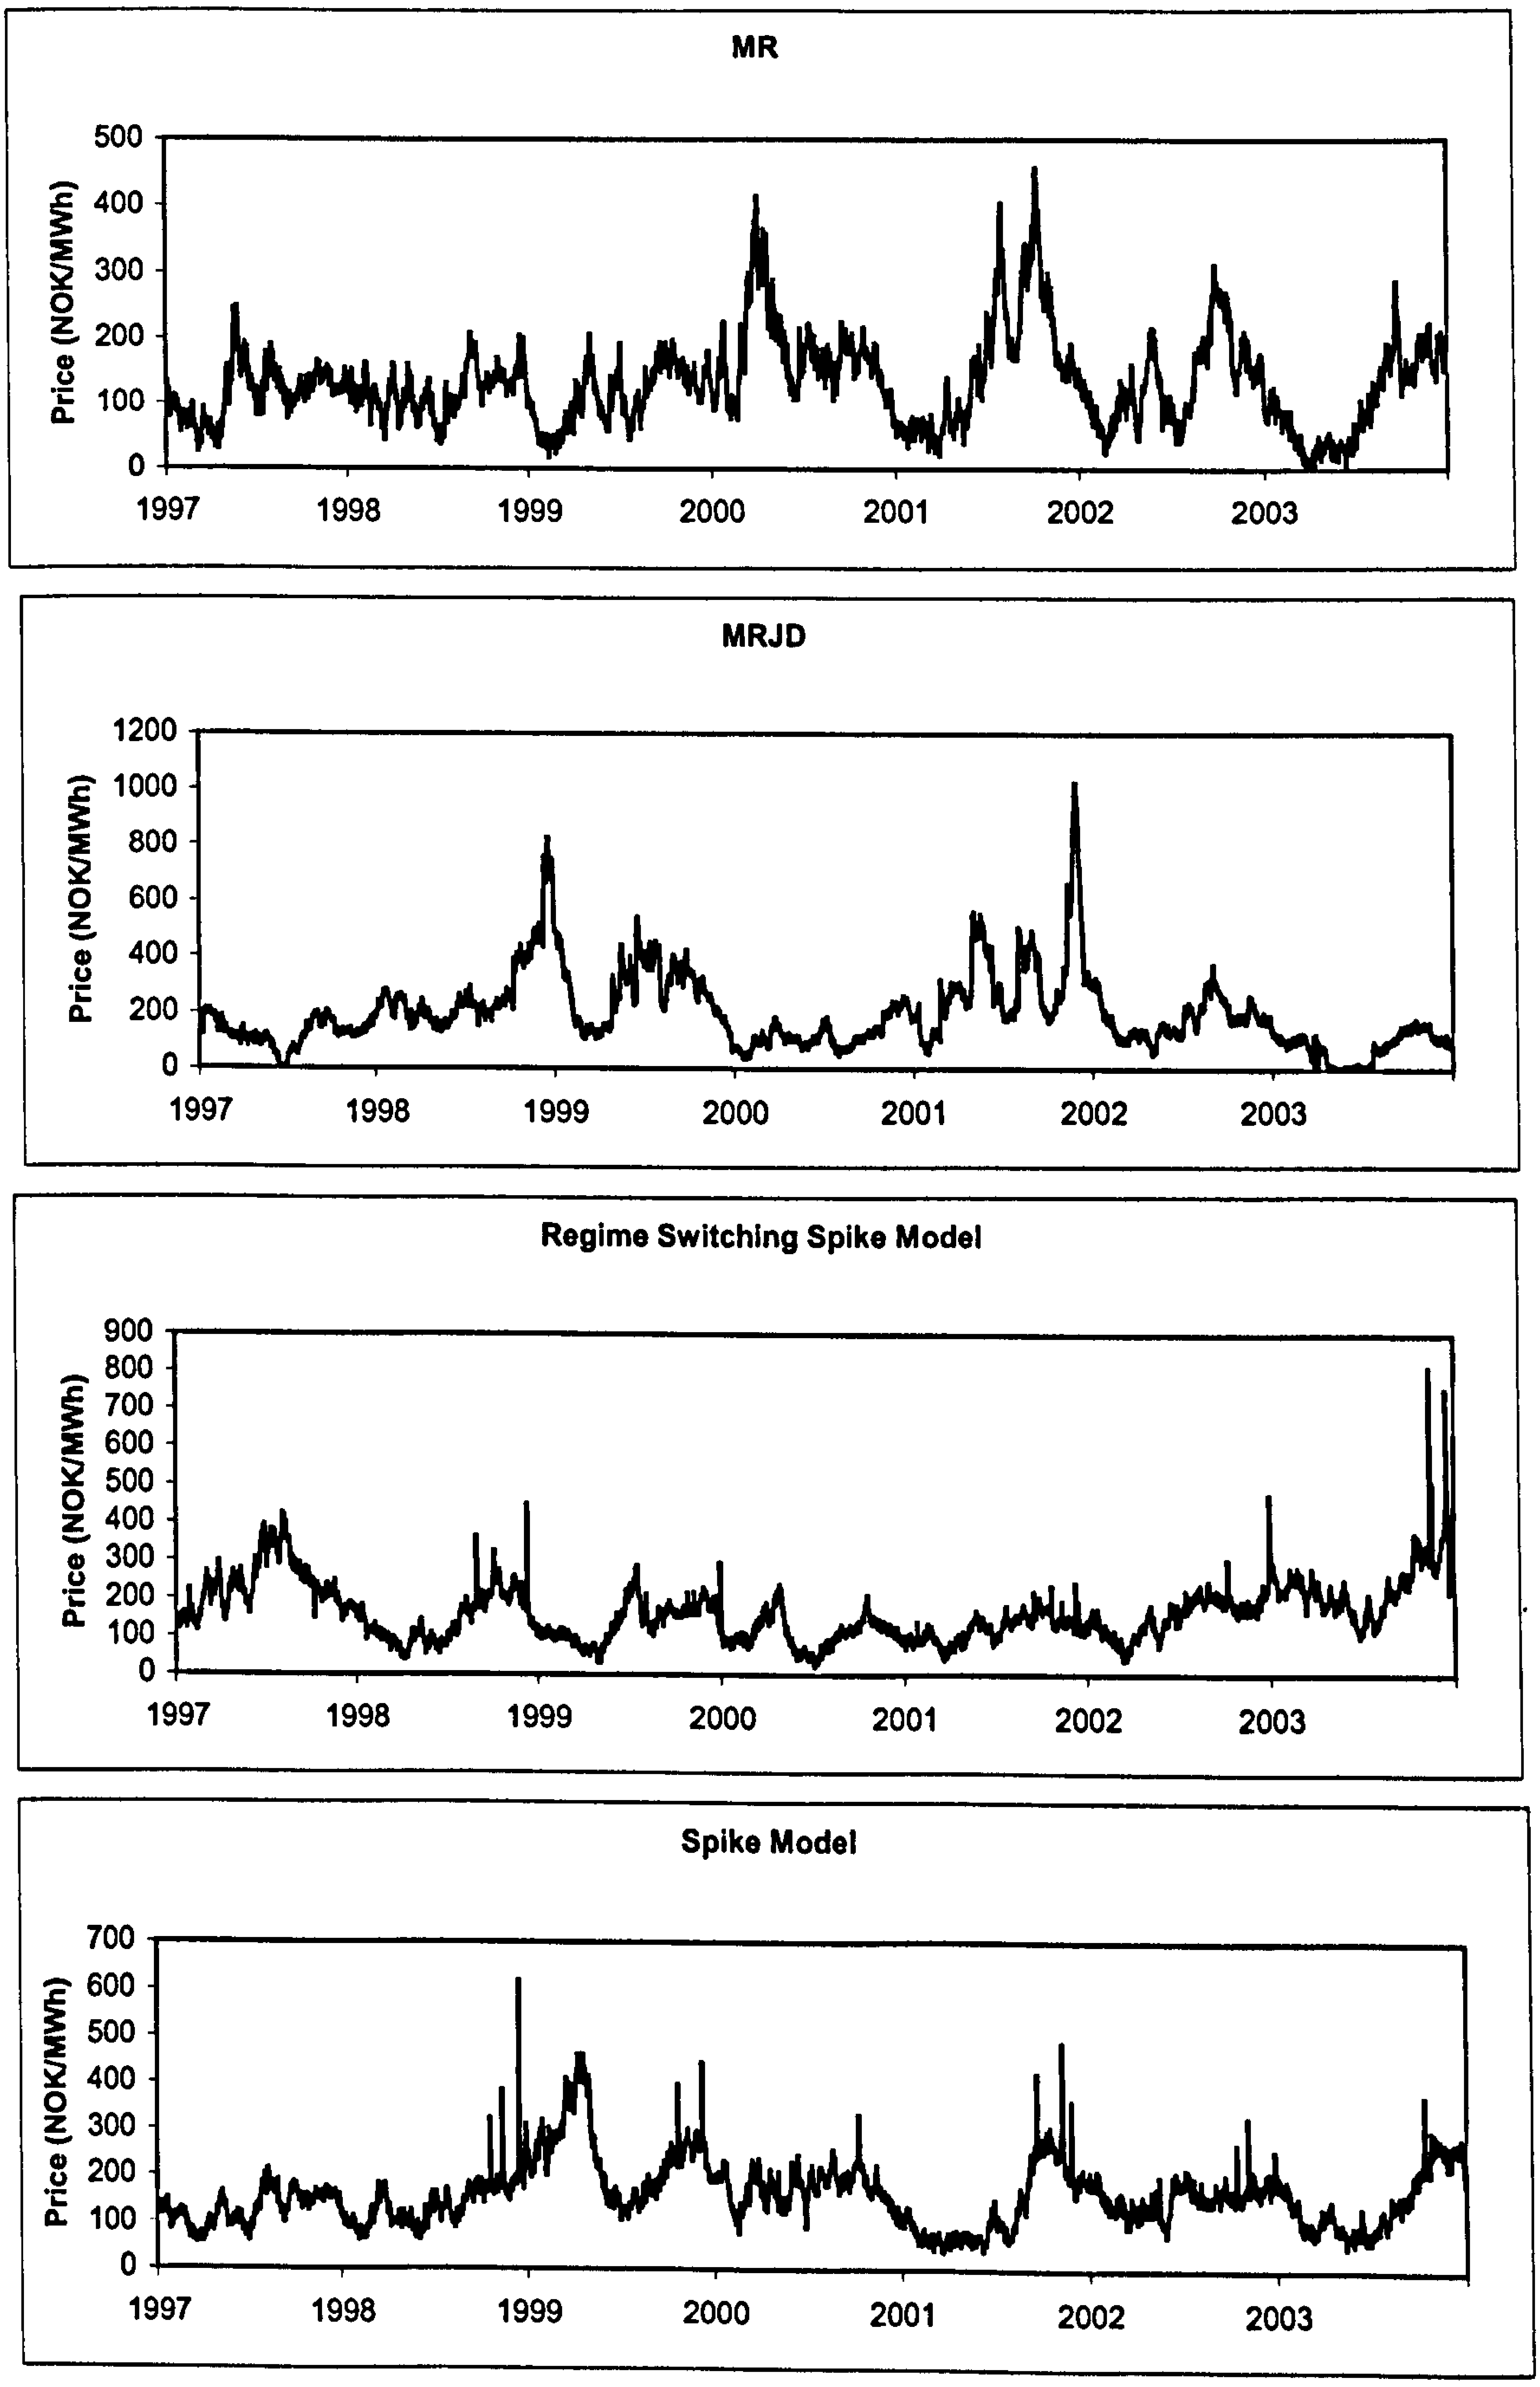

4.4.1 Monte Carlo simulations and the moments of the distribution

Table 4.4 reports the descriptive statistics of the deseasonalised spot price returns, implied by each model and the actual series[25]. In order to get accurate descriptive statistics for all models, we run 60,000 sample paths for each, with 365 time steps each. First thing to note is that all models seem to capture the standard deviation quite well. The MR model does not explain the excess skewness and kurtosis of the actual price, since its implied distribution is normal. On the other hand, all other models capture the excess skewness and kurtosis of the original data. The kurtosis coefficient is slightly better for our models compared to that of the MRJD. However none of the models is able to match exactly the extremely high kurtosis of the actual returns, 66 in Table 4.4. In terms of skewness our models outperform the MRJD, which underestimates the actual skewness considerably. Turing next to the seasonal statistics, our models do a better job in terms of capturing the seasonal distributional characteristics of returns. Compared to the other models, both kurtosis and skewness is much more closer and accurate to that of the actual series. This is expected since in the MRJD model no seasonality on the jumps is induced. Finally our models seem to capture seasonal volatility quite accurately.[26]

The Spike and the Regime-Switching models seem to capture the significant characteristics of the spot prices in an intuitive and more parsimonious way. From all the above results, it seems that the Spike and Regime-Switching Spike models explain the spikes, and especially their very fast mean reversion, quite well. The MRJD on the other hand, is able to explain general jumps in the market, but cannot distinguish between spikes, which are a transitory phenomenon, and long-term shifts in prices. This has important implications for derivatives pricing especially for path dependent options, such as Asian options, because it implies that if for example an upward jump occurs, prices will remain at high levels for some time, making the average price much higher than what is actually the case in the market. Our models however adjust for this, by making prices to revert towards their normal levels at a much faster rate.

---

[25] The results for the Regime Switching spike model are not reported, as they are identical to the Spike model, which is expected since the only difference between the two is in the long run level of the equilibrium to which prices revert. This however does not affect the daily returns levels, which should be similar for both models.

[26] Note that we also ran simulations for log-normally and exponentially distributed jump sizes. However the results showed that there was a high probability that prices would reach levels of more than 1500 NOK/MWh, which has never been observed in the market. Hence care has to be taken when using such distributions, since they have to be truncated (see for instance Weron, 2005).

Table 4.4: Actual Vs Simulated descriptive Statistics

The table reports the descriptive statistics for the actual deseasonalised spot returns, and those implied by the simulations of each model: Spike, Mean reversion Jump Diffusion (MRJD,) and Mean reversion (MR). Panel A reports the results for the entire sample, Panel B for the months of December, January and February and Panel C for the remaining seasons.

	Spike	MRJD	MR	Actual
Panel A: Whole Sample				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.09	0.09	0.09	0.09
Skewness	2.96	1.00	0.00	2.18
Kurtosis	49.0	40.0	3.00	66.0
Panel B: Winter				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.14	0.09	0.09	0.14
Skewness	2.91	1.00	0.00	1.95
Kurtosis	31.0	40.0	3.0	34.0
Panel C: Spring, Summer & Autumn				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.06	0.09	0.09	0.06
Skewness	0.38	1.00	0.00	0.25
Kurtosis	8.54	40.0	3.00	8.64

4.4.2 Closed-form Solutions of the Moments

In the previous section we analysed the moments generated by each model using an Euler discretisation method of the Stochastic Differential Equations for the Monte Carlo Simulations. However, as shown by Das (2001), one can obtain closed-form solutions for the moments for any jump distribution, by first deriving the characteristic function $F(X,Y,t,T;s)$, given that the jump intensity or jump distribution do not depend on the state variables. The solution of the characteristic function can be found using the DPS transform as follows:

$$F(X,T,t,T=t;s) = exp(is(X_t+Y_t)),$$

where $i=\sqrt{-1}$. The above equation is also the boundary condition for the characteristic function. Therefore, the relation between the characteristic function and the transform is given by:

$$\psi^{X}(u,X_{t},t,T)=e^{-r\tau}F\left(X,Y,T;s\right) \tag{4.16}$$

where $X=(X,Y)$ and $u=(is,\ is)$. Therefore the solution of the characteristic function takes the form of the transform and with the above mentioned boundary conditions:

$$F\left(X,Y,t,T;s\right)=\exp(\alpha(t)+\beta_1(t)\bullet X+\beta_2(t)\bullet Y) \tag{4.17}$$

with boundary conditions $a(T)=0$, $\beta_1(T)=\beta_2(T)=is$. Thus, using the results from the transform analysis, and letting $\tau=T\text{-}t$:

$$\frac{\partial\beta_i}{\partial t}=k_i\beta_i \text{ and } \beta_i(T)=is \Rightarrow \beta_i(t)=ise^{-k_i\tau}$$

$$\text{Hence } \beta_1(t)=ise^{-k_1\tau},\ \beta_2(t)=ise^{-k_2\tau}$$

$$a(t)=r\tau+\left(is\right)^2\frac{1}{2}\frac{\sigma_X^2}{2k_1}\left(1\text{-}e^{-2k_1\tau}\right)+is\varepsilon(1\text{-}e^{-k_1\tau})+l\int E\left(\exp(isJe^{-k_2\tau})-1\right)d\tau$$

Now in order to find closed-form solutions for the moments, we need to differentiate the characteristic function successively with respect to s and then find the value of the derivative when $s=0$. Thus if we denote the *nth* moment by m_n, and the *nth* derivative of the characteristic function respectively, by $F_n=\partial^n F/\partial s^n$, then the moments can be derived by $m_n=1/i^n\left[F_n|s=0\right]$. However, since the characteristic function F, depends on α and β, we first have to derive their derivatives with respect to s. First of all the derivatives of β with respect to s are as follows:

$$\frac{d\beta_i}{ds}=ie^{-k_i\tau},\quad \frac{d^2\beta_i}{ds^2}=\frac{d^3\beta_i}{ds^3}=\frac{d^4\beta_i}{ds^4}=0 \tag{4.18}$$

Like wise for α:

$$\frac{1}{i}\left(\frac{d\alpha}{ds}\right)_{s=0}=\varepsilon\left(1-e^{-k_1\tau}\right)+\frac{l}{k_2}E(J)\left(1-e^{-k_2\tau}\right)$$

$$\frac{1}{i^2}\left(\frac{d^2\alpha}{ds^2}\right)_{s=0}=\frac{\sigma_x^2}{2k_1}\left(1-e^{-2k_1\tau}\right)+\frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right)$$

$$\frac{1}{i^3}\left(\frac{d^3\alpha}{ds^3}\right)_{s=0} = \frac{l}{3k_2}E\left(J^3\right)\left(1-e^{-3k_2\tau}\right)$$
$$\frac{1}{i^4}\left(\frac{d^4\alpha}{ds^4}\right)_{s=0} = \frac{l}{4k_2}E\left(J^4\right)\left(1-e^{-4k_2\tau}\right) \tag{4.19}$$

Having derived the above derivatives, and using a lengthy and tedious algebraic manipulation, the moments are given as follows:

$$m_1 = \frac{1}{i}\left[F_1 \middle| s=0\right] = X_t e^{-k_1\tau} + Y_t e^{-k_2\tau} + \varepsilon\left(1-e^{-k_1\tau}\right) + \frac{l}{k_2}E(J)\left(1-e^{-k_2\tau}\right) \tag{4.20}$$

$$m_2 = \frac{1}{i^2}\left[F_2 \middle| s=0\right] = \frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1\tau}\right) + \frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right) + m_1^2 \tag{4.21}$$

$$m_3 = \frac{1}{i^3}\left[F_3 \middle| s=0\right] = \frac{l}{3k_2}E\left(J^3\right)\left(1-e^{-3k_2\tau}\right) + m_1^3$$
$$+3m_1\left(\frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1\tau}\right) + \frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right)\right) \tag{4.22}$$

$$m_4 = \frac{1}{i^4}\left[F_4 \middle| s=0\right] = \frac{l}{4k_2}E\left(J^4\right)\left(1-e^{-4k_2\tau}\right) + 4m_1\left(\frac{l}{3k_2}E\left(J^3\right)\left(1-e^{-3k_2\tau}\right)\right)$$
$$+3\left(\frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1\tau}\right) + \frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right)\right)^2 \tag{4.23}$$
$$+6m_1^2\left(\frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1\tau}\right) + \frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right)\right) + m_1^4$$

where, assuming that the jump size distribution is normal, i.e. $N(\mu_J, \sigma_J^2)$:

$E(J) = \mu_J$, $E\left(J^2\right) = \mu_J^2 + \sigma_J^2$, $E\left(J^3\right) = \mu_J^3 + 3\mu_J\sigma_J^2$ and $E\left(J^4\right) = \mu_J^4 + 6\mu_J^2\sigma_J^2 + 3\sigma_J^4$

Having derived the moments, one can show explicit solutions for the variance, skewness and Kurtosis of the deseasonalised spike process:

$$\text{Variance: } m_2 - m_1^2 = \frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1\tau}\right) + \frac{l}{2k_2}E\left(J^2\right)\left(1-e^{-2k_2\tau}\right) \tag{4.24}$$

Skewness: $\frac{E\left[(z-m_1)^3\right]}{Variance^{3/2}} = \frac{\frac{l}{3k_2}E(J^3)\left(1-e^{-3k_2\tau}\right)}{Variance^{3/2}}$ (4.25)

Kurtosis: $\frac{E\left[(z-m_1)^4\right]}{Variance^2} = 3 + \frac{\frac{l}{4k_2}E(J^4)\left(1-e^{-4k_2\tau}\right)}{Variance^2}$ (4.26)

Likewise for the MR model we can derive the moments and thus find closed-form solutions for the Variance, Skewness and Kurotisis:

Variance: $m_2 - m_1^2 = \frac{\sigma_x^2}{2k_1}\left(1-e^{-2k_1\tau}\right)$ (4.27)

Skewness: $\frac{E\left[(z-m_1)^3\right]}{Variance^{3/2}} = 0$ (4.28)

Kurtosis: $\frac{E\left[(z-m_1)^4\right]}{Variance^2} = 3$ (4.29)

Thus comparing the moments of the spike model against the MR model, the main difference lies in the spike factor Y, which is present in the spike model but not in the MR. Therefore the distribution of MR process is consistent with the normal distribution. However the spike model displays excess Skewness and Kurtosis, whose sign and magnitude depend on the jump parameters l, μ_J and σ_J. This is consistent of course with the results found in the simulation section. For the Mean reversion Jump diffusion, the moments have the same functional form as that of the moments of the spike model; however, this time k_1 and k_2 are equal.

Table 4.5: Actual Vs Model specific descriptive Statistics

The table reports the statistics for the actual deseasonalised spot returns, and those given by the closed-form solutions for each model: Spike, Mean reversion Jump Diffusion (MRJD), and Mean Reversion (MR). Panel A reports the results for the entire sample, Panel B for the months of December, January and February and Panel C for the remaining seasons.

	Spike	MRJD	MR	Actual
Panel A: Whole Sample				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.07	0.09	0.09	0.09
Skewness	3.32	1.00	0.00	2.18
Kurtosis	39.2	40.0	3.00	66.0
Panel B: Winter				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.09	0.09	0.09	0.14
Skewness	1.44	1.00	0.00	1.95
Kurtosis	39.6	40.0	3.0	34.0
Panel C: Spring, Summer & Autumn				
Mean	0.00	0.00	0.00	0.00
St. Deviation	0.06	0.09	0.9	0.06
Skewness	0.15	1.00	0.00	0.25
Kurtosis	5.05	40.0	3.00	8.64

As it was also done in Das (2001), we can compute the moment of the conditional distribution of the deaseasonalised log-spot prices using the estimated parameters for the proposed models. Since our data is daily, the horizon τ is 1/365, given the number of trading days in a year in Nord Pool 's spot market. In order to make a rough comparison as stated by Das (2001), the moments of the changes in the deseasonalised spot prices will correspond to the computed moments at $\tau=1/365$. Table 4.5 reports the descriptive statistics of the deseasonalised spot returns, using the closed-form solutions for each model and the actual. First thing to note is that all models except the spike model, capture the whole sample standard deviation accurately. On the other hand, focusing on the different seasons the spike model fits the standard deviation for the rest of the seasons accurately, but for winter its performance is the same as that of the other models. Why is this happening though? Looking at equation (4.24), we can notice that increasing the speed of mean reversion, for either X or Y, decreases the variance and thus the standard deviation. This directly implies that the high coefficient of the speed of mean reversion of spikes decreases the overall volatility of the spot returns, given the spike model. On the other hand the MRJD model imposes the same mean reversion for both normal and jump shocks, and its coefficient is low; thus the theoretical standard devation implied by the MRJD model is anticipated to be higher than that of the spike model, other things being equal.

The other difference between the two models comes from the definition of jumps and spikes between the two models. In the MRJD model, a jump is defined as any return whose absolute value is greater than 3 standard deviations of all returns. On the other hand in the spike model a jump is any return whose absolute value is greater than three standard deviations of all returns, followed by a reduction or extreme returns of the opposite sign which account for at least half of the original movement within at most 5 days. Therefore, in the spike model if for example we have a positive (negative) jump in the actual series and the next time step we have a negative (positive) jump, the first jump is defined as an actual jump and the second is captured as part of the spike mean reversion. On the other hand, in the MRJD model, both of these shocks will be defined as jumps and they will be taken into account to estimate the jump intensity, l. Therefore we should expect the intensity in the MRJD model to be higher than in the spike model.

Therefore, combining both of the above arguments, since many jumps are captured as part of the spike mean reversion in the spike model and also since the speed of mean reversion is inversely related to the variance, we have this inconsistency. To show how standard deviation, skewness and kurtosis are affected by the spike speed of mean reversion, we have performed a sensitivity analysis, by finding the values of the three moments for different spike speeds of mean reversion, as shown in Figure 4.5. From the analysis it can be clearly seen that by increasing the spike speed of mean reversion, k_2, the values of the descriptive statistics, i.e. standard deviation, skewness and kurtosis, decay exponentially.

Another issue we are facing is the fact that for the spike model, the descriptive statistics given form the closed-form solutions, do not match with those from the simulated state variables, whereas for the MR and MRJD they do. In the simulation example for the spike model, the discretization time step was one day (i.e. $dt=1/365$). Thus using for instance the intensity parameter during winter, l_1, the chance of a jump occurring during one time step is $ldt=1/365*12.17=0.0333$. Thus our simulation is correct up to this point since the probability of one jump is less than one. From the econometric point of view, this simulation approach is correct since the simulation uses daily time steps, which is consistent with the frequency of the data and the estimates for the parameters. Therefore, on this basis we can use the simulations and compare the descriptive statistics generated from the estimates of the models, against those displayed from the data.

However from the continuous time point of view the simulations are not correct. This is mainly caused from the parameter of the speed of mean reversion, which as shown by Clewlow and Strickland (2000), in order for the simulation to approach the continuous distribution of the process, the time step has to be a small fraction of the half-life. In our case the half-life of a spike is 0.88 days, thus the time step has to be less than a day, and therefore we have this inconsistency between the descriptive statistics generated from the simulation using daily time steps and the continuous time formulas for the moments. However simulation experiments showed that our Monte Carlo estimates converged to the continuous time descriptive statistics, when we took hourly time steps, however the paths generated when using hourly or daily time steps are qualitatively similar to those presented in Figure 4.4. A more detailed description of the appropriate Monte Carlo Simulation is shown in the next chapter, where it is used to price Asian Option. As it will be shown there, when a process involves jumps it is not possible to get an "exact" discretisation.

Figure 4.5: Descriptive statistics Vs Spike Mean Reversion k_2

The following figures show how the theoretical standard deviation, skewness and kurtosis (using equations (4.24), (4.25) and (4.26) respectively) change when increasing the speed of mean reversion of spikes.

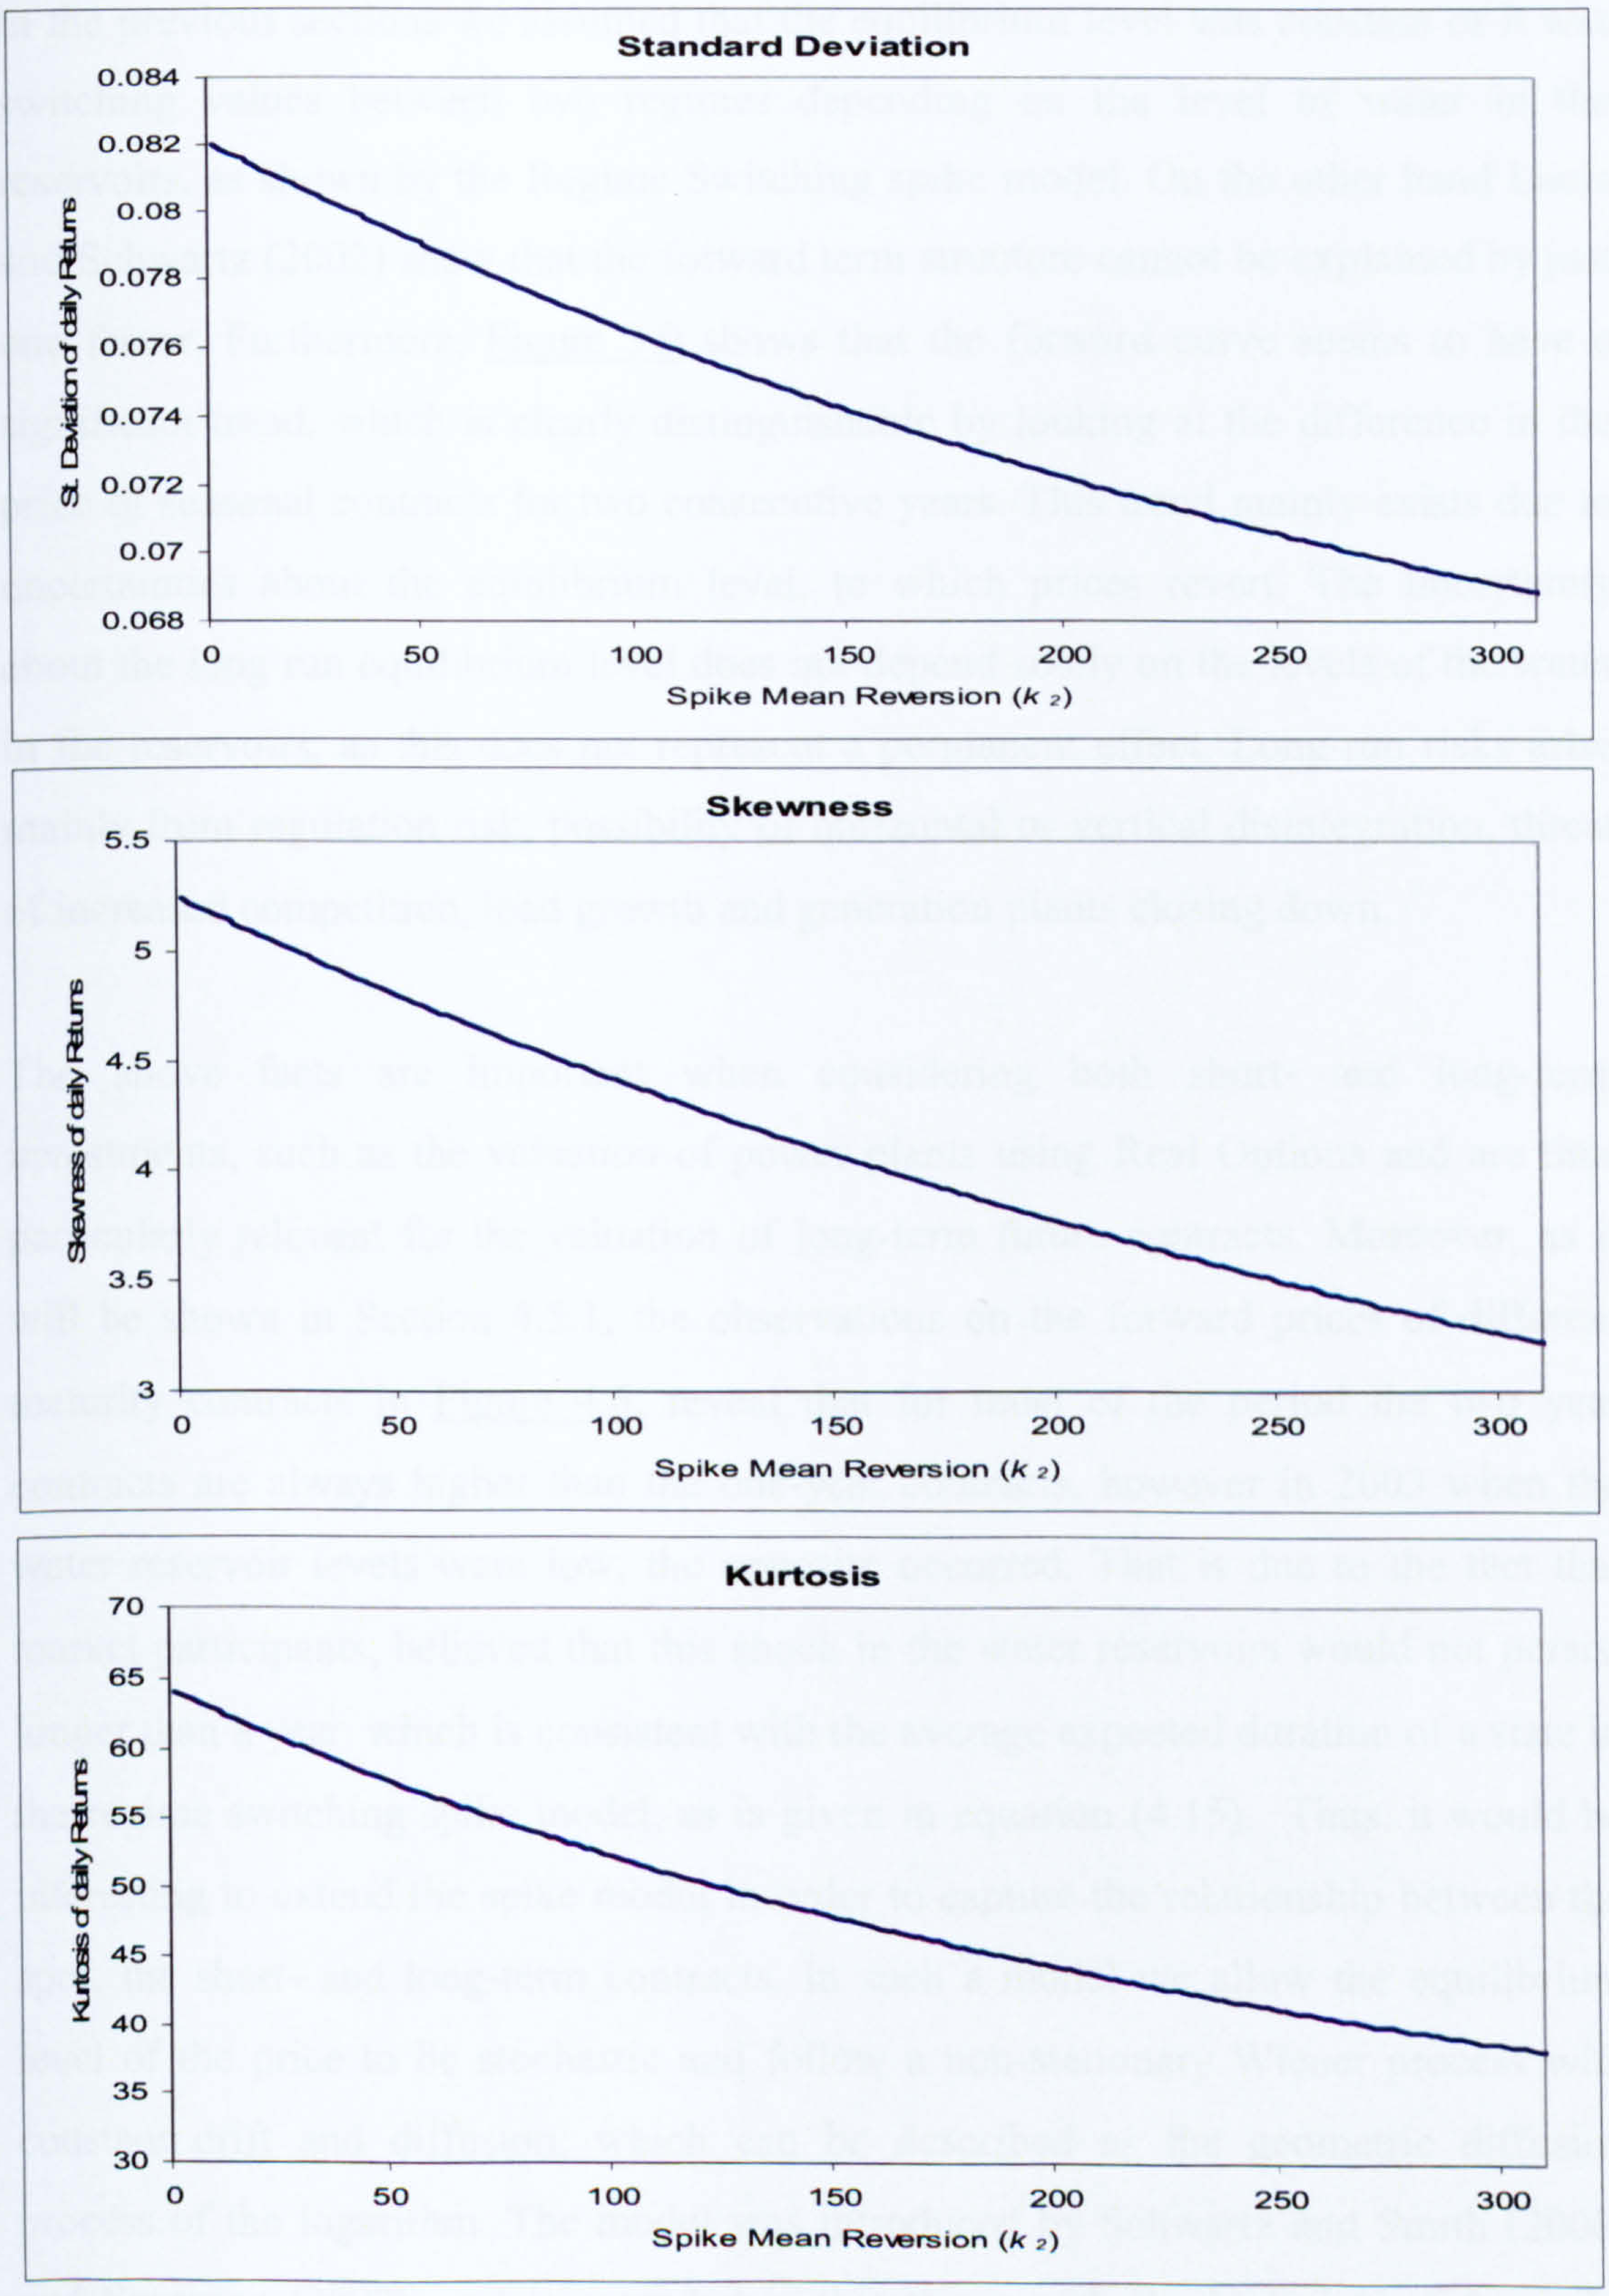

4.5 The Three-factor Model

In the previous sections we assumed that the equilibrium level was constant or it was switching values between two regimes depending on the level of water in the reservoirs, as shown by the Regime Switching spike model. On the other hand Lucia and Schwartz (2002) show that the forward term structure cannot be explained by just one factor. Furthermore, Figure 3.9 shows that the forward curve seems to have a significant trend, which is clearly distinguishable by looking at the difference in the price of seasonal contracts for two consecutive years. This trend mainly exists due to uncertainties about the equilibrium level, to which prices revert. The uncertainty about the long run equilibrium level does not depend solely on the levels of the water in the reservoirs, as this does not represent a permanent effect. Long-run risks arise mainly from regulation risk, possibility of horizontal or vertical disintegration, threat of increased competition, load growth and generation plants closing down.

The above facts are important when considering both short- and long-term investments, such as the valuation of power plants using Real Options and are thus particularly relevant for the valuation of long-term future contracts. Moreover, as it will be shown in Section 4.5.1, the observations on the forward prices of different maturity contracts in Figure 4.6, reveal that for most of the period the two year contracts are always higher than the one-year contracts, however in 2003 when the water reservoir levels were low, the opposite occurred. That is due to the fact that market participants, believed that this shock in the water reservoirs would not persist longer than a year, which is consistent with the average expected duration of a state in the regime switching spike model, as is given in equation (4.15). Thus, it would be interesting to extend the spike model in order to capture the relationship between the spot, the short- and long-term contracts. In such a model we allow the equilibrium level of the price to be stochastic and follow a non-stationary Wiener process with constant drift and diffusion, which can be described as the geometric diffusion process of the logarithm. The model was introduced by Schwartz and Smith (2000) and Sørensen (2002), and is extended in this thesis with the inclusion of the spike component.

We start the analysis by specifying our model. First in equation (4.30) we model prices (P_t) as the sum of a predictable component ($f(t)$) and the exponential sum of three stochastic components. The predictable component takes into account the regularities in the evolution of prices. The first stochastic component (X_t) is assumed to follow a stationary process as in Vasicek (1977), in order to account for the short-term deviations resulting from the difference between the spot and the equilibrium level due to short-term changes in demand resulting from variations in weather or intermittent supply disruptions as well as due to the market behaviour. The second stochastic component (Y_t) represents the spike component and is assumed to follow a mean-reverting process whose shocks are modelled by a compound Poisson process. Thus up until now the first two factors follow exactly the same processes as the ones in the spike and regime switching spike model. However changes in the equilibrium level, ε_t, represent fundamental changes that are expected to persist for the reasons discussed, and thus follows a non-stationary process with constant drift and diffusion parameters:

$$
\begin{aligned}
P_t &= f(t) + \exp(X_t + Y_t + \varepsilon_t) \\
dX_t &= -k_1 X_t dt + \sigma_X dW_X \\
d\varepsilon_t &= \left(\mu_\varepsilon - \frac{1}{2}\sigma_\varepsilon^2\right) dt + \sigma_\varepsilon dW_\varepsilon \\
dY_t &= -k_2 Y_t dt + J(\mu_{J_t}, \sigma_{J_t}^2) dq(l_t)
\end{aligned}
\tag{4.30}
$$

where the Wiener processes dW_X and dW_ε are assumed to be correlated with a constant correlation coefficient ρ.

As before, in equation (4.30), k_1 represents the speed at which X reverts to zero after a shock has occurred. The dW_X represents the increments of the Brownian motion that cause the random shocks in the short-term factor, and it is scaled up by the volatility factor σ_X. For the long-term dynamics, ε is the long-run equilibrium level whose logarithm is assumed to follow a Geometric Brownian motion, with trend μ_ε and volatility σ_ε. As in the spike and the regime switching spike model, k_2 is the speed of mean reversion of the spikes whose shocks are modelled via a compound Poisson process with seasonal intensity l_t. J is a random variable that accounts for the size of the jump, and is assumed to follow a Normal distribution with seasonal mean μ_{J_i} and variance $\sigma^2_{J_i}$.

As shown before in order to price any kind of derivatives, the model should not allow for arbitrage opportunities. Therefore following Harrison and Kreps (1979) and Cox, Ingresoll and Ross (1981) under risk-neutrality the processes become as follows:

$$\begin{aligned} P_t &= f(t) + \exp\left(X_t + \varepsilon_t + Y_t\right) \\ dX_t &= -(\lambda_X + k_1 X_t)dt + \sigma_X dW_X^* \\ d\varepsilon_t &= \left(\mu_\varepsilon^* - \frac{1}{2}\sigma_\varepsilon^2\right)dt + \sigma_\varepsilon dW_\varepsilon^* \qquad \mu_\varepsilon^* = \mu_\varepsilon - \lambda_\varepsilon \\ dY_t &= -k_2 Y_t dt + J(\mu_{J_i}, \sigma^2_{J_i})dq(l_t) \end{aligned} \tag{4.31}$$

Using the results from Appendix 4.8, we can write analytic forms for the distributions of the state variables and spot prices. Since our model is in line with an Affine Jump-Diffusion process, we apply the elementary transform function of Duffie, Pan and Singleton (2000). Hence equation (4.32) shows the expected risk-neutral value of the spot price in the future implied by the model.

$$E_0^*(P_T) = f(T) + \exp\left(e^{-k_1T}X_0 + \varepsilon_0 + Y_0 e^{-k_2T} + \sum_{i=1}^{2} l_i \int_0^T 1_{s\in i}\left(\exp\left(\mu_{J_i}e^{-k_2s} + \frac{1}{2}\sigma^2_{J_i}e^{-2k_2s}\right) - 1\right)ds + A_T\right)$$

$$A_T = \frac{1}{2}\frac{\sigma_X^2}{2k_1}\left(1 - e^{-2k_1T}\right) - \frac{\lambda_X - \rho\sigma_\varepsilon\sigma_X}{k_1}(1 - e^{-k_1T}) + \mu_\varepsilon^* T$$

(4.32)

Using the discussion in the previous section equation (4.32) is the theoretical price of a forward maturing at time *T*. As before several points can be noted about the importance of the state variables to the forward price in equation (4.32). First, as the time to maturity increases, the impact of the state variables *Y* and *X* decreases, and forward price relies more on the long-term equilibrium level. Second, the significance of the state variables in the forward price depends on the mean reversion factor, which acts as a discounting factor. Also this fact shows that a jump in the price, and hence an increase/decrease in the state variable *Y*, will have no significant effect in the forward price, as long as k_2 is very high. This is in line with several empirical findings that a spike in the spot prices, does not translate as a spike in the forward price, since this shock is very short lived. Moreover the higher value of k_2 implies that the value for short-term contracts is mainly governed by the jump component, whereas for longer-term contracts, the premium is governed by the sum of all the risk factors. Finally equation (4.32) shows that in the long run as $t \to \infty$, the spot price does not converge to a constant value, but rather grows at a rate specified by the drift of the long-term factor:

$$f(T) + \exp\left(\frac{1}{2}\frac{\sigma_x^2}{2k_1} - \frac{\lambda_x}{k_1} + \sum_{i=1}^{2} l_i \int_0^T 1_{s \in i}\left(\exp\left(\mu_{J_i} e^{-k_2 s} + \frac{1}{2}\sigma_{J_i}^2 e^{-2k_2 s}\right) - 1\right) ds + \varepsilon_0 + \mu_\varepsilon^* T\right),$$

4.5.1 Model Calibration and Results

4.5.1.1 Data used

For the empirical estimation of the three-factor model, we use the same dataset as in the case of the spike and regime switching spike models. We also use the same constant maturity forward prices, as in the previous sections. Another advantage from the use of this dataset is that it reduces the computational complexity introduced by equation (4.9) for the theoretical forward price, especially for the Kalman filter where we have averages of exponential terms, therefore the state variables do not have a linear relationship with the log-forward prices. Thus the use of constant maturity contracts with no delivery period makes the calibration of our models easier.

The forwards used in this part of the analysis reflect short- to mid-term maturities (1, 2, 4, 7, 12, 16, 20 and 24 weeks, 1, 2 and 3 years to maturity). These maturities are selected to reflect actual traded contracts in the market. Since the model consists of unobserved state variables, both spot and forward prices have to be used simultaneously. The long maturity contracts were chosen so as to give us the most available information about the long-term factor as shown by Schwartz and Smith (2000). However one factor that may affect our results is the fact that the contracts with 1, 2, and 3 years to maturity have a low liquidity, whereas liquidity is high for the first weekly contracts, and the nearest block (month) and seasonal (Quarter) contracts. Figure 4.6 displays the time series of three forward prices used. As already discussed, when compared to the 1-week to maturity contracts, the long-term contracts (2-year) behave in a smoother fashion and seem to not be affected by short-term fluctuations at all.

Figure 4.6: Time series of Forward Prices

The figure shows the time series of the weekly observations for the 1 and 2 weeks to maturity contracts representing shot-run forward prices, as well as for 1 and 2 years two maturity contracts representing long-run forward prices, used in the analysis.

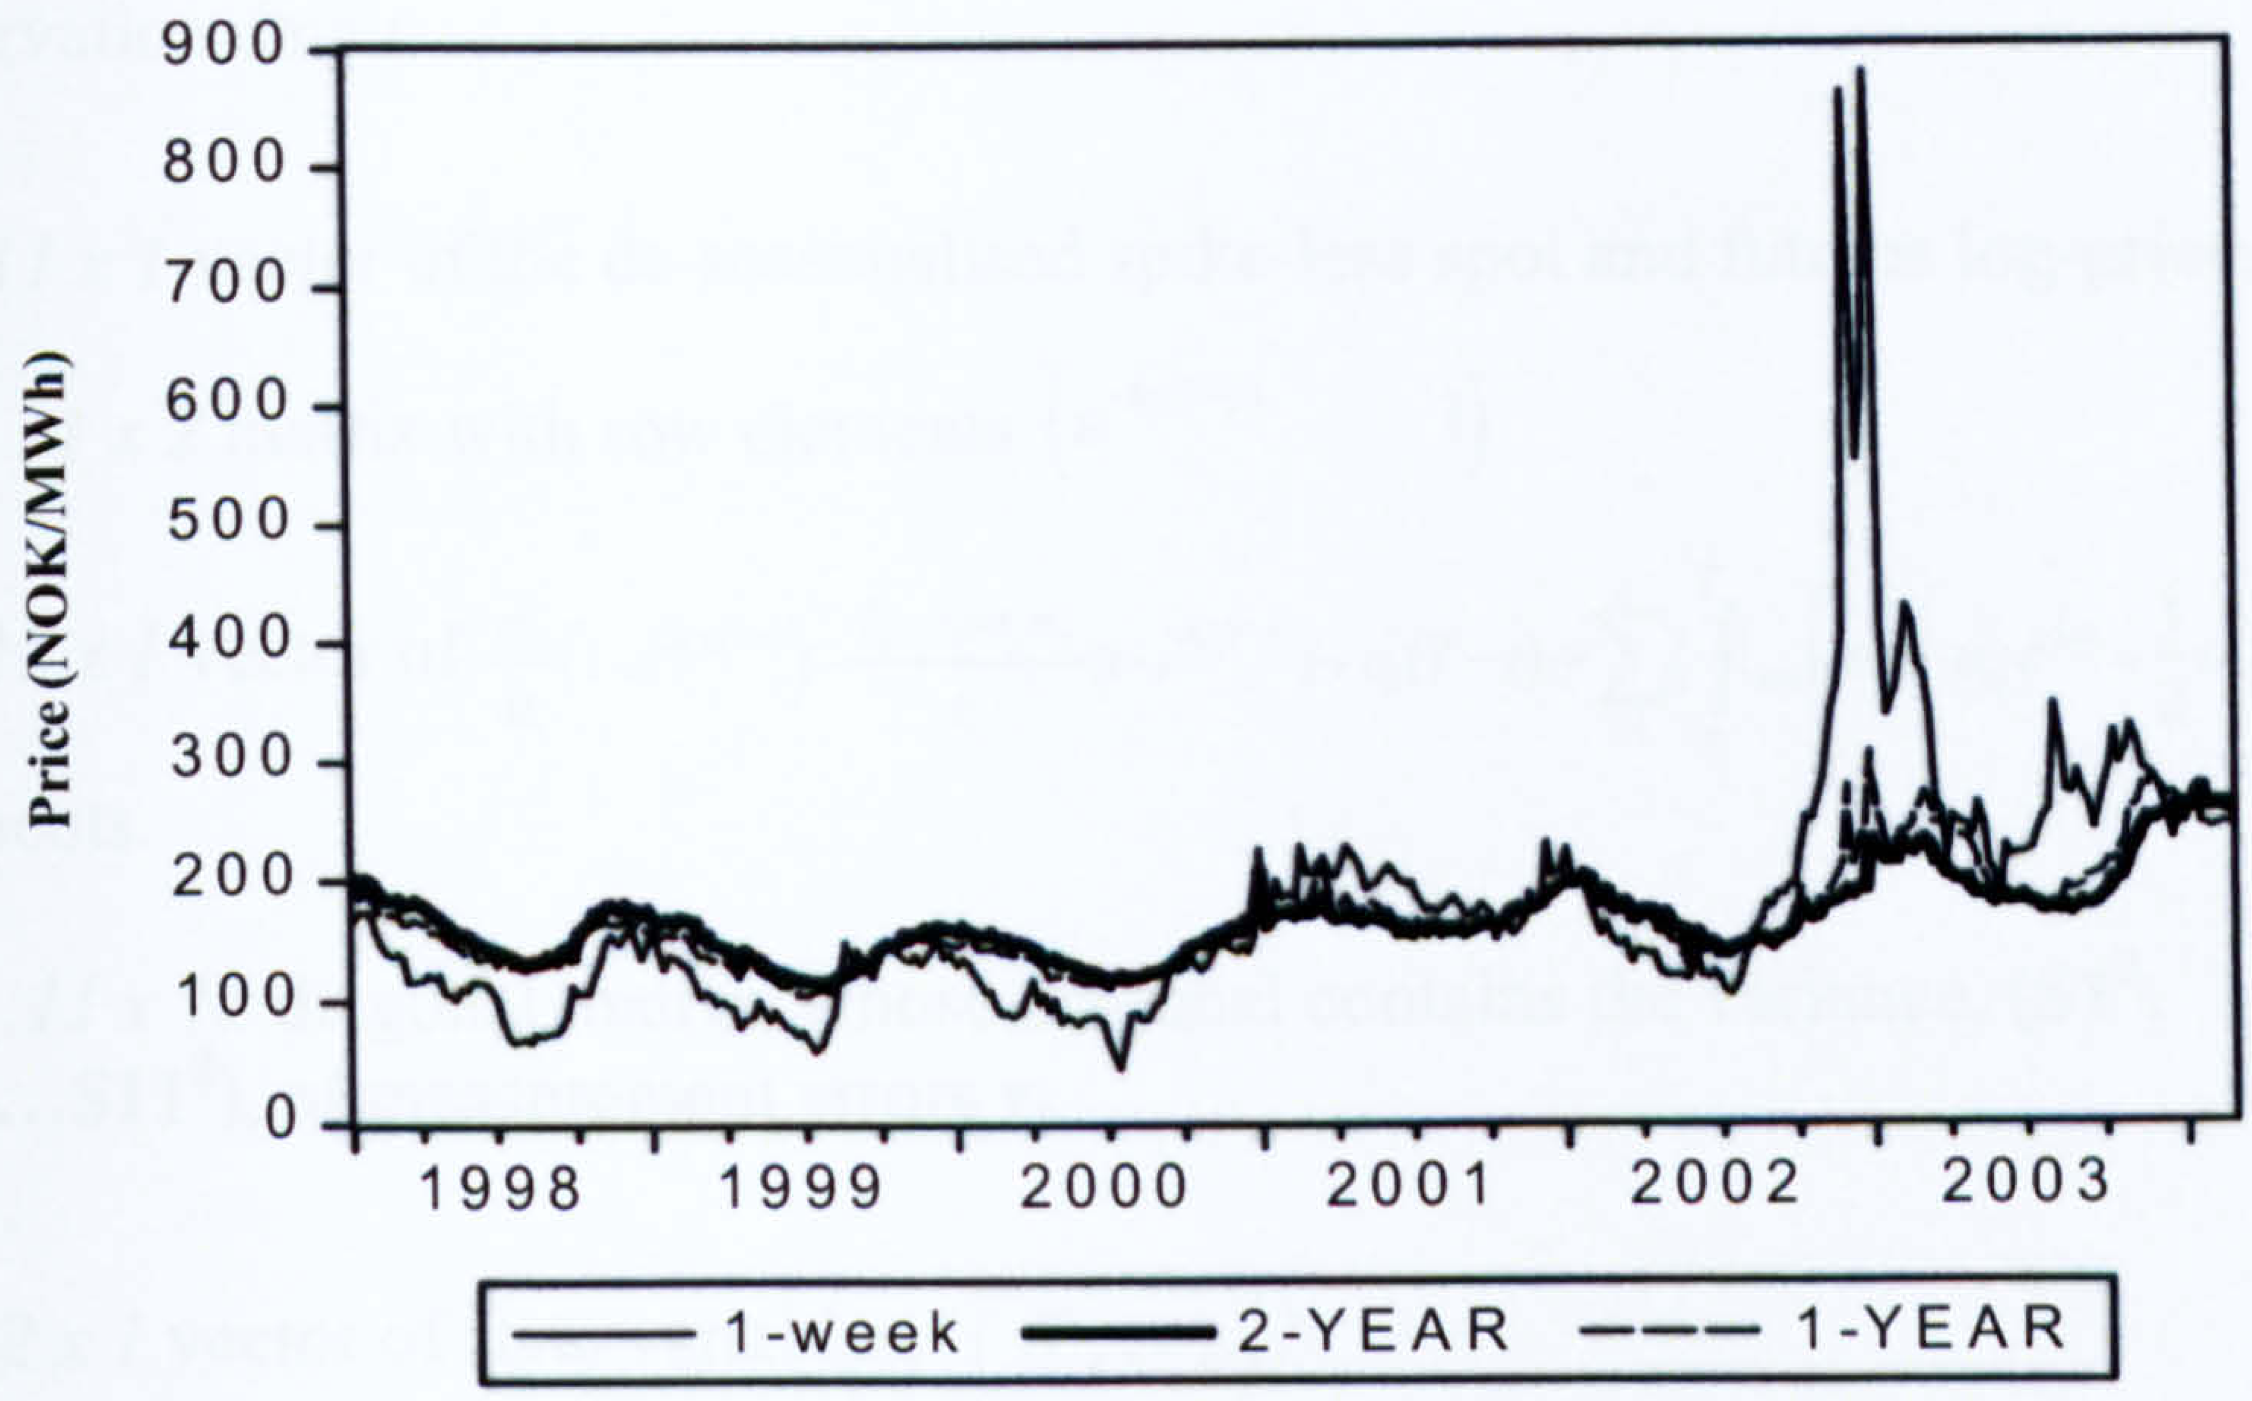

4.5.1.2 Estimation of the Three-factor Model's latent variables and The Kalman Filter:

For the empirical part of the analysis, the estimates of the seasonal and spike component were the same as the ones in the previous section, for the spike and regime switching spike models, respectively. However for the remaining estimates, since the model relies on state variables that are not directly observable, we need a model that "best fits" both spot and forward (futures) prices as well as a filter. Over the years the Kalman filter has been extensively used in finance. In this section we present a general description of the Kalman filter to estimate unobserved state variables. For a more detailed explanation, see for example Harvey (1989), Chapter 3 or Hamilton (1994), Chapter 13.

The Kalman filter may be applied to dynamic models that are in a state-space representation, which include **measurement** and **transition** equations. At each point in time, the **measurement** equation relates a vector of observable variables z_t with a vector of state variables x_t, which in general is not observable:

$$z_t = H_t x_t + d_t + v_t \qquad v_t \sim N(0,R_t) \tag{4.33}$$

Where, given equation (4.32), we can define the following matrices, for each observation date t:

z_t = *11 x 1* vector of the de-seasonalised spike-less spot and futures log-prices at time t

H_t = *11 x 2* matrix with row elements $\left(e^{-k_1(T-t)} \quad 1\right)$

d_t= *11 x 1* vector of $\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1(T-t)}\right)-\frac{\lambda_X-\rho\sigma_\varepsilon\sigma_X}{k_1}\left(1-e^{-k_1(T-t)}\right)+\mu_\varepsilon^*(T-t)+\sum_{i=1}^{2} l_i \int_0^T 1_{s\in i}\left(\exp\left(\mu_{J_i}e^{-k_2 s}+\frac{1}{2}\sigma_{J_i}^2 e^{-2k_2 s}\right)-1\right)ds$ elements.

R_t= *11 x 11* diagonal matrix, whose diagonal contains the variance, (**S1^2, S2^2,…S11^2**), of measurement errors v_t.

x_t= *2 x 1* vector of state variables, $[X_t, \varepsilon_t]'$

The measurement equation assumes the existence of a linear relation between the observed and the state variables. Also, note that the measurement equation contains a disturbance term to allow for measurement errors in the observed data. In other words these measurement errors represent the errors present in the reporting prices or errors in the model's fit to observed prices.

The **transition** equation describes the dynamics of the state variables:

$$x_t = A_t x_{t-1} + c_t + \epsilon_t \qquad \epsilon_t \sim N(0,Q_t) \tag{4.34}$$

Where, given the stochastic differential equations for X and ε in (4.30), we have:

$$c_t \equiv \left[0,\ \mu_\varepsilon - \frac{1}{2}\sigma_\varepsilon^2\Delta t\right],\quad A_t \equiv \begin{pmatrix} e^{-k_1\Delta t} & 0 \\ 0 & 1 \end{pmatrix},\ \text{and}\ \ Q_t = \begin{pmatrix} \frac{\sigma_X^2}{2k_1}(1-e^{-2k_1\Delta t}) & \frac{\rho\sigma_\varepsilon\sigma_X}{k_1}(1-e^{-k_1\Delta t}) \\ \frac{\rho\sigma_\varepsilon\sigma_X}{k_1}(1-e^{-k_1\Delta t}) & \sigma_\varepsilon^2\Delta t \end{pmatrix}$$

Under this representation, the state variables have a multivariate Normal distribution. This assumption can also be relaxed to include non-Gaussian models for the state variables. The measurement and transition equations define what is called the **state space representation.**

The Kalman filter is a recursive procedure for computing the optimal estimator of the state variables at time *t*, based on information at time *t*, and it is continuously updated as new information become available, (Harvey (1989)). The Kalman filter is thus a particular type of Bayesian estimation. Another useful characteristic of the Kalman filter is that it provides consistent model parameter estimates, when maximizing the log-likelihood function of error innovations[27]. Letting F_t denote the conditional variance of z_t conditional on information available at time t-1, the log-likelihood takes the form:

$$\log L(\psi) = -\frac{NT}{2}\log 2\pi - \frac{1}{2}\sum_{t}^{T}\log|F_t| - \frac{1}{2}\sum_{t}^{T}v_t' F_t^{-1} v_t \tag{4.35}$$

where ψ, is the vector of parameters to be estimated.

We start the Kalman filter with a prior mean and covariance matrix based on the observed means and covariance of the data over the entire sample. The estimated state variables and parameters did not appear to be sensitive to different initial values. Moreover since in the electricity market the spot price exists and is reliable and liquid (unlike in the other markets such as oil), and since the main aim of this model is to capture spot dynamics and price contingent claims on the spot, it makes sense to assume a zero measurement error for the spot price. This is implemented by choosing a measurement error covariance matrix with zero variance for the spot price.

4.5.2 Results on the Three-factor model

Table 4.6 shows the parameter estimates of the model. First thing to notice is that although the trend for the long-run factor is significant under the risk neutral world (μ_ε^*), the trend under the objective probability measure is insignificant. Thus we may conclude that although there is always a significant trend in the forward curve, due to

[27] We used the BFGS (Bolgano et al) algorithm to find the optimal estimates of the coefficients that maximise the log-likelihood, in GAUSS 5.0.

the long-run risk faced when trading long-maturity forward contracts in the actual market, there is no significant trend in the spot market. The remaining parameters are significant. Not surprisingly, there is a big difference between the volatility of the short-term factor and that of the long-term factor. The parameter is three times greater than that of the long-term (which is high too). Using this procedure and that kind of data we find that the half-life of the short- term deviations is almost 6 months.

Table 4.6: Annualised Parameter estimates for the 3-Factor Spike model

The table shows the estimated parameters for the three-factor model. (*) and (**) means that the estimate is significant at 5% and 1% significance level respectively.[28]

Three Factor Spike Model			
Parameters of Stochastic Processes		**Measurement Errors**	
μ_ε	0.097 (0.119)	*S1*	0.000
k_1	1.32** (0.041)	*S2*	0.099** (0.004)
k_2	287.61** (52.99)	*S3*	0.117** (0.005)
σ_X	100%** (0.0386)	*S4*	0.153** (0.006)
σ_ε	32.48%** (0.0134)	*S5*	0.195** (0.008)
ρ	-0.42** (0.050)	*S6*	0.242** (0.010)
$l_{(1)}$	12.17	*S7*	0.267** (0.011)
$l_{(2)}$	3.41	*S8*	0.267** (0.011)
$\mu_{(1)}$	41.86%	*S9*	0.267** (0.011)
$\mu_{(2)}$	8.41%	*S10*	0.050** (0.002)
$\sigma_{J(1)}$	38.12%	*S11*	0.000
$\sigma_{J(2)}$	20.40%	*S12*	0.080** (0.0032)
μ_ε^*	0.0623** (0.005)	**Log-Likelihhod** 2798.35	
λ_X	0.5416** (0.062)		

$$P_t = f(t) + \exp(X_t + \varepsilon_t + Y_t)$$

$$f(t) = \beta D_t + \gamma \sin\left((t+\tau)\frac{2\pi}{365}\right) + \delta \sin\left((t+\xi)\frac{4\pi}{365}\right)$$

$$dX_t = -(\lambda_x + k_1 X_t)dt + \sigma_x dW_x^*$$

$$d\varepsilon_t = \left(\mu_\varepsilon^* - \frac{1}{2}\sigma_\varepsilon^2\right)dt + \sigma_\varepsilon dW_\varepsilon^* \qquad \mu_\varepsilon^* = \mu_\varepsilon - \lambda_\varepsilon$$

$$dY_t = -k_2 Y_t dt + J(\mu_{J_i}, \sigma_{J_i}^2)dq(l_i)$$

Where *i=1 for Winter or 2 for Spring, Summer, Autumn*

[28] The deterministic seasonality parameters for *f(t)*, are the same as those for the spike and regime switching spike models shown in Table 4.1.

Figure 4.7: The estimated Equilibrium Price and Spot Price

The figure displays the estimates of the State variables implied by the Kalman Filter as well as the time series of the percentage of water in the reservoirs with its seasonal average. The Equilibrium price = $f(t)+exp(\varepsilon_t)$, where as the Spot price = $f(t)+exp(\varepsilon_t + X_t+Y_t)$.

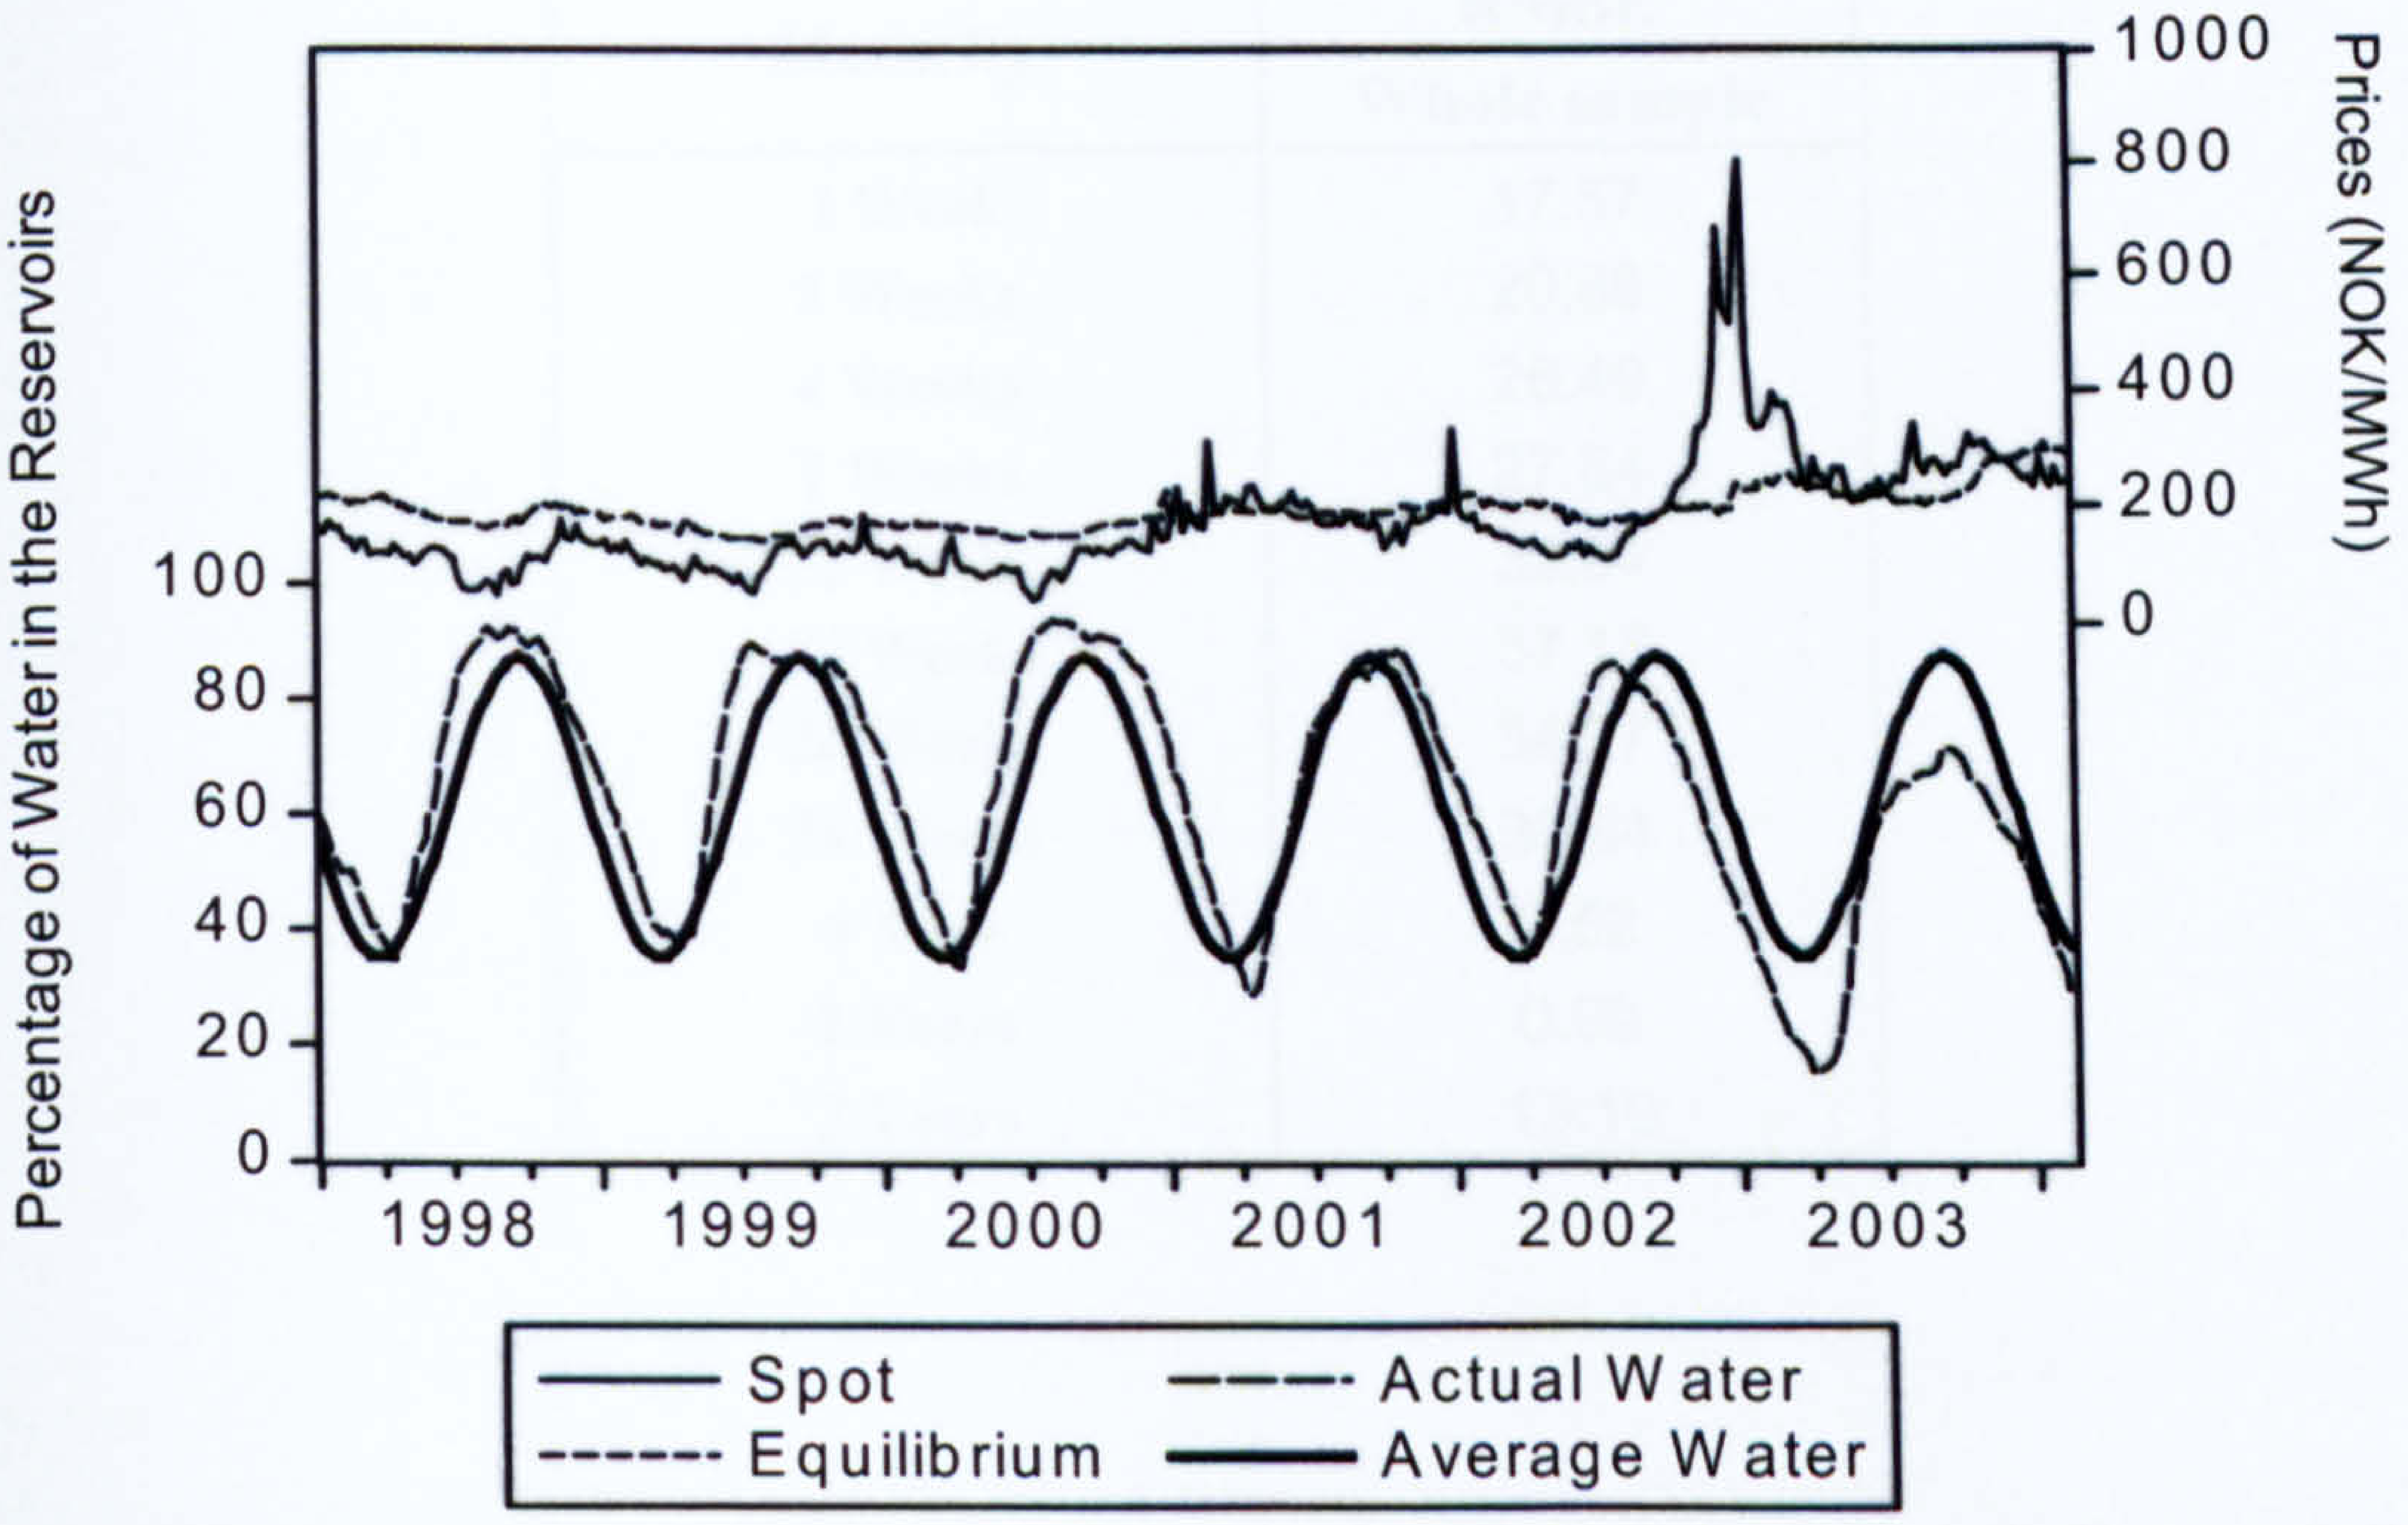

Figure 4.7 shows the estimated time-series of the equilibrium price ($f(t)+\exp(\varepsilon_t)$) and the actual spot price ($P_t = f(t)+\exp(X_t + Y_t + \varepsilon_t)$) at the weekly frequency. The extreme volatility in the spot price is quite distinguishable; however the spot price is sometimes above and sometimes below the equilibrium level. Until 2001, the spot was below the equilibrium due to the high level of water in the reservoirs, providing cheaper energy. In 2001 the water level fell and hence prices increased. At the end of 2002 the water reservoir levels dropped to their lowest level of the past 10 years, hence the prices in Nord Pool increased dramatically. The fact that the equilibrium price did not follow the spot price shows that market participants did not expect this phenomenon to persist. Focusing at the RMSE in Table 4.7, we see that the largest Errors occur for mid-term contracts. Furthermore, the long-term variable seems to give very good fit especially for the prices of 2-year contracts, which is why their measurement error is zero in Table 4.6. Finally, the magnitude of these errors is comparable to those of Lucia and Schwartz (2002).

Table 4.7: Root Mean Square Error for Three-factor model

The table shows the Root of Mean Square Error (RMSE) for the entire sample, between the actual forward price and the theoretical forward price given from the three-factor model, for contracts of different maturities.

Maturity	RMSE
	Whole sample
1 Week	17.57
2 Weeks	20.88
4 Weeks	26.49
7 Weeks	27.84
12 Weeks	32.04
16 Weeks	37.13
20 Weeks	34.27
24 Weeks	30.64
1 Year	2.52
2 Years	0.00
3 Years	13.10

4.6 Three-Factor model Vs Spike and Regime Switching Spike models

Comparing Table 4.2 against Table 4.7, i.e. the RMSEs of the Three-factor spike model, against the RMSEs of the spike and Regime switching spike models respectively, we note that the three-factor spike model is superior to the other two, for all contracts. However there are some disadvantages when relying on this model. First of all as explained earlier, the three-factor model relies on two unobservable variables, X and ε, which have to be calibrated using the Kalman Filter. However, information from the spot is not enough to extract both of the state variables, especially for the long-run factor, which determines the dynamics of the equilibrium level, ε. We thus have to use long-term to maturity contracts to extract information about the long-term dynamics. However as discussed in the data section, long-dated forward contracts have a very low liquidity, which implies as well a large bid-ask spread on the prices. In the Kalman Filter, this is supposed to be captured by the measurement errors. As a result, the estimation of these models may be problematic and the subsequent estimation results may not be reliable.

The second disadvantage of the three-factor spike model stems from the fact that the risk in the reservoirs levels is not captured by the long-run factor, ε, as shown in the analysis, but rather form the short-run factor X.[29] This can be clearly seen from the time series of the state variables in Figure 4.7, which shows especially for 2003, when the reservoir levels reached their historical lowest, it was not the equilibrium level, ε, that increased, but rather the short-term variable, X. However, intuitively this does not make sense, as the short-term factor is supposed to mainly capture the shorter-term risks such as changes in demand resulting from variations in weather as well as market behaviour. Intuitively, the half-life of these shocks should be less than that implied from the three-factor spike model (6 moths). This fact may have serious implications especially for derivatives pricing and more specifically option prices, as it will be shown in the next chapter. This is not to say that the three-factor spike model is not appropriate. One main advantage of the three-factor spike model, except that it provides a better fit, is the fact that it gives a clearer relationship between short- and long-term contracts, and also implies an imperfect correlation between contracts of different maturities. Thus the selection of a model depends on the judgement of a market participant, using his intuition in combination with the empirical findings, and the purpose of its use. For example, as it was mentioned at the beginning of this section, the three-factor spike model might be a very good alternative when valuing long-term contracts such as Power Plants, whose life expectancy is greater than a year, and therefore depend on both long- and short-term risks. On the other hand the regime-switching spike model will capture the risk emanating from fluctuations in the levels of water in the reservoirs, but will not reflect any other long-term risks resulting from changes in supply and demand long-run factors. However if the power plant user, is mainly concerned about the price risk for next year, which is mainly produced by the hydro conditions, the Regime Switching spike model might be more appropriate, as it accounts for changes in the hydro conditions which are the main price determinants in the mid-run.

[29] We also extend the three-factor model by adding another diffusive factor whose dynamics are the same as X, in order to capture the reservoir level risk, however the results gave insignificant estimates.

4.7 Conclusions

In this chapter we examined the major risks involved in the power prices in Nord Pool, and introduced a new model that accounts for the different speeds of mean reversion between normal and spiky shocks in the market. Our proposed model improves significantly the fit between theoretical and observed forward prices as the spikes in the model are discounted at a very fast rate due to their high speed of mean reversion. Consequently, spikes in the spot market do not spillover to the forward market, a finding that is consistent with the actual patterns observed in the market. Furthermore, our proposed model also accounts for seasonality in the risk premium, which reflects the increasing need to hedge against spikes when their probability of occurrence is higher during the colder winter months. These results are in line with previous theoretical and empirical research in other markets, which shows that at these periods the need for hedging is higher and thus forward prices increase in order to entice speculators to trade with market participants. Moreover, we also propose the use of a Regime Switching Spike model that incorporates two separate regimes to distinguish between periods of high and low water levels in the reservoirs. This model is found to be significant in explaining the change in the equilibrium level in the spot market. Finally we extend the spike model, by allowing the equilibrium level not to be dependent just on the level of water in the reservoirs, but also to be affected by longer-run risk, such as load growth, plants shutting down ect. In this model all three stochastic variables are unobservable, and hence two of them have to be estimated using the Kalman Filter; as such we have to rely on long-maturity forward contracts, whose liquidity in the market is very low.

The performance of our proposed models is compared to that of other models proposed in the literature such as a simple mean reversion (MR) and a mean reversion jump diffusion model (MRJD), estimated along the lines suggested by Clewlow and Strickland (2000). This comparison takes place by considering both the accuracy of the theoretical forward prices, implied by each model, in fitting the observed term structure, as well as by observing the trajectorial and distributional properties of the simulated series under the objective probability measure. In particular both the Spike and Regime Switching Spike models provide more realistic simulated price series than the MR and MRJD models, in terms of capturing the jumps and explaining their fast mean reversion. Finally, our models also perform better when the sample is divided up into separate seasons, especially during the winter period, which is expected due to the imposed seasonality in the jump distribution. Overall, it seems that the proposed models' trade-off between complexity and practicality is well behaved, as it accounts for the most significant risks in this extremely volatile market in a parsimonious and intuitive way.

Having discussed the properties of the models in this chapter, the next chapter looks at the implications of these models in terms of options pricing. More specifically we see what kind of volatility shapes we should anticipate from a spike model, by changing each parameter of the stochastic process. We look at the implications of pricing European Options using each of the different models, and give an intuitive explanation on their differences in terms of options pricing and hedging. Finally, we also derive semi-closed form solutions under each of the models.

4.8 Appendix:

Introduction and Application of the Transform Function for the Forward Prices in the Factor Models

A very useful assumption in the finance literature is that the state vector X follows an **affine jump-diffusion** process (AJD). An AJD is a jump-diffusion process for which the drift vector, the "instantaneous" covariance matrix, and jump intensities all have an affine (linear) dependence on the state vector. The affine jump-diffusion processes have been synthesized and extended by Duffie, Pan and Singleton (2000) (henceforth DPS); see also Chacko and Das (2002). Affine diffusions (AD) and affine jump-diffusions (AJD) are quite useful in modelling underlying state variable for several reasons. DPS have shown the close connection between the structure of this kind of models and Fourier transforms, and how from this transform one can obtain derivative prices. The jump diffusion model presented in the main body of the chapter belongs to the class of AJD. Hence, we can use the results provided by DPS in their transform analysis to obtain closed-form solutions implied by our models for forward prices. The DPS transform can be described as follows:

Fix a probability space $\{\Omega, \mathcal{F}, P\}$ and an information filtration $(\mathcal{F}_t) = \{\mathcal{F}_t: t \geq 0\}$, and suppose that X_t is a Markov process in some state space $D \subset \mathbb{R}^n$, following the stochastic differential equation (SDE):

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t + dZ_t \qquad (4.36)$$

Where W is an $(\mathcal{F}_t)$-standard Brownian motion in $\mathbb{R}^n$; $\mu: D \to \mathbb{R}^n, \sigma: D \to \mathbb{R}^{n \times n}$, and Z is a pure jump process whose jump sizes have a fixed probability distribution ν on $\mathbb{R}^n$ and arrive at frequency $\{l(X_t): t \geq 0\}$ for some $l: D \to [0,\infty)$. To be precise, suppose that X is a Markov process whose transition semi-group has an infinitesimal generator

$\mathcal{D}$[30]of the Lèvy type, defined at a bounded C² function $f: D \rightarrow \mathbb{R}$, with bounded first and second derivatives, by

$$\mathcal{D} f(x,t) = f_t(x,t) + f_x(x,t)\mu(x) + \frac{1}{2} tr\left[f_{xx}(x,t)\sigma(x)\sigma(x)^T \right] \\ + l(x) \int_{\mathbf{R}} \left[f(x+z,t) - f(x,t) \right] dv(z) \tag{4.37}$$

Intuitively, $\mu(X_t)$ and $\sigma(X_t)$ are the drift and diffusion terms of the process when no jump occurs, and the jump term captures the discontinuous change of the path with both random arrival of jumps and random jump sizes. That is, conditional on the path of X, the jump times of the jump term are the jumps times of a Poisson process with, possibly, time-varying intensity $\{l(X_s)\colon 0 \le s \le t\}$, and the size of the jump of at a jump time m is independent of $\{X_s : 0 \le s \le m\}$ and has the probability distribution v.

In order for the transform function to work as stated by DPS (2000), the drift, variance-covariance, intensity and discount rate have to be affine functions of the state variables, hence:

$$\begin{aligned} &\mu(x) = K_0 + K_1 \bullet x, \text{ for } K = (K_0, K_1) \in \mathbb{R}^n \times \mathbb{R}^{n\times n}. \\ &\left(\sigma(x)\sigma(x)^T\right)_{ij} = \left(H_0\right)_{ij} + \left(H_1\right)_{ij} \bullet x, \text{ for } \mathrm{H} = (H_0, H_1) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times n\times n}. \\ &l(x) = l_0 + l_1 \bullet x, \text{ for } l = (l_0, l_1) \in \mathbb{R} \times \mathbb{R}^n. \\ &R(x) = \rho_0 + \rho_1 \bullet x, \text{ for } \rho = (\rho_0, \rho_1) \in \mathbb{R} \times \mathbb{R}^n. \end{aligned} \tag{4.38}$$

Let $\theta(c) = \int_{\mathbf{R}^n} \exp\{c \times z\}\ dv(z)$, be the characteristic function of the jump size distribution. The function $\theta(\bullet)$ determines completely the jump size distribution. Also assuming constant interest rates (ρ_0), futures prices are equal to forward prices. Let $\chi \equiv (K, H, l, \rho, \theta)$, it captures both the distribution of the vector process X as well as

---

[30] The generator $\mathcal{D}$ is defined by the property that $\{f(X_t,t) - \int_0^t \mathcal{D} f(X_s,s)ds : t \ge 0\}$ is a martingale for any f in its domain (Ethier and Kurtz, 1986).

the effects of discounting and determines a transform ψ^{χ}: $C^n \times D \times \mathbb{R}_+ \times \mathbb{R}_+ \to C$ of X_T conditional on $\mathcal{F}_t$, when well defined at $t \le T$, by

$$\psi^{\chi}(u, X_t, t, T) = E^{\chi}\left[\exp\left(-\int_t^T R(X_s)ds\right)e^{u \bullet X_T} \mid \mathcal{F}_t\right] \tag{4.39}$$

Where E^{χ} denotes the expectation operator under the distribution of X determined by χ. Hence the difference between the conditional characteristic function of X_T and the transform function ψ^{χ} is the discount factor $R(X_T)$. Therefore under technical regularity conditions DPS (2000) show that

$$\psi^{\chi}(u, X_t, t, T) = e^{\alpha(t) + \beta(t) \bullet x} \tag{4.40}$$

Where α and β satisfy the complex-valued Ordinary-Differential-Equations:

$$\begin{aligned} \dot{\beta}(t) &= \rho_1 - K_1^T \beta(t) - \frac{1}{2}\beta(t)^T H_1 \beta(t) - l_1\left(\theta\left(\beta(t)\right) - 1\right), \\ \dot{\alpha}(t) &= \rho_0 - K_0^T \beta(t) - \frac{1}{2}\beta(t)^T H_0 \beta(t) - l_0\left(\theta\left(\beta(t)\right) - 1\right) \end{aligned} \tag{4.41}$$

With boundary conditions $\alpha(T)=0$ and $\beta(T)=u$.

Application to the Spike Model:

Now for the Spike model of equation (4.1), however with constant jump parameters the connection between the transform function and its use to find the forward prices is as follows:

$$\begin{aligned} P_t &= f(t) + \exp\left(X_t + Y_t\right) \\ dX_t &= \left(-\lambda_X + k_1\left(\varepsilon - X_t\right)\right)dt + \sigma_X dW_X^{\bullet} \\ dY_t &= -k_2 Y_t dt + J(\mu_J, \sigma_J^2)dq(l) \end{aligned} \tag{4.42}$$

Therefore,

$$E_t^*(P_T - f(T)) = e^{r\tau}E_t^*\left(e^{-r\tau}\exp\left(X_T + Y_T\right)\right) = e^{r\tau}\psi^{\chi}(u, X_t, t, T) \\ = e^{r\tau}\exp(\alpha(t) + \beta_1(t)\bullet X + \beta_2(t)\bullet Y) \tag{4.43}$$

Where the vectors $u=(1,1)$, and $X = (X, Y)$ and $\tau=T-t$.

From (4.43) we use (4.41) to reach to the Ordinary Differential Equations and solve them:

$$\frac{\partial\beta_i}{\partial t} = k_i\beta_i \text{ and } \beta_i(T) = 1 \Rightarrow \beta_i(t) = e^{-k_i\tau} \\ \text{Hence } \beta_1(t) = e^{-k_1\tau},\ \beta_2(t) = e^{-k_2\tau} \tag{4.44}$$

$$\theta(\beta_2(t)) = \exp\left(\mu_J e^{-k_2(T-s)} + \frac{1}{2}\sigma_J^2 e^{-2k_2(T-s)}\right)$$

$$\alpha(t) = \int_t^T r + (\lambda_X - k_1\varepsilon)e^{-k_1(T-s)} - \frac{1}{2}e^{-2k_1(T-s)}\sigma_X^2 - l\left(\exp\left(\mu_J e^{-k_2(T-s)} + \frac{1}{2}\sigma_J^2 e^{-2k_2(T-s)}\right) - 1\right)ds$$

Hence using equation (4.43), and the results from (4.44), we have the expected value of the spot, under the risk neutral probability measure:

$$E_0^*(P_t) = f(t) + \exp\left(e^{-k_1 t}X_0 + Y_0 e^{-k_2 t} + l\int_0^t\left(\exp\left(\mu_J e^{-k_2 s} + \frac{1}{2}\sigma_J^2 e^{-2k_2 s}\right) - 1\right)ds + A_t\right)$$

$$A_t = \frac{1}{2}\frac{\sigma_x^2}{2k_1}\left(1 - e^{-2k_1 t}\right) + \left(\varepsilon - \frac{\lambda_x}{k_1}\right)(1 - e^{-k_1 t}) \tag{4.45}$$

Note that for the simple Mean-Reverting model, we exclude Y and its parameters describing its SDE. Similarly, for the Regime Switching spike model, we have two different values for (4.45), depending on the equilibrium price, ε. Equation (4.45) is the general case when the jump parameters remain constant.

Application to the Three-Factor Model:

For the three-factor model, the connection between the transform function and its use to find the forward prices is as follows:

$$
\begin{aligned}
P_t &= f(t) + \exp\left(X_t + \varepsilon_t + Y_t\right) \\
dX_t &= -(\lambda_X + k_1 X_t)dt + \sigma_X dW_X^* \\
d\varepsilon_t &= \left(\mu_\varepsilon^* - \frac{1}{2}\sigma_\varepsilon^2\right)dt + \sigma_\varepsilon dW_\varepsilon^* \qquad \mu_\varepsilon^* = \mu_\varepsilon - \lambda_\varepsilon \\
dY_t &= -k_2 Y_t dt + J(\mu_J, \sigma_J^2)dq(l)
\end{aligned}
\tag{4.46}
$$

where the two Brownian Motions, W_ε and W_X are correlated with a constant correlation coefficient, ρ.

Therefore,

$$
\begin{aligned}
E_t^*\left(P_T - f(T)\right) &= e^{r\tau} E_t^*\left(e^{r\tau} \exp\left(X_T + \varepsilon_T + Y_T\right)\right) = e^{r\tau}\psi^x(u, X_t, t, T) \\
&= e^{r\tau} \exp(\alpha(t) + \beta_1(t) \bullet X + \beta_2(t) \bullet Y + \beta_3(t) \bullet \varepsilon)
\end{aligned}
\tag{4.47}
$$

Where $u=(1,1,1)$, and $X = (X, \varepsilon, Y)$.

Solving the Ordinary Equations we get:

$$
\begin{aligned}
&\frac{\partial \beta_i}{\partial t} = k_i \beta_i \text{ and } \beta_i(T) = 1 \Rightarrow \beta_i(t) = e^{-k_i\tau} \\
&\text{Hence } \beta_1(t) = e^{-k_1\tau},\ \beta_2(t) = e^{-k_2\tau},\ \beta_3(t) = 1 \\
&\theta(\beta_2(t)) = \exp\left(\mu_J e^{-k_2(T-s)} + \frac{1}{2}\sigma_J^2 e^{-2k_2(T-s)}\right)
\end{aligned}
\tag{4.48}
$$

$$
\alpha(t) = \int^T r + \left(\lambda_X - \rho\sigma_\varepsilon\sigma_X\right)e^{-k_1(T-s)} - \left(\mu_\varepsilon^* - \frac{1}{2}\sigma_\varepsilon^2\right) - \frac{1}{2}\left(e^{-2k_1(T-s)}\sigma_X^2 + \sigma_\varepsilon^2\right) - l\left(\exp\left(\mu_J e^{-k_2(T-s)} + \frac{1}{2}\sigma_J^2 e^{-2k_2(T-s)}\right) - 1\right)ds
$$

Hence using equation (4.43), and the results from (4.44), we have the expected value of the spot, under the risk neutral probability measure:

$$
E_0^*(P_t) = f(t) + \exp\left(e^{-k_1 t}X_0 + \varepsilon_0 + Y_0 e^{-k_2 t} + l\int_0^t\left(\exp\left(\mu_J e^{-k_2 s} + \frac{1}{2}\sigma_J^2 e^{-2k_2 s}\right) - 1\right)ds + A_t\right)
$$

$$
A_t = \frac{1}{2}\frac{\sigma_x^2}{2k_1}\left(1 - e^{-2k_1 t}\right) - \frac{\lambda_X - \rho\sigma_\varepsilon\sigma_X}{k_1}(1 - e^{-k_1 t}) + \mu_\varepsilon^* t \tag{4.49}
$$

5. OPTION PRICING ANATOMY OF THE ELECTRICITY MODELS

5.1 Introduction

The previous chapter introduced a spike, a regime switching spike and a Three-Factor spike model that are able to describe the spot price processes of electricity in Nord Pool. We provided closed-from solutions for forward contracts and we then examined their fit on the forward term structure. We also looked at the spot price behaviour using simulations and derived their first four moments that explain the spot price distribution. The obvious next step is therefore to discuss the intuition and implications behind each model in terms of option pricing.

In this chapter we are going to discuss the properties of European and Asian Options using the reduced form models of the previous chapters. In order to do so we first provide semi-closed form solutions, using the transform analysis by Duffie et al (2000) for European Options, which are particularly useful in the option pricing theory for traders and analysts in order to get quick results and make conclusions on how to trade volatility. We discuss the properties of each model in terms of pricing European Options, by looking at the Implied Volatility Skews and performing sensitivity analyses, which we call the pricing anatomy of European option prices.

The pricing anatomy of European option prices is significant, as they have a direct impact on the pricing of products used in the market. In this setting it is important for market practitioners to use option pricing models that are able to price the most significant risks that exist in the market. Thus the aim of this chapter is to provide the potential usefulness of the pricing model introduced in the earlier chapter, in terms of pricing all of the above products. When it comes down to choosing among different models, it is important to have an idea of their similarities and differences particularly in terms of pricing and hedging. What are the implications of choosing the spike model rather than the mean-reverting model? Can we provide some kind of closed-form solutions by making the model more complex and introducing jumps? Which parameters play the most crucial part in order to fit the volatility structure across strike prices and maturities? Does the regime switching spike model give more intuitive prices than the plain spike model? The intuition for European option prices is demonstrated in terms of implied model volatilities, as traders make their decisions in the options market, in terms of Black and Scholes implied volatilities. In this way we can explore the volatility shape implied by the model and thus understand what kind of structure it can capture in the real market. Asian options are then examined in a separate section. The main questions stemming in the Asian option pricing section are to give intuitive explanations on the sensitivity of the products by changing different parameters. This is of paramount importance to market practitioners, who are very much interested on the factors that drive the value and the moneyness of an option, when engaged in trading these products. For example, questions such as, do the jumps affect Asian options traded in the market or are they averaged out, or what does the speed of mean reversion imply, are important for traders in the market. As a priori we would expect the prices of Asian options to increase when increasing the equilibrium level, diffusive volatility, jump intensity, jump size mean and volatility. Interestingly our analysis shows that this is not always the case. The effect of the jump size becomes pronounced only when the jump intensity is high and more particularly when the jump speed of mean reversion is low. Moreover, the effect of the equilibrium price depends strongly on the speed of the mean reversion of the diffusive component. The reasons and the intuition behind these results are examined in the subsequent sections.

Thus the chapter is formed as follows: the following section describes the use of European in the electricity markets. Section 3 reviews the theory of option pricing and derives closed and semi-closed form solutions for the mean-reverting and the spike model. Section 4 explains the theory of implied volatility and model implied volatility. Section 5 displays results on implied volatilities by changing different model parameters and giving intuitive explanations in terms of the moments of the spot price distribution. Section 6 examines the different option prices generated from the four different models examined in Chapter 4. Section 7 explains the use of Asian options, the problems faced in their pricing and conducts a sensitivity analysis based on the spike model. Finally, Section 8 concludes the chapter.

5.2 Use of Options in Power Markets

In power markets, individual options on the spot are rarely used. The market relies more on structured products such as Caps, Floors or Collars. Caps provide price protection for the buyer above a predetermined level, which is called the *cap price*, for a predetermined period of time. For example, within a period $\{T_1 = \tau_1, ..., \tau_n = T_2\}$, the owner of a Cap has each day (hour or any other time interval) the right to exercise a European call option with strike price, K. A Floor on the other hand, guarantees the minimum price that will be paid or received at a predetermined level, called the *floor price*. In other words a Floor is structured such that the holder has the right at each time interval τ_j within the dates T_1 to T_2, to exercise a European Put option with strike price K. A Collar is a combination of a long position in a Cap and a short position in a Floor. All three structured products are settled for a predetermined quantity or volume.

As an example, consider an electricity utility, which needs additional power at times of higher demand e.g. in February when the probability of extreme price movements is high. Since the utility does not know exactly when the load will be high, it may need a structured product to fit its need and hedge the risks. One possibility for the utility is to buy a portfolio of call options, which gives the right but not the obligation, to buy electricity every day (or even hour) within the delivery period, with a capacity of 100MW at a fixed cap price of 100 NOK/MWh. On the other hand the customer demanding protection from high electricity prices might be unwilling to incur the significant costs of a cap. If the customer is willing to sacrifice some possible gain from low prices to pay for the protection, the simple solution is to buy a collar and thus use the premium on the short floor to pay for the cap protection.

Since the cap is a portfolio of call options for a predetermined strike price and dates $\{T_1 = \tau_1, ..., \tau_n = T_2\}$, its valuation is thus the sum of the price of those call options:

$$V_{Cap} = \sum_{j=1}^{n} Call\left(t, \tau_j, P_{\tau_j}, K\right) = \sum_{j=1}^{n} E^{*}\left[e^{-r(\tau_j - t)} \max\left(P_{\tau_j} - K, 0\right)\right] \quad (5.1)$$

where $Call\left(t,\tau_j,P_{\tau_j},K\right)$ is the call price at time t with maturity at τ_j, exercise price K and interest rate r, in the risk-neutral world. Similarly, a floor is a portfolio of European put options $Put\left(t,\tau_j,P_{\tau_j},K\right)$ with strike prices K equal to the floor level and maturity dates equal to the settlement dates τ_j of the floor.

$$V_{Floor}=\sum_{j=1}^{n}Put\left(t,\tau_j,P_{\tau_j},K\right)=\sum_{j=1}^{n}E^{*}\left[\mathrm{e}^{-r(\tau_j-t)}\max\left(K-P_{\tau_j},0\right)\right] \tag{5.2}$$

Moreover many utilities have introduced interruptible or curtailment service contracts which aim to supply electricity at reduced cost by taking advantage of customer's flexibility to manage their load, and thus allow the suppliers to provide electricity to those costumers who are willing to pay the high prices in times of scarcity. Gendra (1994) introduces the concept of Callable Forwards in which the costumer takes a long position in a forward contract and a short position in a call option. Thus the supplier benefits by earning the possibility of calling off the supply in case the price of electricity spikes at expiry. Of course this is done by providing a discount on the forward, which the customer holds, by paying him the price of the option. The strike price of the option, K, is agreed to be equal to the shortage cost a potential user faces in case of a load-curtailment. Furthermore, electricity call and put options are the most effective tools available to merchant power plants and power marketers for hedging price risk because electricity generation capacities can be essentially viewed as call options on electricity, particularly when generation costs are fixed.

5.3 The perfect unrealistic Black & Scholes World

5.3.1 The Black and Scholes Models

In their seminal paper Black and Scholes (1973) (BS) assumed that the process of tradable asset, P, follows a Geometric Brownian Motion with constant Volatility and thus they derived a closed-form solution for European Option Prices under the risk neutral measure:

$$call = P_0 N(d_1) - Ke^{-rT} N(d_2) \tag{5.3}$$

$$put = Ke^{-rT} N(-d_2) - P_0 N(-d_1) \tag{5.4}$$

where

$$d_1 = \frac{\ln(P_0 / K) + (r + \sigma^2 / 2)T}{\sigma\sqrt{T}}$$

$$d_2 = \frac{\ln(P_0 / K) + (r - \sigma^2 / 2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

Note that the first term in (5.3), $P_0 N(d_1)$, denotes the expected value of the discounted spot price at maturity T, in case it is above the discounted strike price K, in the risk-neutral world (Hull 2003). The same of course stands for the second term in (5.4), $P_0 N(-d_1)$, but this time it is for the case the discounted spot price at maturity is below the strike price. $N(d_1)$ and $N(-d_1)$ are also called the deltas and give an indication of how much of the underlying is needed, in order to form a portfolio of the underlying and the option, which is immune to small movements from the spot. On the other hand, the terms $N(d_2)$ and $N(-d_2)$, denote the risk-neutral probability that the options are exercised at maturity. Another important point worth mentioning, is that as the present spot price P_0 tends to infinity, the ratio $ln(P_0/K)$ tends to infinity, and thus $N(d_1)$ and $N(d2)$ equal to one and the call option has a value of $P_0 - e^{-rT}K$, where as the put becomes zero. The volatility parameter has a similar effect, only this time as $\sigma \rightarrow \infty$, then $N(d_1) \rightarrow 1$ and $N(d_2) \rightarrow 0$, thus the value of the call approaches the current value of the spot, and the value of the put tends to the discounted value of the exercise price. Thus we can see that the volatility parameter plays a vital role in the BS equation. However it is an unknown parameter, which sometimes can be extracted from market data on options whenever available and liquidity is satisfactory, and if not, we need to calibrate it from the time series of the spot, as it is the case in the electricity market.

The next obvious step of course is to discuss whether the BS model is appropriate for option pricing in the electricity market. In equations (5.3) and (5.4), the volatility is assumed to be constant across Strike prices (*K*) and Maturities. There are two main disadvantages when using the Geometric Brownian Motion to model Electricity prices, which are mainly caused by the non-storable nature of the commodity. First, electricity prices as discussed in the previous chapters mean revert to a seasonal equilibrium level, and second is their spiky behaviour. In what follows we will explain the consequences in Option pricing when mean reversion and spikes are used to describe the stochastic behaviour of the underlying, instead of the GBM.

5.3.2 Option Pricing under Seasonality and Mean Reversion

As we know electricity prices are mean-reverting to an equilibrium value. Furthermore it has been shown in the market that forward prices become less and less volatile the further their time to maturity is, which is due to mean reversion. Thus the Volatility Term Structure is decreasing with time to maturity, a phenomenon known as the Samuelson effect (1965), since for non-storable commodities such as the electricity any new information in the market will have a more prominent effect on derivative prices that are closer to maturity. This fact can be captured correctly by the seasonal mean-reverting model where the spot price is modelled as follows:

$$P_t = f(t) + \exp(X_t)$$
$$dX_t = \left(k_1\left(\tilde{\varepsilon} - X_t\right)\right)dt + \sigma_X dW_X^* \quad \text{where } \tilde{\varepsilon} = \varepsilon - \lambda_X / k_1 \qquad (5.5)$$

and thus the call and put prices are derived in Appendix 5.9.1:

$$call = e^{-rt}\left[F\left(0,T,X_0\right)N(d_1) - DK\ N(d_2)\right] \qquad (5.6)$$

$$put = e^{-rt}\left[DK\ N(-d_2) - F\left(0,T,X_0\right)N(-d_1)\right] \qquad (5.7)$$

where
$$d_1 = \frac{\ln\left(e^{-rT}\frac{F(0,T,X_0)}{DK}\right) + \left(rT + \frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)\right)}{\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}}$$
$$d_2 = d_1 - \sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}} \qquad (5.8)$$

where $DK=(K-f(T))$ denotes the deseasonalised strike price. The reason we use the deseasonalised strike price follows from the fact that the seasonality component is the deterministic part of the spot and can thus be subtracted directly from the strike price[31]. Thus in a model where there is seasonality the moneyness of the option is directly affected not only by the strike price but also from seasonality. Option pricing formulas (5.6) and (5.7) are very similar to the BS only this time instead of the spot price, we have the discounted deseasonalised forward price $F(0,T,X_0)$ and a standard deviation that converges faster to a constant value with k_1 and T as noted by the standard deviation equation $\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}$, rather than increasing continuously with time to maturity as it is assumed in the BS formula, where we have $\sigma_X\sqrt{T}$. In this way if we calculate the price of an option using equation (5.6), and then use the BS equation (5.3) and find the value of σ such that the BS option price matches that given by equation (5.6), we will see that the implied volatility parameter, σ, will be decreasing with time to maturity, as it will be shown in later sections.

5.3.3 Option Pricing under the Spike Model

As discussed earlier electricity prices are spiking and thus reach levels that are not consistent with a Normal Distribution. That has a direct impact on Option Prices, especially for OTM call options which now have higher probability to end in the money than what the BS Model might predict. This is what causes the implied volatility from market prices in the BS model not to be constant across strike prices but to have a smile or smirk shape as it will be shown in the next section.

[31] See Appendix 5.9.1 for proof.

To calculate the value of the European Option in the risk-neutral world as in (5.3) and (5.4), one has to know the risk neutral distribution. In the BS, the risk neutral density is lognormal, and the expectation of the normalised option payoff can easily be calculated. In more sophisticated models, such as the mean-reverting jump diffusion model of Clewlow and Strickland (1999), there is no closed-form solution for the risk-neutral density. One can use the *Transform Analysis*, by Duffie et al (2000) and implement the Fourier inversion theorem where numerical integration of the imaginary part of complex function is performed. This kind of solution is semi-closed since numerical integration has to be performed to calculate the density function. Note however that even in the BS world, as in any other model that has closed-form solutions for option prices, the calculation of the risk-neutral density function also involves numerical integration of the area under the Normal Distribution, which can nevertheless be easily calculated using software packages and statistical tables. Now recall the spike model discussed in Chapter 4:

$$\begin{aligned} P_t &= f(t) + \exp(X_t + Y_t) \\ dX_t &= \left(k_1\left(\varepsilon - X_t\right) - \lambda_X\right)dt + \sigma_X dW_X^* \\ dY_t &= -k_2 Y_t dt + J(\mu_{J_i}, \sigma_{J_i}^2)dq(l_i) \end{aligned} \tag{5.9}$$

To give an explanation on how one can use the transform analysis, to derive semi-closed form solutions for European option prices, as it is done by Duffie et al (2000), let us start by defining $G_{a,b}(y;)$ as the price of a security paying e^{aX_T} when at maturity $bX_T \leq y$, where X_T is a vector of the state variables that describe the price process and in our example $X_T = \left[X_T, Y_T\right]$. Also note that the price of a European call is given as follows:

$$\begin{aligned} Call &= E^*\left[e^{-rT}\left(P_T - K\right)^+ / \mathcal{F}_0\right] = E^*\left[e^{-rT}\left(e^{1'X_T} - DK\right)^+ / \mathcal{F}_0\right] \\ &= E^*\left[e^{-rT} e^{1'X_T} 1_{1'X_T \geq \ln(DK)} / \mathcal{F}_0\right] - DK\ E^*\left[e^{-rT} 1_{1'X_T \geq \ln(DK)} / \mathcal{F}_0\right] \\ &= G_{1',-1'}\left(-\ln(DK); X_T, T, \chi\right) - DK\ G_{0',-1'}\left(-\ln(DK); X_T, T, \chi\right) \end{aligned} \tag{5.10}$$

where χ captures both the distribution of the vector prices X as well as the effects of discounting and determines the transform function in the DPS (2000), $1'$ and $0'$ denote $2x1$ transpose vectors of ones and zeros respectively. What equation (5.10) implies is that the call option payoff is In-The-Money in case the deseasonalised spot prices $X_T+Y_T>ln(DK)$ or alternatively $-X_T-Y_T<-ln(DK)$. Also denote $\bar{v}=[v,v]$ which is the variable defining the Fourier Transform, Duffie et al (2000) proved that:

$$G_{1',-1'}(-\ln(DK);X,T,\chi)=\psi^{x}(1')-\frac{1}{\pi}\int_{0}^{\infty}\frac{\operatorname{Im}\left[\psi^{x}\left(1'-i\bar{v},X,0,T\right)e^{-i\ln(DK)v}\right]}{v}dv$$
$$=\frac{e^{-rT}F(0,T,X_0)}{2}-\frac{1}{\pi}\int_{0}^{\infty}\frac{\operatorname{Im}\left[\psi^{x}\left(1'-i\bar{v},X,0,T\right)e^{-i\ln(DK)v}\right]}{v}dv \tag{5.11}$$

$$G_{0',-1'}(-\ln(DK);X,T,\chi)=\frac{\psi^{x}(0')}{2}-\frac{1}{\pi}\int_{0}^{\infty}\frac{\operatorname{Im}\left[\psi^{x}\left(-i\bar{v},X,0,T\right)e^{-i\ln(DK)v}\right]}{v}dv$$
$$=\frac{e^{-rT}}{2}-\frac{1}{\pi}\int_{0}^{\infty}\frac{\operatorname{Im}\left[\psi^{x}\left(-i\bar{v},X,0,T\right)e^{-i\ln(DK)v}\right]}{v}dv \tag{5.12}$$

Where $\psi^{x}(u,X,0,T)$ is the transform function derived in the previous chapter, $i=\sqrt{-1}$ and Im[.] is the imaginary part of a complex number. Note that the intuition behind the functions $G_{1',-1'}(ln(-DK);X,T,\chi)$ and $G_{0',-1'}(ln(-DK);X,T,\chi)$ is exactly the same as in the BS model for $P_0N(d_1)$ and $N(d_2)$. Thus $G_{1',-1'}(ln(-DK);X,T,\chi)$ is the expected value of the discounted deaseasonalised spot price, $exp(X+Y)$, at maturity T, in case it is above the discounted deseasonalised strike price DK, in the risk-neutral world. Similarly, $G_{0',-1'}(ln(-DK);X,T,\chi)$ is the risk-neutral probability that the option will be ITM at maturity.

The numerical integration technique used is the Adaptive Simpson Quadrature. The main advantage of the adaptive Simpson Quadrature rule for numerical integration is that it is very accurate and fast, as it divides the area of interest in the integration into smaller areas (or intervals), and uses more points in the areas where they are needed, and less in the areas that are not needed, since the number of intervals needed does not depend on the behaviour of the integrated function everywere, but on the points were the function behaves worst.

5.4 Option Pricing and Moments of the Distribution

In this section we will discuss the impact of the model choice and of its parameters, on the implied volatility smile produced from the BS formula. It turns out that in this context the moments of the risk-neutral distribution play an important role. We will therefore often proceed in two steps, in that we first analyse the impact of the parameters and the model choice on the moments of the risk-neutral distribution, i.e. variance, skewness and kurtosis, and then consider the impact of these moments on the smile. Note that the term "smile" refers to the shape of the implied volatilities, with respect to the moneyness of the option, where ITM and OTM implied volatilities are higher than ATM.

5.4.1 The effect of Moments on Black & Scholes Implied Volatilities

We start by considering the impact of an increase in the variance, while the mean remains constant in the BS model. As a result, the probability mass is shifted from returns close to the centre of the distribution to returns further in the tails, in other words the distribution curve becomes wider. Intuitively, this implies in terms of options pricing that for all options along the strike price axis the probability of positive payoff increases, which leads to an upward shift of the overall level of the smile curve.

Turning now to the skewness, this determines the relation between the prices of OTM puts and OTM calls. Assuming that the mean and the variance of the spot price under the risk neutral measure remain constant, consider a decrease in skewness from normality; the probability mass is then shifted from high prices to very low prices[32]. The prices of OTM calls, which pay off for high spot prices, thus decrease, and the prices of OTM puts, which have a payoff for low spot prices, increase, since there is a greater area under the curve at those points. For a negative skewness, we therefore expect the implied volatility (IV) of OTM puts to be larger than the IV of OTM calls, since the probability of the underlying price reaching low prices is higher than that

[32] In this section we assume constant interest rates, alternatively the interest rate risk could be hedged using the bond market.

implied by the BS model, hence we should expect to see a downward sloping smile as shown by Branger (2004)[33].

Turning now into the kurtosis, in statistical terms it depends on the fourth moment of the distribution, and measures the fatness of the tails. If the kurtosis increases, there is more probability mass in the tails of the distribution, so very low and very high spot prices both have a higher probability of occurring when compared to the normal distribution. In option pricing terms, the kurtosis is the main driver of the curvature of the smile. Similarly to the other cases, we assume that the mean, the variance and the skewness of the distribution are given. Then, increasing the kurtosis directly implies higher probability of extreme prices, and thus higher prices for OTM puts and OTM calls. In order to hedge against these extreme events, one can form a static portfolio of long positions comprising a continuum of calls with strike prices from zero to infinity, as shown by Carr and Madan (2001).

As stated earlier, since the electricity market is non-storable, the sources of risks in the market are more than the number of traded securities, thus our hedging possibilities are very limited and the market is incomplete. Of course one can try and use contingent claims on the same underlying, and thus complete the market. In the spike model, we need one instrument to hedge the diffusive normal risk generated by the normal state variable X, plus a number of other derivatives which has to be equal to the number of possible jumps sizes generated from the spike variable Y. The first is possible since we can use the Forward market. However for Y, the jump sizes are drawn out from a continuous distribution and thus we need a continuum of claims to complete the market. In Appendix 5.9.4, we illustrate how a hedging portfolio may be constructed in the electricity market and the problem of incompleteness when the underlying process involves jumps. From the practical point of view, this is not feasible since the forward options market in Nord Pool is not liquid, especially across all strikes as shown by Hjalmarsson (2003). In addition as noted by Eydeland et al (2003), the liquidity is particularly very low for OTM options in the electricity

[33] Note that the term 'downward sloping smile' does not imply a monotonicity of the implied volatility function in a rigorous mathematical sense.

markets, thus restricting the possibility of trading the volatility skew. This is a surprising finding given the high volatility of the market.

5.4.2 The effect of moments of the Spike Model

As shown in the earlier sections, any jump-diffusion model can generate a skewed risk-neutral distribution with excess kurtosis. Lets first recall the theoretical variance of daily returns (for $\Delta t = 1/365$), derived in Chapter 4, using the characteristic function of the deseasonalised spot process (as in Das, 2001) using the spike model of equation (5.9):

$$Variance = \sigma_x^2 \frac{\left(1 - e^{-2k_1\Delta t}\right)}{2k_1} + l\left(\mu_J^2 + \sigma_J^2\right)\frac{\left(1 - e^{-2k_2\Delta t}\right)}{2k_2} \quad (5.13)$$

The first part of the variance is generated by the normal diffusive variable *X*, and the second from the spike variable *Y*; we call the latter part the *jumpiness*, and it is this part that has the crucial impact on the smile. Also as shown in Chapter 4, the theoretical skewness and kurtosis of the returns implied from the model are:

$$Skewness = \frac{\frac{l}{3k_2}\left(1-e^{-3k_2\Delta t}\right)\left(\mu_J^3 + 3\mu_J\sigma_J^2\right)}{Variance^{3/2}} \quad (5.14)$$

$$Kurtosis = \frac{\frac{l}{4k_2}\left(1-e^{-4k_2\Delta t}\right)\left(\mu_J^4 + 3\sigma_J^4 + 6\sigma_J^2\mu_J^2\right)}{Variance^2} + 3 \quad (5.15)$$

where the *Variance* is given in equation (5.13).

In the previous chapter we discussed that, as the mean of the jump size, μ_J, is the only one that can take negative values, the sign of the skewness of the distribution is determined solely by that parameter. On the other hand we also showed that increasing the speed of mean reversion especially for jumps, k_2, the value of the moments decays. Furthermore, since μ_J is the main determinant of the skewness, the jump size volatility, σ_J and frequency, l, also play another vital role in determining the size of the skewness (not the sign) and kurtosis of the distribution. Now it remains to show what their roles are in option pricing and whether they can actually give derivative traders any intuition with regards to pricing and trading options.

5.5 Theory and Numerical examination of model Implied Volatilities

5.5.1 Model Implied Volatility

The following analysis is based on model's *implied volatilities* from the BS closed-form solution since it is the fundamental tool for option pricing and market participants may use it for pricing and hedging. For example most of the time the traders use the BS model to back out the implied volatilities using prices from the market. The inconsistency of the BS model is the assumption that the volatility is constant across strike prices and time to maturity. Using a more sophisticated model, such as jump diffusion, may yield implied volatilities which are consistent with those implied from the market. However in the current setting, an empirical analysis is not feasible since the options market in the Nord Pool exchange is illiquid and most contracts are traded OTC. Nevertheless the current analysis will give valuable insight behind the different models proposed, the shape of volatility we should anticipate to fit in the market using these specific models, and whether these can price all the risks an option trader faces in the electricity market. Branger (2004) has conducted a similar analysis to ours using Merton's jump diffusion model and Heston's stochastic volatility model, to price options in the stock market. However, to the best of our knowledge, this study has not been implemented before in the power market.

Therefore, despite the shortcomings, the market still relies on BS, so we will be using BS as a benchmark to extract model implied volatilities. However before we start the analysis let us first explain what the term *model implied volatilities* means. This is the anticipated volatility of prices, such that the BS option price in (5.3) and (5.4) matches the option price given by the spike model (or any other model not assuming a GBM process). More specifically, under the log-normality condition, an option price is a value function of the current spot price P, strike K, the current time t in which the option is evaluated, exercise time T, interest rate r, and finally, the volatility parameter σ.

$$V_{Option}^{BS} = f_{BS}\left(P_0, K, t, T, r, \sigma\right) \tag{5.16}$$

The parameter σ reflects the model's consensus on the anticipated random behaviour of prices on the interval *[t, T]*, i.e. from the current date t to the options maturity date T. Now for any given option price from the model V_{option}^{model}, such as the MR, the corresponding implied volatility $\sigma^{implied}$ is defined as the value parameter σ, such that:

$$V_{option}^{model} = f_{BS}\left(P_0, K, t, T, r, \sigma^{implied}\right) \tag{5.17}$$

Applying standard techniques, such as Newton's method, to solve equation (5.17), we obtain the implied volatility

$$\sigma^{implied} = \sigma\left(t, T, X, P_0, r, V_{option}^{model}\right) \tag{5.18}$$

5.5.2 Numerical examination of model implied Volatilities

The numerical examples use a base case where we have chosen the parameters such that we see a pronounced smile for option prices with times to maturity of 0.5, 1 and 2 months, and such that the jumpiness of the daily spot return is equal to 50%. This way we can have a clearer view on the implied volatilities, and thus understand the major contribution of each parameter. Thus for the base case scenario, the parameters chosen are shown in Table 5.1.

In this section we do not want to compare the option prices across different seasons of the year, since we are interested on the implied volatilities, which are affected by the parameters of the stochastic processes, i.e. X and Y. As we see from equation (5.6) and (5.7), the deterministic factor of seasonality $f(T)$ affects directly the moneyness of the option. Thus for the moment we assume that there is no deterministic seasonality, $f(t)$. However, as shown by the closed-form solutions of the moments for the spike model in Chapter 4, the jump parameters l, μ_J and σ_J are the main drivers that affect the distribution of the underlying such as volatility, skewness and kurtosis. The impact of these factors is thus examined in the following section.

Table 5.1: Parameter Values used in the spike model for analysis on European call options Implied Volatilities

The table shows the parameters' value of the spike model in equation (5.9), used for the estimation of European call prices in the base case scenario.

Model Parameters	Parameters values
$X_0 = \left(\varepsilon - \frac{\lambda_X}{k_1}\right)$	5.137
σ_X	0.5
k_1	3
Y_0	0
l	5.5
μ_J	0.1
σ_J	0.283
k_2	290
r	5%

Note: See equation (5.9) for the definition of the variables

Therefore, in this section we analyse the impact that each factor has on Implied Volatilities by considering each factor separately, with the following order. First we examine the effect of mean reversion on implied volatilities, by examining the MR model. Then we turn into the spike model, by first looking at the impact of time to maturity on the implied volatility smile using the base case scenario. Next, we look at the impact of the mean jump size, μ_J, on the volatility skew. Following that, we explore the comparative effect of the jump intensity, l, against the jump volatility σ_J, and explore which parameters affects more significantly the curvature of the volatility smile.

5.5.2.1 Volatility skew in Mean Reversion

Before we start exploring the implied volatilities of the spike model, lets first have an insight on how the implied volatilities of a pure mean-reverting model, such as in equation (5.5), look like when the diffusive volatility, σ_X is 100%, all jump parameters set to zero and the remaining parameters are as in Table 5.1. In this analysis we consider short-term to maturity European options (15 days), and three different equilibrium levels of X; The first case is when the risk neutral equilibrium level is equal to the current spot price i.e. $\exp\left(\tilde{\varepsilon}\right) = P_0$; the second is the case where the equilibrium level is below the current spot price, and finally in the third case the equilibrium level is above the spot price[34].

Figure 5.1 and Figure 5.2 present the implied volatilities of the MR model for call and put options, respectively, where moneyness is defined as the ratio K/P. Thus looking at both figures from left to right, we are moving from ITM to OTM call options, and the opposite is true for put options (i.e. from OTM to ITM). Starting with the call option, we can note that even a pure mean-reverting model displays some skewness. This is due to the mean reversion of prices, since if prices are mean-reverting they do not deviate substantially from their mean. This implies that the volatility must be decreasing when moving from ITM to OTM call options. In the Base and High cases where the risk-neutral mean is equal and above the current spot price, mean reversion is good for ITM call options, since it pulls prices towards the mean and thus prices will not be able to fall to very low levels, as there is a force pulling the prices above the strike most of the time. However the opposite is true for OTM call options; the mean reversion coefficient does not allow prices to deviate from their mean substantially and this has the implication that, if the mean is below the strike price, the probability of the option ending ITM is reduced. This results in the skew evident in the graph. Turning next to the Low case, where the equilibrium level is below the current spot price, we see an increasing volatility as strike prices increase. This suggests that prices are now pulled towards lower levels due to mean reversion, and thus ITM calls are cheaper than what the BS model would have predicted. However

[34] The second case corresponds to the risk-neutral equilibrium level in the regime switching spike model when the water level in the reservoirs is above its seasonal mean, and thus the risk-neutral equilibrium level is 4.91. Conversely, the third case corresponds to the case when the water level in the reservoirs is below its seasonal mean, and thus the risk-neutral equilibrium level in the regime switching spike model is 5.51.

OTM options are not really affected by this since they have already a low probability of ending ITM due to the speed of mean reversion as the previous results showed. Similar results are observed in the case of put options, which are shown in Figure 5.2, where the implied model volatilities display exactly the opposite image from that of call options. Thus the speed of mean reversion and the equilibrium level play a critical part on the implied volatility skew.

Figure 5.1: Implied volatility of mean-reverting model for European Call Options

The figure shows the implied volatilities, for European Call Options, of the MR model across different strike prices (Moneyness), using different equilibrium levels. The Base case corresponds to the cases when the risk neutral equilibrium level of X is equal to the current value of X, the Low and High equilibrium cases correspond to cases when the risk neutral equilibrium level is below and above the current value of X.

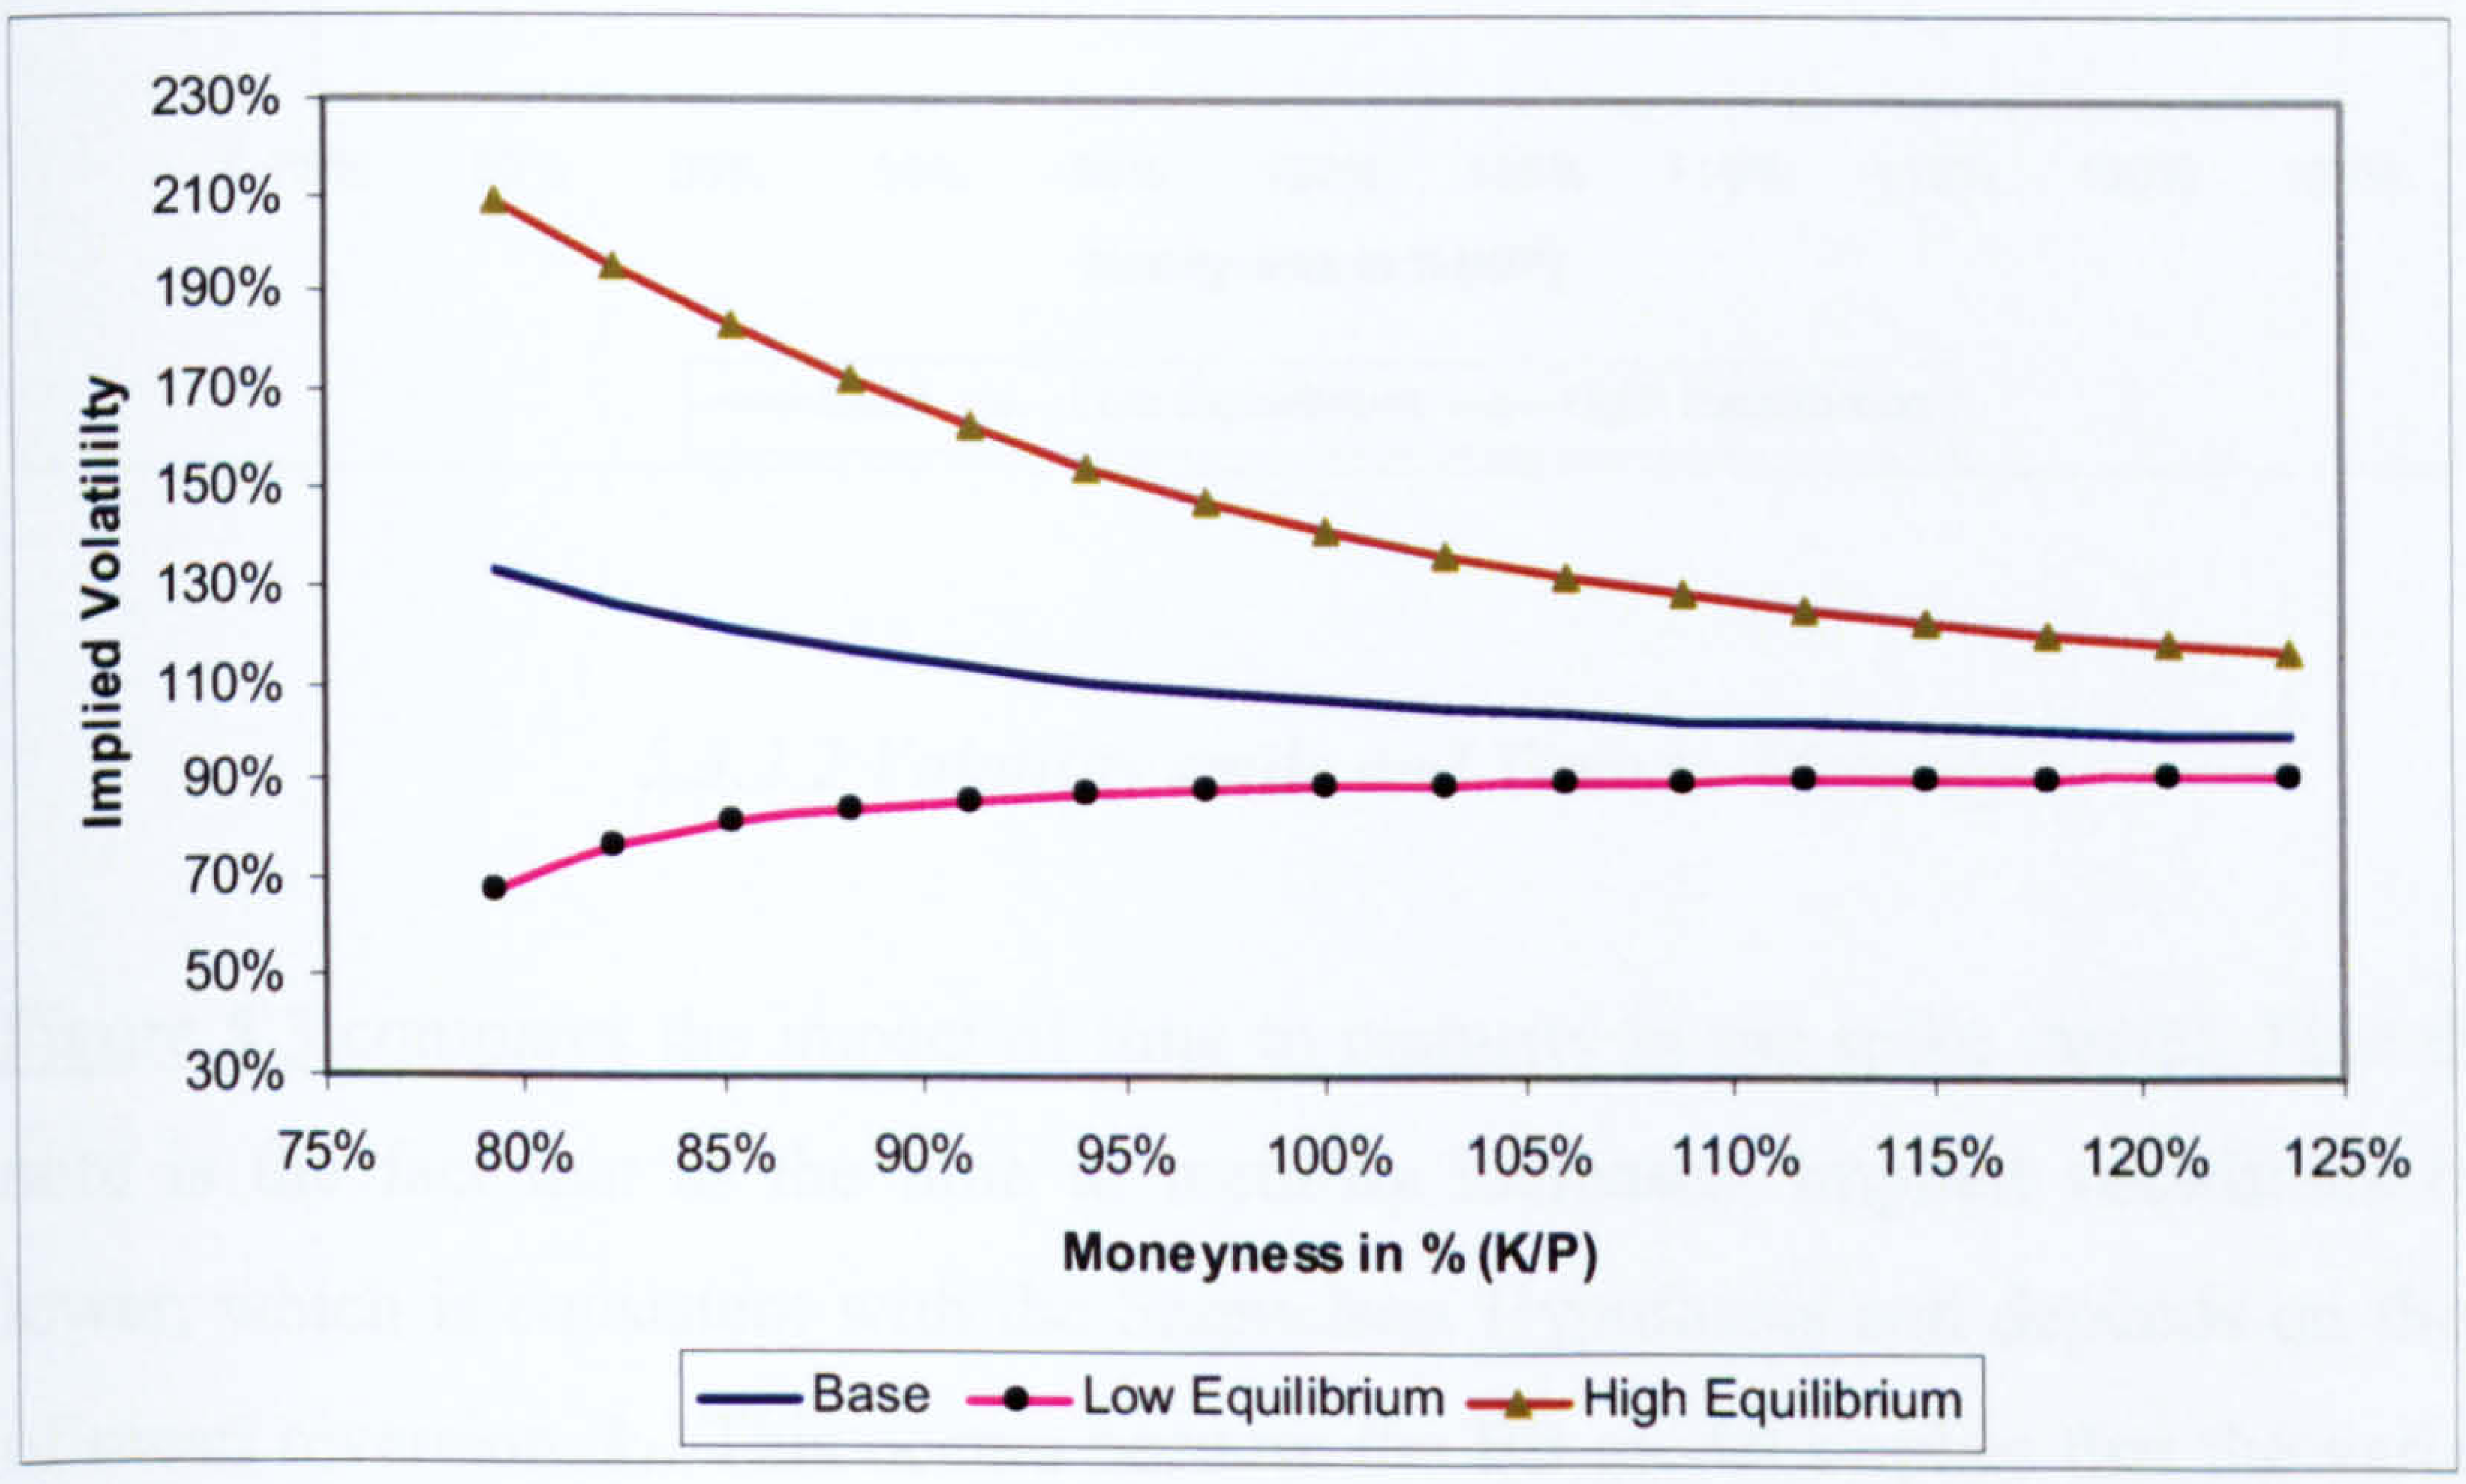

Figure 5.2: Implied volatility of mean-reverting model for European Put Options

The figure shows the implied volatilities, for European Put Options, of the MR model across different strike prices (Moneyness), using different equilibrium levels. The Base case corresponds to the cases when the risk neutral equilibrium level of X is equal to the current value of X, the Low and High equilibrium cases correspond to cases when the risk neutral equilibrium level is below and above the current value of X.

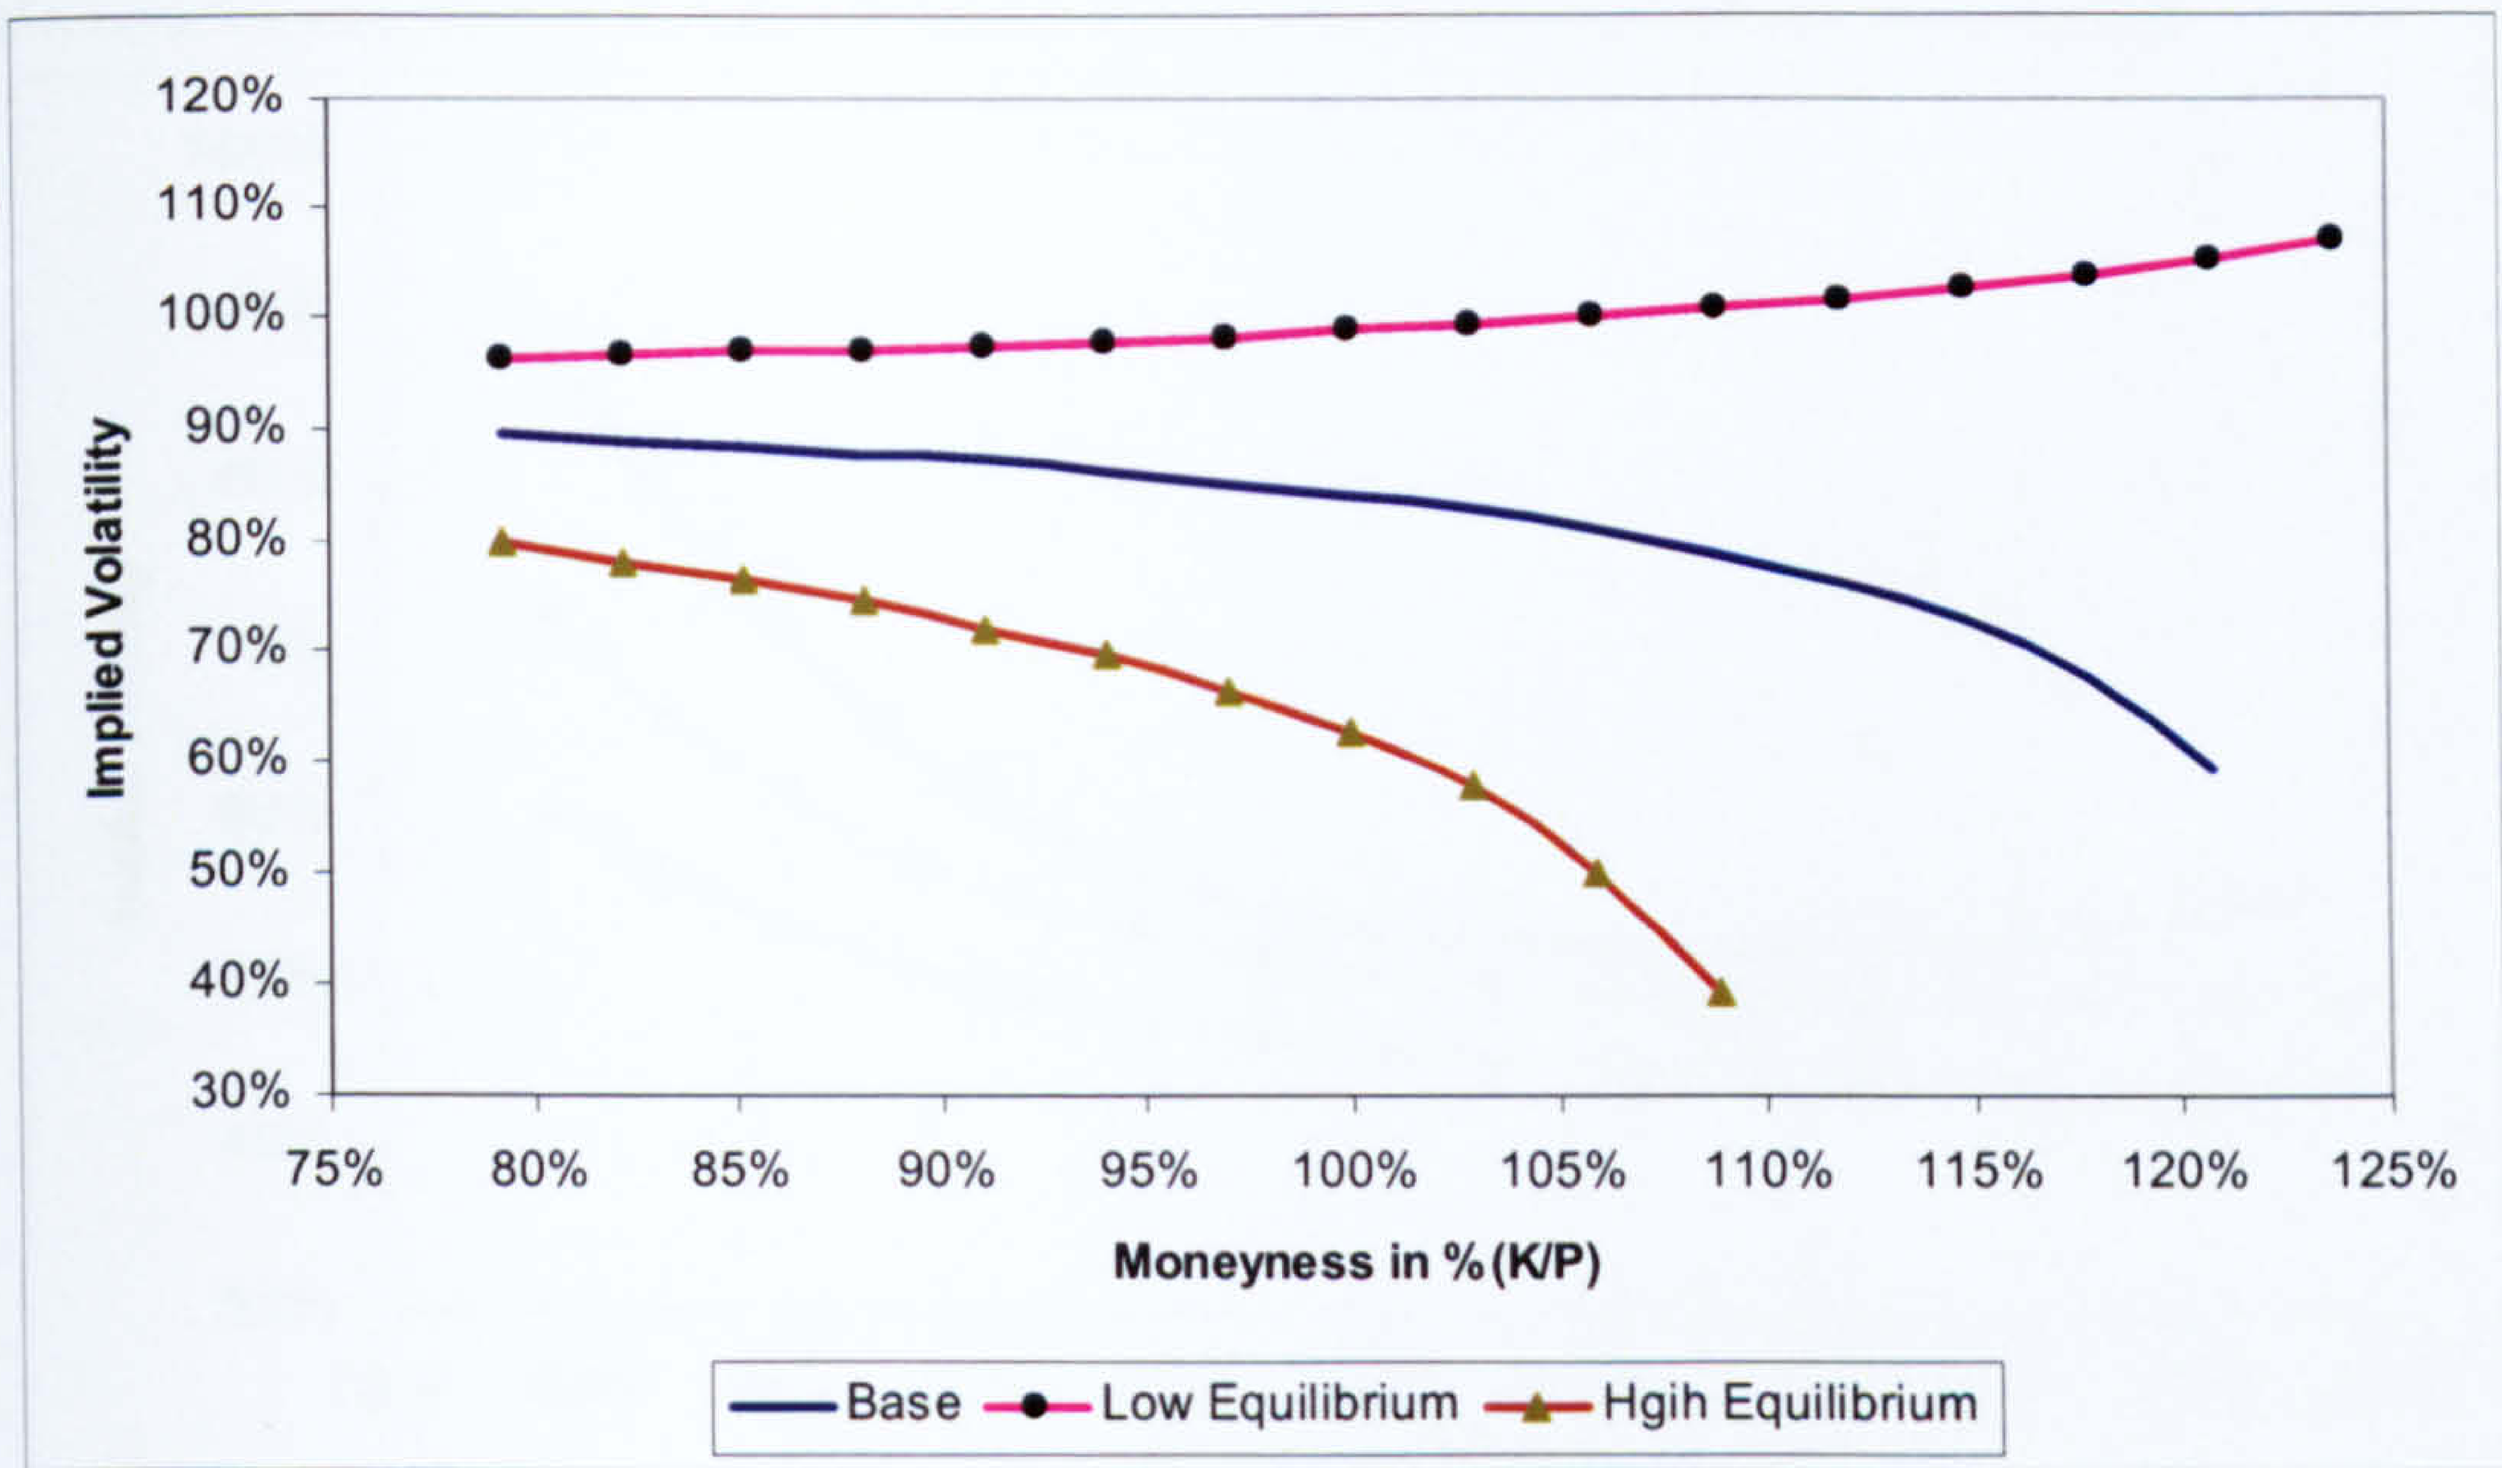

5.5.2.2 Volatility smile and Time to Maturity

Figure 5.3 compares the impact of time to maturity in the spike model. First thing to note is the fact that as the time to maturity increases, implied volatilities become lower, which is consistent with the Samuelson Hypothesis and depends on the speed of mean reversion, k_1. This occurs because the BS model implies that the variance of the spot prices should increase with time to maturity, where as the variance equation in (5.13) implies that the variability of the spot prices in the future converges to a constant value. Another way to explain this is to say that the volatility of the forward prices decreases with time to maturity. This is also an intuitive result for the contribution of the spikes in the model. We know that the speed of mean reversion in the spikes is much higher than that in the normal factor X. Thus the contribution to the volatility coming from the spikes is significantly lower than that of the normal process, the longer the maturity of the option. On the other hand, for short dated options (i.e. 15 days to maturity) the impact of jumps is more pronounced. From this point and onwards the graphs refer to call options, since similar conclusions can be made for both put and call options as we saw in the previous analysis.

Figure 5.3: Base case scenario of implied volatility skew from the spike model with time to maturity of 15 days, 1 and 2 months for European Call Options.

The figure shows the implied volatilities, for European Call Options, of the Spike model across different strike prices (Moneyness), using different maturities (15 days, 1 and 2 months), and using the Base Case parameters.

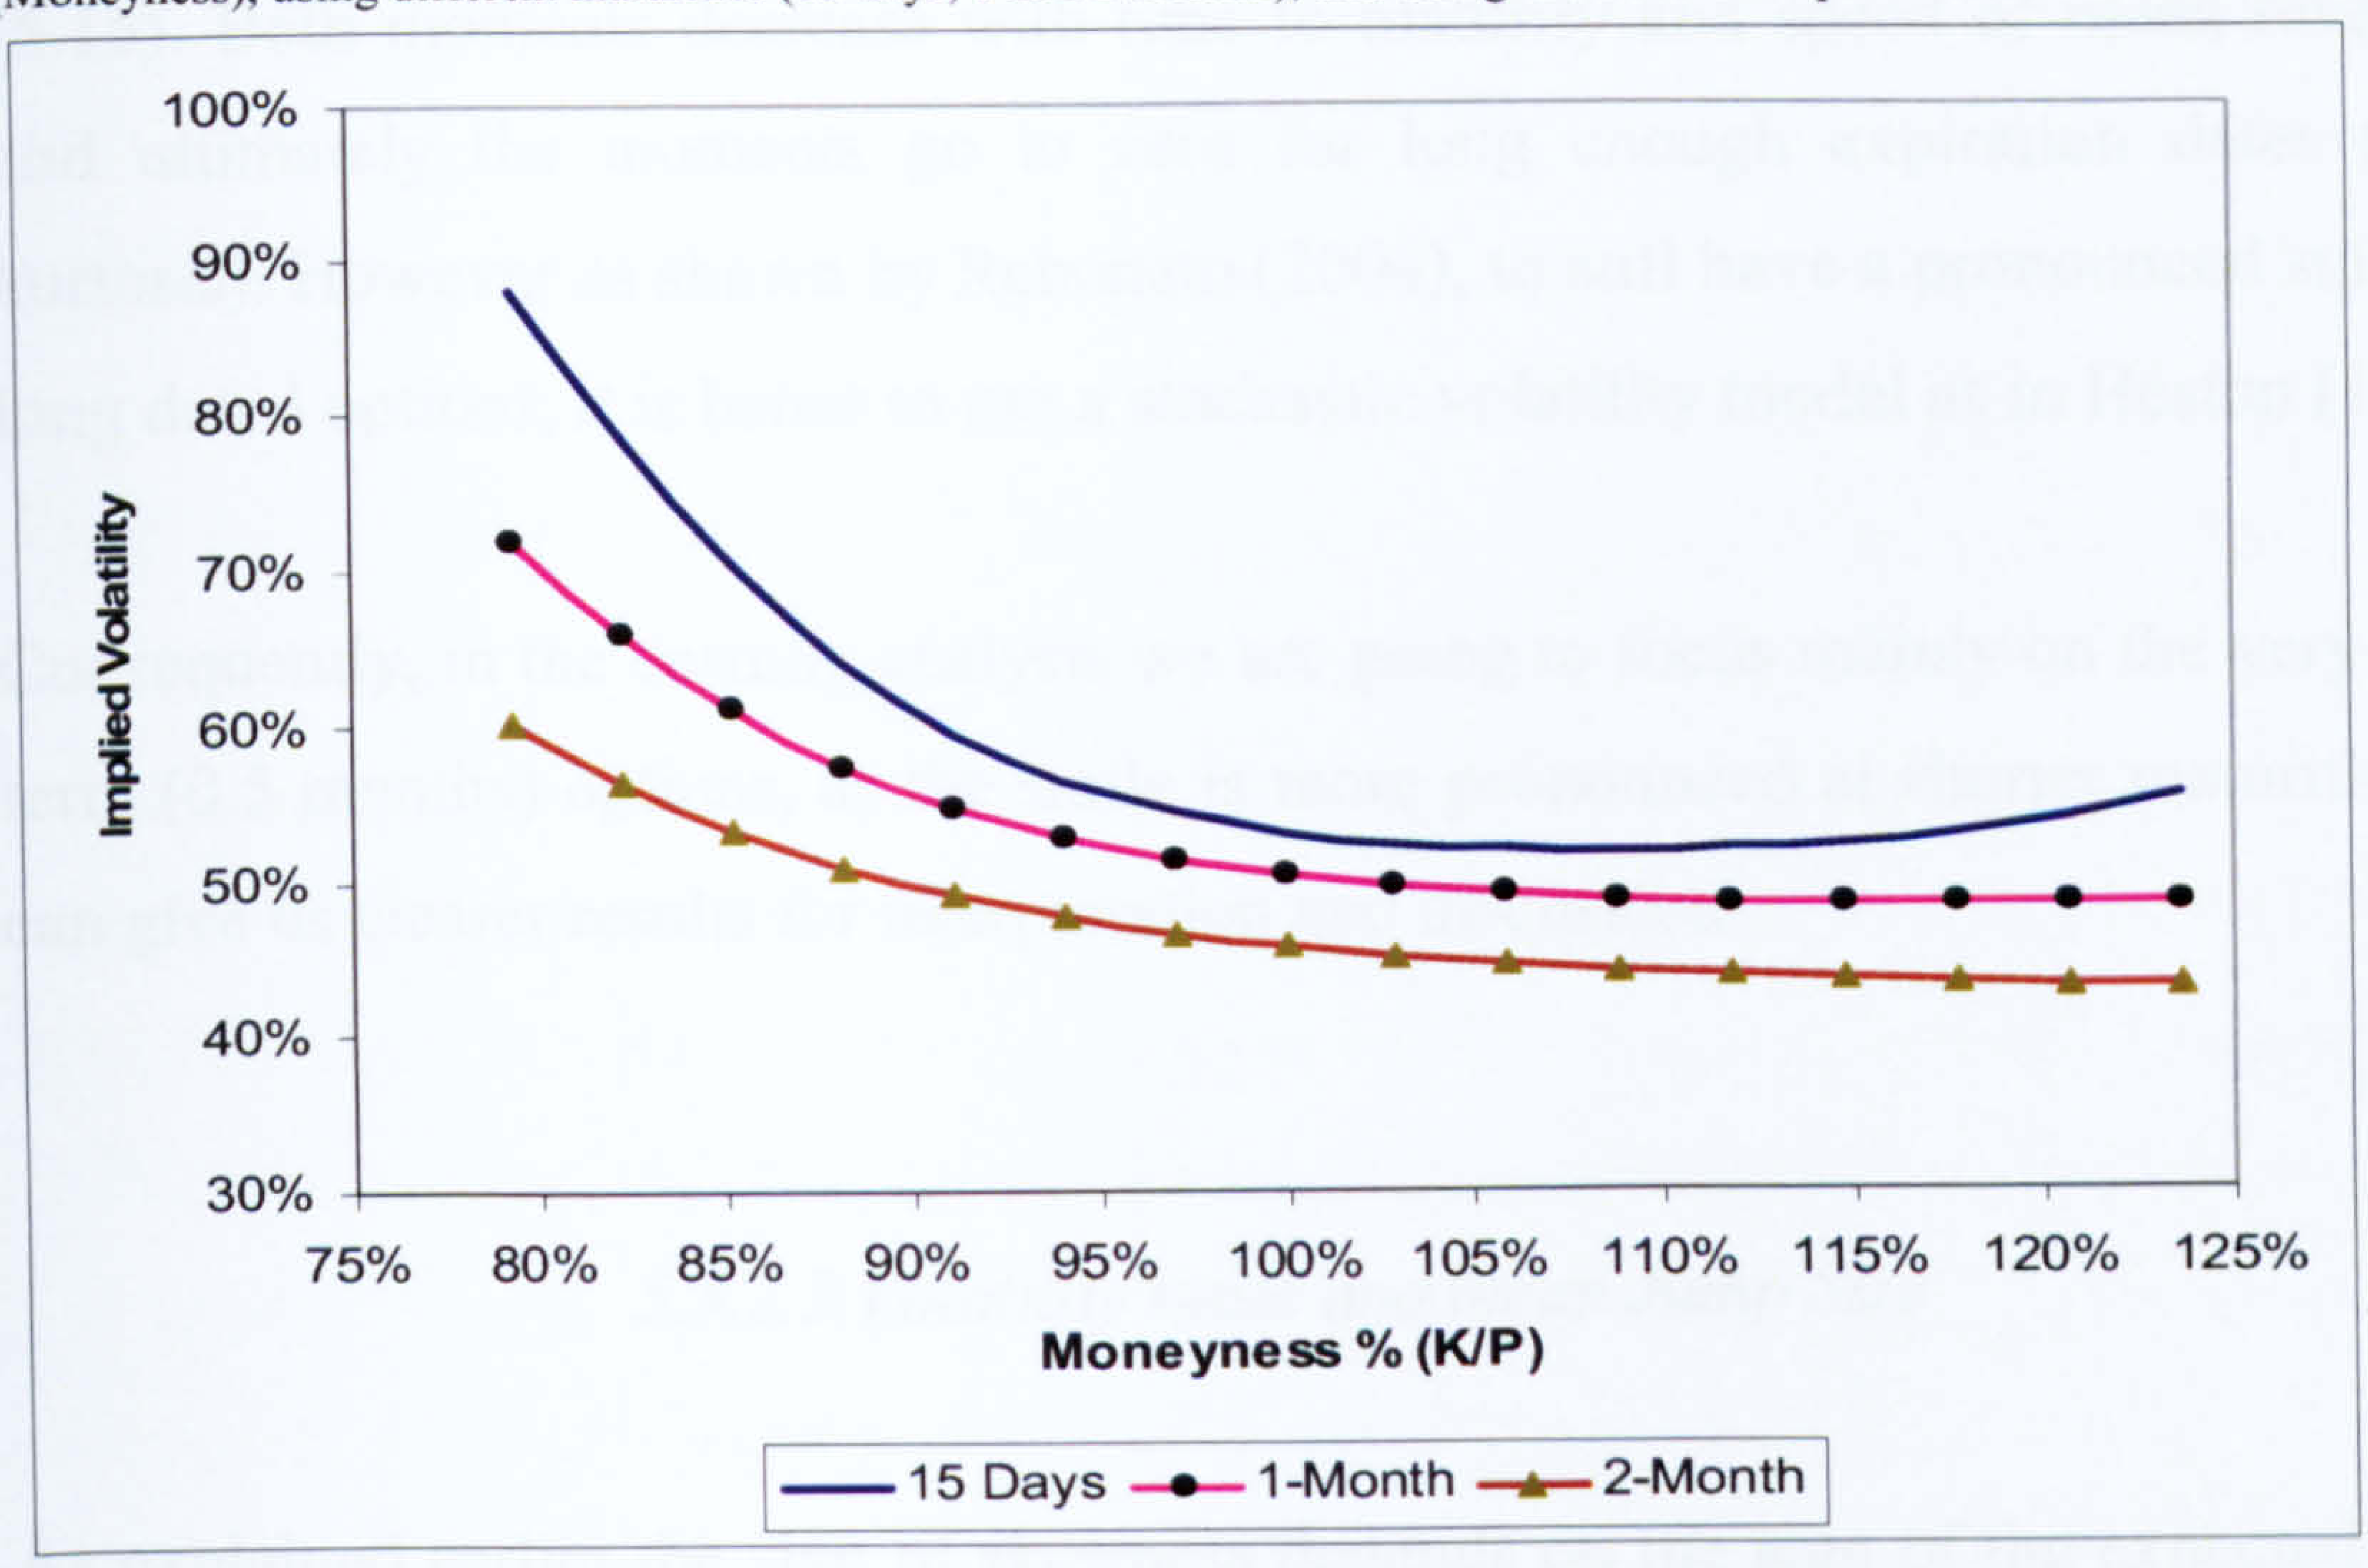

In a jump diffusion model such as Merton (1976), with positive mean jump size we would expect OTM call options to have higher implied volatilities than ATM and ITM options. However this is not generally the case when mean reversion is included in the model. The speed of mean reversion does not allow prices to deviate from the equilibrium level of 170.20 NOK/MWh[35]. Thus the probability of OTM calls ending up ITM is significantly reduced the stronger the speed of mean reversion, compared to the GBM model. On the other hand, prices cannot go to very low levels as well due to mean reversion again, thus ITM call options have a higher probability of ending way ITM than what the GBM might allow. Thus in a classical mean reversion model we should see implied volatilities, which decrease in an exponential fashion with strike prices, thus producing a negative skew as shown earlier.

Focusing now on the skewness and kurtosis, it is evident from Figure 5.3 that there is a pronounced smile for short-term options, however as the time to maturity increases the smile flattens out particularly for OTM options[36]. This pattern of large differences in the smile for different times to maturity and very steep smile for short times to

[35] Since the risk-neutral equilibrium level of X is 5.137, thus $exp(5.137)=170.20$.

[36] See Branger (2004) for evidence on how the smile flattens for long maturity options, when using Merton (1976) model.

maturity is consistent with a jump diffusion model. The intuitive reason for the significant impact of time to maturity in the spike model can be found by looking at the skewness and kurtosis of the risk neutral distribution in equations (5.14) and (5.15). Both moments decrease with time to maturity and speed of mean reversion, and ultimately the moments go to zero for long enough expiration dates (3 for kurtosis). However as shown by Rebonato (2004), to still have a pronounced smile for long dated options, it is better to use a stochastic volatility model as in Heston (1993).

Consequently, in the ensuing analysis we are going to focus mainly on the very short-term (0.5 months) options, as the smile is more pronounced at shorter maturities and can give us clearer results for interpretation and discussion.

5.5.2.3 Volatility smile and mean Jump Size

As explained earlier the sign of skewness depends on the sign of the expected jumps $E\left[Jdq(l)\right]$ in log-returns, a fact that is depicted in Figure 5.4. If μ_J is negative, the volatility smile seems to be upward slopping when moving from ITM to ATM call options and then becomes relatively flat. The effect seems to be very similar to that produced by the case of low equilibrium level in Figure 5.1. However in the current scenario option prices are not affected by the fact that the mean reversion pulls spot prices to lower equilibrium values but rather by the fact that there is a higher probability, than that predicted from a normal distribution, for prices to reach extreme low levels. It is worth noting that the pattern evidenced here is the opposite of what would have been predicted using Merton's (1976) model. To illustrate this, consider the effect the equilibrium level has on call option prices. The speed of mean reversion pulls prices towards the equilibrium level which, in the base case considered here, is higher than the exercise price for ITM call options and lower than the exercise price for OTM call options. Thus prices cannot deviate from their mean substantially and that benefits mainly ITM call options, since OTM options have less probability of being exercised. Now if we allow for extreme negative shocks, OTM options are not so much affected since mean reversion has already reduced their probability of being exercised and there is enough certainty that they will end OTM. For ITM options on the other hand, options from the point where they benefited from mean reversion and there was high probability that they would end ITM, the negative jumps reduce substantially this likelihood, and thus there is a higher probability that they will end OTM.

Therefore in the case of negative jumps, the probability mass is shifted to much lower spot prices, thus reducing the probability of the call option ending ITM. Of course when the expected jump size is positive, we see a more pronounced smile. Here the speed of mean reversion creates a pronounced skewness for ITM call options, whose value increases with the positive jumps. In addition, positive jumps also allow more extreme values to be reached, thus giving a greater probability of OTM options to end ITM than what BS might predict, making the smile more pronounced. Comparing the two curves, we can note that the implied volatilities from positive expected jump sizes are higher due to the greater likelihood of high spot prices, which increases probability of the call options ending ITM.

Figure 5.4: Implied volatility Skew for 15-Day Call Options with respect to the sign of mean jump size μ_J

The figure shows the implied volatilities, for European Call Options with 15 days to maturity, of the Spike model across different strike prices (Moneyness), using positive and negative mean jump sizes, μ_J, and the remaining parameters being the same as in the Bases Case.

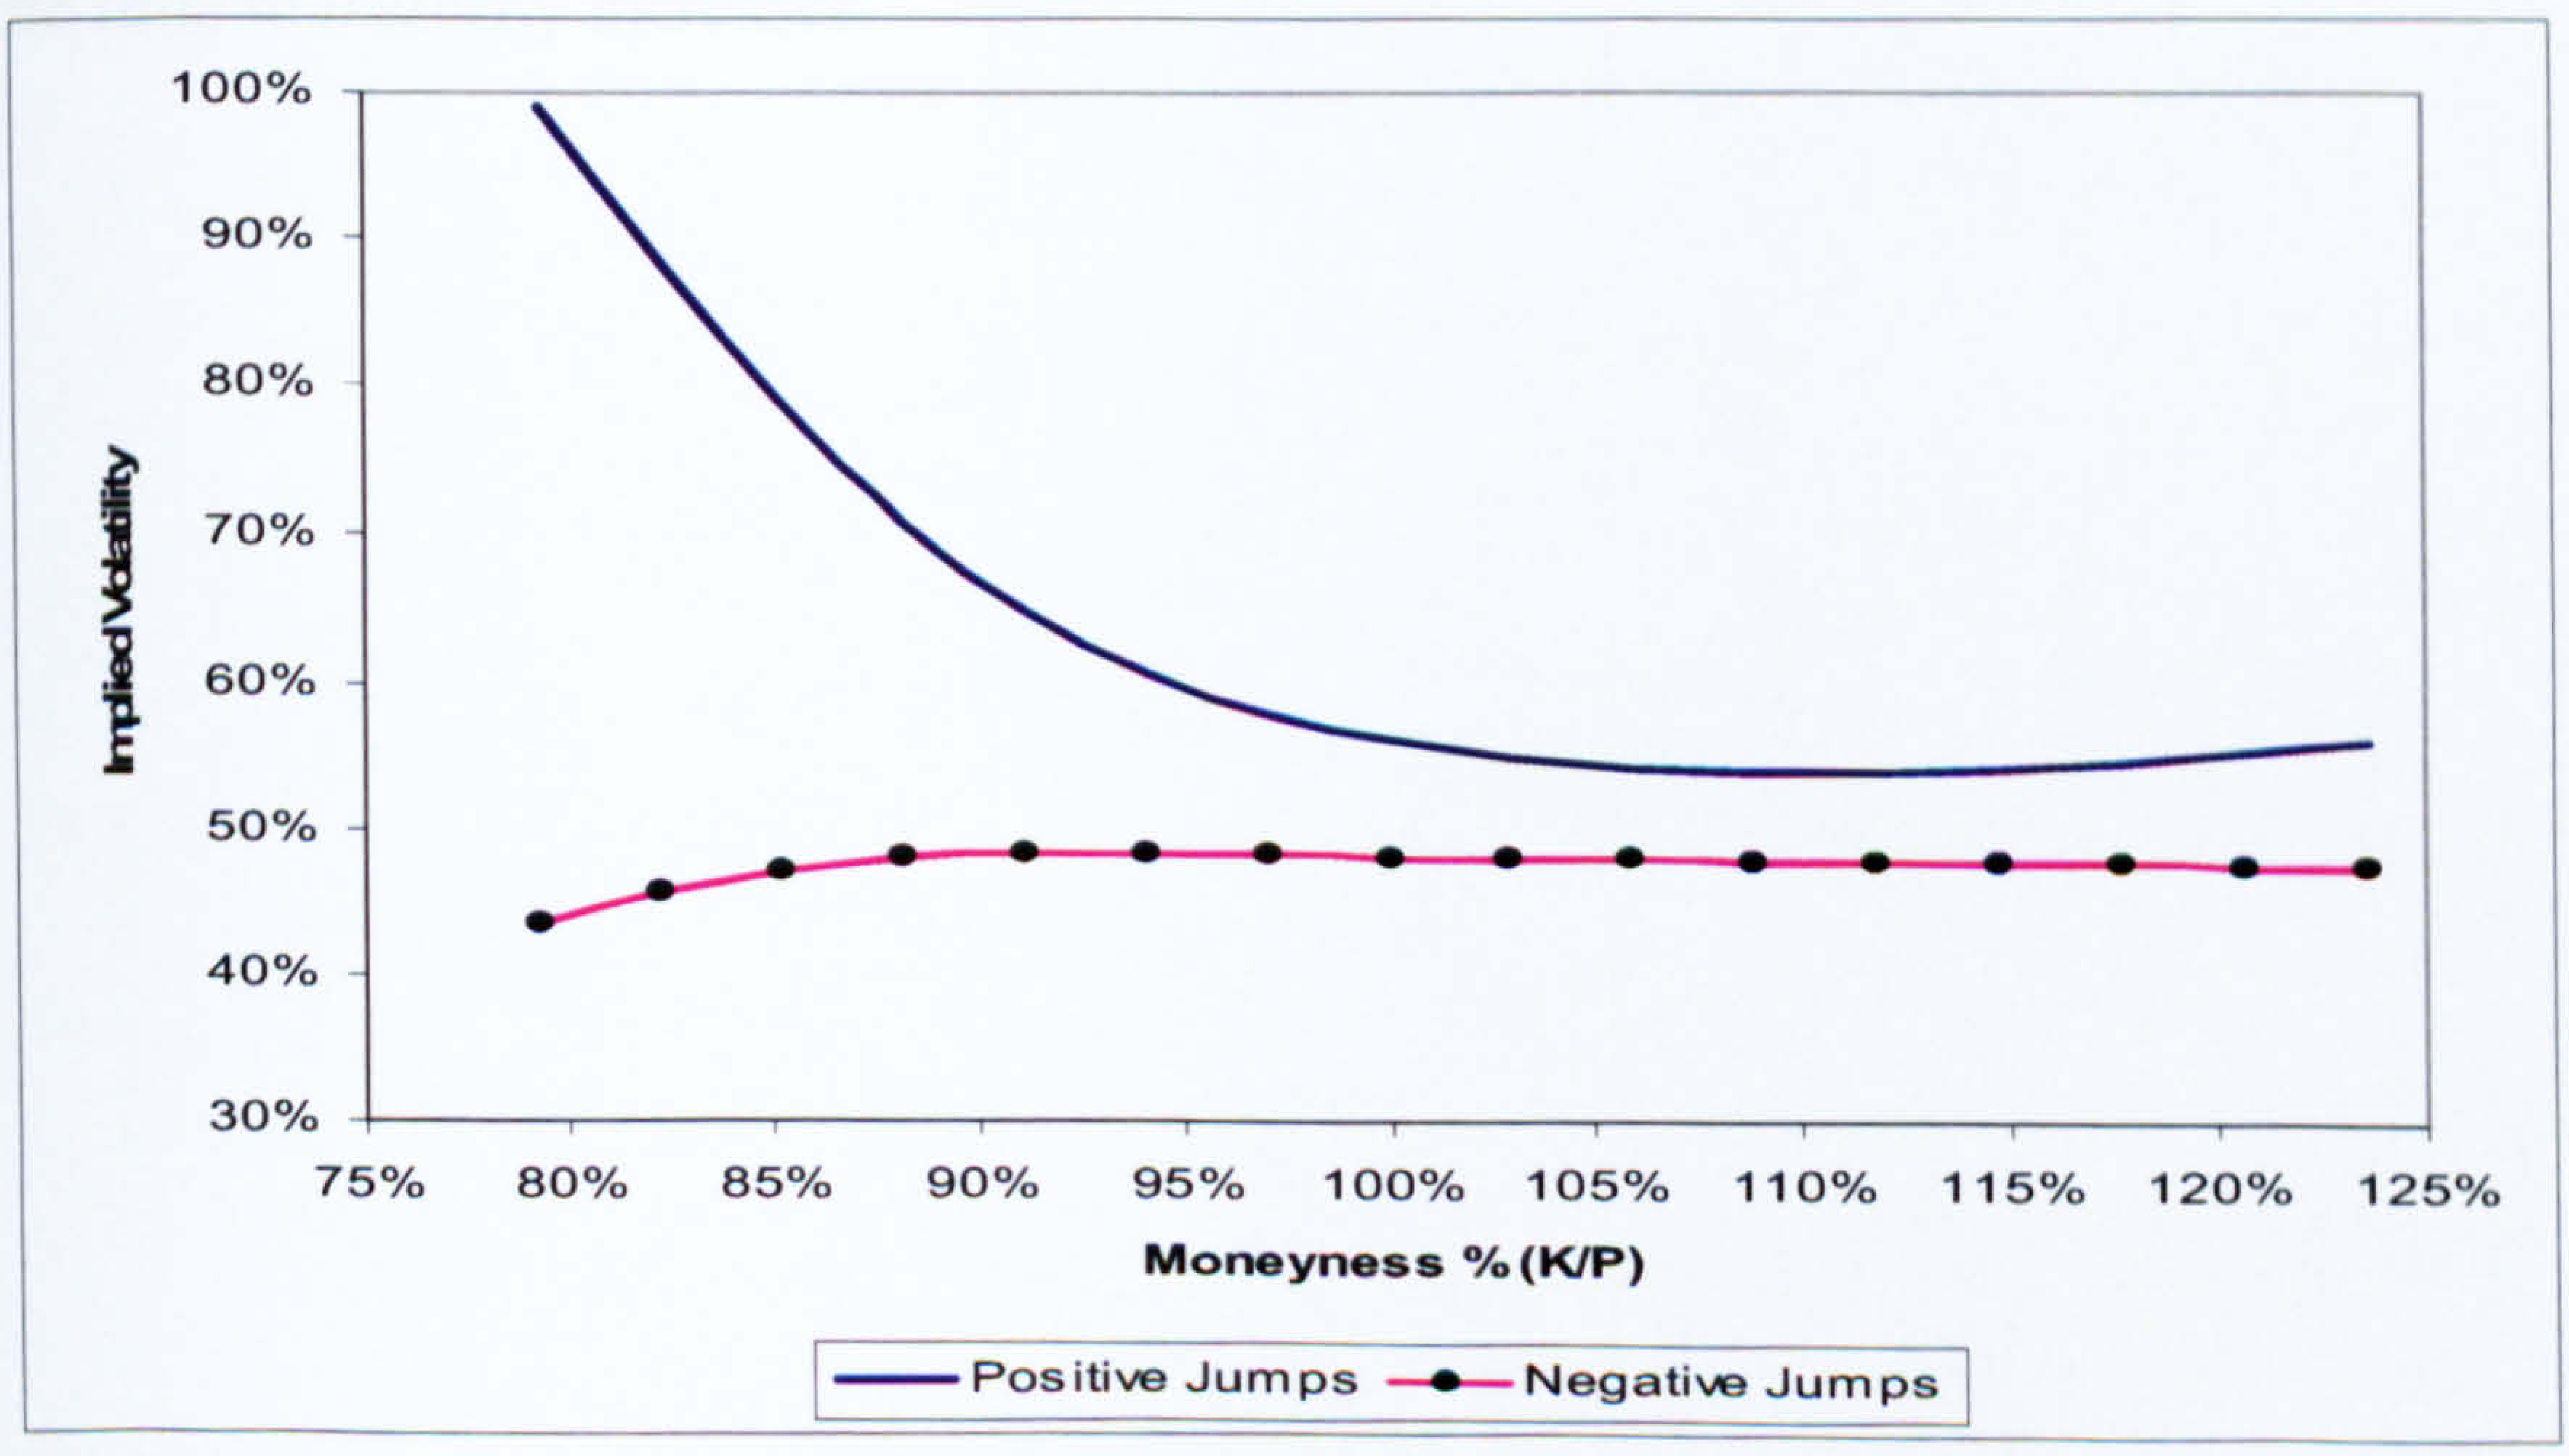

5.5.2.4 Volatility smile curvature; Jump Intensity versus Jump size volatility

Our last point is to analyse the impact of jumpiness, which is the contribution of the jump component in the total variance as shown in equation (5.13). A jump diffusion model can be interpreted as a mixture of a pure jump process without any diffusive component, as in the spike variable Y, and a diffusive process such as a GBM or a mean reversion. The issue is then to examine the impact of jumpiness on the smile, and whether this is affected differently by changing each of the constituent components of jumpiness, namely σ_J and l. The two extreme cases are when the jumpiness is zero and the implied volatilities for OTM options becomes flat, and when the jumpiness is one, and the smile has the maximum curvature. Figure 5.5 and Figure 5.6 display the implied model volatilities where the jumpiness, as a ratio of the total volatility given by equation (5.13), is increased to levels of 10%, 50%, 75% and 90%, by increasing the jump size volatility, σ_J, for options that are 0.5 and 2 months to maturity. In both cases we see that the smile becomes more pronounced as we increase jumpiness and the implied volatilities of course are higher, since the overall volatility increases (due the increases in jumpiness). The same can be observed in Figure 5.7 and Figure 5.8 when instead we increase the jump intensity, l, to increase the jumpiness. In all cases again the smile becomes significantly less pronounced as the time to maturity increases.

Figure 5.5: Implied Volatility Skew for 15-Day European Call Options for different levels of jumpiness due to changes in Jump size volatility σ_J

The figure shows the implied volatilities, for European Call Options with 15 days to maturity, of the Spike model across different strike prices (Moneyness), by changing the contribution of jumpiness to the total variance, with respect to σ_J and all other parameters staying the same as in the Base Case.

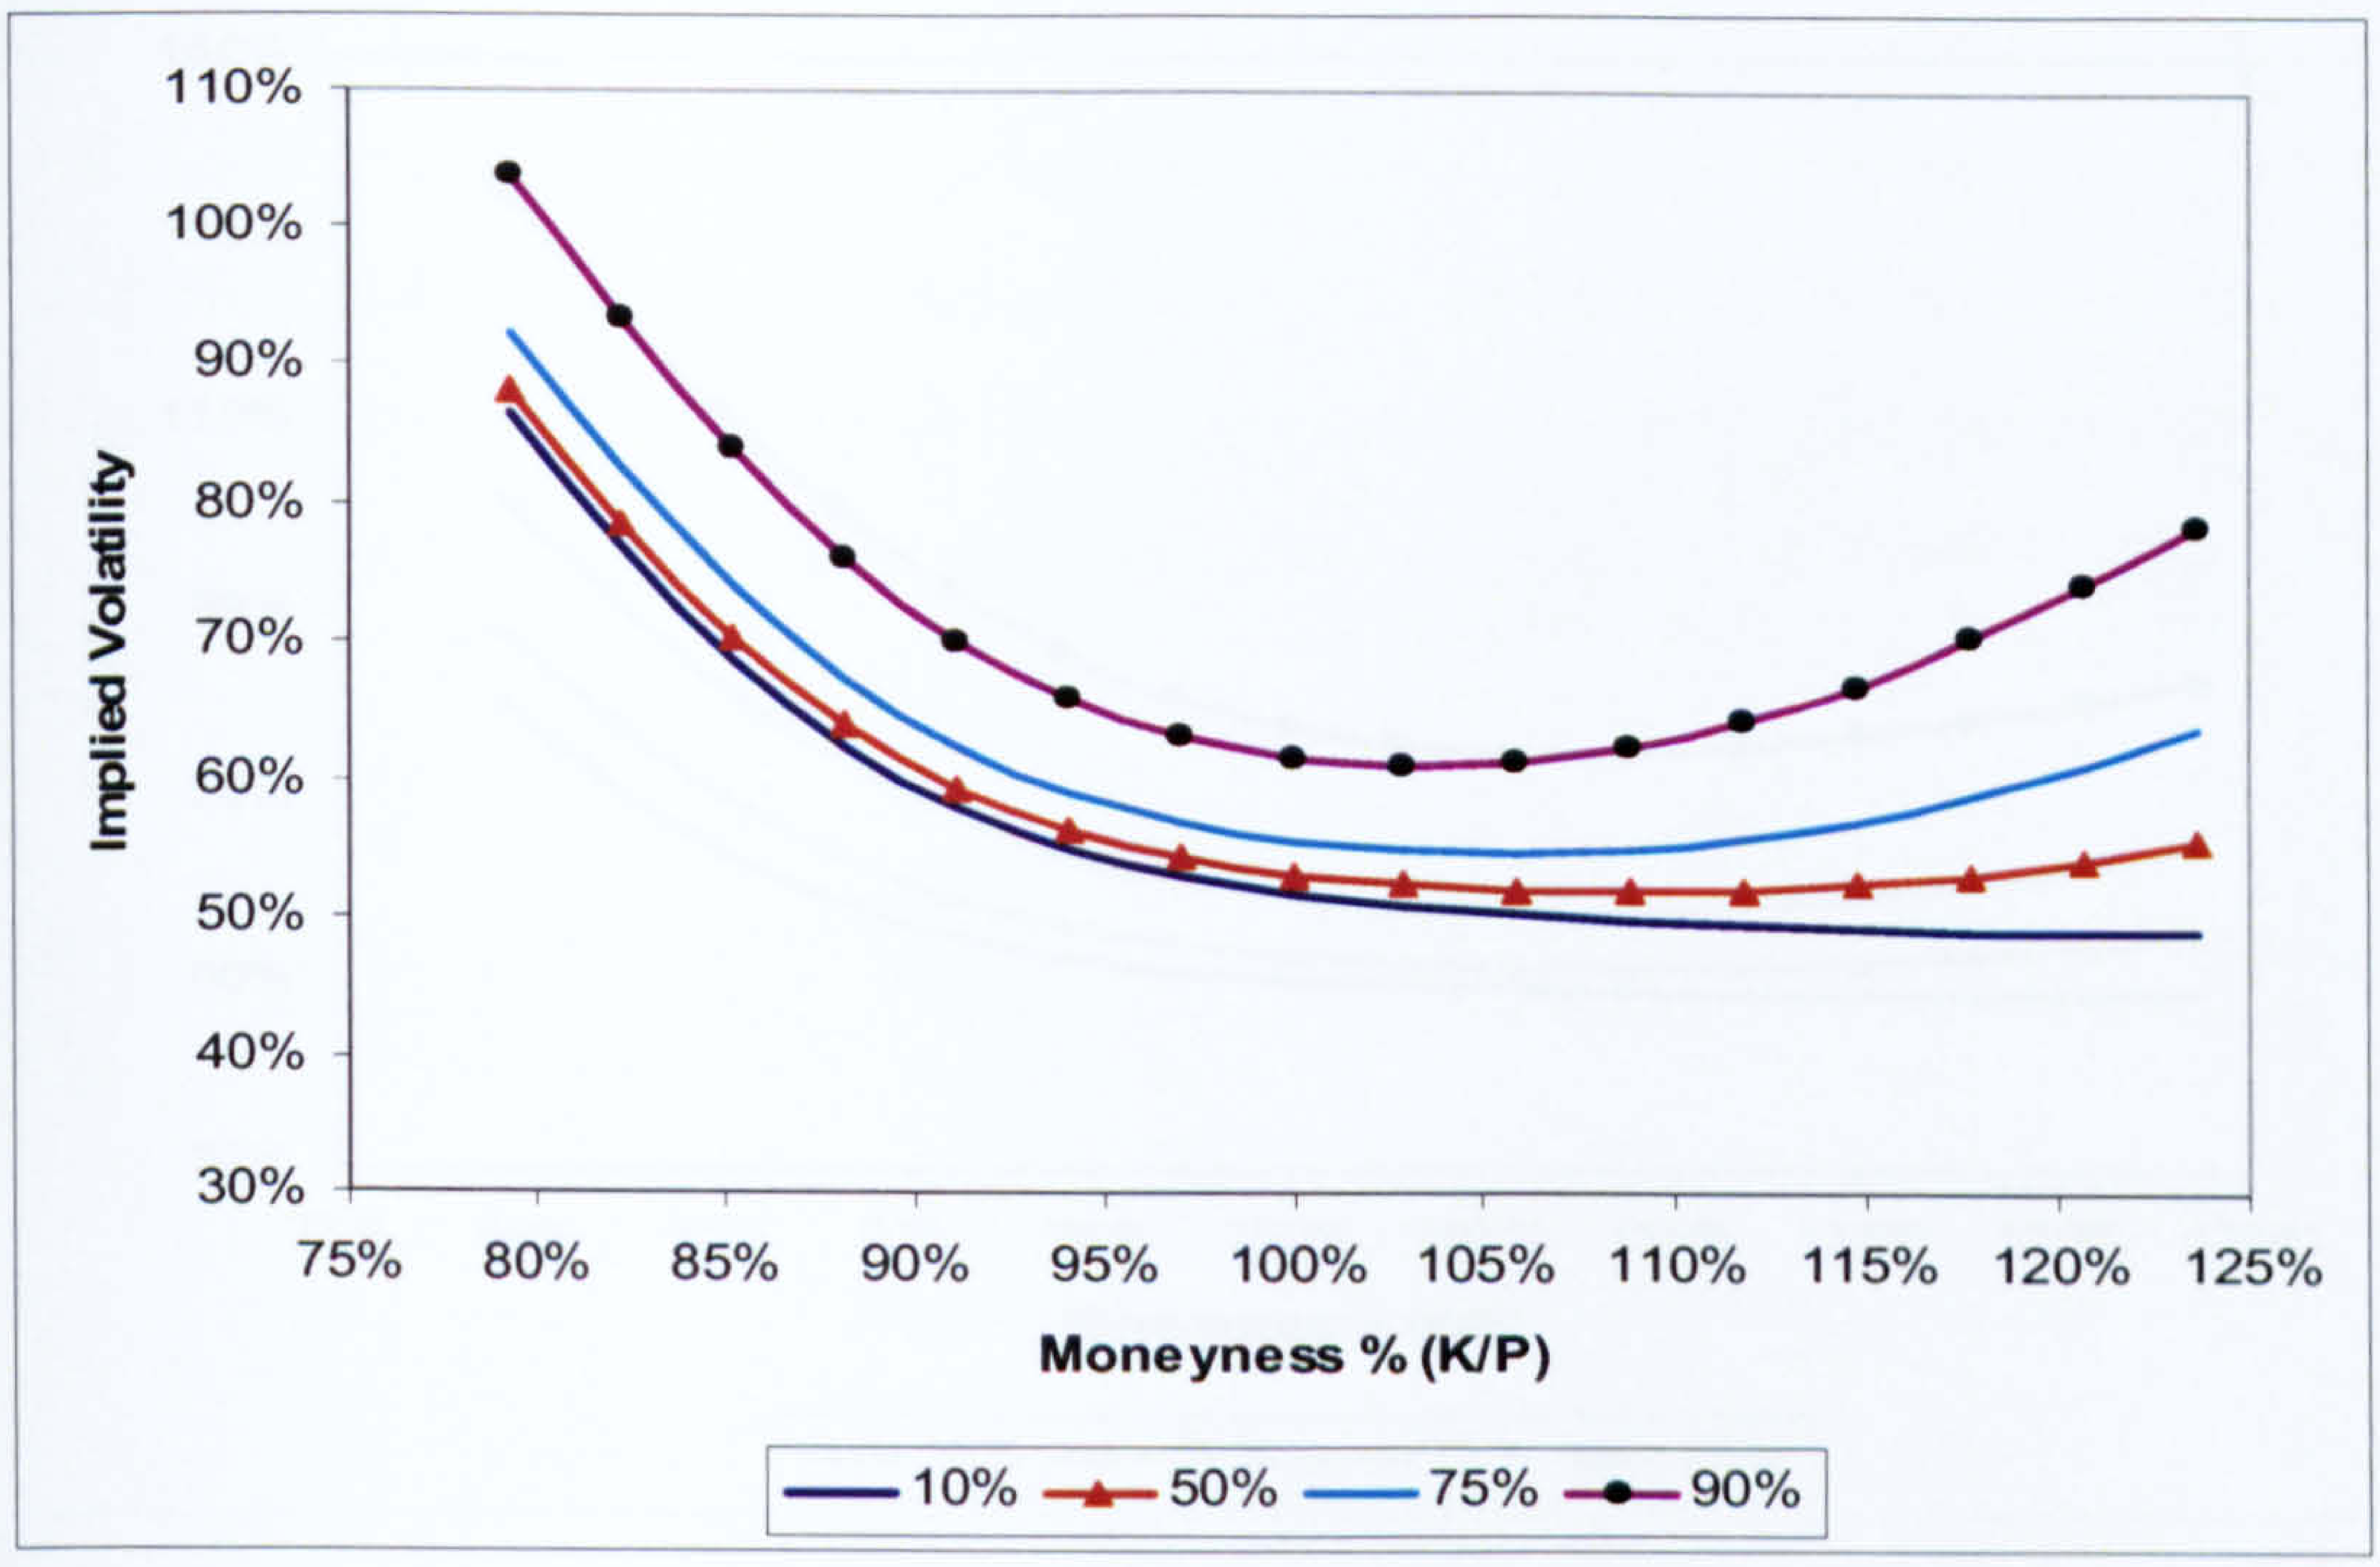

Figure 5.6: Implied Volatility Skew for 2-Months European Call Options for different levels of jumpiness due to changes in Jump size volatility σ_J

The figure shows the implied volatilities, for European Call Options with 2 months to maturity, of the Spike model across different strike prices (Moneyness), by changing the contribution of jumpiness to the total variance, with respect to σ_J and all other parameters staying the same as in the Base Case.

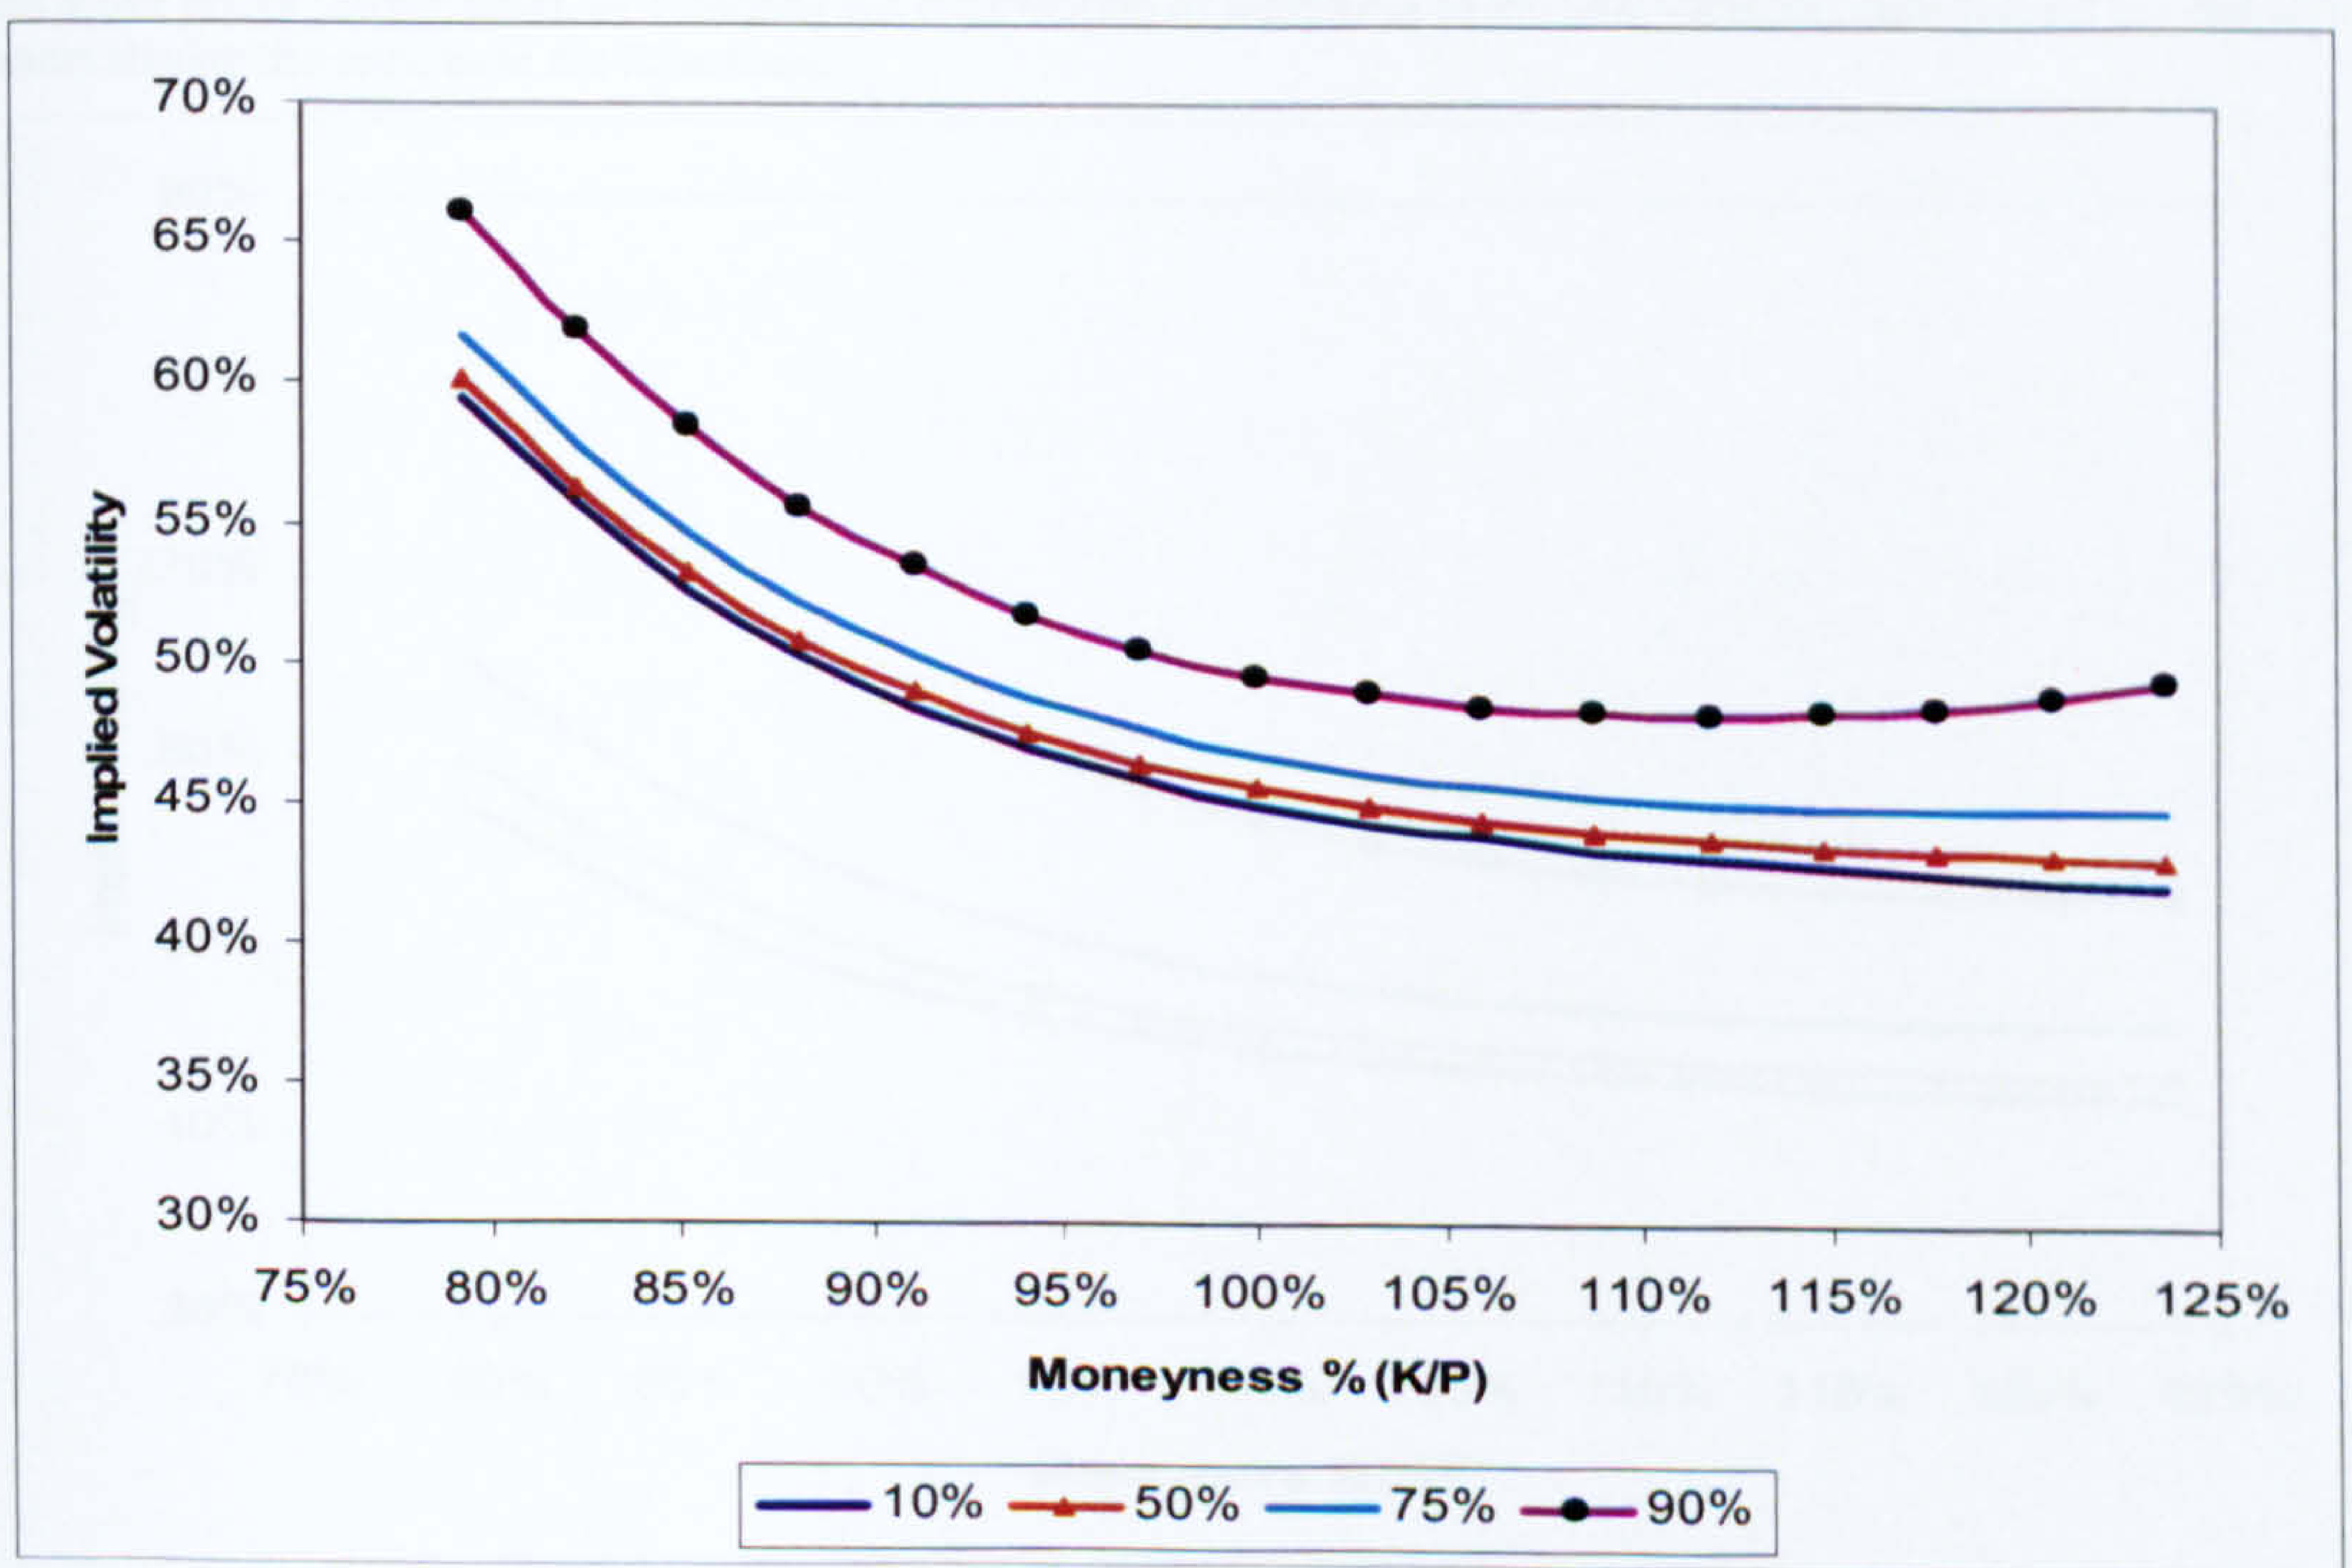

Figure 5.7: Implied Volatility Skew for 15-Day European Call Options for different levels of jumpiness due to changes in Jump intensity *l*

The figure shows the implied volatilities, for European Call Options with 15 days to maturity, of the Spike model across different strike prices (Moneyness), by changing the contribution of jumpiness to the total variance, with respect to *l* and all other parameters staying the same as in the Base Case.

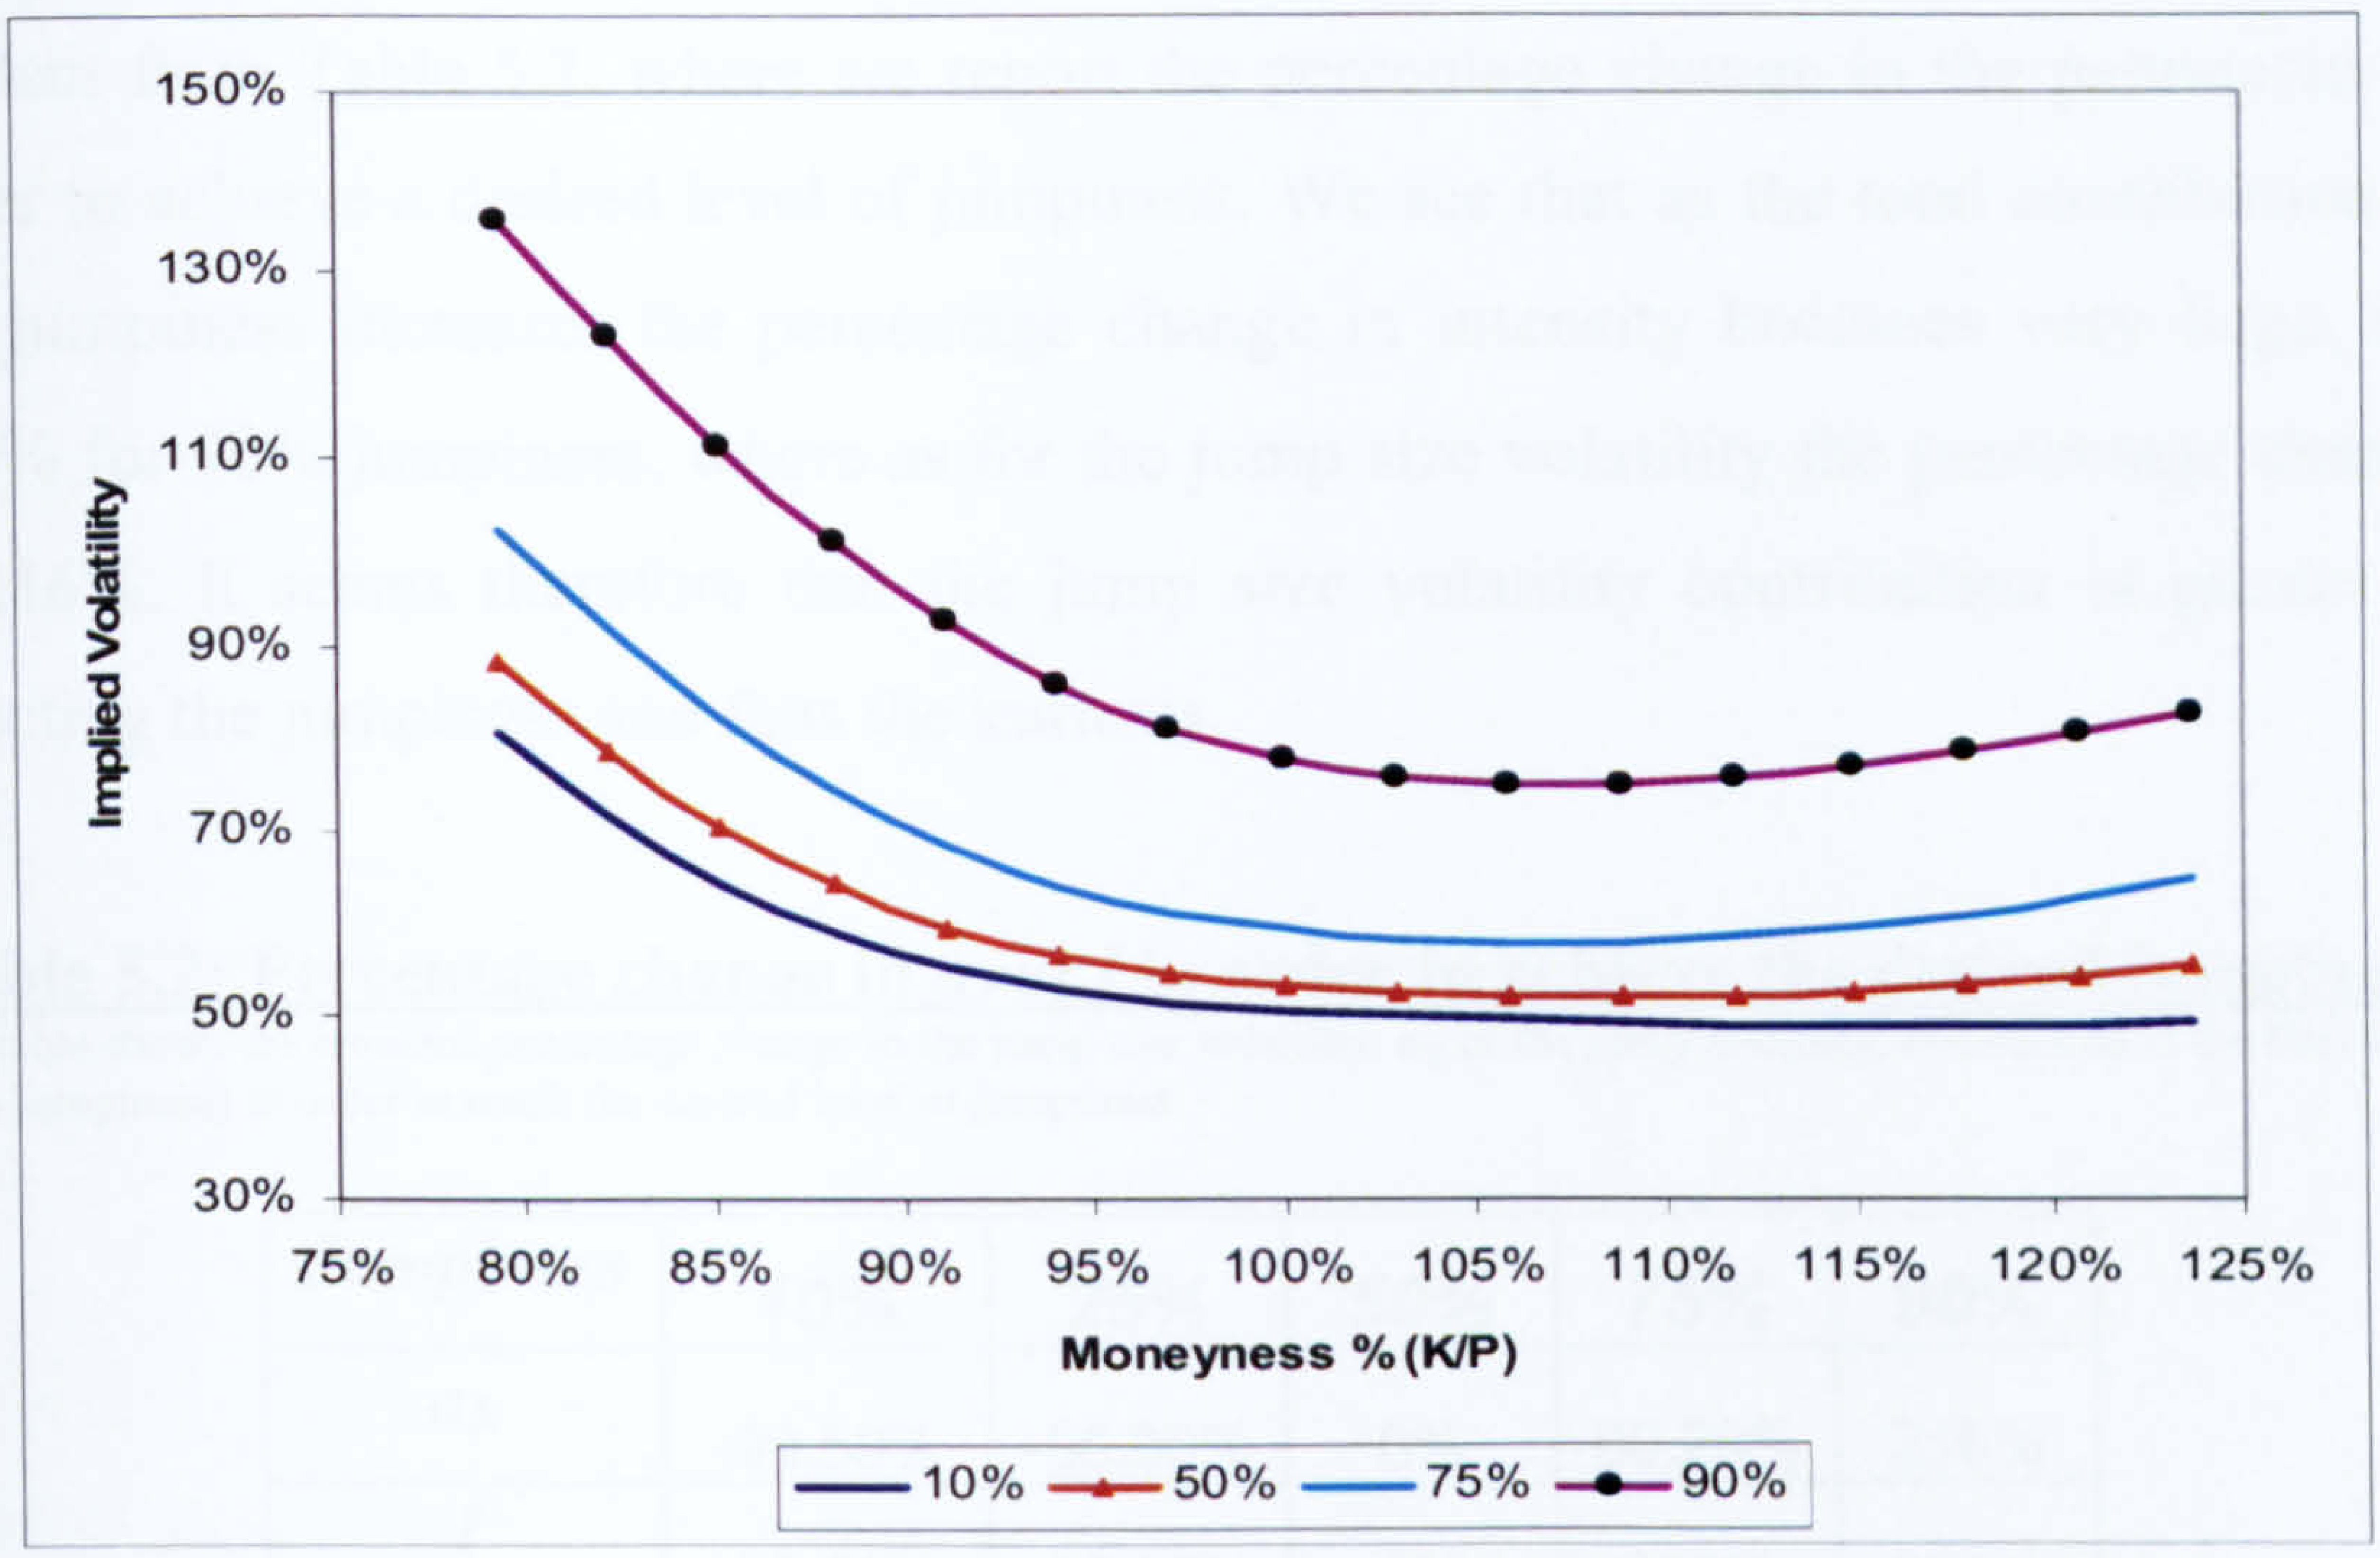

Figure 5.8: Implied Volatility Skew for 2-Month European Call Options for different levels of jumpiness due to changes in Jump intensity *l*

The figure shows the implied volatilities, for European Call Options with 2 months to maturity, of the Spike model across different strike prices (Moneyness), by changing the contribution of jumpiness to the total variance, with respect to *l* and all other parameters staying the same as in the Base Case.

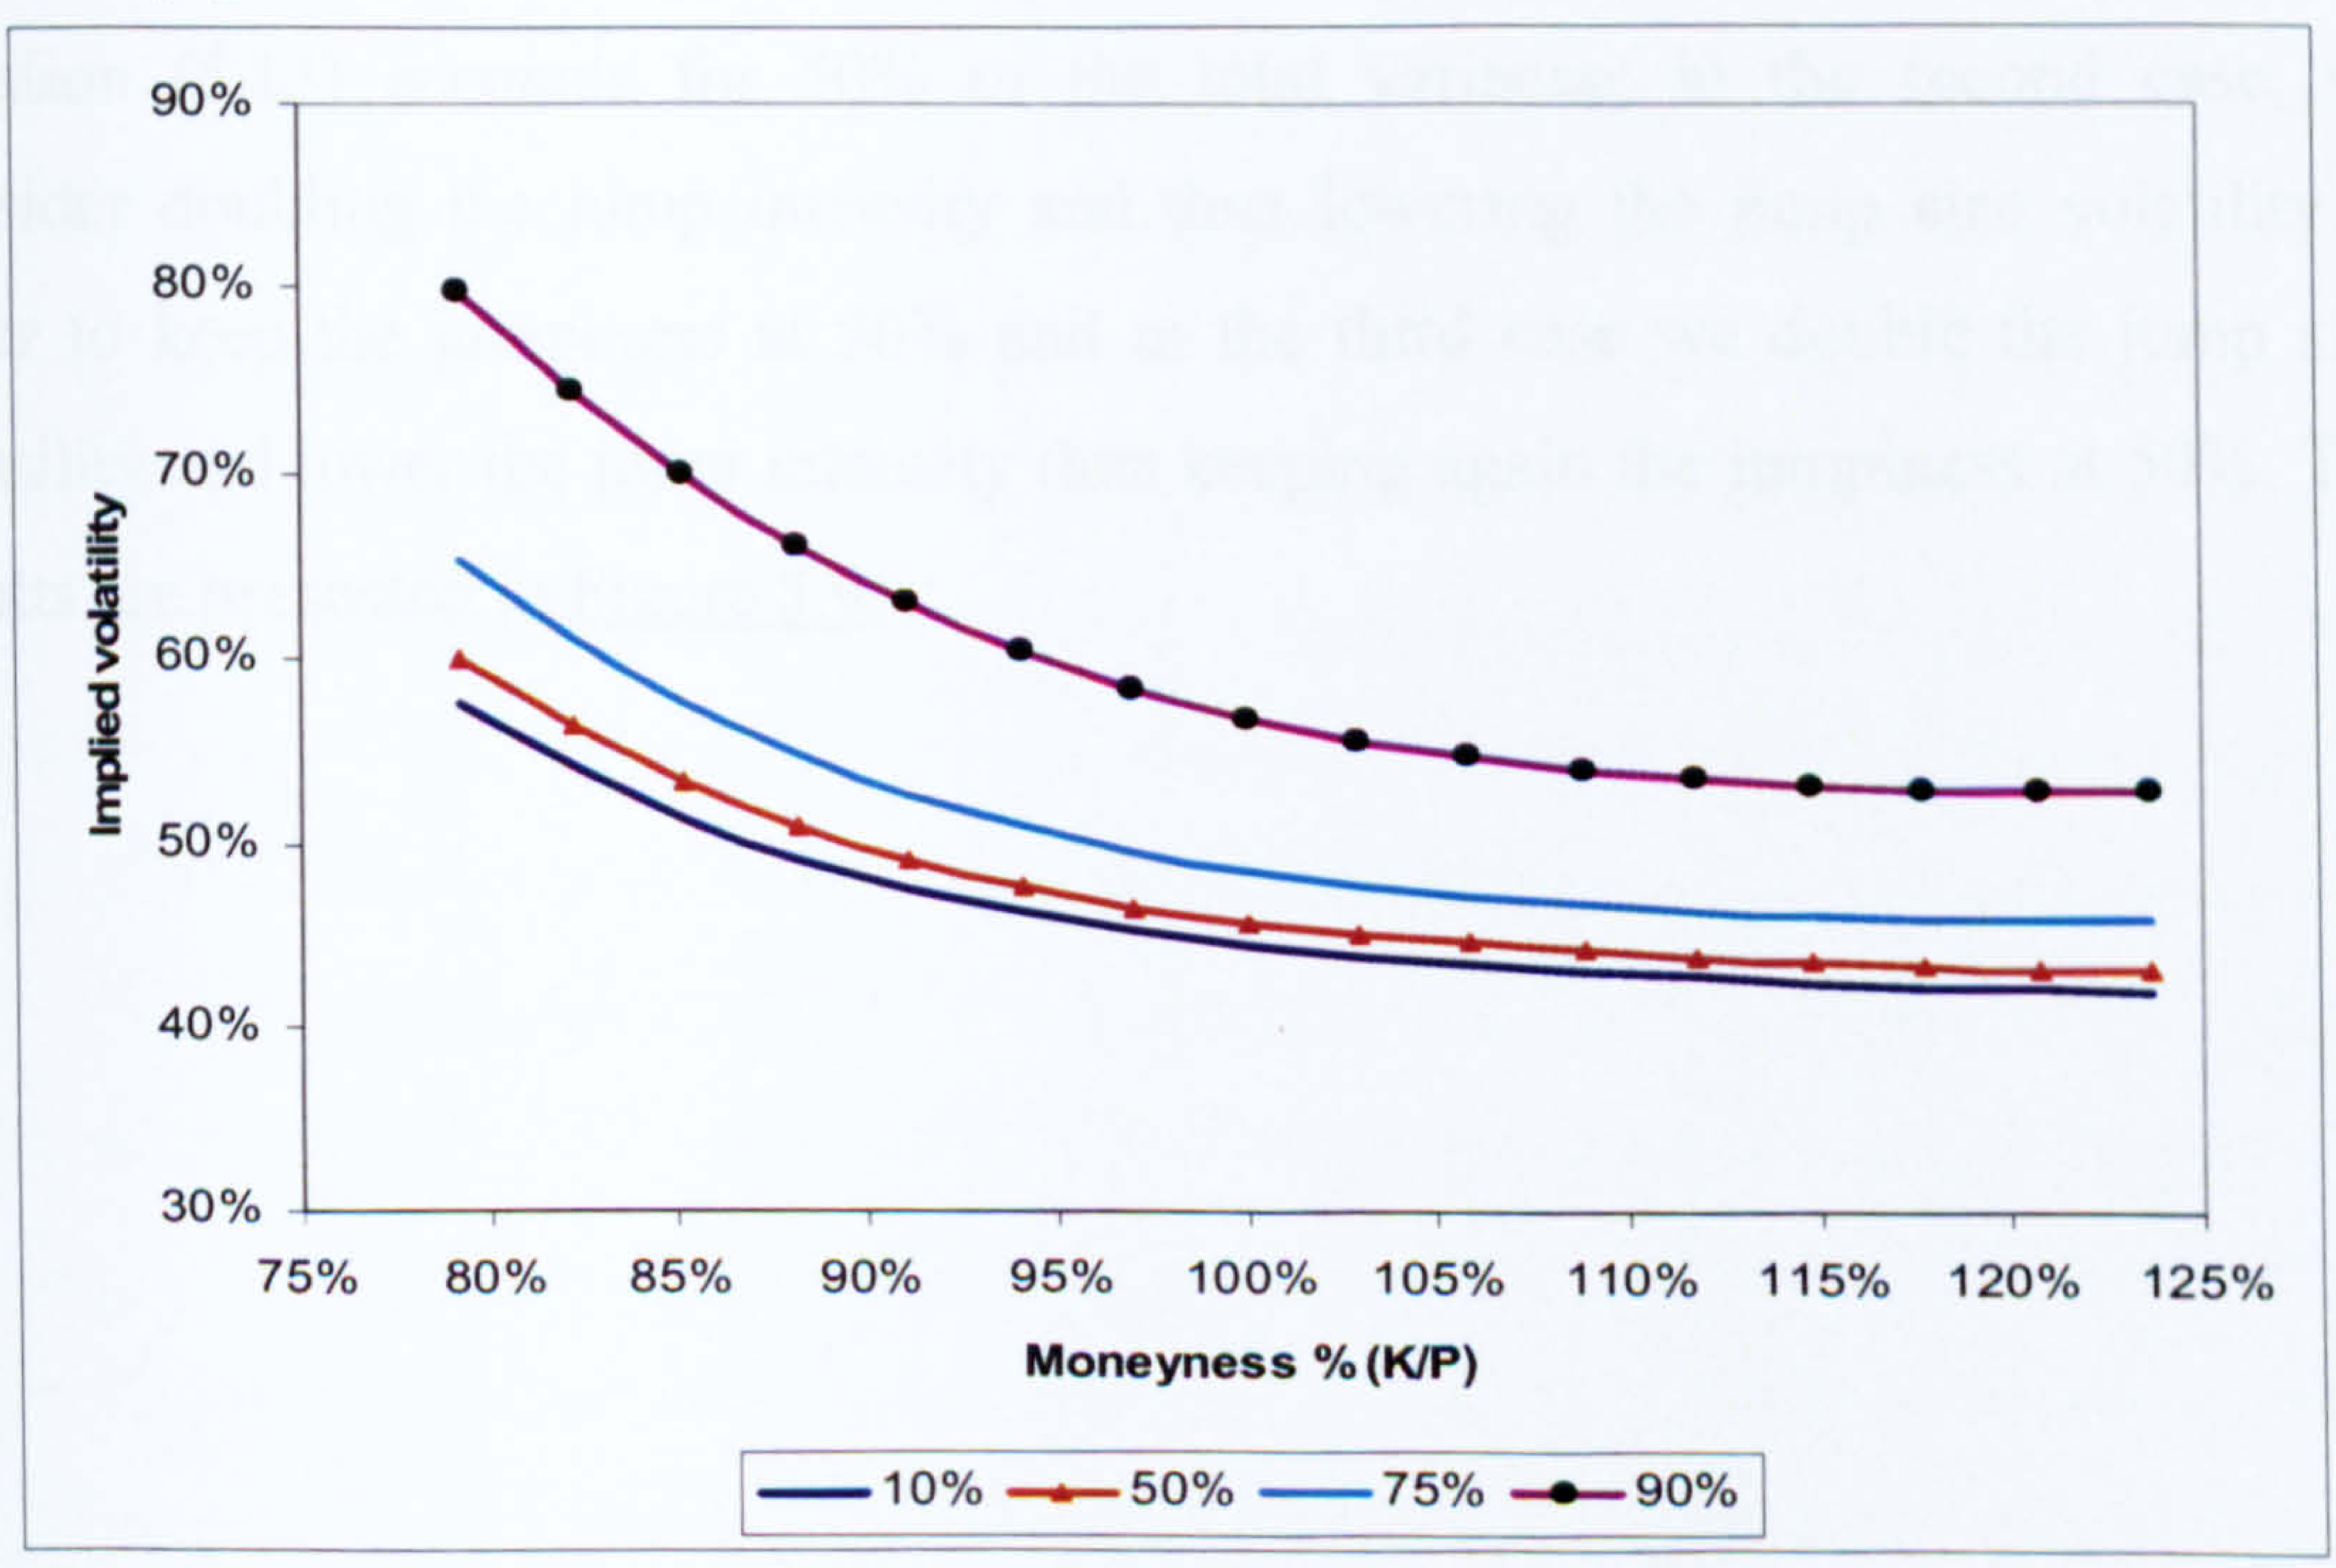

In this example we also note, that when the increase in the contribution of the jumpiness is due to increases in jump size volatility, the smile appears to be more pronounced than when it is due to the increases in jump arrivals. This fact is also evident from Table 5.2, where we report the percentage change in the parameters in order to achieve a desired level of jumpiness. We see that as the total contribution of the jumpiness increases the percentage change in intensity becomes very large, i.e. 800% for 90% jumpiness, where as for the jump size volatility the percentage change is 216%. It seems therefore that the jump size volatility contribution is greater in affecting the jumpiness and thus the kurtosis.

Table 5.2: Percentage change in σ_J or l in order to achieve the desired jumpiness

This table shows the required percentage change in the jump size volatility, σ_J, or the jump intensity, l compared to the Base Case (50% jumpiness) in order to reach the desired level of jumpiness.

Jumpiness	*10%*	*25%*	*50%*	*75%*	*90%*
σ_X	-99.59%	-50.00%	0%	80.28%	216%
l	-88.89%	-66.67%	0%	200%	800%

A more appropriate method for assessing whether the smile curvature is affected more from the jump size volatility or intensity is to employ the following three scenarios; one is the base case using the same parameters as in Table 5.1 where the jumpiness in equation (5.13) accounts for 50% of the total variance; in the second case, we consider doubling the jump intensity and thus lowering the jump size volatility in order to keep the jumpiness at 50% and in the third case we double the jump size volatility and lower the jump intensity thus keeping again the jumpiness at 50%. The results are presented in Figure 5.9.

Figure 5.9: Implied volatility Skew for 15-Days European Call Options with respect to Jump Intensity l and Jump Volatility σ_J

The figure shows the implied volatilities generated from the spike model, when doubling either the jump size volatility or the jump intensity, but always keeping the level of jumpiness at 50%.

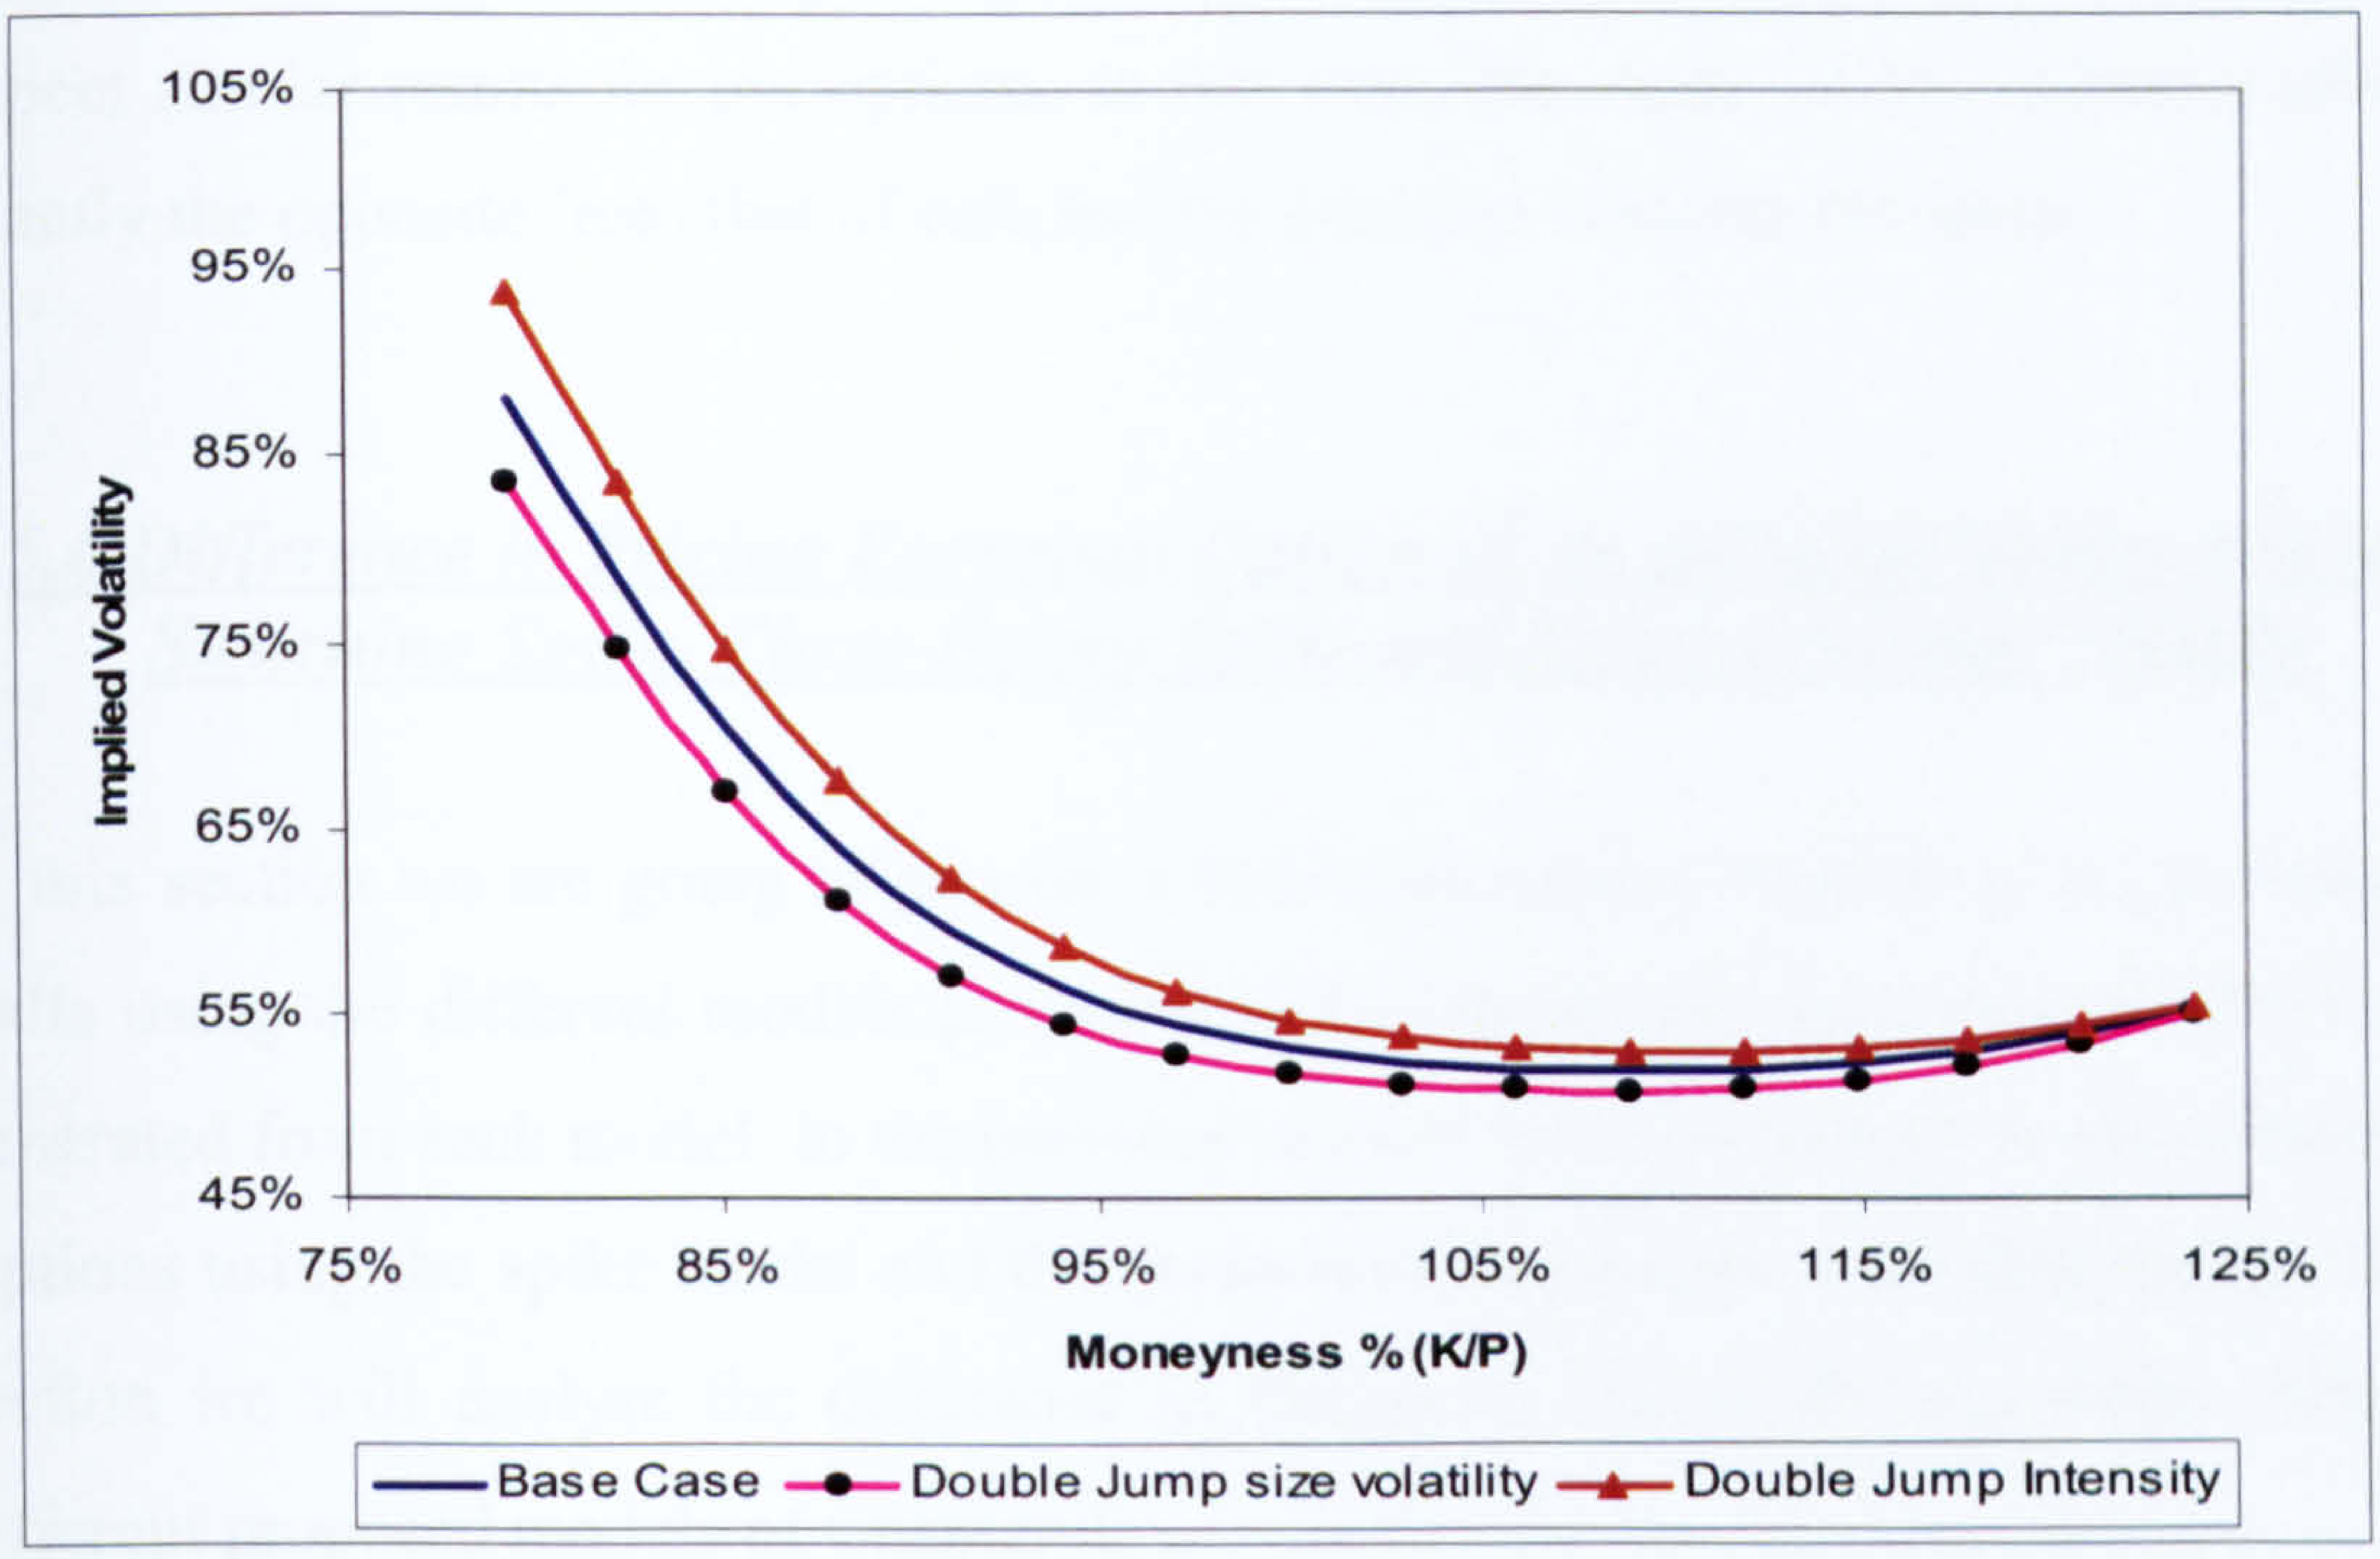

We can see that by doubling the jump intensity l, the overall level of the implied volatilities increases and there is still a pronounced smile, but the slope of the implied volatility curves (mainly for OTM call options) is lower than that when doubling σ_J, since for ITM call options the implied volatilities for the double intensity is at higher levels, and as we get closer to OTM the implied volatilities from doubling σ_J, converge to those of doubling l. This implies that for an increase in the jump size volatility, implied volatilities of OTM options increase more than the implied volatilities of ATM and the volatility curve is steeper than when increasing the jump intensity. This is also consistent with the analysis shown in Figure 5.5 to Figure 5.8, which show a more pronounced smile when the jumpiness level is changed with respect to the jump size volatility. Thus by keeping the spot returns variance constant, if the jump volatility is reduced and this is offset by an increase in the jump intensity, then the excess kurtosis of the distribution is reduced, making the smile flatter. On the other hand if an increase in the jump size volatility goes together with a decrease in the jump arrivals, the smile becomes more pronounced due to the increase in the kurtosis. Thus from this result we can see how important it is for a model to capture the low intensity high volatility jumps, as it is done using the filtering method of estimation, rather than the opposite as the Maximum Likelihood methods do. In all cases ITM volatilities are higher again due to mean reversion. Of course we would expect similar results for put options; in this case, the shape of the volatility smile is exactly the opposite from that of call, but the intuition remains the same.

5.6 Difference in Pricing European Option when using the Spike, Regime Switching Spike, Three Factor Spike and Mean-reverting Models

In this section we are going to provide a brief description on how to price European Calls using the different models proposed in this thesis and also compare the results generated from each model. In the previous section we showed how to price European Options using the spike model and the mean-reverting model with seasonality. In this section we will analyse the difference in European Option Prcing, when using the different proposed models of Chapter 4.

In the Regime-Switching Spike model, the equilibrium levels switch between two different regimes; Regime w is when water reservoir levels are above their seasonal average and the equilibrium price is low, and Regime d is when water reservoir levels are below their seasonal mean in which case we have to rely on more expensive fuel generators and thus the equilibrium level is higher. The pricing of European Options under the regime-switching spike model has the same intuition as that of the pricing of forwards. That is we first calculate two values for the option; one using the equilibrium level in Regime d when the water reservoir levels are low, $Call_d$, and another using the equilibrium value when we are in Regime w, $Call_w$. The formula for calculating each of these two separate option values is exactly the same as the one we used for the spike model in equation (5.10). Thus under the regime switching model the value of the option is the weighted average of the two separate options prices using a different equilibrium level, where the weights denote the probability of being in their respective states at that specific time in the future when the option expires. Thus, a call option for example is given by:

$$call = p_w call_w + p_d call_d \tag{5.19}$$

On the other hand when using the Three-Factor Spike model, the pricing formula is exactly the same as (5.11) to (5.12), however this time $X = [X, Y, \varepsilon]$, $1'$ and $0'$ denote $3x1$ transpose vectors of ones and zeros respectively and $\bar{v} = [v, v, v]$ and we used the derived transformed function from the previous chapter the same way we did for the spike model.

Returning to the option pricing analysis Figure 5.10 shows the values of European call options using the spike model of equation (5.10), for options starting in February with a time to maturity of up to 2 years across different strike prices. Seasonality is induced both by the deterministic variable $f(T)$ as well as by the jump parameters using the same ones calibrated in the previous chapter. From the graph it is very clear that seasonality affects directly the option price. For instance, call option prices peak at 0.92 and 1.92 years to maturity, which correspond to the months of January when electricity demand is at its highest. On the other hand, call option values have some sharp drops at maturities of 0.25 and 1.42 years, which correspond to non-working days. Excluding seasonality, we see that the price of the option generally increases with time to maturity, however seasonality includes some cyclical effects as well as some non-smoothing patterns due to the change of the jump parameters between different seasons or due to the fact that sometimes the option price expires during a holiday or a weekend[37].

[37] In such cases we saw in the previous chapter that spot prices fall significantly, due to lower demand, and we incorporated this in our model using dummy variables.

Figure 5.10: European Call Option Price for the Spike Model Across Time to maturity and Moneyness (including seasonality in both constant and jump parameters)

The figure shows the prices of European Call Options generated using the spike model in equation (5.10), across different levels of moneyness (*K/P*) and time to maturity.

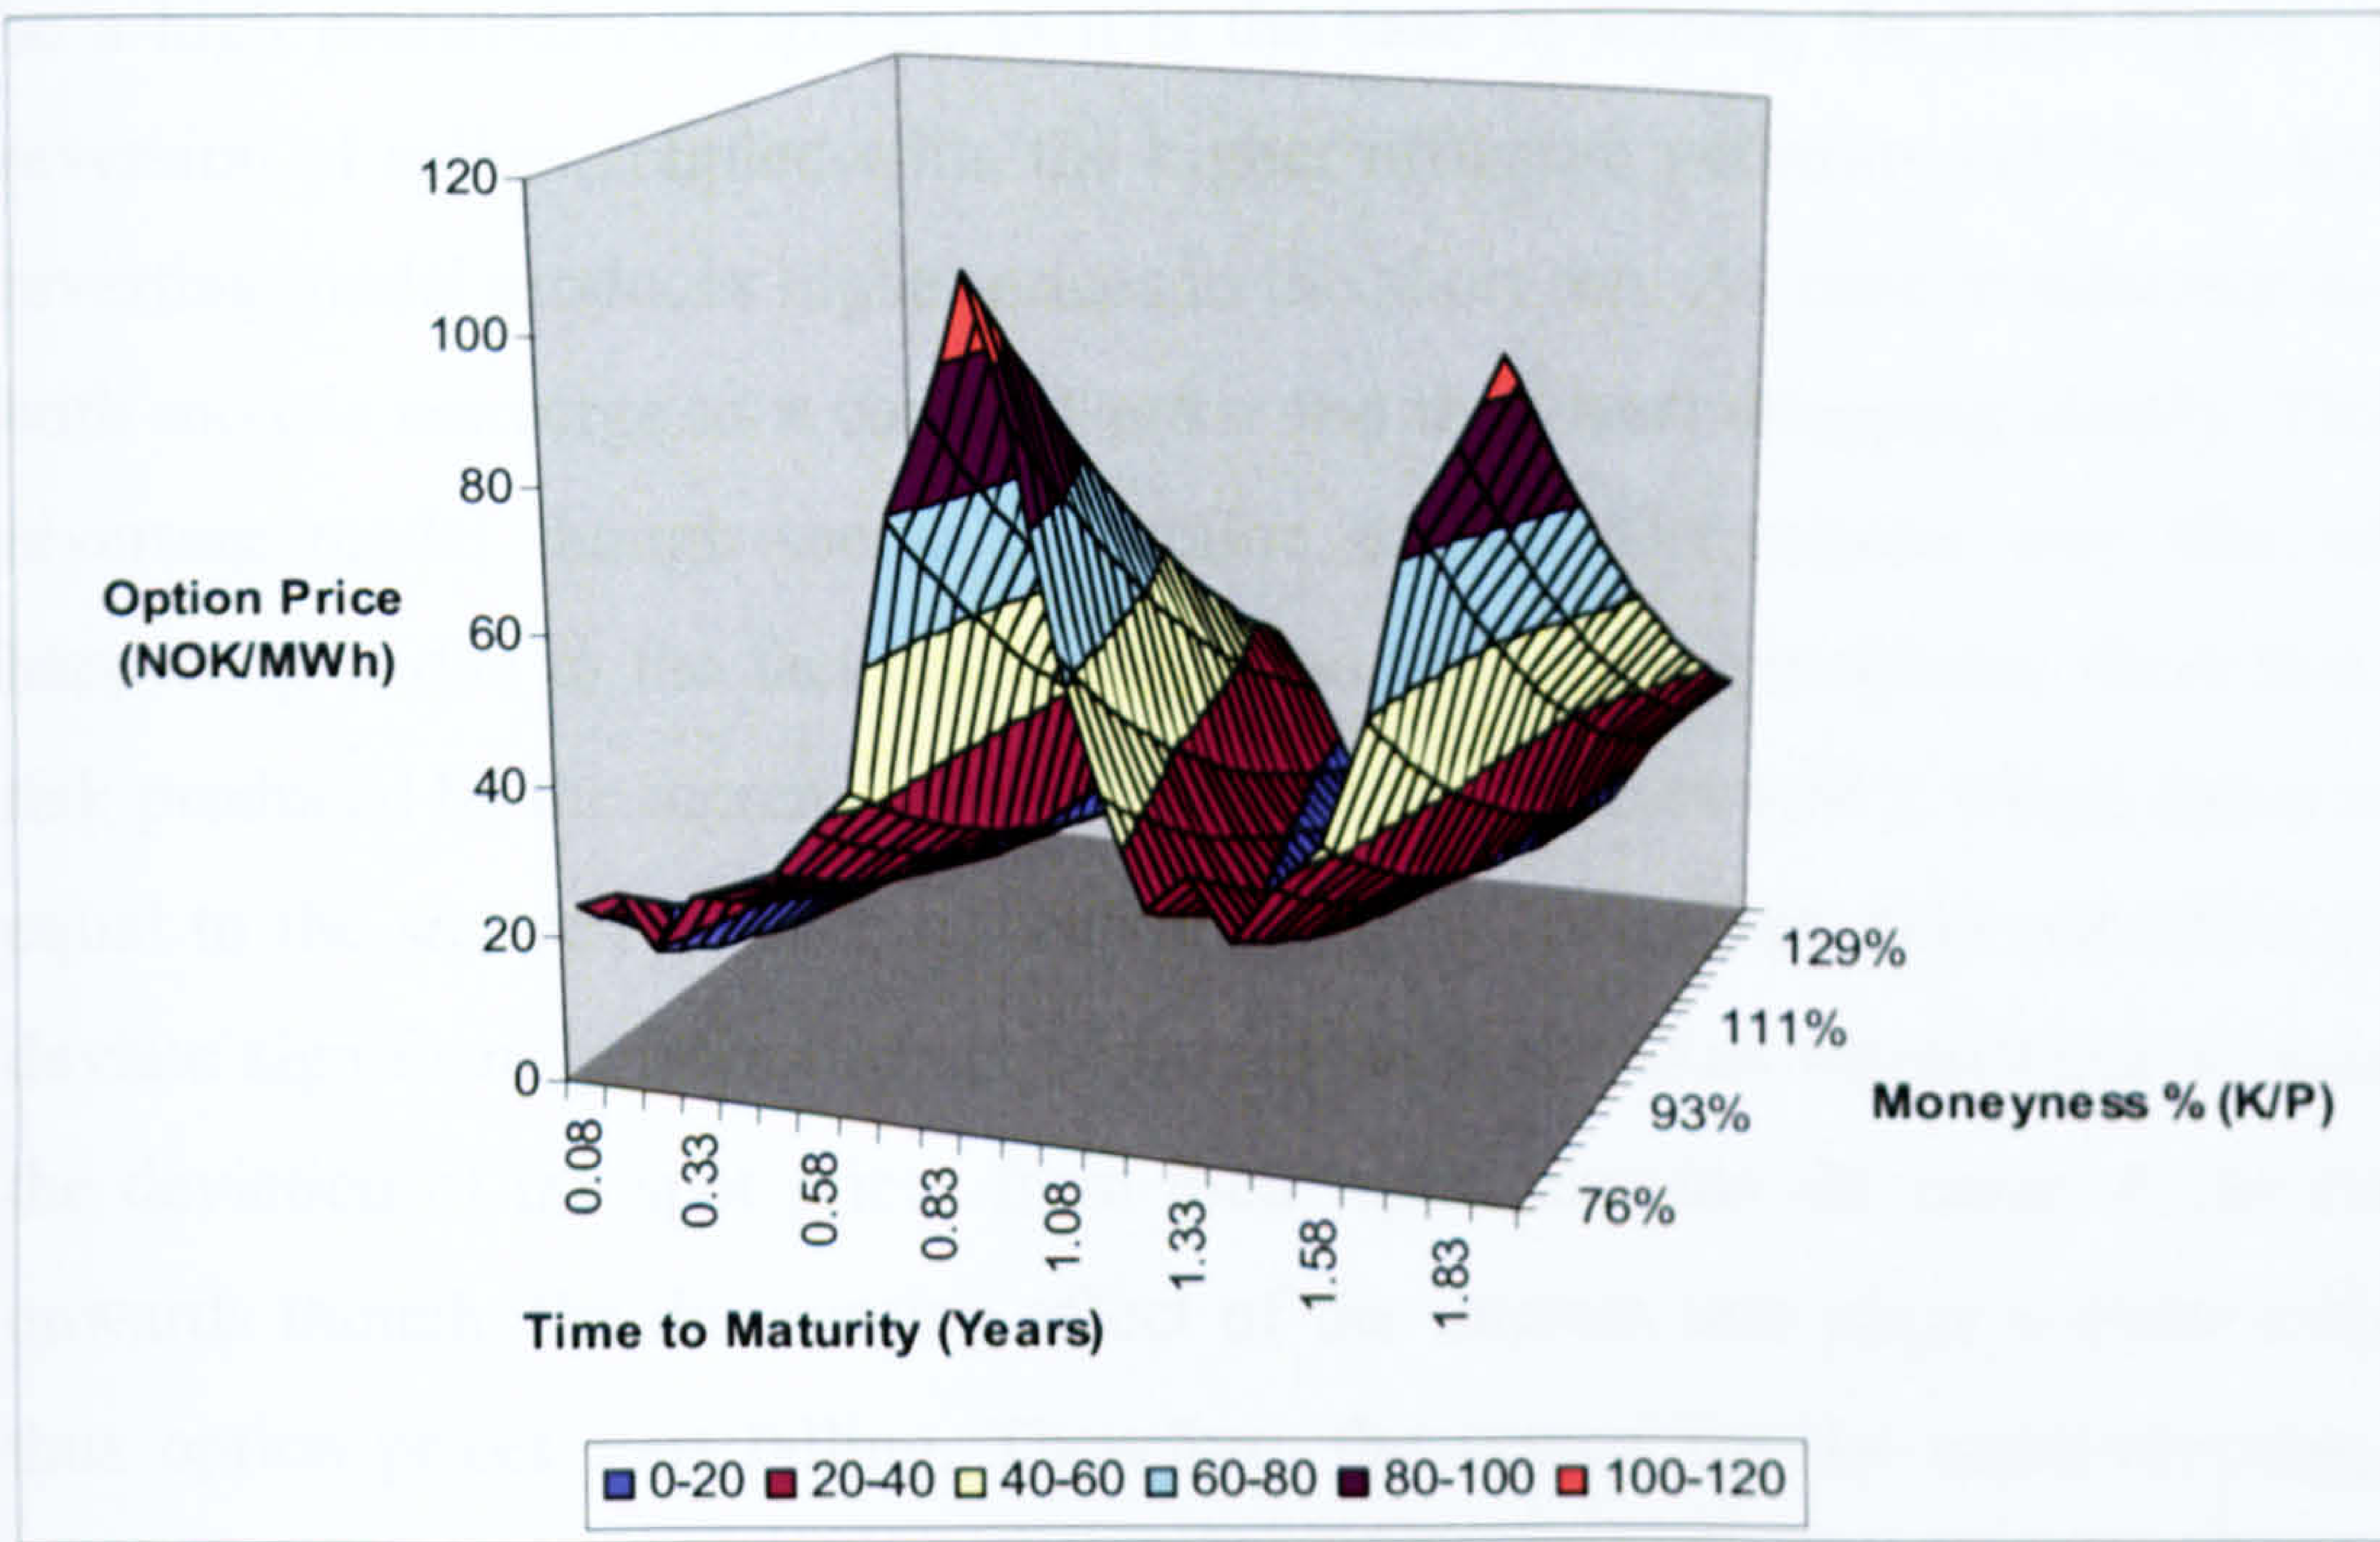

Figure 5.11 shows the call option prices across different times to maturity using the four different models, for both winter jump parameters (B) and jump parameters for the other seasons (A), respectively. In this analysis we do not consider the deterministic seasonal variable *f(T)*, in order to have a clearer picture of the effect of time to maturity on options prices, and the current spot price is assumed to be equal to the strike price and the risk neutral equilibrium level, i.e. the option is ATM. We can see that the Option prices between panels (A) and (B) are similar. This means that even though the jump parameters are higher during winter than the rest of the seasons, these low frequency larger price movements do not affect substantially the option prices. This is due to the fact that the probability of a jump occurring at a specific date is low, and thus in order for the jumps to play a significant role their frequencies have to be much higher, or their speed of mean reversion has to be lower, so that jumps can stay at high levels for longer periods. This is going to be analysed in greater depth at the next section where we discuss about Asian options, as well as the next chapter where we look at the extra swing premium.

Comparing the properties of the two short-run models, namely the spike and the mean-reverting model, the latter produces higher prices for short-term options (i.e. $T<0.25$ years) than the spike model in both cases. So, despite the fact that there may be a high probability of spikes, as it is the case in winter, the high degree of mean reversion of spikes coupled with, the higher diffusive volatility (167%) in the mean-reverting model produces higher prices in the short run. As time to maturity increases, both models converge to a constant price and then start dropping slowly. The mean-reverting model though seems to decline earlier. The reason why this might be happening is due to the fact that as the time to maturity increases, there is a greater risk produced by the increments of the Brownian motion (dW), which has a volatility equal to the square root of time. However, mean reversion does not allow prices to deviate significantly from their equilibrium level and thus we arrive at a stage where the deviation of the spot prices from their mean remains the same. From that point onwards though, the discounting effect of the interest rate plays a more critical role thus option prices start falling. Therefore, the reason for the mean-reverting model declining at a sooner time than the spike model is that the speed of mean reversion in the MR model (8.03) is much greater than that of the diffusive component in the spike model (2.98).

Figure 5.11: European Call Option Prices For Different models, across time to maturity with Spike parameters for Spring, Summer, Autumn (A) and Winter (B)

The figures show the option prices generated from the four different models across time to maturity using equations (5.6), (5.10) and (5.19), for spike parameters during Spring, Summer, Autumn (A) and Winter (B) respectively.

(A)

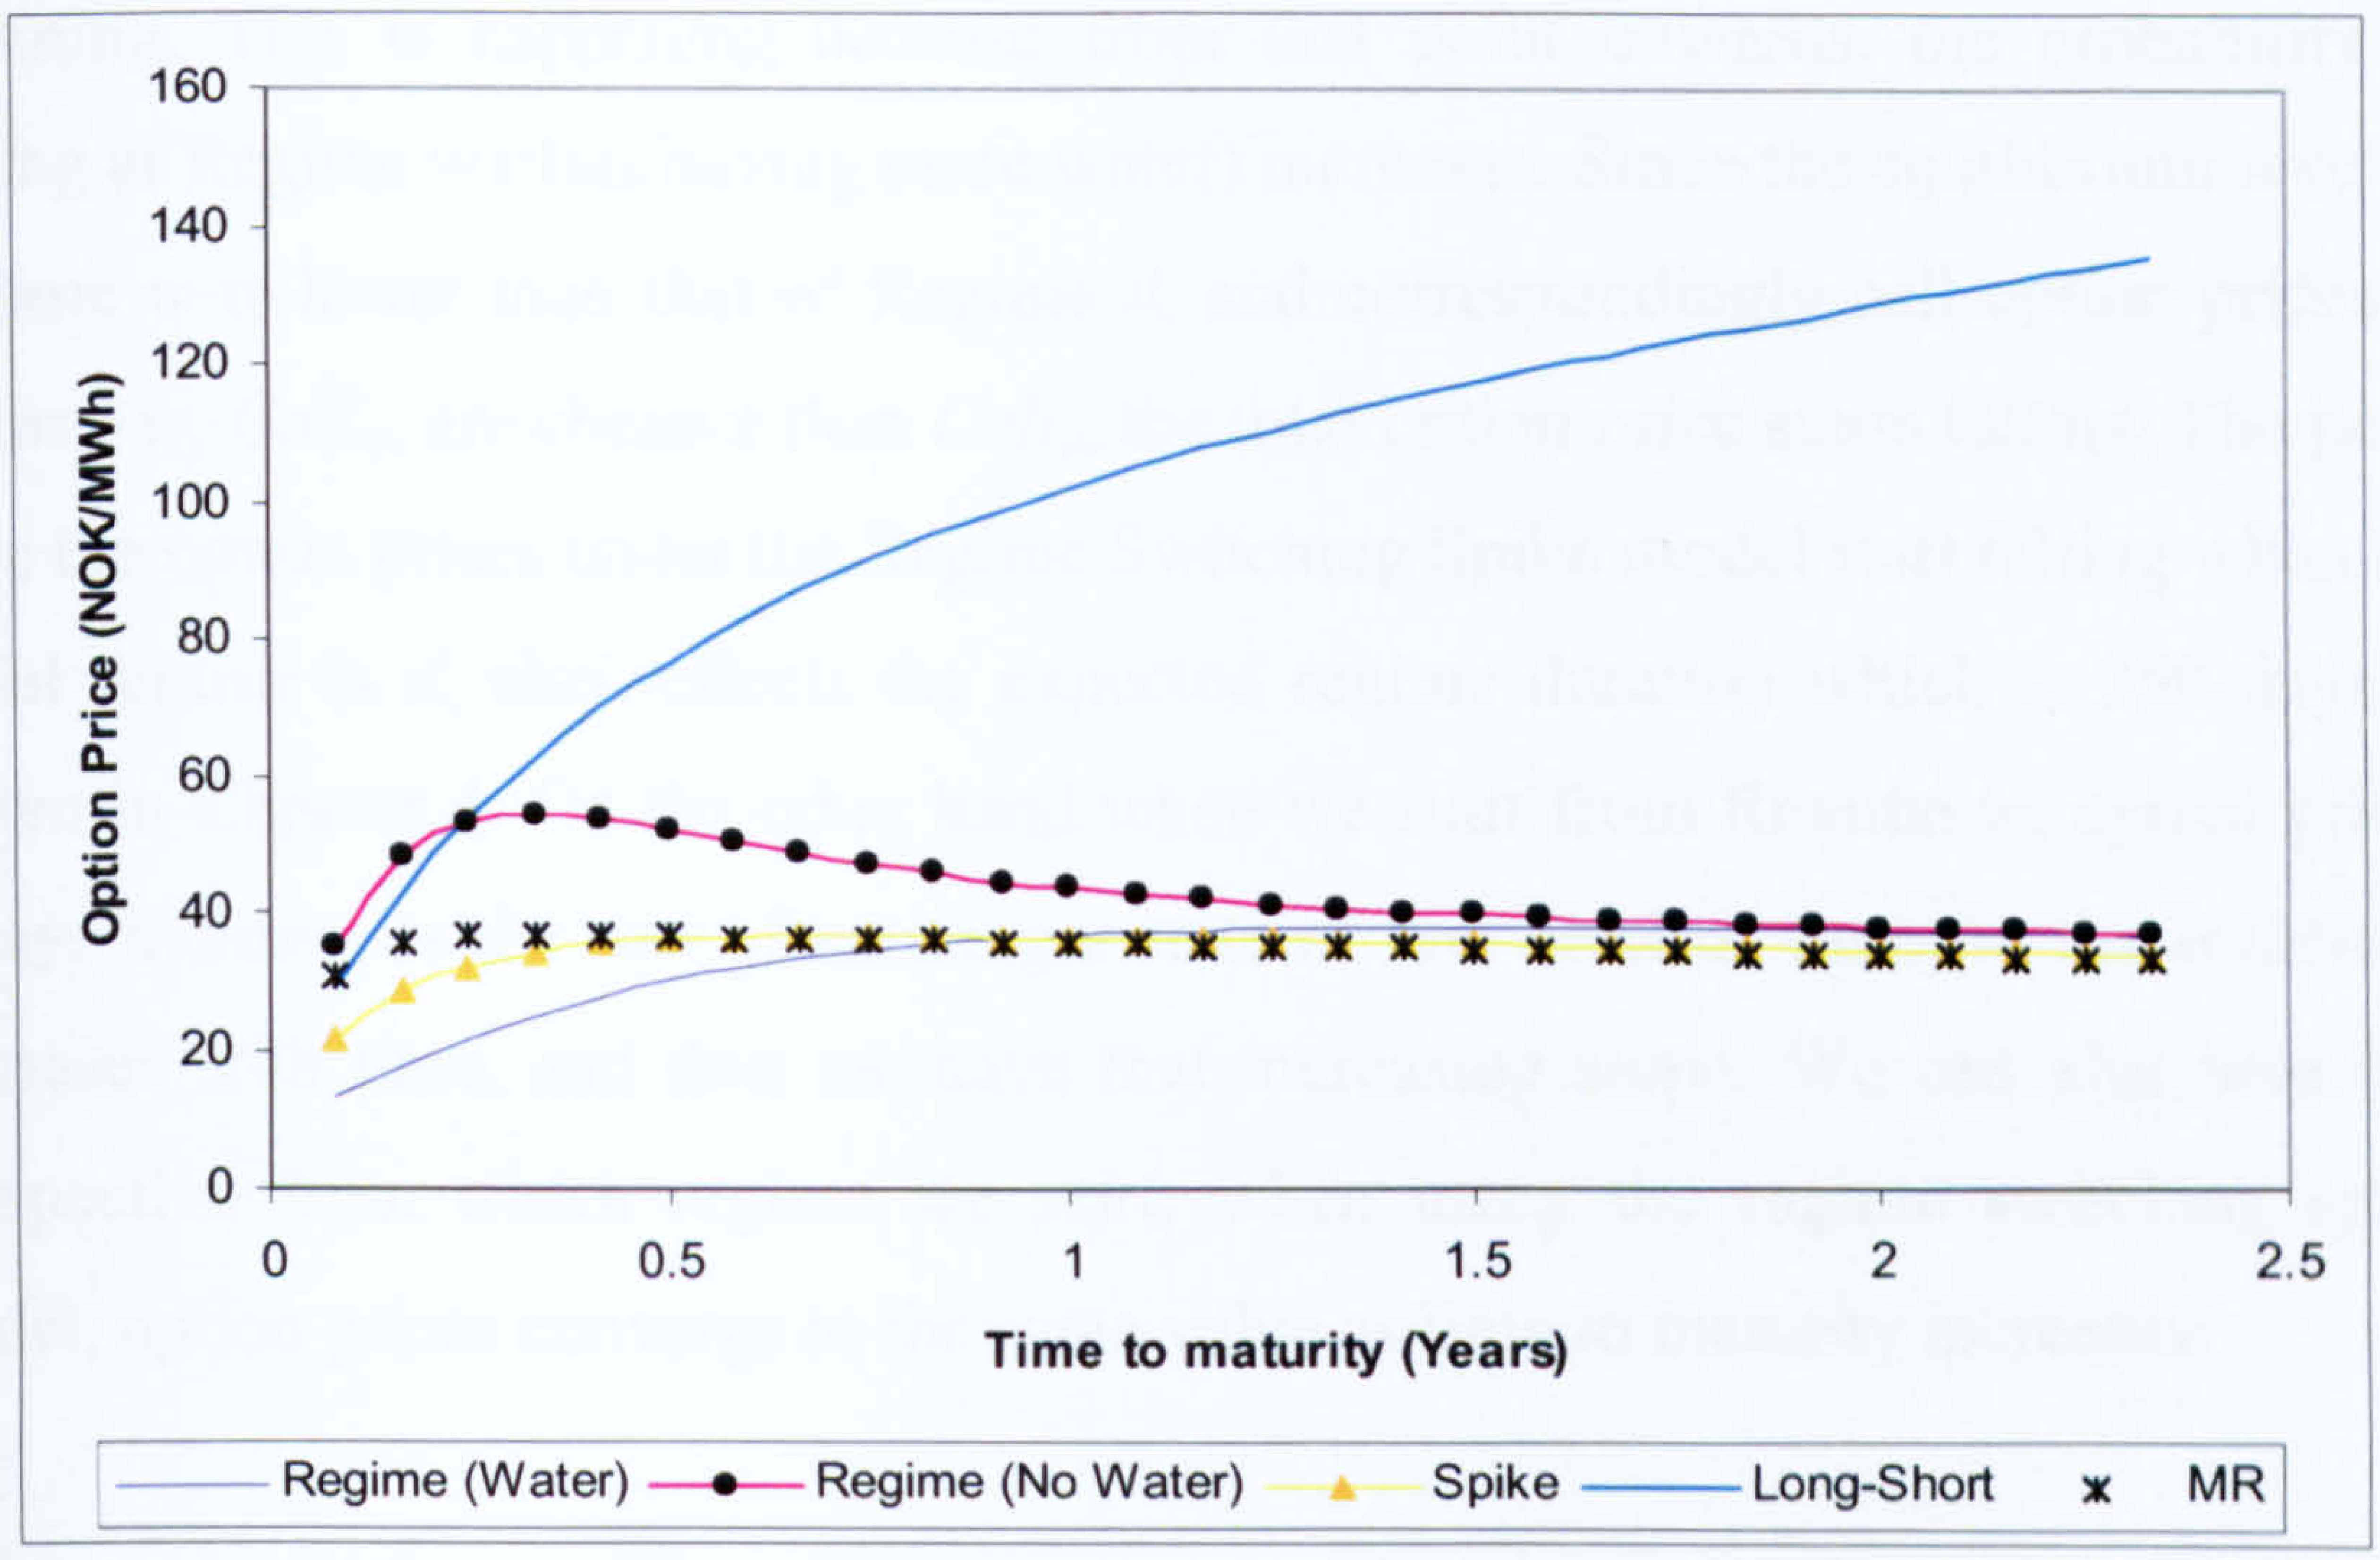

(B)

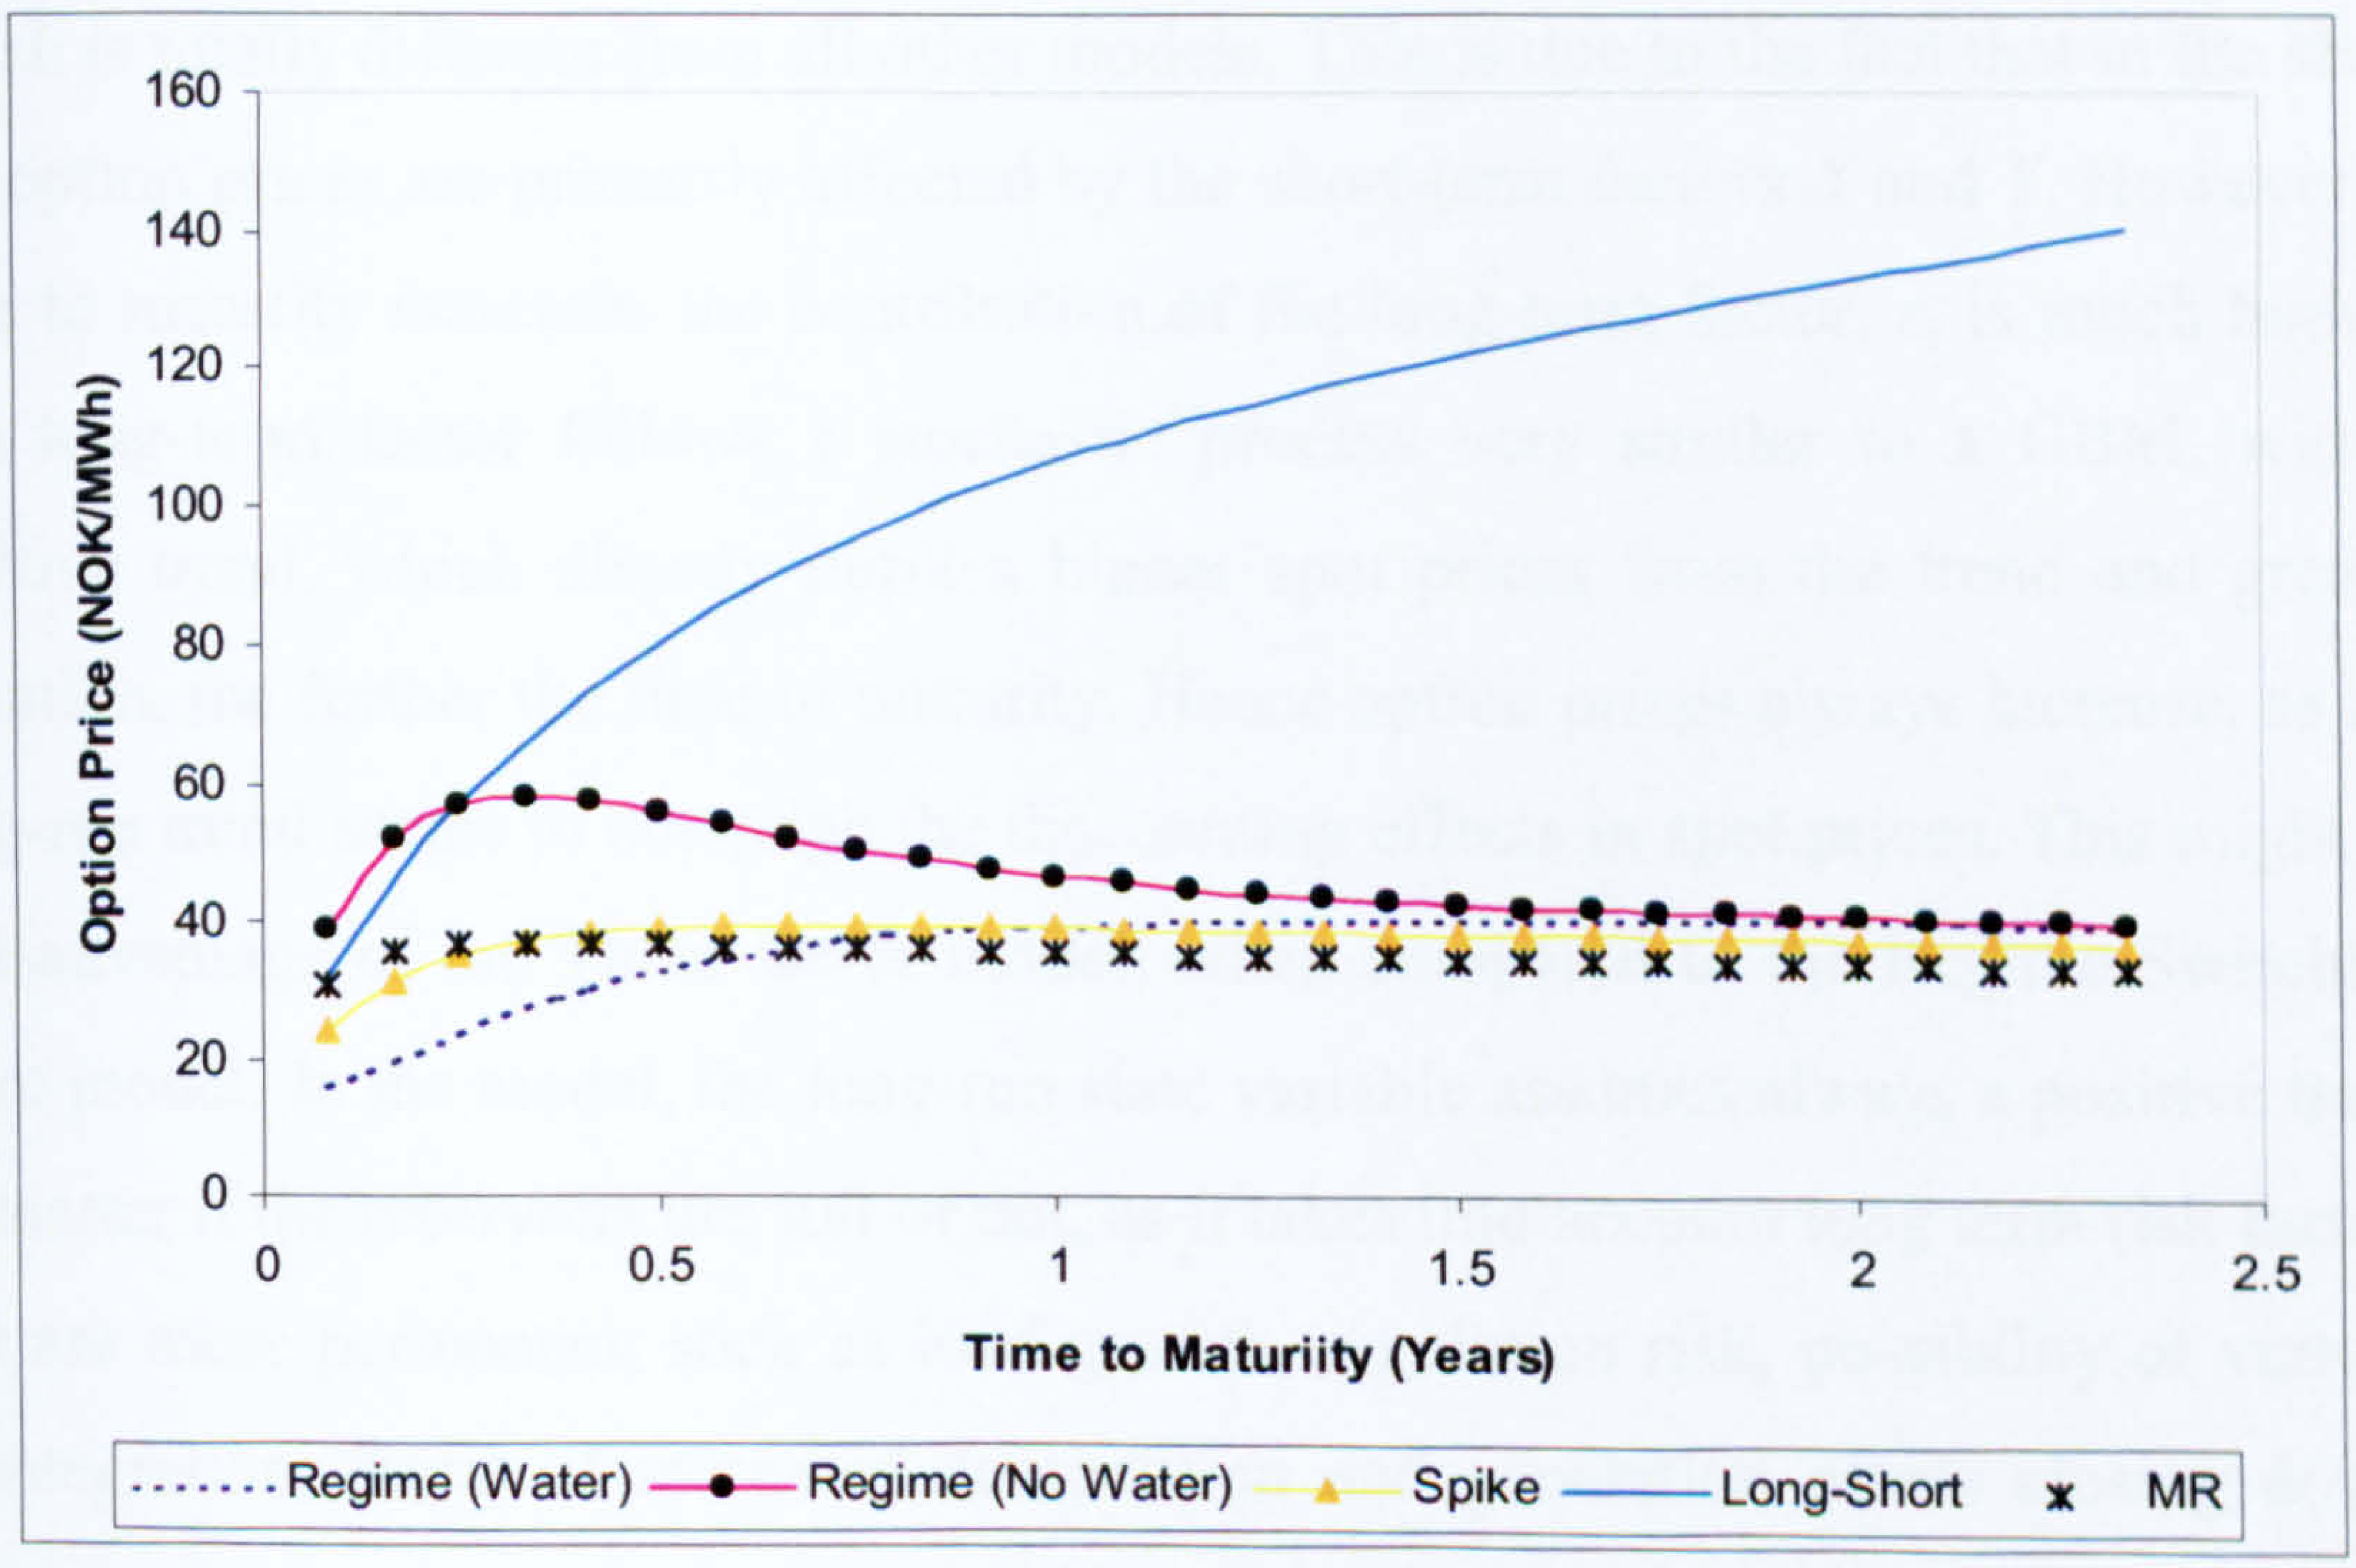

On the other hand looking at the regime-switching model, we see a very intuitive result. When we start from Regime *d* where the water reservoirs are below their seasonal level, as time to maturity increases option prices start increasing, up to maturities of half a year. From that point onwards, option prices clearly start declining. This is happening because from that point onwards, the probability of ending at Regime *w* (thus having more water) increases. Since the equilibrium level of Regime *w* is lower than that of Regime *d*, and correspondingly call option prices in Regime *w*, $Call_w$, are cheaper than $Call_d$, the total option price starts falling. The point were the option prices under the Regime Switching Spike model start falling when the initial regime is *d*, also reflects the expected regime duration which is 250 days as shown in Chpater 4. On the other hand when we start from Regime *w*, option prices always increase as the risk of entering a state of low level of water in the reservoirs increases with time, and thus we have that increasing slope. We can also note that irrespective from which regime we start, when using the regime switching spike model, option prices converge to the same value as time to maturity increases.

Finally, when using the Three-Factor spike model option prices increase with time to maturity, as the trend increases the further the maturity of the option, a characteristic which is totally different from all other models. This is due to the fact that in the short run option prices are primarily affected by the short-term factors X and Y. However as time to maturity increases the contribution of the long-term factor, ε, is much higher. The long-term factor follows a stochastic process very similar to a GBM, with a positive trend, which directly implies higher spot prices from the trend and greater variation, the further the time to maturity. Hence option prices always increase, as the long-run trend seems to outweigh the discounting effects in spot prices. This might be a disadvantage of the Three-factor model, when compared to the Regime Switching spike model. In the model, the long-run state variable assumes always a positive trend no matter if the reservoirs are full or not, as it takes into account long term risk factors that are more permanent, such as load growth, regulation risk, possibility of vertical disintegration, threat of increased competition and generation plants closing down. However this variable does not accurately reflect the risk of the level of water in the reservoirs, since water levels change from year to year depending on the rainfall, thus inducing some degree of mean reversion or regime switching from year to year, rather than a constant change. As discussed in the previous chapter, the reservoir risk in the

Three-Factor spike model seems to be captured by the short-term variable X, which is not very intuitive as reservoir risk is a type of risk that should affect mainly mid- to long- run behaviour.

Finally, Figure 5.12 shows the call option prices from the Three-Factor spike model, across different correlation parameters, $\rho_{X,\varepsilon}$ between X and ε, and different maturities. We can see that option prices increase as the correlation coefficient increases, reflecting the fact that the volatility of the sum of two random variables increases directly with correlation, as follows:

$$Var[X, \varepsilon]=Var[X]+Var[\varepsilon]+2\,\rho_{X,\varepsilon}\sqrt{Var[X]Var[\varepsilon]}$$

Figure 5.12: European Call Option Prices produced by the Three-Factor model, across different correlation coefficients and time to maturity.

The figure shows how European Call option prices are affected in the Tree-Factor spike model for different values of the correlation coefficient $\rho_{X,\varepsilon}$.

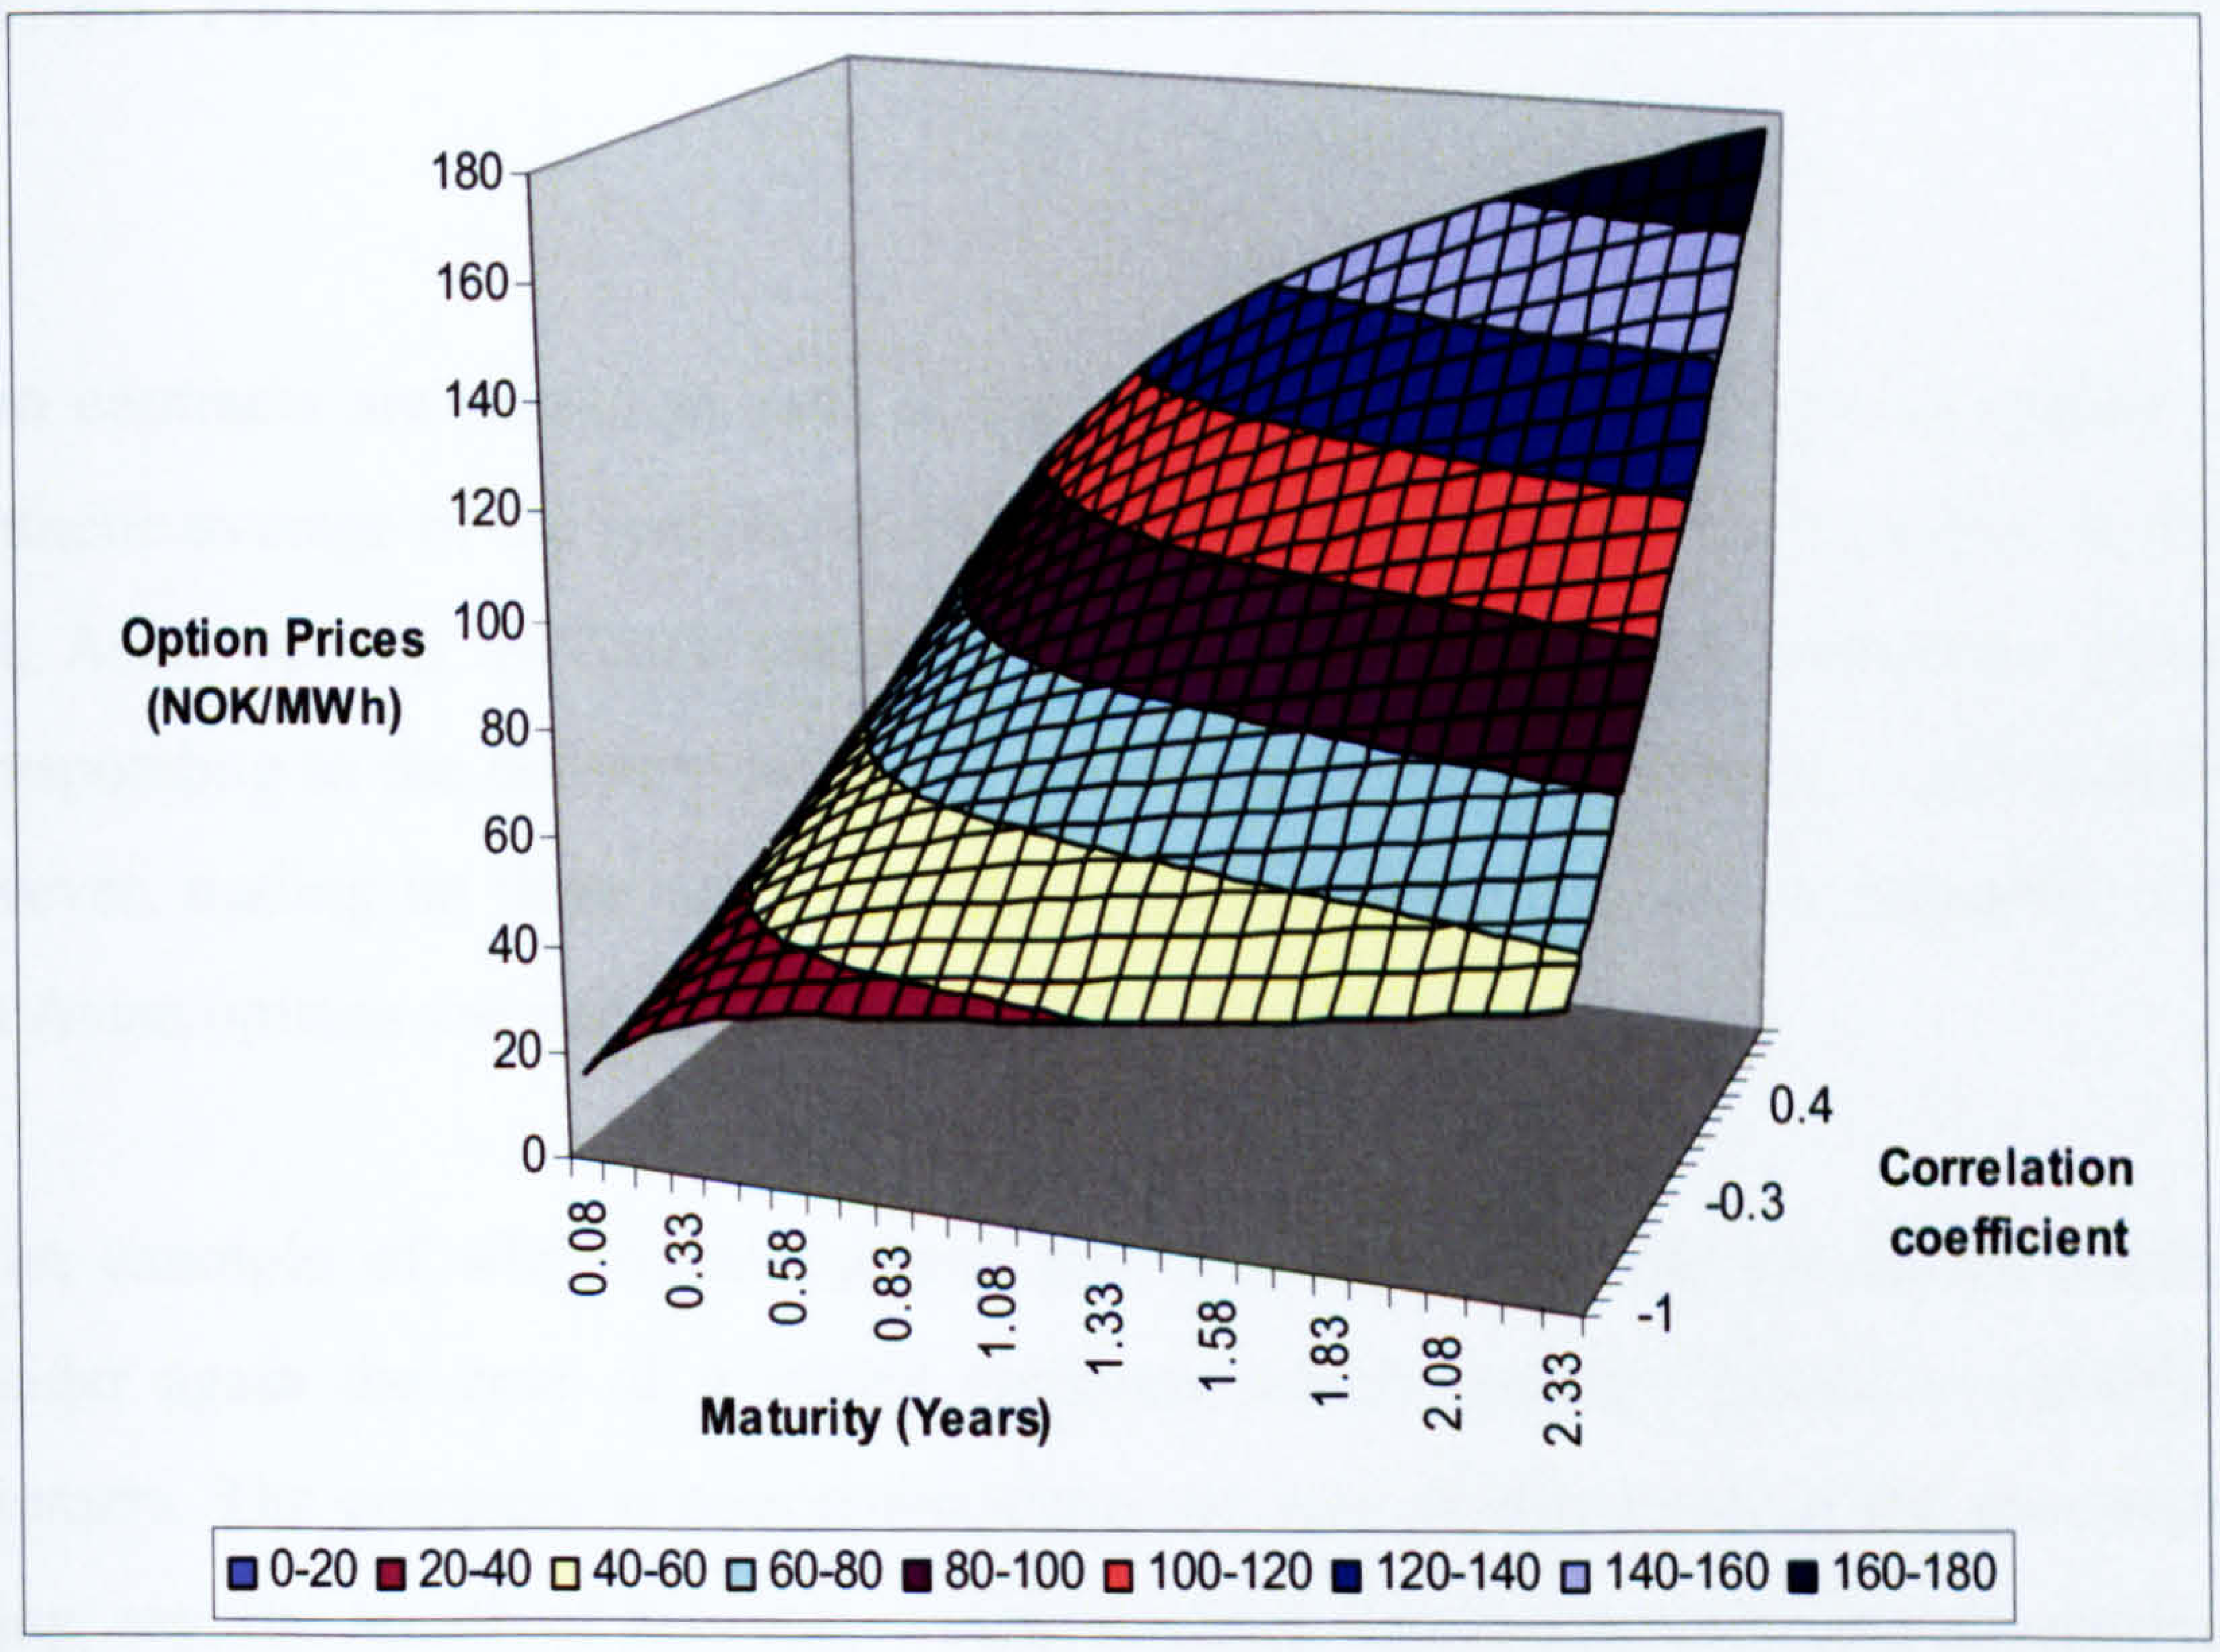

5.7 Asian Options

Asian options are options whose final payoff is based in some way on the average level of the electricity price during some or all of the life of the option. The payoff from an *average price call* is $max(P_{ave} - K, 0)$, where P_{ave} is the average value of the underlying asset calculated over a predetermined averaging period $\{T_1 = \tau_1, ..., \tau_n = T_2\}$[38]. In the same manner the payoff of an average price put Asian Option is $max(K - P_{ave}, 0)$. Thus the value of call and put average price Asian Options are respectively:

$$Asian\ Call = E^{*}\left[e^{-r(T_2-t)}\max\left(\frac{\sum_{i=1}^{n} P_{\tau_i}}{n} - K, 0\right)\right] \tag{5.20}$$

$$Asian\ Put = E^{*}\left[e^{-r(T_2-t)}\max\left(K - \frac{\sum_{i=1}^{n} P_{\tau_i}}{n}, 0\right)\right] \tag{5.21}$$

Asian contracts are settled in cash at the end of each settlement period against the arithmetic average of the system price in Elspot during the settlement period. In Nord Pool, Asian options on block (four weeks) futures contract with settlement periods corresponding to the delivery period of the reference futures contract were available. However, trading on these contracts was ceased in 2001 due to low liquidity. Since then Asian options are only traded in the OTC market.

As an example of why Asian options are very helpful in the electricity markets, consider again the case of a utility company which supplies power to residential customers. The company is concerned about the cost of electricity in the spot market during, say, the month of February which is a high demand period. One alternative is for the company to use a forward contract to fix the price of electricity. However, the utility believes that there is a good chance that, if there are no major supply disruptions in the system or some days the weather is not that cold as February is the last month of winter, the spot prices will be much lower than currently suggested by

[38] There are also average strike options, where the strike price is the average of the spot price during the delivery period, and at expiry it is settled against the spot price, P_T. However as these are not traded in Nord Pool, they are not considered in this thesis.

the forward market. Consequently, the company wants to benefit from the lower prices, but at the same time wants a measure of protection in case prices increase, contrary to expectations. One solution might be to buy Caps, however the cost involved is very high due to the number of call options involved, as the company needs to buy one caplet for each day of the delivery month. A more cost effective solution is to buy a call option on the average spot price in February, struck at the current level of the forward price. In this way, when prices fall as expected, the utility enjoys the benefits of cheap delivery. The main advantages of holding Asian call options are the following: First, if prices increase the utility will have to buy in the expensive market, but its higher cost will be offset by the payoff from the option it bought for protection. Second, it will provide a very good hedge against the overall risk in that certain month, for example in case the water reservoir levels for that month are extremely low and demand is high. Third, Asian options are also very helpful for the writers of the options, since market manipulation in the short-run or high volatility in any particular date will not affect the average significantly. Moreover, Asian Options are cheaper than regular European options, since the volatility of the average is lower than the volatility of the actual underlying. However, as it will be shown latter, Asian option do not provide perfect hedge against spikes.

5.7.1 Pricing Asian Options

The put-call parity gives the relationship between the prices of European put and call options. As shown by Vehviläinen (2001), by considering Asian options as European options on the average spot price, one can also derive a similar put-call parity relation for Asian Options. Assume that there is an electricity forward contract with current forward price F_{t,τ_1,τ_n} (such as a block contract), for which the delivery period is the same as the averaging period of the Asian option i.e. from τ_1 to τ_n. Next, consider two portfolios:

1. Portfolio A: a long position in an Asian call option with strike price K and a cash amount of $e^{-r(\tau_n-t)}\left(K-F_{t,\tau_1,\tau_n}\right)$.
2. Portfolio B: a long position in an Asian put option with strike price K and an electricity forward contract with forward price F_{t,τ_1,τ_n}.

At the end of the delivery period of the options, the call option in Portfolio A has a payoff of $max(P_{ave} - K, 0)$, where P_{ave} is the average spot price over the delivery period, and the cash amount is worth $(K - F_{t,\tau_1,\tau_n})$. In Portfolio B, the put option has a payoff of $max(K - P_{ave}, 0)$ and the electricity forward payoff since it is a block contract is $(P_{ave} - F_{t,\tau_1,\tau_n})$. Thus at the end of the delivery period, both portfolios have the same total payoff of $max(P_{ave} - F_{t,\tau_1,\tau_n}, K - F_{t,\tau_1,\tau_n})$. Thus, the no arbitrage assumption states that the present values of the portfolios must be equal, and since there is no cost in entering a forward contract:

$$e^{-r(\tau_n - t)}\left(K - F_{t,\tau_1,\tau_n}\right) = Asian\ Put - Asian\ Call$$

However it is not possible to create hedging portfolios for Asian options with electricity forwards, despite the connection of the payoffs. Trading with an electricity forward contract ceases before the averaging period of the corresponding Asian option but the payoff of the Asian option depends on the events of the delivery period. Thus after the delivery period has started, dynamic adjustment of the hedging portfolio is not possible.

As stated in the introduction of the chapter, the pricing of Asian options does not exist in closed-form, if a model follows a lognormal or any other log-distribution. This is because the distribution of the arithmetic average of a set of e.g. lognormal distributions does not have any analytically tractable properties. On the other hand the geometric average of a set of e.g. lognormal distributed variables is also lognormal. Thus closed-form solutions exist for the geometric average, e.g. Kemna and Vorst (1990), Turnbull and Wakeman (1991) and Curran (1992). However in the electricity markets Asian options depend on the arithmetic average. Thus we have to rely on numerical methods, such as Monte Carlo.

Monte Carlo methods are very popular in empirical finance, since they are in general easy to implement and allow the treatment of problems with high dimensionality. In particular, when there are multiple stochastic factors, numerical methods like finite–differences or binomial techniques become impractical while Monte Carlo is still appropriate. Moreover due to the structure of the payoff of the Asian options, which is path dependent, the Monte Carlo method becomes more practical as it describes the price path generated for every simulation.

Since the convergence of Monte Carlo is fairly slow, other methods may be preferred as long as the underlying stochastic process is simple and the number of risk factors is small. However, in the case of energy derivatives appropriate stochastic processes for the underlying, e. g. the electricity price, are in general considerably complicated since they have to account for seasonality, mean reversion, spikes etc. Therefore, calibration of a model for the energy prices to real market data usually requires at least two stochastic factors, frequently in conjunction with a jump diffusion process. Consequently Monte Carlo methods may be the best choice if the stochastic processes under consideration are expected to be close to reality, particularly when pricing path dependent derivatives such as Asian Options. The use of Monte Carlo to price Asian options is described next.

5.7.1.1 Pricing Options using Monte Carlo Simulation methods

Before discussing the discretization of the Stochastic processes for the electricity price recall our general Spike Model for the electricity spot price (P_t) in the risk neutral world (as in equation (5.9)):

$$\begin{aligned} P_t &= f(t) + \exp(X_t + Y_t) \\ dX_t &= \left(k_1\left(\varepsilon - X_t\right) - \lambda_X\right)dt + \sigma_X dW_X^* \\ dY_t &= -k_2 Y_t dt + J(\mu_{J_i}, \sigma_{J_i}^2) dq(l_t) \end{aligned} \tag{5.22}$$

f(t) is the deterministic seasonal component of the model, and as such it does not have an error from discretization or random sampling since it involves no stochastic terms. On the other hand X_t and Y_t, are the stochastic components of the model, which explain the behaviour of the short-term prices under the regular conditions and when spikes occur. Both stochastic factors follow a mean reversion to an equilibrium level, however X_t is affected by shocks generated by a Normal distribution, where as the spike factor Y_t, is affected by shocks occurring at discrete points in time with probability generated by a Poisson distribution, assumed to be independent from dW_X^*; in other words, the *spike* process is independent of the Normal short term process X_t, and whenever spikes occur, their amplitude is generated from a Normal distribution, with mean and volatility which vary depending on the season of the year.

In order to demonstrate the discretisation process, we consider first the case where the Spot price depends only on X_t, that is to say that $P_t = e^{X_t}$. First of all the stochastic differential equation for X_t is *explicitly solvable* and has the following solution in terms of stochastic integral (Itô's integral):

$$X_t = e^{-k_1\Delta t} X_{t-1} + \left(\varepsilon - \frac{\lambda_X}{k_1}\right)(1 - e^{-k_1\Delta t}) + \sigma_X e^{-k_1\Delta t} \int_0^{\Delta t} e^{k_1 s} dW_s^* \quad (5.23)$$

Where Δt is the disctetisation step used; for example, if the Average needs to be taken over the days of a certain month, Δt can be 1 day. Using the above equation it can be shown that X_t, follows a Normal Distribution with Mean and Variance, as shown in the derivation of the moments in Chapter 4:

$$E[X_t] = e^{-k_1\Delta t} X_{t-1} + \left(\varepsilon - \frac{\lambda_X}{k_1}\right)(1 - e^{-k_1\Delta t}) \quad (5.24)$$

$$Var[X_t] = \sigma_X^2 \frac{(1 - e^{-2k_1\Delta t})}{2k_1} \quad (5.25)$$

In this equation again we can note that the expected value of X_t is a weighted average between the previous period's value of X_t and the risk neural equilibrium level $\left(\varepsilon - \frac{\lambda_X}{k_1}\right)$, weighted according to $e^{-k_1\Delta t}$, which also acts as a discounting factor of previous period's value to the current. In addition, as the discretisation step approaches infinity ($\Delta t \rightarrow \infty$), the mean converges to the risk neutral equilibrium level, $\left(\varepsilon - \frac{\lambda_X}{k_1}\right)$ and the variance converges to $\frac{\sigma_X^2}{2k_1}$.

In order to perform the simulation it is necessary to get the discrete-time equation for the process of X_t. Dixit and Pindyck (1994) show that the correct discrete-time format for the continuous-time process of mean reversion is the stationary first-order autoregressive process, AR(1). So the sample path ***simulation equation*** for X_t is performed by using the *exact*[39] (i.e. valid for large Δt) discrete-time expression:

$$X_t = e^{-k_1 \Delta t} X_{t-1} + \left(\varepsilon - \frac{\lambda_X}{k_1} \right)(1 - e^{-k_1 \Delta t}) + \sigma_X \sqrt{\frac{\left(1 - e^{-2k_1 \Delta t}\right)}{2k_1}} N(0,1) \quad (5.26)$$

where N(0,1) is a standard normally distributed variable. Note that the equation above can be viewed as the sum of the expected value equation with a random term generated from a standard Normal distribution times the standard deviation scaling the Brownian Motion. Equation (5.26) gives the discrete-time equation for an Arithmetic Ornstein-Uhlenbeck process and permits an exact discretisation in the sense that its accuracy doesn't decrease at all if a larger time-step (Δt) is used. Because the error from the discretisation is eliminated using the above exact discretisation equation, the only error that remains in the simulation is the error from the random numbers. On the other hand, alternative discretisation procedures such as the first-order Euler's approximation introduce discretisation errors into the simulation and have higher computational cost because they require small Δt.

In the same manner lets now focus on the spike part of the model and in this section lets assume that the spot price is a function of Y_t only, i.e. $P_t = e^{Y_t}$. Jumps are represented by the term $J(\mu_{J_i}, \sigma^2_{J_i})dq(l_i)$ in equation (5.22), which most of the time is zero and when they occur their arrival rate is l_i, and size J which is Normally distributed with mean μ_{J_i} and variance $\sigma^2_{J_i}$. Therefore:

$Jdq(l_i) = 0$ with probability $\boldsymbol{1 - l_i\, dt}$

$Jdq(l_i) = J$ with probability $\boldsymbol{l_i\, dt}$

[39] Where by "exact", we refer to the solution of the Stochastic Differential Equation of X_t, in order to give us accurate simulations no matter how big the time-step used.

The introduction of the *jump* component in the model adds a significant new feature. The *jump* term consists of two random processes. First of all the jump size can be easily simulated using a Normal distribution as follows:

$$\mu_{J_i} + \sigma_{J_i} N(0,1) \qquad (5.27)$$

On the other hand the timing of a jump is distributed as a Poison process and as such $dq(l_i) = 0$ most of the time and takes the value of 1 when a randomly timed jump occurs. Following Clewlow and Strickland (2000), we can simulate the timing of the jump by a random sample u generated from a uniform (0,1) and then checking the condition that $(u < l\,\Delta t)$; if the condition is satisfied then it takes the value of 1 and, thus a jump occurs; otherwise it is zero. This procedure on average generates jumps randomly at the correct average frequency. The spot process thus now, assuming still that it depends solely on Y_t, becomes:

$$P_t = e^{Y_t} = \exp\left(e^{-k_2 \Delta t} Y_{t-1} + \left(\mu_J + \sigma_J N(0,1)\right) 1_{u < l\Delta t}\right) \qquad (5.28)^{40}$$

However the discretisation error in this case is not the same as for a normally distributed random variable, since jumps do not occur continuously in time. Although the discretisation for the mean-reverting part of the equation of X_t is exact (valid for very large Δt), the presence of jumps together with their mean reversion is problematic when using large values of Δt, thus there is no exact way of discretising the equation for simulation purposes. Intuitively if we have on average 1 jump per day, and we take a time step of 7 days, the simulation would fail the criterion of the "exact discretisation", as it accounts for 1 jump for every 7 days. In particular, if the time interval Δt is *sufficiently small*, the probability of two jumps occurring at the same time step, $(l\Delta t)^2$, is much lower than $(l\Delta t)$. In addition, the speed of mean reversion for the spike factor is less than a day, hence a finer discretisation step may be required to capture the fast mean reversion of spikes.

To illustrate this, we performed various simulations for pricing European Options using $\Delta t = 1day$, and we found out that when the spot price was dependent only on X_t, using the exact discretisation method without jumps, the Monte Carlo price was the

[40] Note that $e^{-k_2\Delta t} Y_{t-1}$ is the exact solution of Y_t if we have no random terms such as the compound Poisson process, as it was the case for X_t.

same as the one produced in the closed-form solution. However, when the spike factor was included, convergence with the semi-closed form solution of equation (5.10), occurred when the time steps used were smaller than the speed of mean reversion. As a result the discretisation process requires finer time steps than one day, and as such for the simulation of Y we used hourly price time-steps (i.e. $\Delta t=1/(365*24)$). For European Options, the results were consistent with the semi-closed form solution results with an average error of 0.4%. The prices generated from the Monte Carlo algorithm, thus have the following form:

$$
\begin{aligned}
X_t &= e^{-k_1\Delta t}X_{t-1} + \left(\varepsilon - \frac{\lambda_X}{k_1}\right)(1 - e^{-k_1\Delta t}) + \sigma_X\sqrt{\frac{\left(1 - e^{-2k_1\Delta t}\right)}{2k_1}}N(0,1) \\
Y_t &= e^{-k_1\Delta t}Y_{t-1} + \left(\mu_J + \sigma_J N(0,1)\right)\left(u < l\Delta t\right) \\
P_t &= e^{X_t + Y_t + f(t)}
\end{aligned}
\tag{5.29}
$$

where for X_t, $\Delta t=1$day (i.e. $1/365$) and for Y_t, $\Delta t=1$ hour (i.e. $1/(365 * 24)$)

In order to price an Asian Option with Monte Carlo, for each sample Path i, we first take the average of the spot prices generated during that specific sample path for the averaging period specified in the contract, AV. Then at maturity T, we take the pay-off of the derivative for each specific sample path, and thus find their average:

$$
\begin{aligned}
AV_i &= \frac{1}{n}\sum_{i=1}^{n} P_{\tau_{j,i}} \\
Asian\ Put &= \frac{1}{m}\sum_{i=1}^{m} e^{-r(T_n - t)}\max\left(AV_i - K, 0\right) \\
Asian\ Call &= \frac{1}{m}\sum_{i=1}^{m} e^{-r(T_n - t)}\max\left(K - AV_i, 0\right)
\end{aligned}
\tag{5.30}
$$

where $P_{\tau_{j,i}}$ is the spot price generated form the Monte Carlo algorithm in equation (5.29), at time τ_j for sample path i, n is the number of days in the delivery period, and m is the number of sample paths generated by the Monte Carlo algorithm.

5.7.2 Asian Options Sensitivity Analysis Results

In the current setting we are looking at a sensitivity analysis of an Asian call option using our spike model and varying its parameters. The parameter values for the base case scenario can be found in Table 5.3, where the jump parameters used here are not seasonal, but the total average from the data. The reason why we do not consider the seasonal parameters and seasonality *f(t)* is that we do not want our results to be specific to any particular season, but rather we want to see how the different parameters from the stochastic factors affect the prices of Asian Options and in particular the impact of volatility factors during the averaging period. The Asian Option contract priced here is similar to that used to be traded in Nord Pool, i.e. the delivery date is 4 weeks. At the current setting we are pricing the Asian option one day before the averaging period starts.

Table 5.3: Parameter Values used in the spike model for the analysis on Asian call options

The table shows the value of the parameters of the spike model in equation (5.9), used for the Monte Carlo estimation of Asian call prices in the base case scenario.

Model Parameters	Parameters values
$X_0 = \left(\varepsilon - \frac{\lambda_X}{k_1}\right) = K$	5.137
σ_x	1.03
k_1	2.98
Y_0	0
l	5.57
μ_J	0.26
σ_J	0.35
k_2	287.61
r	5%
Time to maturity = Averaging period	4 weeks

We implement the described method of Monte Carlo Simulation by using the "antithetic" variables, technique. The antithetic variate approach uses pairs of negatively correlated random numbers that in return tend to produce pairs of negatively correlated simulation results. For example, if we have in one draw from a normal distribution a number *z*, then for another trial we use *–z*. Similarly for a uniform distribution, if in one draw the number is *u*, then for another trial we use *1-u*. This technique has been shown to produce results which are significantly more precise than the ones produced by the ordinary random sampling, since the results from each pair are averaged therefore making them less variable (since the correlation is perfectly negative) and thus reducing significantly the standard error of the results (see Boyle, Broadie and Glasserman 1997). Figure 5.13 presents the Asian option values for different numbers of trials, using the base case parameters. From the graph, it can be seen that convergence in prices occurs at approximately after 50,000 simulations. From that point onwards, the marginal change is decreasing significantly. In our analysis we used 200,000 simulation paths in order to be precise even though we had convergence for lower number of simulations most of the time. The reason is that after some tests we realized that changing some parameters, especially the jump parameters and volatility, the rate of convergence was lower. For reasons of consistency, in each sensitivity analysis we used the same random numbers generated from the normal distributions for *X* and the jump size and the uniform distribution for the jump arrivals.

Figure 5.13: Monte Carlo Convergence for Asian Options

The figure shows how Asian Call Options prices converge, as we increase the number of Monte Carlo simulation paths.

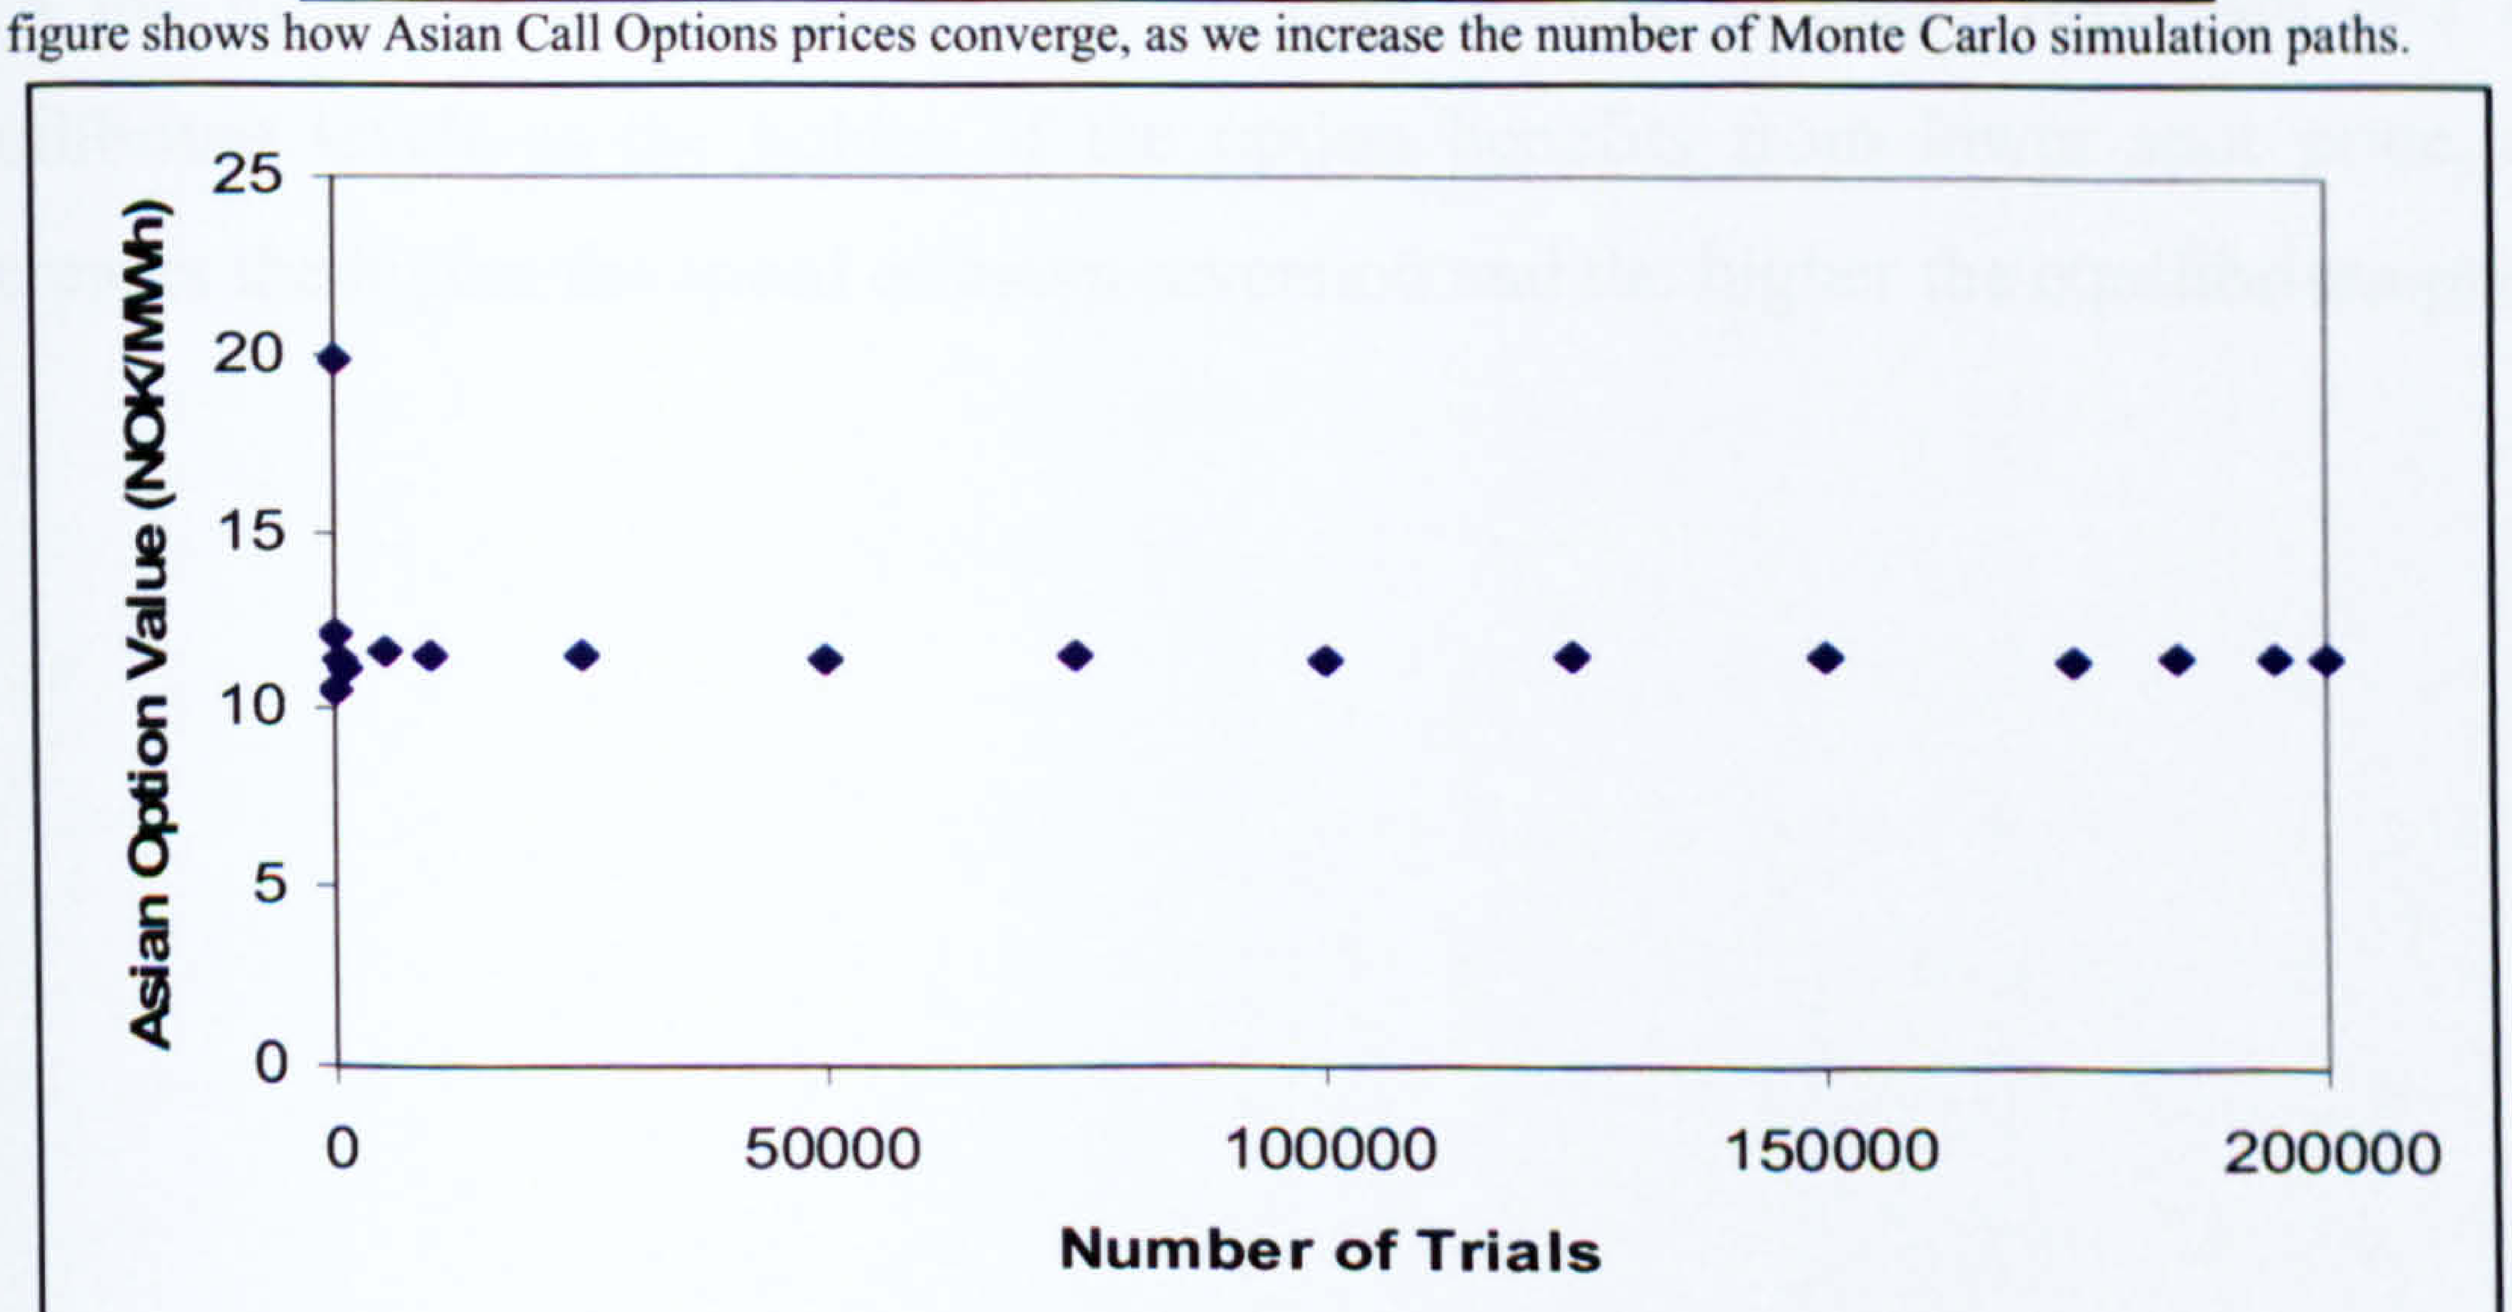

Thus, next we conduct a sensitivity analysis on Asian option prices, with respect to the parameters in the spike model, i.e k_1, k_2, $\tilde{\varepsilon}$ (risk-neutral equilibrium level), σ_X, σ_J, μ_J and l. In every case, we change two of the parameters and keep the remaining constant as in Table 5.3. In all cases we consider an ATM Asian option, meaning that the strike price $K=exp(X_0+Y_0)$=170.20 NOK/MWh.

Figure 5.14 shows the sensitivity analysis for different values of k_1 and $\left(\varepsilon-\frac{\lambda_X}{k_1}\right)$; that is we keep all other parameters constant, and we start varying the parameter of the speed of mean reversion of X and the risk-neutral equilibrium price. For the equilibrium level, it has to be noted that option prices have been estimated for equilibrium levels that are lower, equal and higher than the current spot price. The results are very intuitive indeed. What we actually see is that at low risk-neutral equilibrium levels, option prices decrease the higher the speed of mean reversion. This is because as the speed of mean reversion increases, shocks in the spot price die out much faster, and thus the prices mean revert to the lower equilibrium level, thus decreasing the average level of the spot even more and making the option less probable to end ITM. On the other hand, when the equilibrium level is higher, the value of the Asian option increases with the speed of mean reversion, since the prices mean revert to the high equilibrium level faster, thus the average increases and the option is more probable of ending ITM. Now in the case of Put Asian options we should expect exactly the opposite due to the structure of the pay-off. Hence the value of a put option is always increasing for high speeds of mean reversion and low equilibrium levels as the holder of the option benefits from lower spot price, and decreases the higher the speed of mean reversion and the higher the equilibrium price.

Figure 5.14: Asian Call Option Prices with respect to the speed of mean reversion k_1 and the risk-neutral equilibrium level of X, $\left(\varepsilon - \frac{\lambda_X}{k_1}\right)$

The graph displays a sensitivity analysis, where the prices of Asian Call Options are estimated using the Monte Carlo Simulation explained in the chapter, by changing the parameters of the speed of Mean Reversion, k_1, and the Risk Neutral Equilibrium Level to which the Diffusive component reverts to.

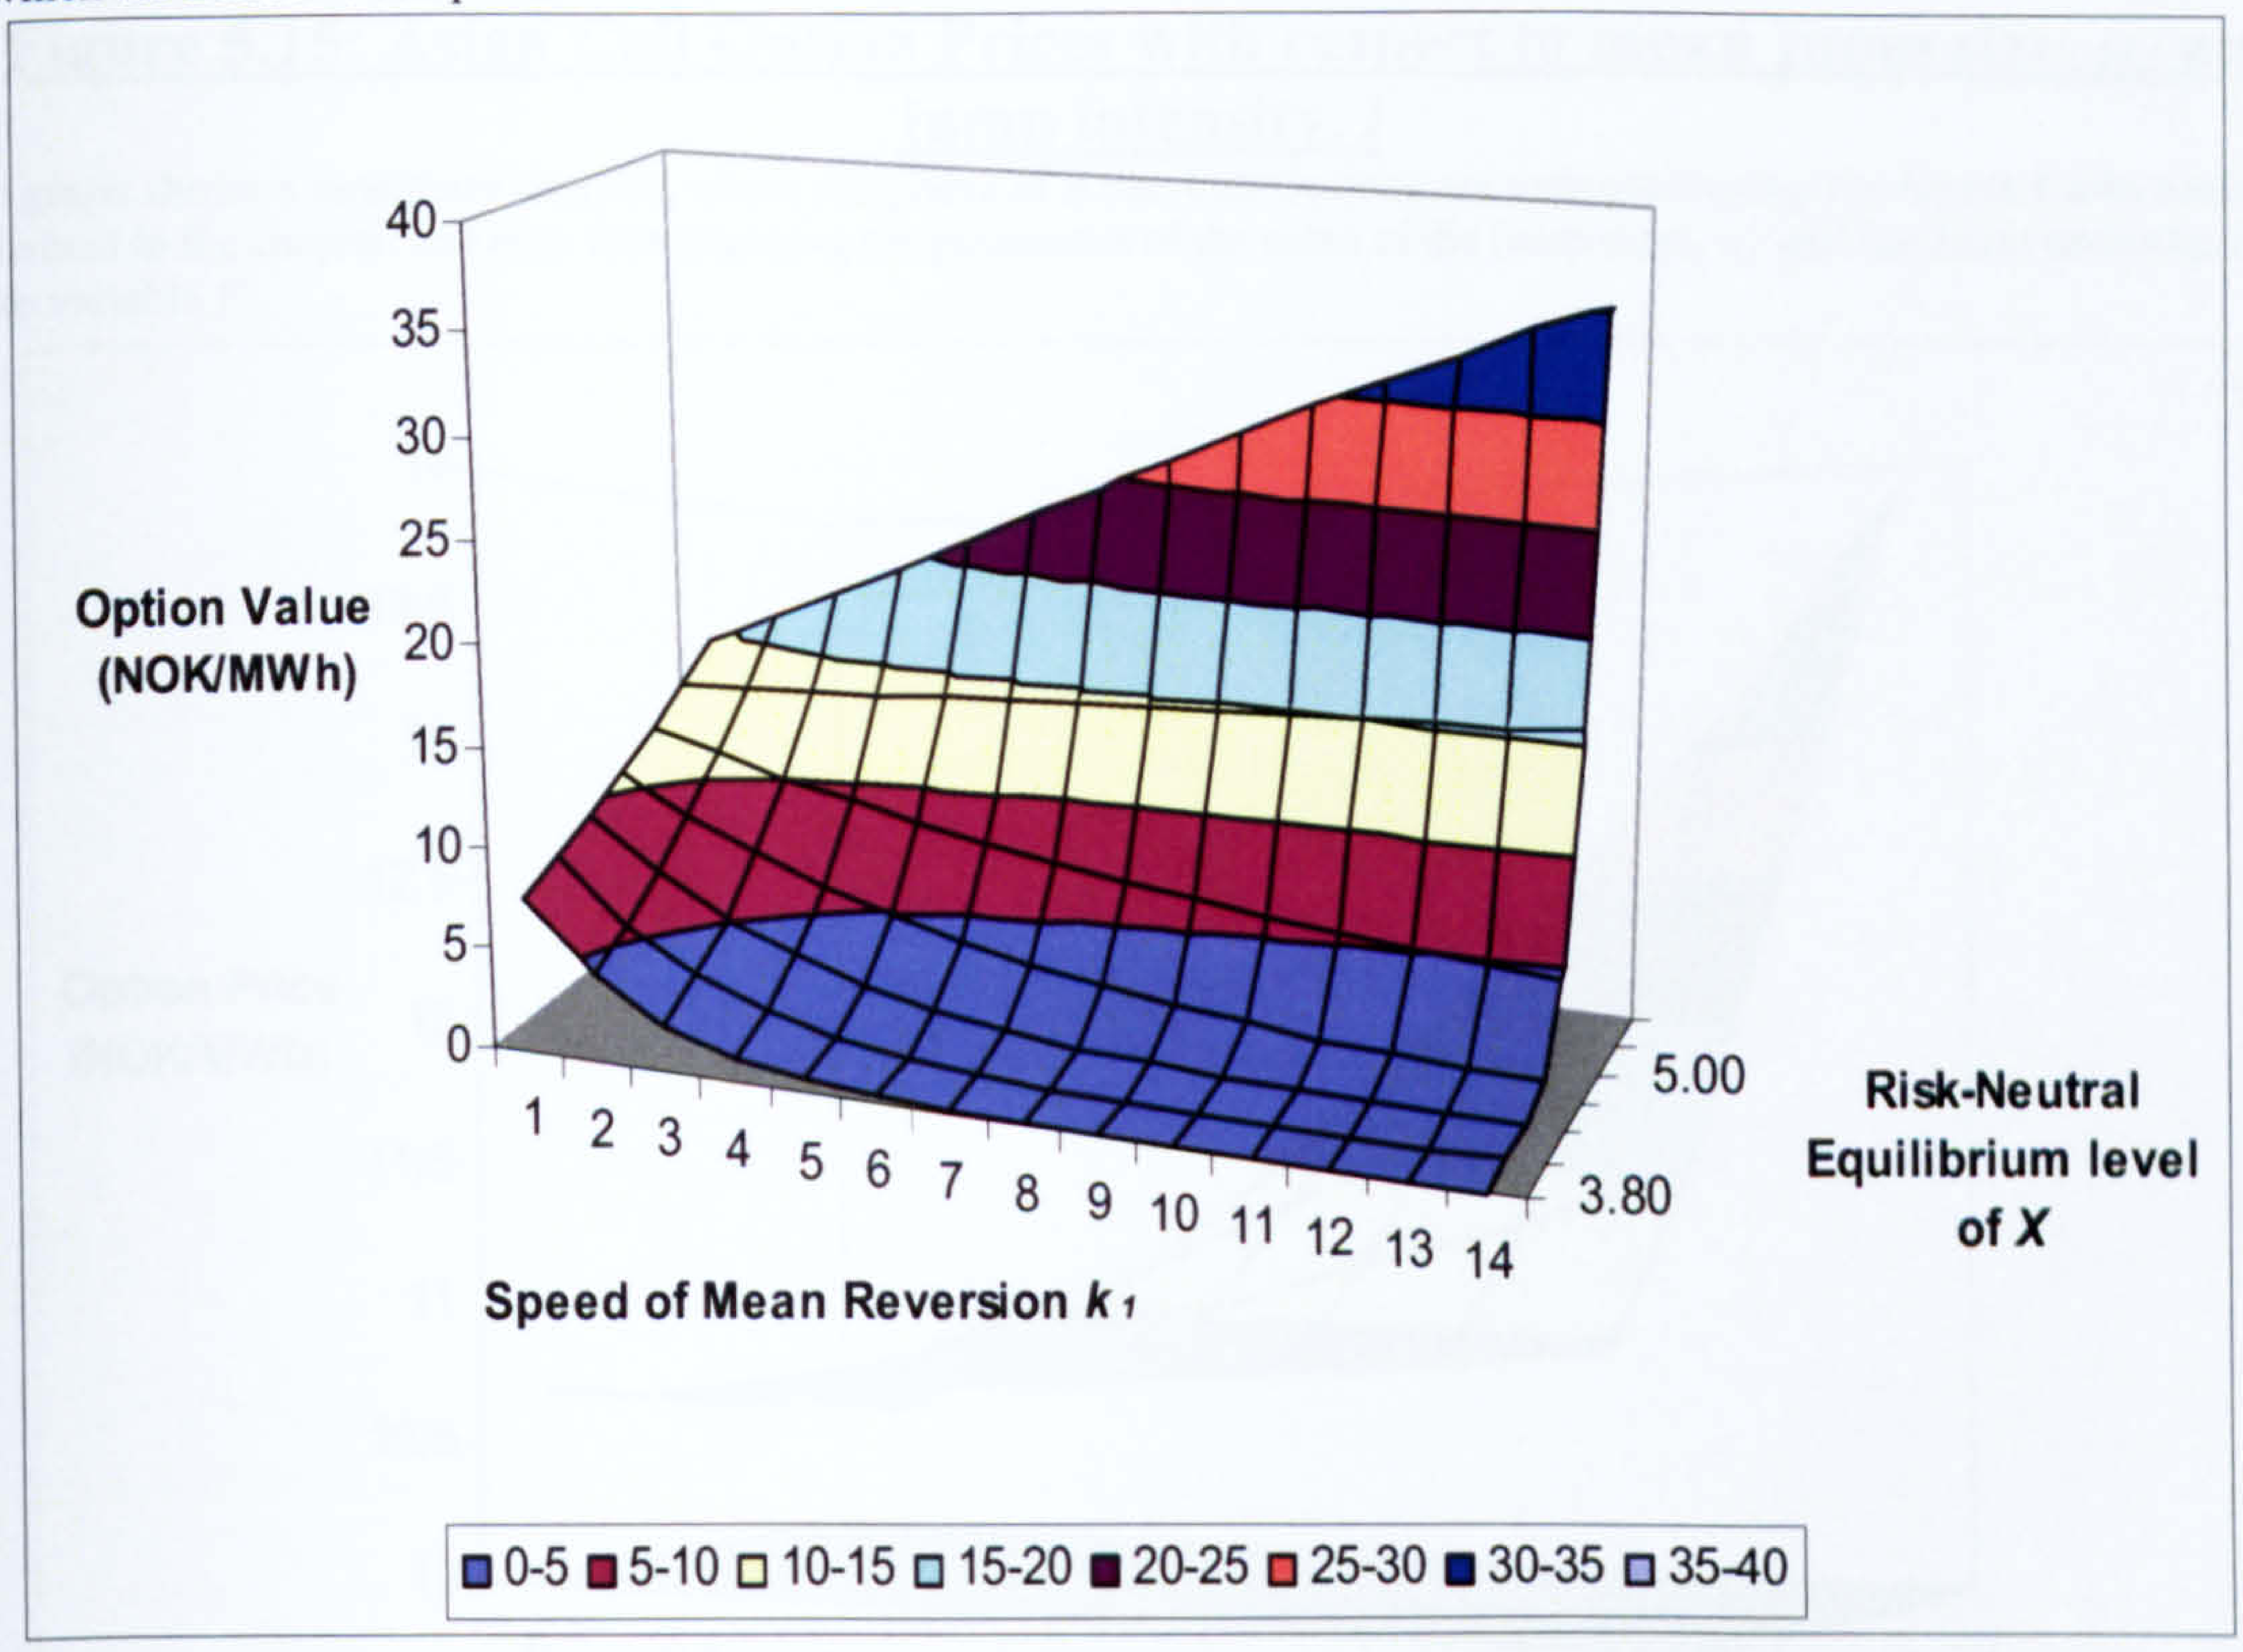

In order to see the impact of some of the jump components and on Asian option prices, Figure 5.15 shows the sensitivity with respect to the expected jump size return and jump arrivals. As expected, option prices increase in both the jump size and the arrival. However this occurs in a non-linear fashion. The contribution of the size increases at higher intensity levels. This is because if the jump intensity is low, the jump size does not play such a significant role, since its probability of occurrence is low, and thus the limited numbers of spikes that occur in the market are averaged out due to the specification of the Asian option contract. However, as the jump intensity increases jumps start to play a much more important role in the average price and, as their size increases, the total average of the spot price increases thus giving a higher probability of the option ending ITM. Another interesting fact is that with negative jump sizes, increasing the jump intensity the Asian call options decrease slightly, whereas with positive, it increases the price. However it seems that the absolute change in the Asian option price is small when the jumps become more negative, compared to the cases when the jumps become more positive. Similarly for Asian put options, we should see that option prices increase the larger the jump intensity, and the more negative the mean jump size, since the option holder now benefits more from the higher probability of negative jump sizes.

Figure 5.15: Asian Call Option Prices with respect to mean jump size, μ_J and jump intensity, l

The graph shows a sensitivity analysis, where the prices of Asian Call Options are estimated using the Monte Carlo Simulation explained in the chapter, and each time changing the parameters of the mean of the jump sizes, μ_J and the Jump intensity, l of the spike variable Y.

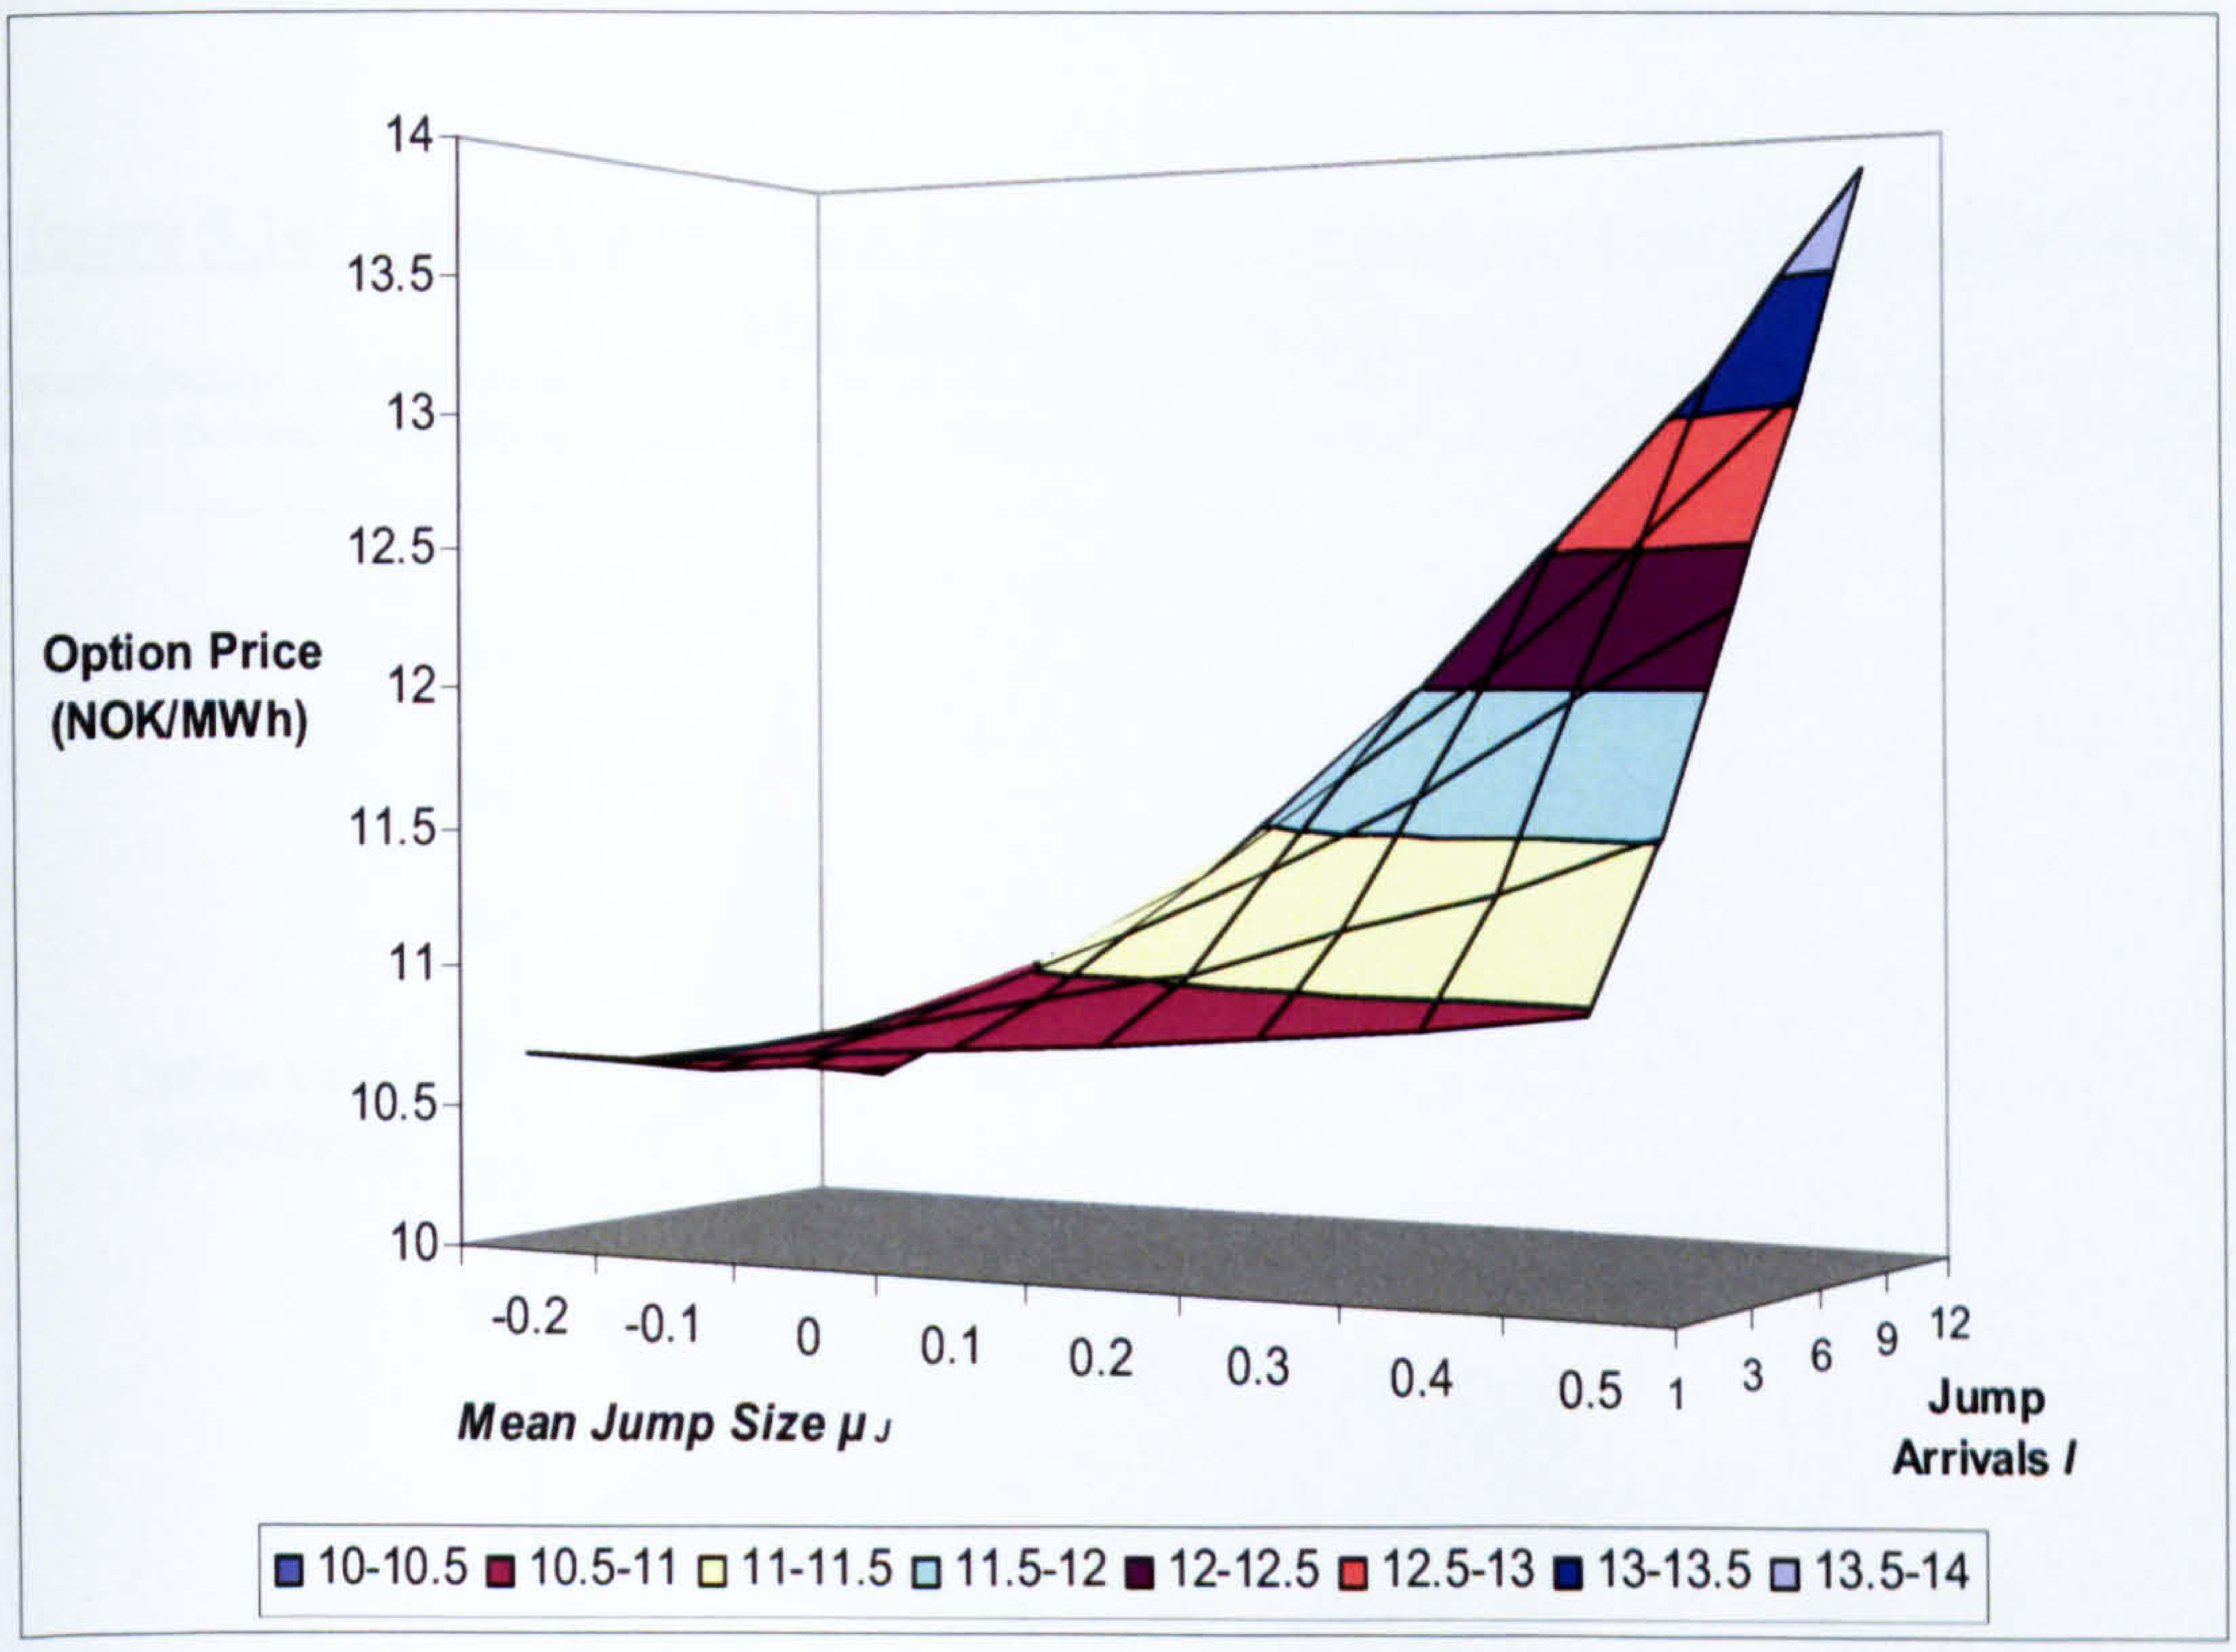

Regarding the jump size volatility and the spike mean reversion, in Figure 5.16 we see that Asian call options prices are inversely related to the spike mean reversion, but increase with the jump size volatility. This result is very interesting, as it shows the significance of the speed of jump mean reversion. The difference is large for mean reversion levels of approximately 50 (half-life of 5 days) and less. This is due to the fact that at low levels of mean reversion, if a jump occurs in the market, the spot price will stay at extreme price levels for a long period of time. Therefore this has a significant impact on the four-week average spot price, thus giving very high probabilities for the Asian option to end ITM. It is at those low levels of spike speed of mean reversion where the jump size volatility plays a very critical role, since if we increase the jump volatility, the jump sizes become much larger and prices stay at high levels for a long period of time thus increasing the value of the Asian call options. However as the speed of mean reversion increases, the marginal difference by increasing the jump size volatility is becoming less significant, thus the value of the Asian call option changes much less. This is because any extreme event that may occur, will affect the price levels for about a day, and thus the 28-day average of the spot is not significantly affected. For Asian Put options we should also expect the same impact since again the holder will benefit from big variations on the price that will last longer, given of course that the mean jump size is negative this time.

Figure 5.16: Asian Call Option Prices with respect to spike mean reversion, k_2 and jump size volatility, σ_J

The graph displays a sensitivity analysis, where the prices of Asian Call Options are estimated using the Monte Carlo Simulation explained in the chapter, by changing the volatility of the jump sizes, σ_J , and also its speed of mean reversion, k_2 of the Spike Variable *Y*.

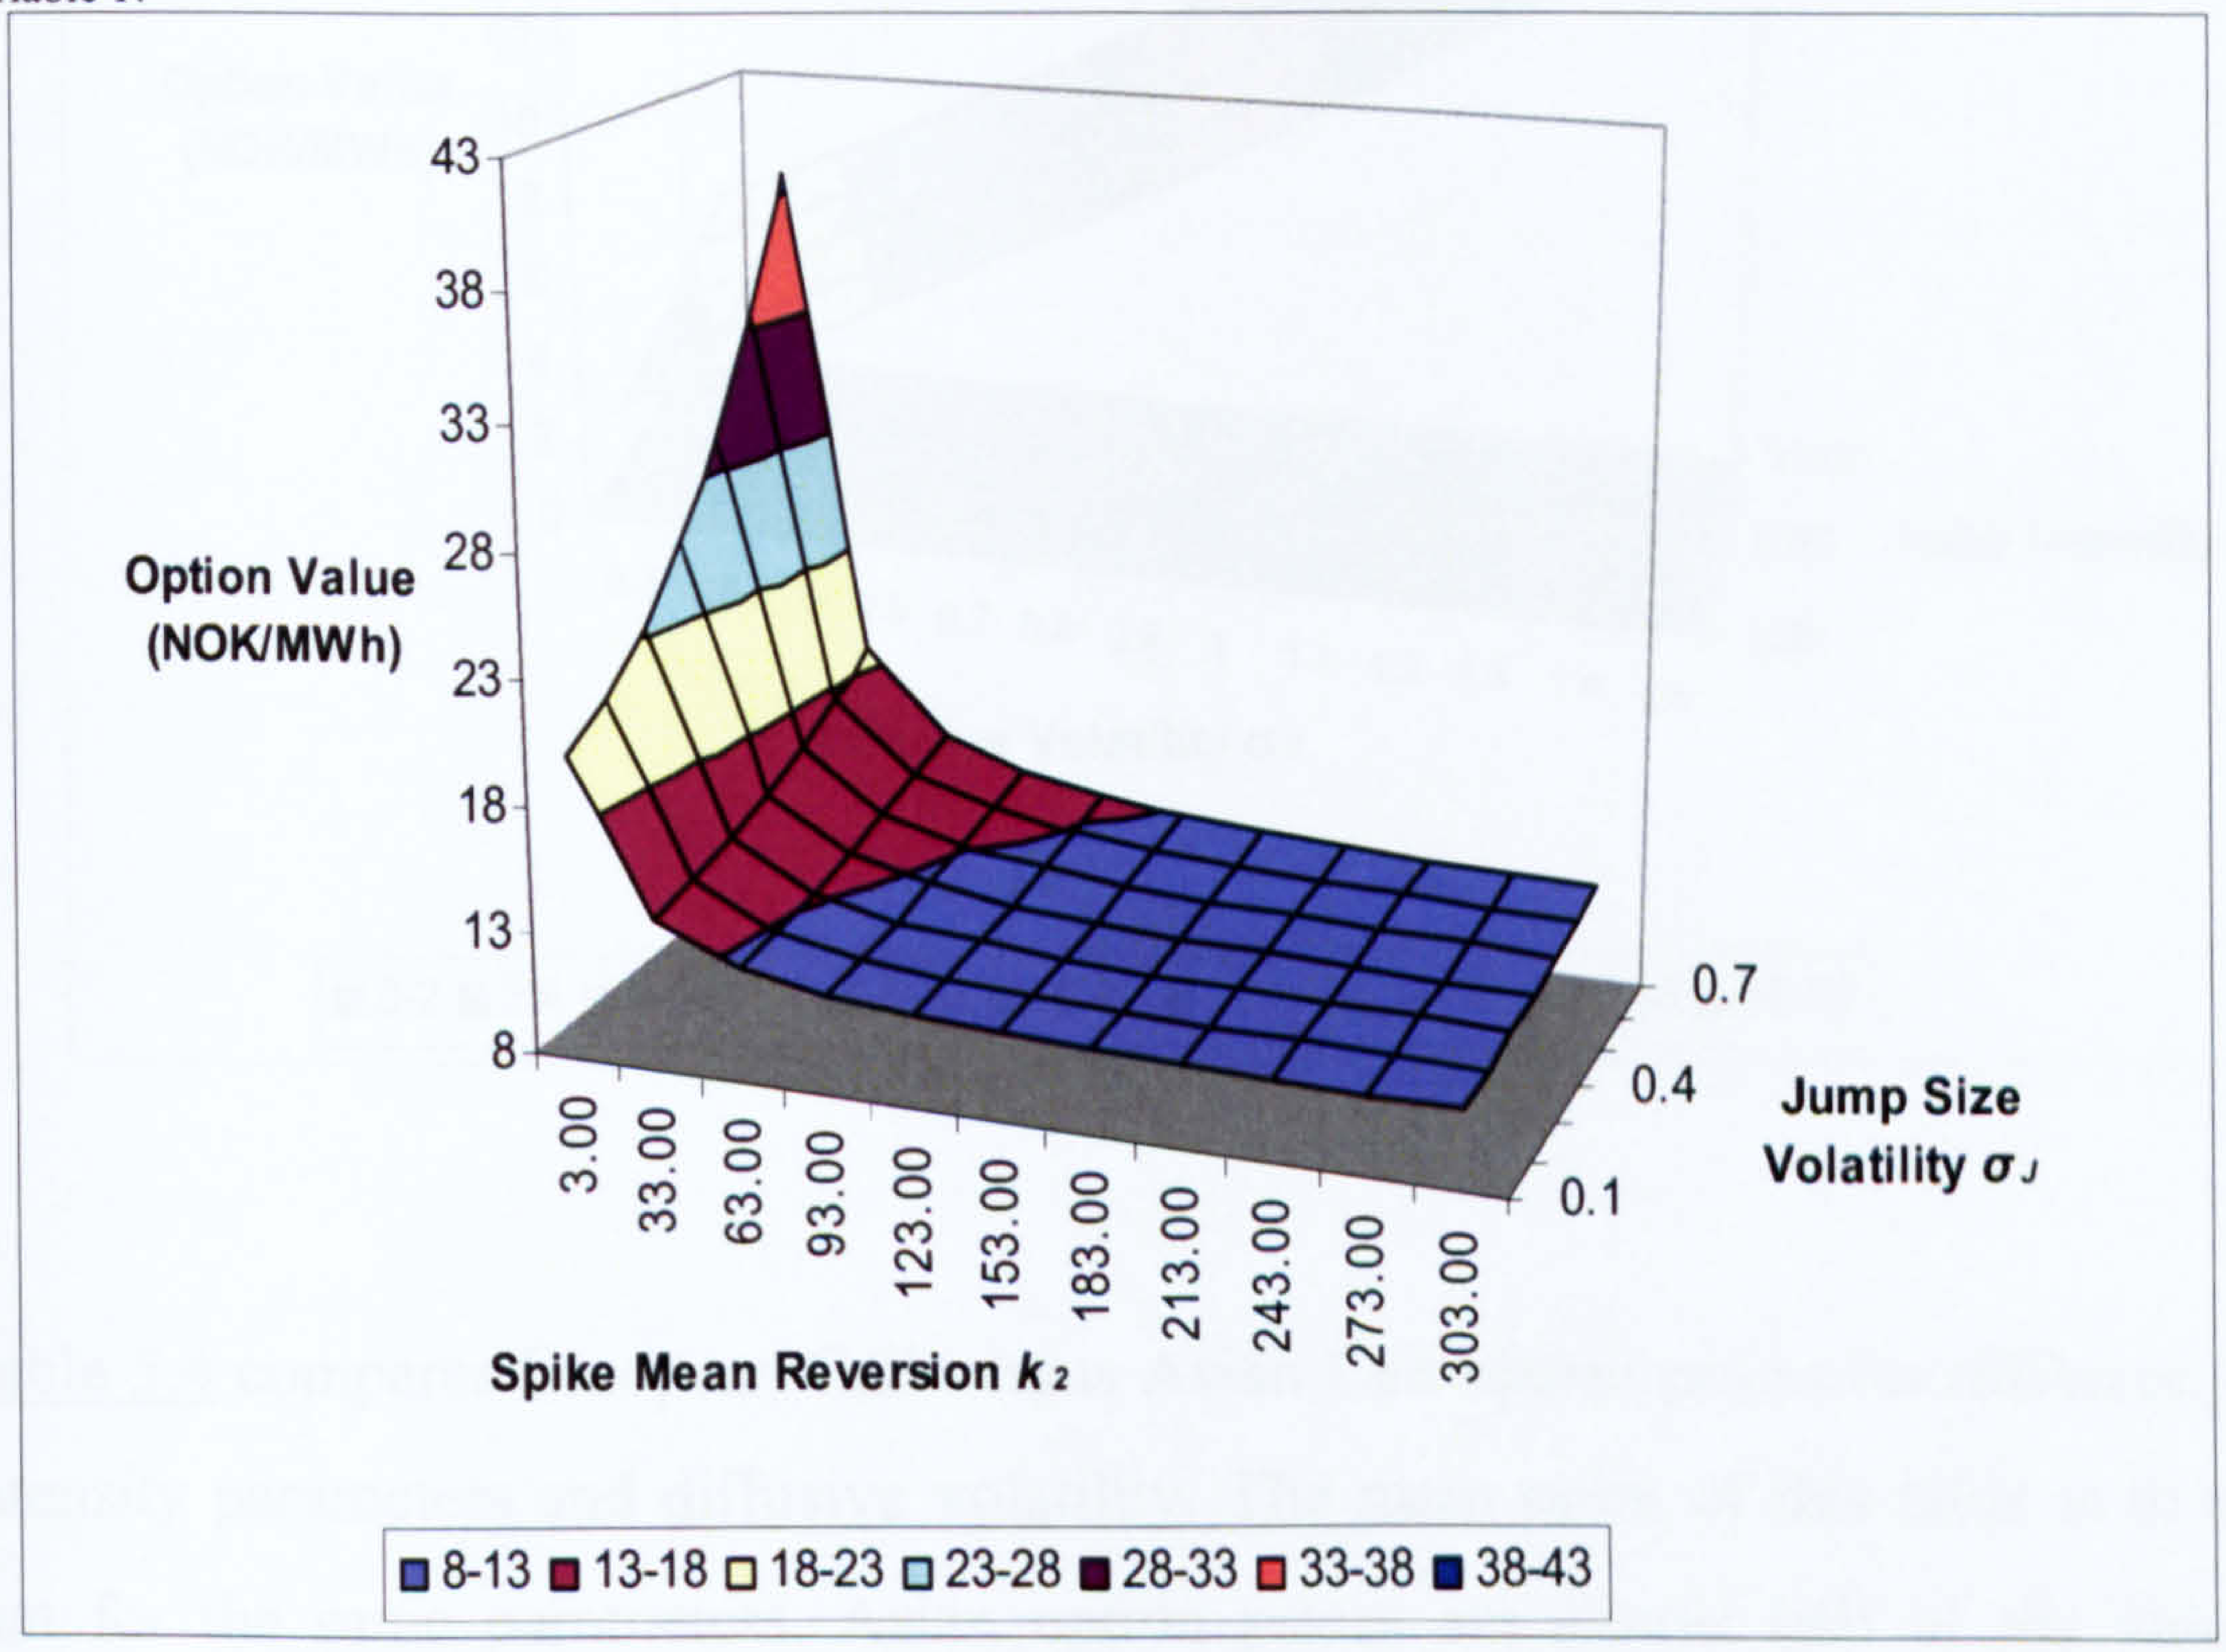

Figure 5.17 shows the sensitivity of the Asian call option prices with respect to both jump intensity and the diffusive volatility. In this graph we see that both the jump intensity and the diffusive volatility increase the option value. The reasons for the jump intensity were explained earlier. For the diffusive volatility, the higher it is, the more extreme a shock in the market is, and since diffusive shocks are continuous and generated from a Brownian motion, the greater the scaling the more probable higher prices are and thus the greater the average spot price. However the effect would be

less steep than that observed in the case of a European Option due to the impact of the averaging period.

Figure 5.17: Asian Call Options Prices with respect to jump intensity, *l* and diffusive volatility, σ_X

The graph displays a sensitivity analysis, where the prices of Asian Call Options are estimated using the Monte Carlo Simulation explained in the chapter, by changing the parameters of the volatility, σ_X, of diffusive variable X, and also the jump intensity, l, of the spike variable Y.

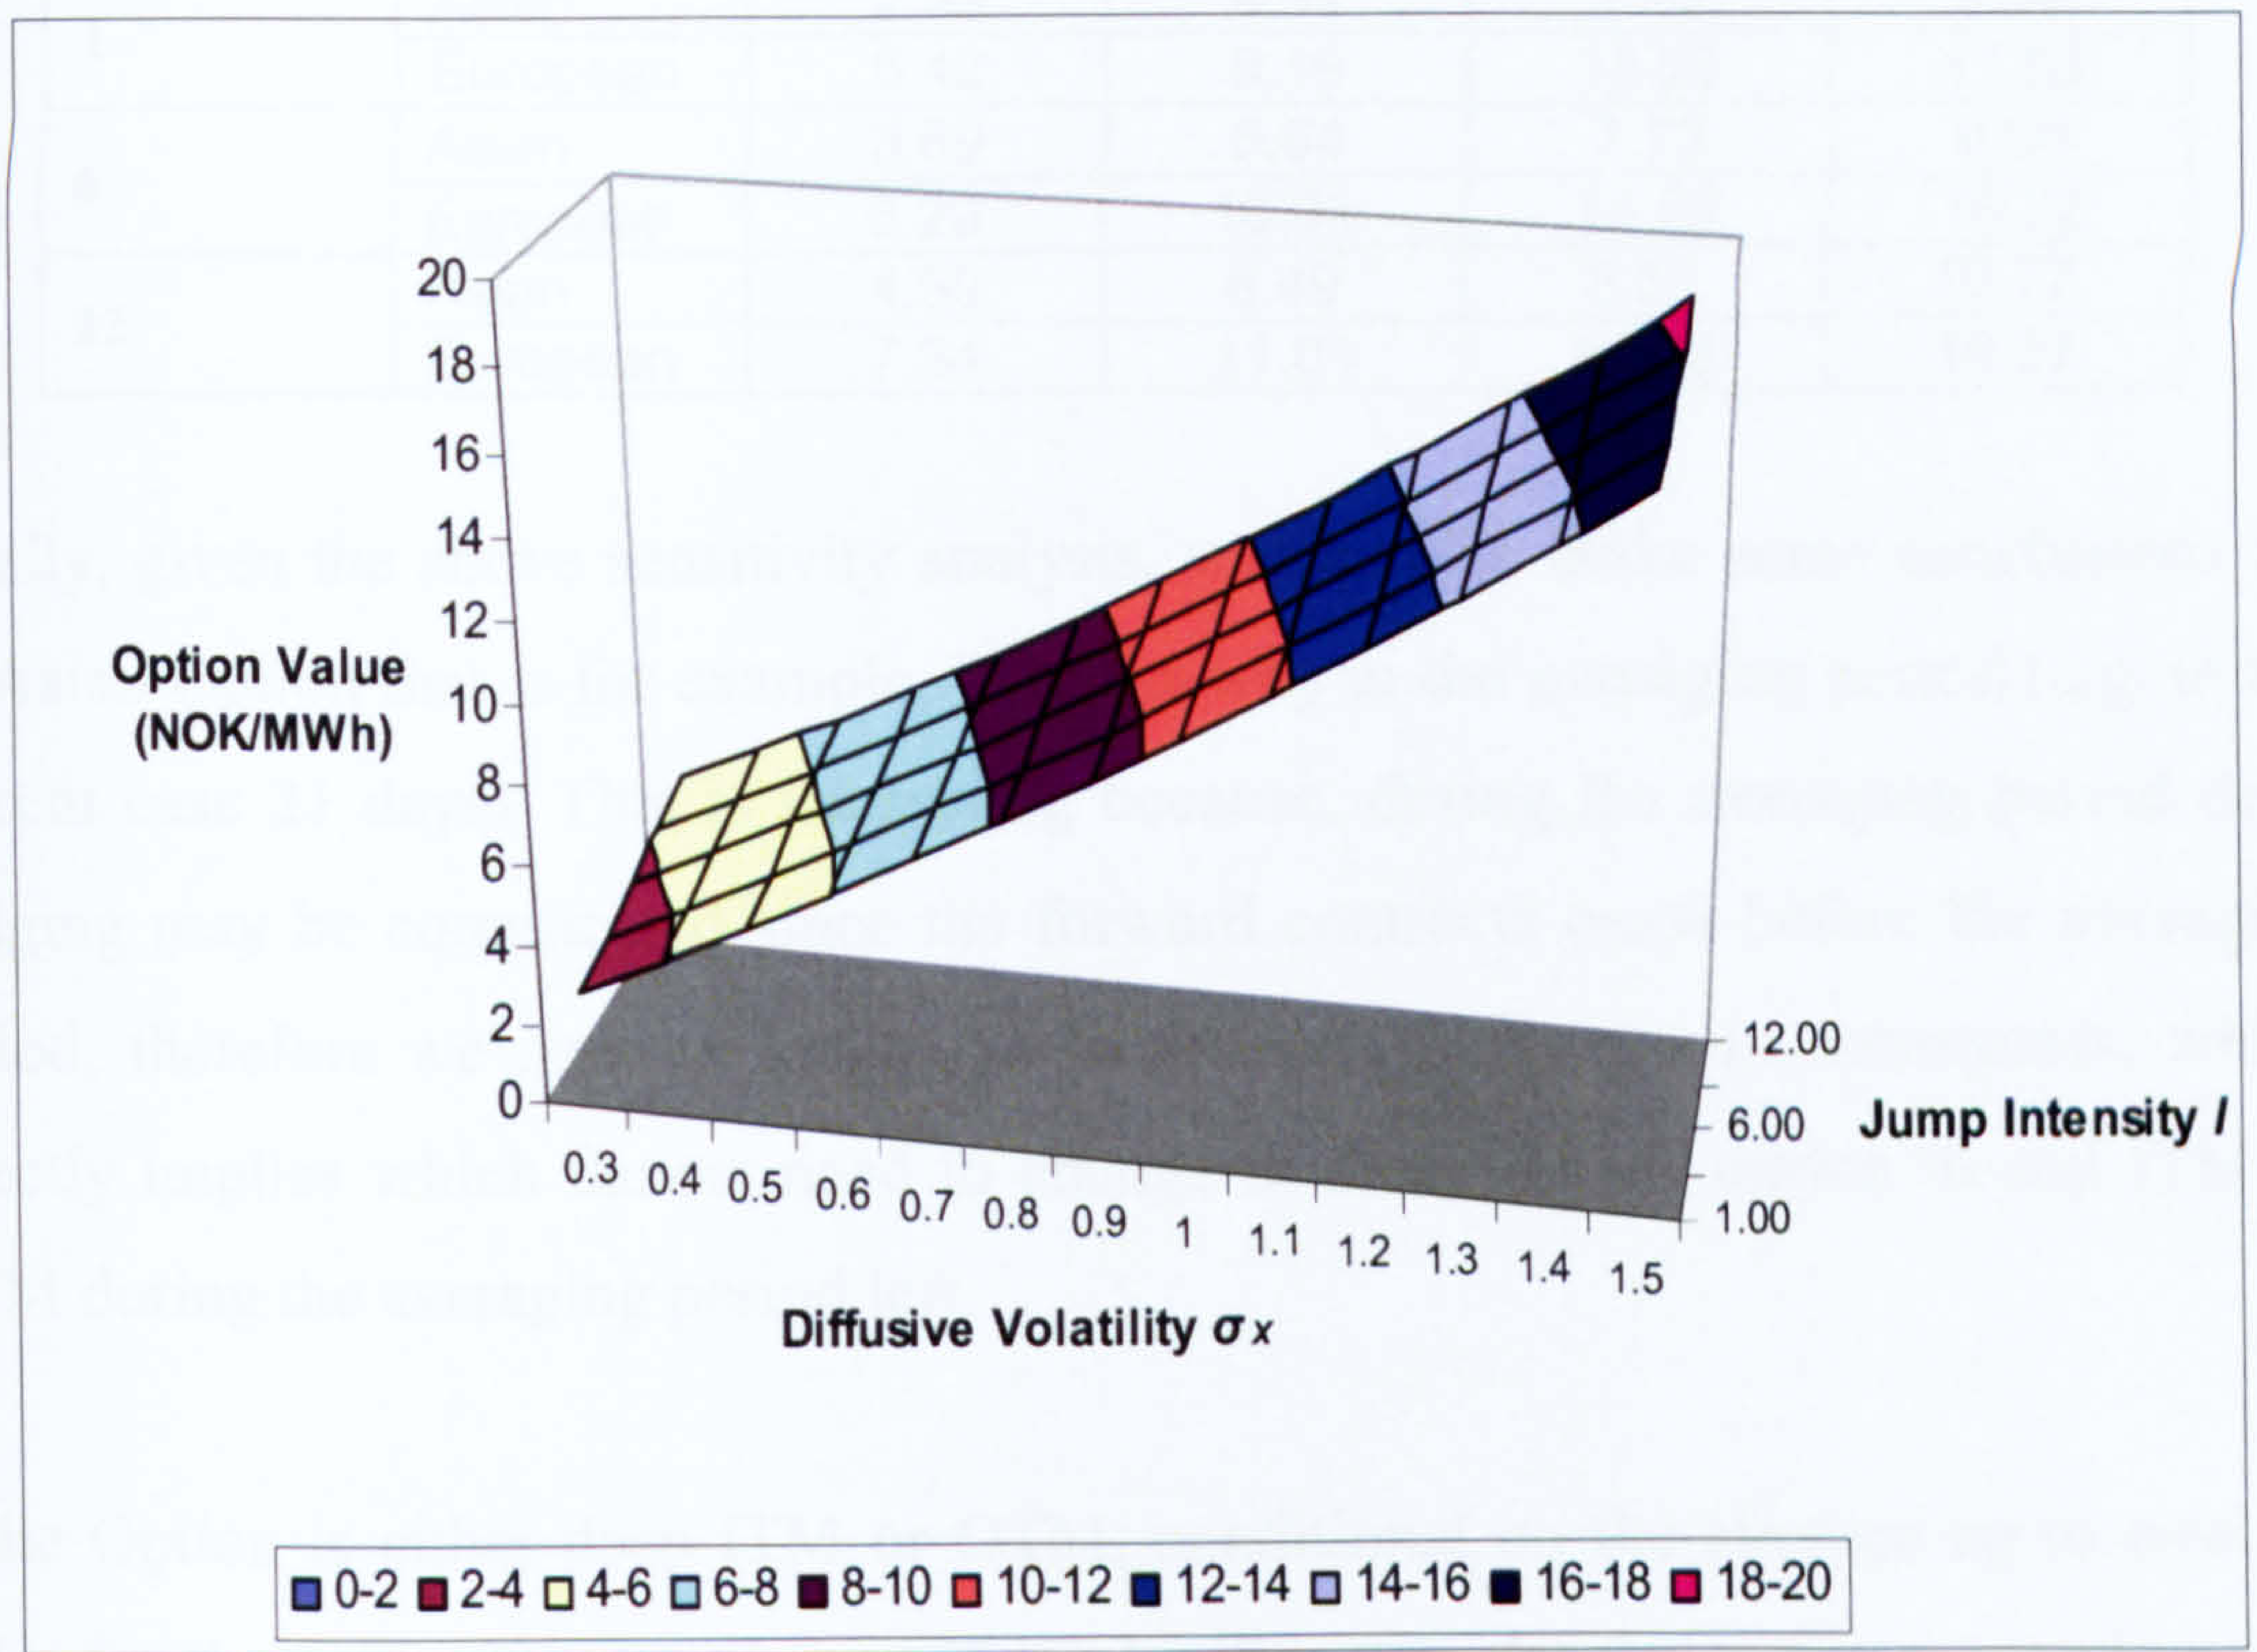

Table 5.4 compares European Call versus Asian Call option prices for different jump intensity parameters and diffusive volatility. The main point of this table is to show that for the same parameters, Asian option prices are almost half of the value of European call prices. This is because the actual spot price is much more volatile than the average spot price. This implies that the value of caps is much higher than the value of Asian options. Thus Asian options still provide protection, at a cheaper price. The problem however comes when short-lived extreme price movements occur, which, as shown earlier, are averaged out in the Asian options. On the other hand, if one has a portfolio of call options on different dates, he is protected from those extreme events, albeit at a much higher price.

Table 5.4: European Versus Asian Call Option Prices (in NOK/MWh) for different Jump Intensity, l, and Diffusive Volatility Parameters, σ_X

		Diffusive Volatility σ_X			
Jump Intensity, l		0.3	0.5	0.7	0.9
1	Asian	2.93	4.91	7.02	9.29
	European	5.42	9.19	13.22	17.53
6	Asian	3.69	5.64	7.73	9.96
	European	6.29	10.01	14.02	18.32
12	Asian	4.59	6.49	8.53	10.77
	European	7.34	11.01	14.98	19.27

Finally, given the above sensitivity analysis, we can also make some conclusions for an Asian Option that is for example, ¾ of the way in the averaging period (e.g. in the current case 21 days). This is interesting because, during the averaging period delta hedging may be complicated since the forward contracts cease before the averaging period, therefore we need to know the likelihood of changes in moneyness, which directly implies which factors need to change in order for the option to end ITM or OTM during the averaging period left.

If the Option is either deep ITM or OTM, conditional on the average up to now, in order for the moneyness to change state significantly, the following parameters will play a significant role. First of all the jump intensity has to increase significantly, so that there is a high probability of a spike. Next, the absolute value of the mean has to increase and together the volatility. However the above factors will not have any significant effect if the spikes do not last. Thus we believe the most important factor that can change the moneyness of the option is the speed of mean reversion of the spike. If the speed of mean reversion is low, any jump that will occur in the market, will affect the average significantly more. Alternatively we need a change in the equilibrium level, but in order for that to affect the moneyness of the option, the change has to be significant and also the speed of mean reversion of the normal factor has to be very fast, a fact that is not evident in the Nord Pool market.

5.8 Conclusions

In this chapter we started by explaining the importance of Caps, Floors and Asian Options for hedging price risk in the electricity markets. Even though Caps and Floors can give protection for every day, the premium may be too large, thus an alternative cheaper approach would be to use Asian options and get protection from the overall level of price in a given delivery period, rather than just a single day event.

Next we derived closed-form solutions under the mean-reverting and the spike models. The benefits of the mean-reverting model rely on the Samuelson hypothesis, which is consistent for any non-storable asset, and thus provides decreasing volatility term structure. The option pricing formula is very similar to the Black & Scholes (1973), but instead the underlying seems to be the forward price, as it is the expectation of the spot. However when introducing jump processes, only semi-closed form solutions are available. Using the transform analysis by Duffie et al (2001), the semi-closed form solutions for the spike model are derived, and we provide an intuitive explanation on each of the terms of the pricing equation. Numerical integration in the semi-closed form solution can be easily achieved using advanced methods such as the Adaptive Simpson Quadrature.

To get a further insight on the spike model, we use it to calculate options prices and then we use the prices as an input in the BS model to back out the implied volatilities. A similar analysis is also implemented for the mean-reverting model. We find that even an Ornstein-Uhlenbeck process displays volatility skews but always for ITM options, which is a result of mean reversion. However the presence of jumps displays volatility skews for OTM options depending on the sign of the mean jump size. We also show that mean reversion reduces the volatility smile as time to maturity increases. Furthermore, the jump size volatility and jump intensity mainly affect the kurtosis and thus the curvature of the smile. Moreover, we also provide evidence that for a desired jumpiness level, the jump size volatility affects the volatility smile more than the jump intensity.

Next we explored the option prices generated by the four proposed models in the previous chapters. The mean-reverting model seems to converge to an option price very fast as time to maturity increases, and from a point onwards, the discounting effect of the interest rate starts playing a more important role. On the other hand, the regime switching spike model gives the following results; if we are in a state where there is enough water in the reservoirs, we see that call option prices increase with time to maturity consistently. In this case, the premium paid is derived by the risk that we may end up in a state where the reservoir level is low. The opposite effect happens when we start from a state when the reservoir levels are low, and thus we have a downward sloping option price. The same analysis was performed using the Three-Factor model, however in this case the long-run factor, increases the longer the time to maturity. Also increasing the correlation between the two normal state variables increases the volatility and thus the option price.

Finally we performed a sensitivity analysis using the spike model on the Asian call options. Since no closed-form solution is available and also the model displays jumps, the most appropriate method is the Monte Carlo simulation. The sensitivity analysis showed that Asian call option prices increase with jump intensity and jump size, as the averaging becomes more affected by the intensity. However the effect of the speed of mean reversion has a dual effect depending on the equilibrium level. On the other hand, the spike mean reversion plays a vital role especially at low levels, as it may give very high Asian option prices.

Having examined these issues, the next chapter is devoted on Swing Options, which are a very useful structured product in the energy markets particularly for hedging both volumetric and price risk at any time of convenience of the option holder. One of the main difference from the products examined in the current chapter, is that swing options allow for early exercise, therefore it will be very interesting to examine the factors driving their early exercise features. We will show how the pricing is performed both theoretically and practically, conduct a sensitivity analysis mainly for the extra swing premium i.e. what is the extra price we have to pay to use this particular product rather than a more simple one such as Caps and Floors, to the different parameters, define the boundary conditions, as well as exploring the effect of penalties on the options.

5.9 Appendix:

5.9.1 Derivation of Option pricing formula using the Seasonal Mean-Revering Model

Recall the electricity model when we have only mean reversion in the risk-neutral world:

$$P_t = f(t) + \exp(X_t)$$
$$dX_t = \left(k_1\left(\tilde{\varepsilon} - X_t\right)\right)dt + \sigma_X dW_X^* \quad \text{where } \tilde{\varepsilon} = \varepsilon - \lambda_X / k_1 \tag{5.31}$$

Thus the Option price is now equal to the expected pay-off in risk-neutrality:

$$Call = E^*\left[e^{-rT}\left(P_T - K\right)^+ / \mathcal{F}_0\right]$$
$$= e^{-rT}E^*\left[\left(\exp\left(X_0 e^{-k_1T} + \tilde{\varepsilon}\left(1-e^{-k_1T}\right) + \sigma_X \int_0^T e^{-k_1(T-s)}dW_X^*\right) - (K - f(T))\right)^+ / \mathcal{F}_0\right]$$

Now we use the following facts:

1. The seasonal component is deterministic and thus we can think of a new deseasonalised strike price equal to: *DK=(K-f(T))*
2. The deseasonalised forward price of electricity under the model is a forward price on the deseasonalised spot, thus depending only on X:
$$F(0,T,X_0) = \exp\left(X_0 e^{-k_1T} + \tilde{\varepsilon}(1-e^{-k_1T}) + \frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)\right)$$
3. Due to the quadratic variation of a Brownian Motion,
$$E^Q\left[\left(\sigma_X \int_0^T e^{-k_1(T-s)}dW_X^*\right)^2 / \mathcal{F}_0\right] = \frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1T}\right)$$
4. $\sigma\int_0^T e^{-k_1(T-s)}dW_X^* \sim N\left(0, \frac{\sigma_X^2}{2k_1}\left(1-e^{-2k_1T}\right)\right)$ has the same distribution as

$\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}y$, where *y~N(0,1)*

Thus the option pricing formula for the call now becomes:

$$Call = e^{-rT}\int_{-\infty}^{\infty}\left(F(0,T,X_0)e^{-\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)+\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}y} - DK\right)^+ e^{-\frac{y^2}{2}}\frac{1}{\sqrt{2\pi}}dy \tag{5.32}$$

So we need to find the value from which the integration should start:

$$F(0,T,X_0)\mathrm{e}^{-\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)+\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}y} > DK$$

$$y > \frac{\ln\left(\frac{DK}{F(0,T,X_0)}\right)+\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)}{\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}} = d \tag{5.33}$$

So what we are actually saying here is that if $P_T>K$, we are integrating d, where as if $P_T<K$, we are integrating zero. Hence, equation (5.32) now becomes

$$Call = \overbrace{e^{-rT}\int_d^{\infty} F(0,T,X_0)\mathrm{e}^{-\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)+\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}y}\, e^{-\frac{y^2}{2}}\frac{1}{\sqrt{2\pi}}dy}^{a} - \underbrace{e^{-rT}\int_d^{\infty} DKe^{-\frac{y^2}{2}}\frac{1}{\sqrt{2\pi}}dy}_{b} \tag{5.34}$$

Now take each term of the integration (a and b) separately. First for a, start by completing the square:

$$a = \int_d^{\infty} F(0,T,X_0)\mathrm{e}^{-\frac{\left(\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}-y\right)^2}{2}}\frac{1}{\sqrt{2\pi}}dy$$

Now substitute $\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}} - y = z \Rightarrow y = z + \sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}} > d$ and thus $z > d - \sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}$.

$$a = \int_{d-\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}}^{\infty} F(0,T,X_0)\mathrm{e}^{-\frac{(z)^2}{2}}\frac{1}{\sqrt{2\pi}}dz = F(0,T,X_0)N\left(-d+\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}\right)$$

Therefore

$$-d+\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}} = \frac{\ln\left(\frac{F(0,T,X_0)}{DK}\right)+\frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)}{\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}} = d_1 \tag{5.35}$$

In the same way for b:

$$b = \int_d^{\infty} DKe^{-\frac{y^2}{2}} \frac{1}{\sqrt{2\pi}} dy = DK \ N(-d)$$

$$-d = \frac{\ln\left(\frac{F(0,T,X_0)}{DK}\right) - \frac{\sigma_X^2}{4k_1}\left(1-e^{-2k_1T}\right)}{\sigma_X\sqrt{\frac{\left(1-e^{-2k_1T}\right)}{2k_1}}} = d_2 \quad (5.36)$$

Hence using (5.35) and (5.36), and substituting in equation (5.34), we can recover the value of a call option:

$$call = e^{-rt}\left[F(0,T,X_0)N(d_1) - DK \ N(d_2)\right] \quad (5.37)$$

5.9.2 Put Call Parity in Electricity Market

Now we can also use the put-call parity relation to retrieve the value of a Put option, however since electricity is non-storable, our portfolios have to be hedged with futures. Consider that you have European call and put options on the underlying electricity, with the same strike price, K, and expiry date T. Form two separate portfolios:

A: a European call option and a cash amount of Ke^{-rT} invested in the risk-free rate
B: a European put option, a long position in a futures contract with the same expiry date and an amount of cash equal to the discounted price you locked for the futures position i.e. $F(0,T,P_0)e^{-rT}$ invested in the risk free rate

We can clearly see that the value of both portfolios at expiry will be equal to $max(F(T,T,P_0),K)=max(P_T,K)$, assuming that the futures converge with the spot at their expiration to diminish any arbitrage opportunities. Since the value of both portfolios is equal we have the following relationship:

$$Call + Ke^{-rT} = Put + F(0,T,P_0)e^{-rT}$$
$$e^{-rt}\left[F(0,T,X_0)N(d_1) - DK \ N(d_2) + K\right] = put + \left(f(T) + F(0,T,X_0)\right)e^{-rT} \quad (5.38)$$
$$put = e^{-rt}\left[DK \ N(-d_2) - F(0,T,X_0)N(-d_1)\right]$$

5.9.3 Proof of the formula for the security price $G_{a,b}(y)$

For $0 < \tau < \infty$ and a fixed $y \in \mathbb{R}$,

$$\frac{1}{2\pi}\int_{-\tau}^{\tau}\frac{e^{ivy}\psi^{\chi}(a-ivb,X,0,T)-e^{-ivy}\psi^{\chi}(a+ivb,X,0,T)}{iv}dv$$
$$=\frac{1}{2\pi}\int_{-\tau}^{\tau}\int_{\mathbb{R}}\frac{e^{-iv(z-y)}-e^{iv(y-z)}}{iv}dG_{a,b}(z;X,T,\chi)dv$$
$$=-\frac{1}{2\pi}\int_{\mathbb{R}}\int_{-\tau}^{\tau}\frac{e^{-iv(z-y)}-e^{iv(y-z)}}{iv}dvdG_{a,b}(z;X,T,\chi)$$

where *Fubini* is applicable because

$$\lim_{y\to\infty} G_{a,b}(y;X,T,\chi)=\psi^{\chi}(a,X,0,T)<\infty$$

given that χ is well defined at (a,T). Next we note that, for $\tau>0$,

$$\int_{-\tau}^{\tau}\frac{e^{-iv(z-y)}-e^{iv(y-z)}}{iv}dv=-\frac{\operatorname{sgn}(z-y)}{\pi}\int_{-\tau}^{\tau}\frac{\sin(v|z-y|)}{v}dv$$

is bounded simultaneously in z and τ, for each fixed y. By the bounded convergence theorem,

$$\lim_{\tau\to\infty}\frac{1}{2\pi}\int_{-\tau}^{\tau}\frac{e^{ivy}\psi^{\chi}(a-ivb,X,0,T)-e^{-ivy}\psi^{\chi}(a+ivb,X,0,T)}{iv}dv$$
$$=-\int_{\mathbb{R}}\operatorname{sgn}(z-y)dG_{a,b}(z;X,T,\chi)$$
$$=-\psi^{\chi}(a,X,0,T)+\left(G_{a,b}(y;X,T,\chi)+G_{a,b}(y-;X,T,\chi)\right)$$

where $G_{a,b}(y-;X,T,\chi)=\lim_{z\to y,z\le y}G_{a,b}(z;X,T,\chi)$. Using the integrability condition

$\int_{\mathbb{R}}\left|\psi^{\chi}(a+ivb,X,0,T)\right|dv<\infty$, and by the dominated convergence theorem we have

$$G_{a,b}(y;X,T,\chi)=\frac{\psi^{\chi}(a,X,0,T)}{2}+\frac{1}{4\pi}\int_{-\infty}^{\infty}\frac{e^{ivy}\psi^{\chi}(a-ivb,X,0,T)-e^{-ivy}\psi^{\chi}(a+ivb,X,0,T)}{iv}dv$$

Because $\psi^{\chi}(a-ivb,X,0,T)$ is the complex conjugate of $\psi^{\chi}(a+ivb,X,0,T)$ we have

$$G_{a,b}(y;X,T,\chi)=\frac{\psi^{\chi}(a,X,0,T)}{2}+\frac{1}{\pi}\int_{0}^{\infty}\frac{e^{-ivy}\psi^{\chi}(a+ivb,X,0,T)}{iv}dv$$

5.9.4 Option Replication in the electricity market

In this section, we are describing how a hedging portfolio may be constructed in the electricity market, to replicate a derivative contract. Since electricity is not storable and thus cannot be used for the usual delta hedging strategies, the obvious replacement would be a forward. Lets assume that the spot price evolves as in the spike model. For this example let's assume that we have no seasonality in order to get a clear picture, and the jump distribution is assumed to be free, i.e. we are not restricting the distribution of the jump size, *J*, to be Normal. The spot price under the spike model is described as follows:

$$
\begin{aligned}
P_t &= \exp(X_t + Y_t) \\
dX_t &= k_1\left(\varepsilon - X_t\right)dt + \sigma_X dW_X \\
dY_t &= -k_2 Y_t dt + J dq(l)
\end{aligned}
$$

As it may be seen, we have two sources of risk: one arising from the wiener process of the normal diffusive variable, X and the second arising from the jump variable dq. Suppose now that we need to hedge a derivative on the electricity price, for example a call, *C*. In order to do so, we create a portfolio consisting of one call, to be replicated by α units of a forward, $F(t,T)=F_t$, β units of a bond B, used to finance the portfolio, and γ units of a second ("hedging") option, V. Thus the value of the portfolio at time t is as follows:

$$\Pi_t = C_t + \alpha F_t + \beta B_t + \gamma V_t$$

The main aim of the portfolio is to make it riskless and thus have a deterministic value. If we do so, we may chose an amount of the discount bond and our investment today should have zero cost, and thus zero growth. Before describing the evolution of our portfolio Π, we need to show the dynamics of the derivatives. Under the risk neutrality measure, the dynamics of C are as follows:

$$dC = \left(\frac{\partial C}{\partial t} + k_1(\varepsilon - X)\frac{\partial C}{\partial X} - k_2 Y \frac{\partial C}{\partial Y} + \frac{\sigma_x^2}{2}\frac{\partial^2 C}{\partial X^2} \right) dt$$
$$+ \left(C(Y+J) - C(Y) \right) dq - lE\left(C(Y+J) - C(Y) \right) dt$$
$$\sigma_x \frac{\partial C}{\partial X} dW_x$$

the first dt terms, and the last terms dW_X, are the usual ones described by Ito's Lemma. However the middle terms arise from the extended Ito's Lemma. $\left(C(Y+J) - C(Y) \right) dq$ is the change in the derivative's price if a jump occurs, which is discontinuous. However, the extra term $lE\left(C(Y+J) - C(Y) \right) dt$, is subtracted, in order to compensate for the jump process in the derivative's price, and thus make it a martingale. V follows the same kind of process as C. On the other hand, the forward price has a zero cost of entering, and thus the first dt terms should not exist, as there should be no return when entering a forward contract. Thus the forward price evolves as follows:

$$\frac{dF_t}{F_t} = e^{-k_1(T-t)}\sigma_x dW_x + (e^{Je^{-k_2(T-t)}} - 1)dq - lE\left(e^{Je^{-k_2(T-t)}} - 1 \right) dt$$

Thus the portfolio evolves as follows:

$$d\Pi_t = \left(\frac{\partial C}{\partial t} + \frac{\partial V}{\partial t} + k_1(\varepsilon - X)\left(\frac{\partial C}{\partial X} + \gamma\frac{\partial V}{\partial X} \right) - k_2 Y\left(\frac{\partial C}{\partial Y} + \gamma\frac{\partial V}{\partial Y} \right) + \frac{\sigma_x^2}{2}\left(\frac{\partial^2 C}{\partial X^2} + \gamma\frac{\partial^2 V}{\partial X^2} \right) + \beta r B_t \right) dt$$
$$+ \left(\left(C(Y+J) - C(Y) \right) + \gamma\left(V(Y+J) - V(Y) \right) + \alpha F_t\left(e^{Je^{-k_2(T-t)}} - 1 \right) \right) dq$$
$$- l\left(E\left(C(Y+J) - C(Y) \right) + \gamma E\left(V(Y+J) - V(Y) \right) + \alpha F_t E\left(e^{Je^{-k_2(T-t)}} - 1 \right) \right) dt$$
$$+ \left(\frac{\partial C}{\partial X} + \gamma\frac{\partial V}{\partial X} + \alpha F e^{-k_1(T-t)} \right) \sigma_x dW_x$$

In order to make the portfolio immune from changes in the Wiener process, the amount of forwards, α, needed in the portfolio must equal to:

$$a = -\frac{\frac{\partial C}{dX} + \gamma\frac{\partial V}{dX}}{F_t e^{-k_1(T-t)}}$$

By analogy with the stochastic-diffusion treatment, one would be tempted to choose the amount of the second derivative, γ, in such a way that all dq terms should be identically equal to zero:

$$\gamma = -\frac{\frac{\partial C}{\partial X}\left(e^{Je^{-k_2(T-t)}} - 1\right) - e^{-k_1(T-t)}\left(C(Y+J) - C(Y)\right)}{\frac{\partial V}{\partial X}\left(e^{Je^{-k_2(T-t)}} - 1\right) - e^{-k_1(T-t)}\left(V(Y+J) - V(Y)\right)}$$

In the above expression, the hedging amount of option V, depends on J, which is the jump size over the next time interval if a jump occurs. However, we have a problem since J is a random variable and can take any value generated from its distribution. If we knew the magnitude and sign of the jump, J, before it happened it would be possible to choose a value of γ, such that the change in the value of the option V, due to the jump could perfectly compensate for the jump component in the option C. However from the same expression, it is clear that no single value of γ can ensure that for any possible realizations of the random variable J, the amount γ of the option V would be just right to ensure that the change in the value of the rest of the portfolio due to the jump would be exactly compensated by the change in V, also due to the jump. Therefore after this simple illustrative discussion, we may conclude that perfect replication is impossible when the dynamics of the underlying involve jumps.

6. SWING OPTIONS AND THE EXTRA SWING PREMIUM

6.1 Introduction

In order to hedge themselves against extreme price fluctuations of certain commodities, many consumers enter into forward contracts, which give them the right and the obligation to purchase a fixed amount of the commodity for a predetermined price. On the other hand, as we saw in Chapter 5, market participants may also use Caps, Floors or Asian options in order to hedge against price movements. Our results show that the cost of Caps and Floors is very high since their price is the sum of European Call or Put options for each day over the hedging period. We thus proposed the use of Asian options which are particularly useful for hedging against the average price over the desired period, although this strategy may not be effective for hedging risk over a specific day, especially when it comes to hedging against spikes that occur on a given day.

Moreover, the above proposed strategies for risk reduction may not be sufficient for some market participants, since they may not know their exact future need of the commodity, in terms of quantity. This is particularly the case for commodities that cannot be stored or for which storage is very expensive, in which case inventories cannot be used to carry forward the underlying asset. As a result, so-called Swing contracts have been developed in order to give their holder flexibility with respect to the amount to be purchased in the future. These contracts are mainly used in energy markets since energy is difficult (or expensive) to store and exhibits extreme price fluctuations. This is especially the case in electricity, although Swing contracts appear also in coal (as in Jaskow 1985 and 1987) and gas markets (Clewlow and Strickland 2000). The main characteristic properties of Swing contracts, such as the multiple early exercise features, are the same for all underlying commodities. Only the choice of suitable a stochastic process depends strongly on the type of underlying asset. In the following we concentrate mainly on Swing options on electricity.

The popularity of these instruments is due to the fact that a large number of banks and financial institutions are active participants in the commodity markets and the electricity markets, in particular. Many market participants do not own the physical assets (i.e. generators) to make delivery of electricity a viable option and hence must focus on financial hedging if they are to participate in the market. Thus, if a bank wishes to enter into a transaction with a client in which the client would like to buy a financial (cash settled) swing contract to hedge the cost of their electricity supply, the bank needs to have in place a hedging strategy in case it cannot find another party to take the opposite position and thus benefit from the bid-ask spread.

However, as noted in previous chapters, hedging contingent claims in electricity is more difficult compared to other markets. Delta hedging in electricity markets can be performed using forward contracts, but that only hedges the diffusive risk coming from the Brownian motion. On the hand, it is impossible to hedge the jump risk completely, while illiquidity in the options market does not allow Gamma or Vega hedging. Since hedging is therefore difficult to implement, the sellers of these contracts should require a premium that is consistent with the risks of the market that cannot be hedged and takes into consideration the optionality features provided by this structured product. In addition, market participants would like to know the behaviour of the contract for different parameters in the model and also compare it to a benchmark or another financial product with which they are more familiar and is easier to price. Such contracts are strips of European options, which as it will be discussed later, provide the lower bounds for the price of swing options.

Hence, in this chapter we focus on the pricing of swing contracts using the different reduced form models described in Chapter 4 and explore their sensitivities by comparing them against strips of European Option. To our knowledge, only Eydeland and Wolyniec (2003) have given an example on the dependence of the incremental premium relative to strips of European options with respect to convenience yield and volatility, however no detailed work has been done on the sensitivity analysis of swing options contracts, and their incremental premium relative to regular European Option strips with respect to other factors such as jumps and mean reversion. This analysis is both interesting and also complements the analysis that was conducted in the previous chapter. Therefore in this chapter we are going to investigate, the additional premium one has to pay in order to have the optionality of the times to exercise relative to European options. This analysis will give an insight and a more in depth understanding of what actually drives the incremental premium and how it is affected by the model's parameters. Thus, as discussed earlier, market practitioners can have a benchmark and get a better insight on the premia exchanged for the trading of such complicated structured products. In addition, since we are dealing with short-term contracts, we examine the swing option prices produced by the different short-term models proposed in this thesis, namely the spike and the mean-reverting, for different seasons in the year and give the lower and upper bounds in terms of American and strips of European Options.

Our findings are very intuitive and in summary we find that the incremental swing option premium is positively affected by increases in jump size volatility, jump mean and jump intensity. Furthermore, we see that in terms of spike mean reversion there is a cut-off point after which increasing the speed of mean reversion decreases the incremental premium which nevertheless always remains positive. Interestingly, we find that the incremental swing option premium decreases as diffusive volatility increases, as there are more spot exercise possibilities and thus the swing option becomes similar to a strip of European Options. Finally, another interesting point is the fact that the additional swing premium becomes larger the smaller the equilibrium level and the higher speed of mean reversion.

The remainder of this chapter is organised as follows. Section 2 describes the characteristics of a typical swing option. Section 3 presents the literature review on swing options. Section 4 provides an in-depth description of general swing contracts and shows how they can be priced using stochastic optimal control. Section 5 reviews the Least-Squares algorithm and describes its implementation for swing option pricing. Section 6 provides the results of the analysis for the incremental swing option premium, under the different models and examines the impact of increasing the swing rights and penalties. Finally section 7 concludes the analysis.

6.2 Description of Swing Contracts

A typical swing contract will contain an agreement to take delivery of a fixed volume of a commodity for a fixed price over a fixed period of time. The swing element is an option to take delivery of a further volume of the commodity on certain dates in the future. The maximum and minimum volumes available at each date are specified, and the purchaser will be given a number of exercise rights to exercise at his discretion on any of a number of dates chosen in the future. The basis of calculation of the strike price is agreed in advance, as is a penalty to be charged to the purchaser if he does not consume the minimum quantity or does not exercise a minimum number of rights that is agreed at the outset of the contract. The minimum volume may be expressed in percentage or units of the underlying commodity as a take or pay amount, and the penalty may also be set in percentage of the final swing date strike price.

Swing contracts are very flexible with regards to the contract specifications, and the structure of this product is such that it can have favourable terms for both power producers and suppliers. For example, their American style feature that allow for early exercise makes them a very effective tool to hedge against extreme short-lived price movements such as spikes. On the other hand, by prescribing the maximum number of up- and down-swings as well as the penalty function, swing contracts reduce the short-term uncertainty in power demand experienced by the generators.

In this chapter we consider a swing option with 4 swing rights, maturing in 14 days. The first swing opportunity starts a day after the contract is agreed and the last 14 days later. The option provides the holder with daily swing rights- i.e. only one option can be exercised at a particular day- where the notification should occur within a short time interval after the announcement of the following day's power price by Nord Pool[41]. The incremental amount will then be delivered on the following day. It will be assumed that the holder of the option has another forward contract with the obligation to take one unit of power per day, where the price is fixed to the forward price for the same delivery period as that of the swing option. For simplicity, the unit will be set to 1 MWh[42]. Therefore, the swing contract will give the holder the right to double the

[41] Since Nord Pool is a day ahead market as explained in Chapter 4.
[42] As the problem scales with the amount, every amount can be used here.

amount four times during the options life within the two weeks of interest where the price is the same as in the forward contract.

The reason for specifying such a contract with small time steps between the swing opportunities, rather than having larger time steps, e.g. monthly, is due to the fact that in electricity hedgers are mainly concerned about daily variations in prices and volumes caused primarily by the spikes. In other markets, such as natural gas or coal, swing opportunities have time steps of a month (like the ones examined by Jaillet et al 2004) and hedgers are mainly concerned about mid- to longer-term variations in prices. Of course a similar contract can also be traded in electricity, since they are tailor made, but discussion with market practitioners in Nord Pool showed that they are mainly interested in shorter time contracts, as they provide a better fit to their requirement in the market. Nevertheless, the qualitative properties of the results would have been the same if we had chosen larger time steps.

6.3 Literature Review on Swing Contracts

There is a wealth of literature in swing option pricing. As is the case for path-dependent derivatives with American optionality inherent to them, the main methods of pricing are divided into trees/mesh methods and Monte Carlo integration.

Starting from the mesh methods, Lari-Lavassani et al (2001) consider a general swing contract in the gas market, entitling the owner to a number of up (purchase gas) and down (sell gas) swings that can be exercised on certain discrete dates. A penalty is levied to the owner of the contract, if he does not make net purchases within a specified volume over the life of the contract. The pricing of the contract is done using a discrete binomial forest method and considers both one and two factor models. They then use the one factor model to estimate the hedge parameters. They further suggest that the swing contract can be considered as a combination of Bermudan put and call options, corresponding to the number of up and down swing rights. They further prove convergence between the discrete binomial forest model and a partial differential equation model based on Black and Scholes (1973) which is used to value the swing contract as a series of European options.

Jaillet et al (2004) introduce a trinomial forest pricing method for swing options; using a one-factor lognormal mean-reverting forward curve model, the trees are constructed along the lines of Hull and White (1994) and also incorporate seasonality in the underlying process. They find that the swing option is exercised in a "bang-bang" strategy; that is to say that whenever the expected spot price deviates from its equilibrium level, the option is exercised, but this is mainly the case when no penalties are imposed.

The price behaviour of swing contracts is also investigated by Clewlow, Strickland and Kaminiski (2001). They propose that a swing contract behaves as a string of European options or vanilla swap under certain circumstances. More specifically, they find that when a contract is close enough to maturity and the penalty is sufficiently high then the maximum local volume should be purchased at every possible swing date in order to avoid a penalty if the global minimum has not been reached by expiration, hence the contract is similar to a vanilla swap[43]. Similarly, when the contract is sufficiently in the money, or the penalty is sufficiently low, the contract is similar to a strip of vanilla options. They also note that there exists a critical spot price, which changes over time, above which the swing contract holder will exercise if further volume needs to be purchased to avoid a penalty. This critical spot price may be below the strike price if the penalty is sufficiently punitive. The analysis is performed for a swing contract where the number of swing rights is equal to the number of possible swing dates and the pricing is performed using a trinomial forest as in Jaillet et al (2004).

Turning now to the Monte Carlo methods, in recent years stochastic mesh-methods and regression-based methods have been developed for the pricing of American options. These methods are suitable for problems with higher dimensions (such as having more than two state variables) and can handle stochastic parameters as well as jumps. In particular the method by Longstaff and Schwartz (2001) estimates the optimal exercise boundary problem by performing a set of least squares regressions on the continuation value, and thus early exercise takes place when the intrinsic value is higher than the estimated value function. The method's application to swing options has been performed by Dörr (2003) to value swing options using a simple mean-reverting model for electricity prices. He then monitors the expected payoffs of Swing options for different (sub-optimal) exercise strategies and compares with the Least Squares optimal exercise strategy.

Ibanez and Zapatero (2004) propose a Monte Carlo simulation method for pricing Bermudan options based on the computation of the optimal exercise frontier. The optimal exercise frontier for each exercise date is the price of the underlying or the locus of points at which the exercise (intrinsic) value equals the non-exercise (live) value. Ibanez (2004) extends Ibanez and Zapatero (2004) Monte Carlo algorithm to price swing options. He calculates an optimal exercise frontier for each possible number of remaining swing rights. A recursive algorithm is used to calculate each point on the optimal frontier at each swing date, which uses a forward tree to calculate the intrinsic value. He then simulates paths of the forward curve (using again a one-

[43] By local minimum or maximum, we referrer to the minimum or maximum volume that can be purchased at each swing date. By global minimum or maximum we refer to the total maximum or minimum volume that can be purchased during the whole lifetime of the swing option. In their paper, Clewlow et al (2001) impose a penalty only for a global minimum.

factor mean-reverting or log-normal model) and decides at each swing date whether to exercise or not. Exercise occurs if the spot price is above the optimal frontier at the swing date. Exercise occurs at each date dependent upon the optimal frontier because the expected path of the forward would not provide greater financial profit from exercise at later swing dates. The simulation is repeated a number of times and the mean of the sum of the discounted payoffs in each path, gives the contract's value. The forward path increments are the time steps between swing dates in this model, in contrast to the many time steps in the model of Lari-Lavassani et al (2001), and in line with Jaillet et al (2004) and Clewlow et al (2001).

Finally Keppo (2004) considers the cases of one-swing right per contract and full swings (where the number of swing rights matches the number of swing dates), but not the case where the number of swing rights is smaller than the possible swing dates. He assumes a complete market for the electricity, which is not a robust assumption, and no arbitrage so that there is a selection of European options and forwards with the same strike prices as the swing contract to replicate swing strategies. He suggests that the one-swing contract can be replicated using an option purchased for the maturity at which the swing contract holder is expected to exercise their right. Keppo (2004) recognises that in his one-swing case the hedging strategy will be dynamic, depending on which date is expected to be the optimal exercise date. He provides a lower bound for the swing option valuation by considering the value of the option/forward needed to purchase the minimum required amount. Thus his strategies rely on a bang-bang exercise fashion.

Following the discussion on the literature review, we can see that no analysis has been done to see how the presence of jumps in the market affects the swing option premium. However if someone uses a jump diffusion model, hedging becomes almost impossible especially when the jump size distribution is continuous and thus a continuum set of other derivatives is required to hedge the jump risk, as discussed in the previous chapter. We also note that due to the American-type features of swing options, no closed form solutions are available thus one has to rely on numerical methods. Finite difference schemes and trees, as discussed in the previous chapter, are well suited for low dimensional problems and standard dynamics, which do not incorporate jumps. In case when the state variables of the underlying asset of the derivative contract are more than two, finite difference schemes and trees become impractical as the number of branches explodes. As proposed in the literature, to overcome this problem Monte Carlo techniques are the best alternative (see e.g. Ibanez and Zapatero 2004). The two methods used in the literature are the ones based on either estimating the value function based on the Least Squares Monte Carlo method by Longstaff and Schwartz (2001) or by the estimation of the optimal exercise frontier proposed by Ibanez and Zapatero (2004). As reported by Ibanez and Zapatero (2004), both methods provide low-biased estimators, however the advantage of Ibanez and Zapatero (2004) is that it provides the holder of the option with an optimal exercise rule. On the other hand as reported on their paper, their method is much more computationally intensive. Since the main focus of this paper is on the sensitivities of swing options to changes in the parameters that affect their value and we thus need the most accurate and efficient algorithm (since sensitivities involves hundreds of prices to be estimated), we will use the LSM method by Longstaff and Schwartz (2001). Moreover, as shown by Dörr (2003), one can also identify the exercise boundary when using the LSM algorithm. Finally another contribution to the literature is the exploration of the incremental swing premium sensitivity with respect to strips of European options, which can serve as a benchmark to market participants and provides further intuition on how much they have to pay for the extra optionality.

6.4 Swing Options Pricing

6.4.1 Description

In this section, we will explain more in depth the characteristics of swing option contracts, along the lines of Jaillet et al (2004). Swing options usually are base-load forward contracts with embedded optionality. A base-load contract is the obligation to deliver or take a certain (daily) amount of the underlying (e.g. MWh) commodity for a predetermined price (e.g. NOK/MWh). The swing part of the contract allows for flexibility to change the delivery amount that is agreed in the base-load forward contract, during its lifetime. There is a great variety of swing options traded in the energy markets, however they all share some common characteristics.

In the current setting let 0 denote the time when the contract is written, T_1 and T_2 the first and the last day of the exercise period of the swing contract that in general coincides with the first and the last day of the delivery period in the forward contract, N is the maximum number of rights the holder is allowed to exercise, and $\{T_1 \leq \tau_1, \ldots, \tau_n \leq T_2\}$ are the discrete sets of dates at which the exercise of one right out of the N is allowed for each τ_j, where $N \leq n$. Note that if N exercise rights were allowed on a given date, then the contract would be equivalent to N Bermudan options. Generally if a right is exercised on a given date, there is also a refraction time Δt_R, which is the time lay before the next exercise can take place, and in most cases equals to $(\tau_{j+1} - \tau_j)$.[44]

The swing contract may provide the holder with the right to decide whether to, for example, take a positive incremental amount $v^+(\tau_j)$ on top of the volume specified in the base-load contract and thus receive an increased amount of the underlying commodity, or a negative incremental amount and thus deliver a volume of $v^-(\tau_j)$ of the underlying commodity, on a specific exercise date τ_j, on which a swing opportunity is allowed. Hence for a specific exercise date τ_j, the exercise decision can be summarised by the following variable:

[44] The notation largely follows Jaillet et al (2004).

$$\chi^{+}(\tau_j)=\begin{cases}1 \; \textit{if the swing holder exercises} \\ \quad \textit{for more volume on date } \tau_j \\ 0 \; \textit{otherwise}\end{cases}$$
$$\chi^{-}(\tau_j)=\begin{cases}1 \; \textit{if the swing holder exercises} \\ \quad \textit{for less volume on date } \tau_j \\ 0 \; \textit{otherwise}\end{cases} \tag{6.1}$$

Since by logic, a simultaneous exercise for additional or less volume does not make sense and no more than N exercises rights are allowed, the following constraints are associated with the contract:

$$0 \le \chi^{+}(\tau_j)+\chi^{-}(\tau_j) \le 1 \text{ for all } 1 \le j \le n$$
$$0 \le \sum_{i=j}^{n}\left(\chi^{+}(\tau_j)+\chi^{-}(\tau_j)\right) \le N \tag{6.2}$$

In addition we define the total number of times the holder is allowed to consume extra volume by N^{+}, and the total number of times he is allowed to decrease the base-load volume by N^{-}, therefore N^{+} might be different from N^{-}. Thus, we have:

$$0 \le \sum_{i=j}^{n}\left(\chi^{+}(\tau_j)\right) \le N^{+} \text{ and } 0 \le \sum_{i=j}^{n}\left(\chi^{-}(\tau_j)\right) \le N^{-} \text{ for all } 1 \le j \le n \tag{6.3}$$

Moreover, there are also constraints on the swing volume. As stated in Jaillet et al (2004) there are in general two main kinds of volume specifications in swing contracts that are categorised according to the duration of the exercise decision. One category is given by *local effects*, where the exercise of a right modifies the delivery volume only on the date of exercise and thus the delivery reverts to the level specified in the base-load contract thereafter. In these kind of contracts, the incremental volumes $v^{+}(\tau_j)$ and $v^{-}(\tau_j)$ are bounded from above (by $v^{+}_{\max}$ and $v^{-}_{\max}$) and below (by $v^{+}_{\min}$ and $v^{-}_{\min}$) due to physical constraints in the market (e.g. there is a maximum MWh that a transmission cable in a grid can carry), thus for each exercise date τ_j, $1 \le j \le n$:

$$\chi_j^- v_{\min}^- \le v^-(\tau_j) \le \chi_j^- v_{\max}^-$$
$$\chi_j^+ v_{\min}^+ \le v^+(\tau_j) \le \chi_j^+ v_{\max}^+ \tag{6.4}$$

Therefore if the volume agreed in the bases-load contract is represented by $V_{base\text{-}load}$, then the total daily volume delivered to the holder of a swing option is bounded from above and below as follows:

$$V_{base-load} - v_{\max}^- \le V_{base-load} - v_{\min}^- \le V_{base-load} \le V_{base-load} + v_{\min}^+ \le V_{base-load} + v_{\max}^+ \tag{6.5}$$

The second type of specification volume is given by the *global effect,* since in this category the exercise right influences all future volumes and thus the volume remains at the new level until the next swing right is exercised. Thus the overall "swing" volume until the last exercise date, τ_n is given by[45]:

$$V_{\tau_n} = \sum_{j=1}^{n} \left(\chi^+(\tau_j) v^+(\tau_j) - \chi^-(\tau_j) v^-(\tau_j) \right) \tag{6.6}$$

In the current analysis we will use the first category, i.e. the *local effects.* The price of the underlying for which the holder of the option can perform up- (take extra volume) and down-swings (deliver volume), so to speak the strike price K, may depend on time, meaning that strike prices can be a function of forward prices and do not necessarily have to be the same for both up and down swings. Instead, the strikes can contain for instance bid-ask spreads. For simplicity in the current setting we assume that they are both equal.

Another common restriction in swing contracts is the agreement of a penalty payment $\varphi(V)$, that is due when the total cumulative quantity $V = V_{\tau_n} + nV_{base-load}$ taken during the contract's term is below a minimum value V_{min} or above a maximum value V_{max} respectively. Note that the total volume delivered to the holder during the whole delivery period, is the sum of the total base-load volume delivered from the forward contract and the total "swing" volume. Most of the time the penalty is paid at the maturity of the contract T_2. There are several ways to define the penalty amount as a function of the quantity taken. For example (as in Jaillet et al 2004), the penalty can be constructed to be a constant amount C NOK, if the lower limit is more than the

[45] "Swing" volume refers to the total excess volume that is consumed due to the optionality of the contract.

total consumption V, and a per unit penalty of P_{T_2} if the total consumption is above the volume allowed:

$$\varphi(V)=\begin{cases} C & if\ V<V_{\min} \\ 0 & if\ V^{-}<V<V_{\max} \\ (V-V_{\max})P_{T_2} & if\ V_{\max}<V \end{cases} \tag{6.7}$$

In another situation the upper and lower limits may be dictated by technical constraints which cannot be violated. Thus we have:

$$\varphi(V)=\begin{cases} \infty & if\ V<V_{\min} \\ 0 & if\ V^{-}<V<V_{\max} \\ \infty & if\ V_{\max}<V \end{cases} \tag{6.8}$$

The penalty functions that will be used in this chapter belong to the first category.

6.4.2 Upper and Lower Bounds of Swing Options

As shown by Jaillet et al (2004) when there are no penalties the following boundary conditions exist for swing options:

1. A swing option with one exercise right (*N=1*) is equivalent to a Bermudan option with the same strike price, *K*, and exercise dates, $\{T_1 \le \tau_1,....,\tau_n \le T_2\}$ given the fact that a Bermudan option can be exercised (like swing options) at discrete points in time.

2. A swing Option with one up- ($N^{+}=1$) and one down-swing ($N^{-}=1$) opportunity is equivalent to the sum of one Bermudan put and one Bermudan call with the same strike price, *K*, and exercise dates, $\{T_1 \le \tau_1,....,\tau_n \le T_2\}$.

3. A swing option with exercise rights equal to the exercise opportunities (i.e. *N=n*) is the same as a strip of European Options, with maturities equal to

the dates of the swing opportunities. Therefore if we have $N=n$ and only up-swings (down-swings) are allowed, this is equivalent to Caps (Floors).

4. For a swing option with N^+ upswing rights and N^- downswing rights, the upper bound equals to the sum of N^+ Bermudan call options and N^- Bermudan put options. Because the Bermudan options can be exercised altogether at once when the optimal exercise condition is satisfied, contrary to the corresponding swing option which allows only to exercise a single swing per exercise opportunity, the value of the Bermudan option bundle is always greater than or at least equal to the swing option's value.

5. For a swing option with N^+ up-swing rights and N^- down-swing rights, the lower bound equals to the sum of the N^+ most expensive European calls and the N^- most expensive European puts among the strip of European options covering all exercise opportunities. This is based on a set of the best exercise dates seen at the valuation time. Thus, the exercise dates are already chosen at the valuation time. On the other hand, the swing options allow the holder to postpone the exercise decision to another swing opportunity. Consequently, the value of the European option collection is always less than or at most equal to the value of the swing option.

6. When there are not any penalties for overall consumption, the highest or the lowest level allowed by the local constraints should be taken or delivered when a swing opportunity is exercised on a given date.

6.4.3 Intuitive Valuation of Swing Options

In this section we provide a valuation example of a swing option with the following terms: First of all let's assume the amount of commodity that can be delivered at an exercised swing date is constant. That is to say that, $v^+(\tau_j) \equiv v^+$ and $v^-(\tau_j) \equiv v^-$, thus the total "swing" volume allowed is given by $V_{\tau_n} = v^+ N^+ - v^- N^-$. Let also $u(\tau_j)$ and $d(\tau_j)$, denote the number of up- and down-swings performed up to an exercise date τ_j, including the exercise decision taken at τ_j. Thus $u(\tau_n) \leq N^+$ and $d(\tau_n) \leq N^-$, but also

$u(\tau_{j-1}) \leq u(\tau_j)$ and $d(\tau_{j-1}) \leq d(\tau_j)$. Moreover we assume that the exercise price is constant, K, and is the same for both up- and down-swings.

The value of the swing option for the underlying price P_{τ_j} at the swing opportunity date τ_j, with $u(\tau_j)$ up-swings and $d(\tau_j)$, down-swings exercised, is denoted by ${}_{swing}V_{\tau_j}^{u(\tau_j),d(\tau_j)}$. Having in mind that it is only allowed to exercise a single swing right at every opportunity.

Now, when the holder exercises an up-swing, which is only possible if the maximum number of up-swings has not been reached so far, he receives the incremental volume v^+ for a price of K and thus adds v^+ $(P_{\tau_j} - K)$ to his Profit&Loss. Moreover, the holder has to account for the fact that the swing contract is left with one up-swing right less for the future, i.e. the value now becomes ${}_{swing}V_{\tau_j}^{u(\tau_{j-1})+1,d(\tau_{j-1})}$. The same, of course applies to down-swings.

Given the above notation, and assuming that the holder tries to maximise the option value, we can say that the valuation of a swing option becomes a stochastic optimal control problem, since the aim of the holder is to find a suitable maximisation strategy. It is very similar to the valuation of American and Bermudan options but with more dimensions. If neither the maximum number of up-swings nor the number of down-swings has been reached, i.e. $u(\tau_n) < N^+$ and $d(\tau_n) < N^-$, the value of the option will be:

$${}_{swing}V_{\tau_j}^{u(\tau_j),d(\tau_j)} = \max\left({}_{swing}V_{\tau_j}^{u(\tau_{j-1})+1,d(\tau_{j-1})} + v^+\left(P_{\tau_j} - K\right),\ {}_{swing}V_{\tau_j}^{u(\tau_{j-1}),d(\tau_{j-1})+1} + v^-\left(K - P_{\tau_j}\right), {}_{swing}V_{\tau_j}^{u(\tau_{j-1}),d(\tau_{j-1})} \right) \quad (6.9)$$

Thus the value of the swing option is the maximum out of three strategies: The first strategy is to exercise an up-swing, receive a profit of v^+ $(P_{\tau_j} - K)$ and thus remain with one upswing less from what we had after the previous exercise decision was taken at τ_{j-1}. The second strategy is to exercise a downswing, receive a profit of v^- $(K-P_{\tau_j})$ and thus remain with one down-swing less from what we had after the previous exercise decision was taken at τ_{j-1}. And the final strategy is to perform no exercise at all and remain with the same rights we had after the exercise decision at τ_{j-1}.

On the other hand if the maximum number of up-swings, N^+, has been reached, the value of the swing becomes:

$$ {}_{swing}V_{\tau_j}^{N^+,d(\tau_j)} = \max\left({}_{swing}V_{\tau_j}^{N^+,d(\tau_{j-1})+1} + v^-\left(K - P_{\tau_j}\right),\ {}_{swing}V_{\tau_j}^{N^+,d(\tau_{j-1})} \right) \tag{6.10} $$

whereas if the maximum number of down-swings, N^-, has been reached, the value of the swing becomes:

$$ {}_{swing}V_{\tau_j}^{u(\tau_j),N^-} = \max\left({}_{swing}V_{\tau_j}^{u(\tau_{j-1})+1,N^-} + v^+\left(P_{\tau_j} - K\right),\ {}_{swing}V_{\tau_j}^{u(\tau_{j-1}),N^-} \right) \tag{6.11} $$

The above equations (6.9) - (6.11) are valid only if $j<n$, i.e. before the last swing opportunity when there is still time to exercise more rights. At the last exercise opportunity penalty payments are due if any of the restrictions are violated. That is to say that if the penalty is a function of the total volume consumed, e.g. $\phi^{u(\tau_n),d(\tau_n)} = \phi\left(u(\tau_n)v^+ - d(\tau_n)v^-\right)$, then the value of the swing option at the last exercise opportunity date is as follows:

$$ {}_{swing}V_{\tau_n}^{u(\tau_n),d(\tau_n)} = \max\left(v^+\left(P_{\tau_n} - K\right) - \phi^{u(\tau_{n-1})+1,d(\tau_{n-1})},\ v^-\left(K - P_{\tau_n}\right) - \phi^{u(\tau_{n-1}),d(\tau_{n-1})+1},\ -\phi^{u(\tau_{n-1}),d(\tau_{n-1})}\right) \tag{6.12} $$

$$ {}_{swing}V_{\tau_n}^{N^+,d(\tau_n)} = \max\left(v^-\left(K - P_{\tau_n}\right) - \phi^{N^+,d(\tau_{n-1})+1},\ -\phi^{N^+,d(\tau_{n-1})} \right) \tag{6.13} $$

$$ {}_{swing}V_{\tau_n}^{u(\tau_n),N^-} = \max\left(v^+\left(P_{\tau_n} - K\right) - \phi^{u(\tau_{n-1})+1,N^-},\ -\phi^{u(\tau_{n-1}),N^-} \right) \tag{6.14} $$

However we have to note that the above solution is based on the perspective of profit maximisation, and thus assuming that there exists a liquid spot market from which the holder of the structured product can buy the underlying or sell it, when feasible. Of course there exist cases when the holder is obligated to exercise a right, not because it maximises his balance sheet but for demand needs, whenever the spot market is not liquid at all. For example if the contract specifies delivery at a specific location, which itself is not liquid market, it might not be possible to sell the unneeded volume to a liquid grid. Moreover, when the demand increases substantially and the spot market is illiquid at the current location, the utility might exercise a swing right even though it might not be the best strategy. However whenever the exercise strategy is driven solely by the commodity demand the customer faces, we are no longer dealing with an option anymore, but with a *load-serving contract.* In its basic form a *load-serving contract* involves serving a load at a given location under the assumption that the supplier controls the physical dispatch rights of the load, ensuring that the amount of load to serve is not directly determined by the exercise policy of the demand owner. In this way we are making sure that we are serving the actual physical load and not the net load after the owner of the load has exhausted cheaper sources of delivery. Thus for example the owner of the load in some circumstances might be obligated to call on delivery of power when the true demand is low and the price is very high. The load-serving contracts are not considered here. For more on this look at Eydeland and Wolyniec (2003).

6.5 Least Square Monte Carlo for Swing Options pricing

6.5.1 Monte Carlo Simulation for American Options valuation

As discussed in the previous chapter, Monte–Carlo methods are very popular in practical finance, since they are – in general – easy to implement and allow the treatment of problems with high dimensionality. In particular, when there are multiple stochastic factors, or the process involves jumps, numerical methods like finite–differences or binomial techniques become impractical while Monte Carlo is still appropriate.

In the case of electricity derivatives, appropriate stochastic processes for the underlying are in general considerably complicated, since they have to account for seasonality, mean reversion, spikes etc. Calibration of a model for the energy price to real market data usually requires at least two stochastic factors, frequently in conjunction with jump diffusion. Therefore Monte–Carlo methods may be the best choice.

However, the treatment of early exercise features is a great challenge for Monte–Carlo methods. For American and Bermudan options several approaches to this problem have been discussed in the literature section. As we explained among the most popular methods is the one by Ibanez and Zapatero (2004), which focuses on the computation of the optimal exercise frontier, which is the locus of points where the exercise value matches the continuation value. Another direction in the literature, focuses directly on the conditional expectation function involved in the iterations of dynamic programming and used least squares regression to estimate the conditional expectations. These conditional expectations, i.e. the continuation values under the assumption that the option is not exercised at a particular opportunity (iteration step) are estimated by least squares regression. One example is the Least Squares Monte Carlo (LSM) method, an algorithm proposed by Longstaff and Schwartz (2001), which seems to have become more and more popular among practitioners. In the current setting we will explain how to use the LSM algorithm to value Swing Options, as it was presented by Dörr (2003); the alternative algorithm introduced by Ibanez (2004) provides the same degrees of low-biasness in the pricing but at the same time is more complex to develop and much more computationally intensive as discussed in the section of the literature review, therefore it will take much longer to estimate the hundreds of swing option prices that need to be estimated for our analysis.

The basic idea of the Longstaff and Schwartz algorithm, described in detail in Clement et al (2002), is to use least squares regression on a finite set of functions as a proxy for conditional expectation estimates. In a first step, the time axis has to be discretized, i.e. if the American option is alive within the time horizon $[T_1, T_2]$ early exercise is only allowed at discrete times $\{T_1 = \tau_1, ..., \tau_n = T_2\}$. The American option is thus approximated by a Bermudan option.

For a particular exercise date τ_k, early exercise is performed if the payoff from immediate exercise exceeds the continuation value, i.e. the value of the (remaining) option if it is not exercised at τ_k. This continuation value can be expressed as the conditional expectation of the option payoff with respect to the risk neutral pricing measure. The expectation is taken conditional on the information set $\mathcal{F}_{\tau_k}$ which is available at τ_k. Representing the continuation value by $F(\omega;\tau_k)$ we can write:

$$F(\omega;\tau_k) = E^*\left(\sum_{j=k+1}^{n} e^{-r(\tau_j-\tau_k)} C\left(\omega,\tau_j;\tau_k,T_2\right) / \mathcal{F}_{\tau_k}\right)$$

where r is the continuous risk-free rate of interest and $C(\omega, \tau_j; \tau_k, T_2)$ denotes the remaining cash flows, conditional on the option not being exercised at or prior to time τ_k and the holder following the optimal exercise strategy for all remaining opportunities τ_j between τ_{k+1} and $\tau_n=T_2$. Note that for each path ω there may be at most one j with $C(\omega, \tau_j; \tau_k, T_2) > 0$, since the Bermudan option has only one exercise right, otherwise there is no exercise and $C(\omega\ \tau_j; \tau_k, T_2) = 0$.

Starting at $\tau_n=T_2$ we calculate the pay-off at expiry and then work backwards through time. The early exercise decision at time τ_{n-1} is made by comparing $F(\omega; \tau_{n-1})$ with the immediate payoff $Pay\left(P_{\tau_{n-1},\omega}\right)$, e.g. in the case of a call option $Pay\left(P_{\tau_{n-1},\omega}\right) = \max\left(P_{\tau_{n-1},\omega} - K, 0\right)$, where $P_{\tau_{n-1},\omega}$ is the value of the underlying at time

τ_{n-1} in path ω. While $Pay\left(P_{\tau_{n-1},\omega}\right)$ is known, the functional form of $F(\omega;\ \tau_{n-1})$ is unknown, but can be represented as a linear combination of a countable set of $\mathcal{F}_{\tau_{n-1}}$ - measurable *basis functions*. A *basis* is a set of vectors in a linear combination of which can represent every vector in a given vector space, and such that no element of the set can be represented as a linear combination of the others. A *basis function* is an element of the basis for a fuction space, meaning that each function in the function space can be represented as a linear combination of the basis functions.Using this definition, Longstaff and Schwartz (2001) show that the continuation value can be represented by a set of basis functions B_z as:

$$F(\omega;\tau_{n-1}) = \sum_{z=0}^{\infty} a_z(\tau_{n-1}) B_z\left(P_{\tau_{n-1}}\right) \tag{6.15}$$

Where $B_z\left(P_{\tau_{n-1}}\right)$ are the basis functions and $a_z(\tau_{n-1})$are the corresponding coefficients. Since we consider a Markov process for the state variable P, only current realizations of it can be included in the basis functions B_z. For practical purposes $F(\omega;\ \tau_{n-1})$ has to be approximated by using the first $L < \infty$ basis functions, and denote this approximation by $F_L(\omega;\tau_{n-1})$. For example, Longstaff and Scewartz (2001) use the first three polynomials i.e. $P^0_{\tau_{n-1}}, P^1_{\tau_{n-1}}, P^2_{\tau_{n-1}}, P^3_{\tau_{n-1}}$, as basis functions.

At this stage the crucial step is to estimate $F_L(\omega;\tau_{n-1})$ by regressing the discounted values of $C(\omega,\ \tau_j;\ \tau_{n-1},\ T_2)$, i.e. the cash flows, which occur at τ_n, onto the basis functions. Since the early exercise decision is only relevant for those paths where the option is in the money at τ_{n-1} the regression is restricted to these paths. In this way we limit the region over which the continuation value must be estimated and thus far fewer basis functions are needed to obtain an accurate approximation of the continuation value. The fitted values from this regression $\hat{F}_L(\omega;\tau_{n-1})$, are an (unbiased) estimator of $F_L(\omega;\tau_{n-1})$as shown by Longstaff and Schwartz (2001). Now the exercise decision is made by comparing $Pay\left(P_{\tau_{n-1},\omega}\right)$ with $\hat{F}_L(\omega;\tau_{n-1})$, and thus determining whether early exercise is optimal for an In-the-Money path ω at τ_{n-1}.

Once the exercise decision has been identified, the option cash-flow paths $C(\omega, \tau_j; \tau_{n-2}, T_2)$, can then be approximated. The recursion proceeds by rolling back to time τ_{n-2}, and repeating the procedure until the exercise decisions at each exercise time along each path have been determined. The American option then is priced by starting at time zero moving forward at each path ω until the first stopping time occurs and discounting back the cash flow to time zero. At the end the result obtained by Least–Squares Monte Carlo, V_{LSM}, is the average over the discounted cash-flows from each path (note that there is at most one cash flow per path)

$$V_{LSM} = \frac{1}{m}\sum_{i=1}^{m} C_{LSM}(\omega_i) \tag{6.16}$$

where $C_{LSM}(\omega_i)$ denotes the discounted cash flows, which result from following the LSM strategy described above.

Since the LSM method represents one particular *strategy* the "real" option value V (which represents the optimal strategy) must be greater than or equal to V_{LSM}. The convergence of V_{LSM} to V, i.e. the proposition that for any $\epsilon > 0$ there exists an $L < \infty$, is such that

$$\lim_{m\to\infty} \Pr\left[\left|V - V_{LSM}\right| > \in\right] = 0 \tag{6.17}$$

for which the proof is provided by Longstaff and Schwartz (2001).

6.5.2 Extension of LSM for Swing Options

For the valuation of Swing options the basic concepts of Least–Squares Monte Carlo can be directly adopted, as it is shown by Dörr (2003). Since we now have more than one exercise right, however, we have to deal with an additional "dimension", i.e. the number of exercises left. As discussed at the beginning of the chapter, in this chapter we are considering swing options with 4 up-swing rights and 14 opportunity dates. The reason why we choose up-swings, is due to the fact that we are examining the case of a utility company that wants to hedge against adverse price movements that would affect its highly volatile costs. At the end of the section, we also look at down-swings and compare the different models proposed in this thesis. Thus in this section we will describe the methodology for pricing these swing options using Monte Carlo simulation.

We consider a Swing option with exercise opportunities at times τ_1 to τ_{14}, i.e. with 14 exercise opportunities. There are 4 up-swings, and the strike price at each opportunity is K. Sampling m paths yields the $m \times 14$-spot price matrix P. The main difficulties arising from the presence of more than one exercise rights are the following:

- The benefit from immediate exercise is not only the payoff, but the payoff plus the value of the remaining Swing option, which has one up-swing less than the original one as stated earlier;
- When early exercise is performed at τ_j, rearranging the cash flows at later opportunities requires the cash flow matrix of the Swing option with one up-swing less than the original one.

The generalized cash flow matrix of our algorithm must therefore have three dimensions:

1. Number of paths (m)
2. Number of time-steps (exercise opportunities=14)
3. Number of exercise rights (upswings) left

We thus consider the following steps in the implementation of the LSM algorithm for swing options:

STEP 1:

We denote the cash flow matrix for u upswings left as C^u. In our example, there are thus four $m \times 14$ matrices C^1, C^2, C^3 and C^4. At the initial step, we can calculate at the expiration date cash flows, and up to $14-(u-1)$ time steps before for each one of the matrices C^u, since only one exercise right is allowed at each time step. For example if we have 4- upswing left and there are 4 exercise rights remaining, and we are at time τ_{11}, we may exercise one right at each time-step i.e. τ_{11}, τ_{12}, τ_{13} and τ_{14} without having to look at the continuation value. Thus the contract becomes equivalent to a portfolio of 4 European call options with expiry dates τ_{11}, τ_{12}, τ_{13} and τ_{14} respectively. Thus after the initial step in the procedure, the cash flow matrices for each C^u look like the following:

$$C^4=\begin{pmatrix} \% & \% & \% & \cdots & Pay(P_{11,1}) & Pay(P_{12,1}) & Pay(P_{13,1}) & Pay(P_{14,1}) \\ \% & \% & \% & \cdots & Pay(P_{11,2}) & Pay(P_{12,2}) & Pay(P_{13,2}) & Pay(P_{14,2}) \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \vdots \\ \% & \% & \% & \cdots & Pay(P_{11,M}) & Pay(P_{12,M}) & Pay(P_{13,M}) & Pay(P_{14,M}) \end{pmatrix}$$

$$C^3=\begin{pmatrix} \% & \% & \% & \cdots & \% & Pay(P_{12,1}) & Pay(P_{13,1}) & Pay(P_{14,1}) \\ \% & \% & \% & \cdots & \% & Pay(P_{12,2}) & Pay(P_{13,2}) & Pay(P_{14,2}) \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \vdots \\ \% & \% & \% & \cdots & \% & Pay(P_{12,M}) & Pay(P_{13,M}) & Pay(P_{14,M}) \end{pmatrix}$$

$$C^2=\begin{pmatrix} \% & \% & \% & \cdots & \% & \% & Pay(P_{13,1}) & Pay(P_{14,1}) \\ \% & \% & \% & \cdots & \% & \% & Pay(P_{13,2}) & Pay(P_{14,2}) \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \vdots \\ \% & \% & \% & \cdots & \% & \% & Pay(P_{13,M}) & Pay(P_{14,M}) \end{pmatrix}$$

$$C^1=\begin{pmatrix} \% & \% & \% & \cdots & \% & \% & \% & Pay(P_{14,1}) \\ \% & \% & \% & \cdots & \% & \% & \% & Pay(P_{14,2}) \\ \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \vdots \\ \% & \% & \% & \cdots & \% & \% & \% & Pay(P_{14,M}) \end{pmatrix}$$

where $Pay\left(P_{j,i}\right)=\max\left(P_{\tau_j,i}-K,0\right)$ is the pay-off of an up-swing at the simulated path i and time step τ_j. The $\%$ means that the cash flows are undefined at the corresponding stage.

For C^4 we can combine the last four time steps in the initial step of the algorithm since it is obvious that early exercise takes place at τ_{11} whenever the payoff at this time step is positive. Note that this is the fourth time step from the last. Similarly, when three up-swings are left immediate early exercise is performed at τ_{12} and thus we can combine the last three time steps for C^3. The matrix C^2 corresponds to two up-swings left and C^1 corresponds to the cash flow matrix in the original Longstaff and Schwartz (2001) algorithm for Bermudan options. Thus at this stage all pay-offs have a known functional form and are equivalent to European call options, where as at earlier time steps, the continuation value needs to be estimated. As an example, if we have eight paths, after the initial step the matrices might look as follows:

$$
\begin{aligned}
C^{4} &= \begin{pmatrix}
\% & \% & \% & \cdots & Pay_{11,1} & Pay_{12,1} & Pay_{13,1} & Pay_{14,1} \\
\% & \% & \% & \cdots & Pay_{11,2} & Pay_{12,,2} & Pay_{13,2} & 0 \\
\% & \% & \% & \cdots & Pay_{11,3} & Pay_{12,3} & Pay_{13,3} & Pay_{14,3} \\
\% & \% & \% & \cdots & Pay_{11,4} & Pay_{12,4} & 0 & Pay_{14,4} \\
\% & \% & \% & \cdots & Pay_{11,5} & Pay_{12,5} & Pay_{13,5} & Pay_{14,5} \\
\% & \% & \% & \cdots & Pay_{11,6} & Pay_{12,6} & Pay_{13,6} & 0 \\
\% & \% & \% & \cdots & Pay_{11,7} & Pay_{12,7} & Pay_{13,7} & Pay_{14,7} \\
\% & \% & \% & \cdots & Pay_{11,8} & Pay_{12,8} & 0 & Pay_{14,8}
\end{pmatrix} \\
C^{3} &= \begin{pmatrix}
\% & \% & \% & \cdots & \% & Pay_{12,1} & Pay_{13,1} & Pay_{14,1} \\
\% & \% & \% & \cdots & \% & Pay_{12,2} & Pay_{13,2} & 0 \\
\% & \% & \% & \cdots & \% & Pay_{12,3} & Pay_{13,3} & Pay_{14,3} \\
\% & \% & \% & \cdots & \% & Pay_{12,4} & 0 & Pay_{14,4} \\
\% & \% & \% & \cdots & \% & Pay_{12,5} & Pay_{13,5} & Pay_{14,5} \\
\% & \% & \% & \cdots & \% & Pay_{12,6} & Pay_{13,6} & 0 \\
\% & \% & \% & \cdots & \% & Pay_{12,7} & Pay_{13,7} & Pay_{14,7} \\
\% & \% & \% & \cdots & \% & Pay_{12,8} & 0 & Pay_{14,8}
\end{pmatrix} \\
C^{2} &= \begin{pmatrix}
\% & \% & \% & \cdots & \% & \% & Pay_{13,1} & Pay_{14,1} \\
\% & \% & \% & \cdots & \% & \% & Pay_{13,2} & 0 \\
\% & \% & \% & \cdots & \% & \% & Pay_{13,3} & Pay_{14,3} \\
\% & \% & \% & \cdots & \% & \% & 0 & Pay_{14,4} \\
\% & \% & \% & \cdots & \% & \% & Pay_{13,5} & Pay_{14,5} \\
\% & \% & \% & \cdots & \% & \% & Pay_{13,6} & 0 \\
\% & \% & \% & \cdots & \% & \% & Pay_{13,7} & Pay_{14,7} \\
\% & \% & \% & \cdots & \% & \% & 0 & Pay_{14,8}
\end{pmatrix} \\
C^{1} &= \begin{pmatrix}
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,1} \\
\% & \% & \% & \cdots & \% & \% & \% & 0 \\
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,3} \\
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,4} \\
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,5} \\
\% & \% & \% & \cdots & \% & \% & \% & 0 \\
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,7} \\
\% & \% & \% & \cdots & \% & \% & \% & Pay_{14,8}
\end{pmatrix}
\end{aligned}
\tag{6.18}
$$

Where $Pay_{j,i}$ stands for a positive pay-off (i.e. $P>K$), and *0* means that the option is Out-of-the-Money at that specific point in the path.

STEP 2:

We now start stepping backwards in time. Let C_j^u denote the column vector j of matrix C^u, i.e. all cash flows at period τ_j. Thus C_{14}^1 denotes the cash flows at τ_{14} of Matrix C^1. On the hand $\hat{C}_{14}^1$ contains the elements of C_{14}^1 that belong to the paths that are In-The-Money at τ_{13} discounted back to τ_{13}. For τ_{13} we calculate the continuation values for one upswing left by least squares regression of the cash flow vector $\hat{C}_{14}^1$, onto the basis functions. As suggested by Longstaff and Schwartz (2001) the LSM method is very robust against the choice of basis functions, so the choice of a constant, $P_{\tau_j}^1, P_{\tau_j}^2$ and $P_{\tau_j}^3$ as basis functions should be appropriate. By considering only paths that are In-the-Money, we limit the region over which the conditional expectation must be estimated, and far fewer basis functions are needed to obtain an accurate approximation to the conditional expectation function as shown in Longstaff and Schwartz (2001). The continuation value $Cont_{13}^1(i)$ for a specific path i at time-step 13 with one up-swing left, is then estimated using the estimates of the coefficients of the basis function $a_z(\tau_{n-1})$ from the regression, as follows:

$$Cont_{13}^1(i) = \hat{a}_0(\tau_{13}) + \hat{a}_1(\tau_{13}) P_{\tau_{13,i}} + \hat{a}_2(\tau_{13}) P_{\tau_{13,i}}^2 + \hat{a}_3(\tau_{13}) P_{\tau_{13,i}}^3$$

Thus early exercise takes place for each element, i, that satisfies the following condition at τ_{13}:

$$Pay_{13,i} > Cont_{13}^1(i) \qquad (6.19)$$

Having determined where early exercise is optimal for the Bermudan option at τ_{13}, by satisfying condition (6.19), the cash flows for column 13 at matrix C^1 are now determined and the matrix may look like as follows:

$$C^1 = \begin{pmatrix} \% & \% & \% & \cdots & \% & \% & Pay_{13,1} & 0 \\ \% & \% & \% & \cdots & \% & \% & Pay_{13,2} & 0 \\ \% & \% & \% & \cdots & \% & \% & Pay_{13,3} & 0 \\ \% & \% & \% & \cdots & \% & \% & 0 & Pay_{14,4} \\ \% & \% & \% & \cdots & \% & \% & Pay_{13,5} & 0 \\ \% & \% & \% & \cdots & \% & \% & Pay_{13,6} & 0 \\ \% & \% & \% & \cdots & \% & \% & Pay_{13,7} & 0 \\ \% & \% & \% & \cdots & \% & \% & 0 & Pay_{14,8} \end{pmatrix} \quad (6.20)$$

Thus, as we can see from (6.20), we assume that early exercise occurs at τ_{13} for paths 1 to 3 and 5 to 7 in this example. Note however, that for paths 1, 3, 5 and 7 the option was In-the-Money at τ_{14}, however since we assumed that for these paths, the continuation value at τ_{13} is lower than the immediate exercise, the cash flow *(P-K)* is placed at τ_{13}, and the corresponding rows at column 14 have been replaced with zero, since this swing option is like a Bermudan and hence allows for one exercise right. At this stage, matrices 2 to 4 remain unchanged.

STEP 3:

Moving now to τ_{12}, we need to estimate the continuation value for the second matrix, $Cont_{12}^2$ (i.e. the value of a swing option with 2 exercise rights at τ_{12}, by not exercising early at τ_{12}), since early exercise possibilities might exist. First of all we add up the cash flows for C^2 for τ_{14} and τ_{13}, i.e. we add up columns 14 and 13 of C^2. This vector sum is denoted as C_{13+14}^2, but in this vector we omit all paths where $Pay_{12,i}$ is zero and the remaining are discounted back to τ_{12}, and thus we get $\hat{C}_{13+14}^2$. The continuation vector $Cont_{12}^2$ is then obtained by linear regression of $\hat{C}_{13+14}^2$ on the basis function, which comprise a constant, $P_{\tau_{12,i}}^1, P_{\tau_{12,i}}^2$ and $P_{\tau_{12,i}}^3$. Thus the early exercise condition now reads as follows

$$Pay_{12,i} + Cont_{12}^1(i) > Cont_{12}^2(i) \quad (6.21)$$

This means that we also have to estimate $Cont_{12}^1$ before early exercise is performed for C^2. In other words condition (6.21) implies that early exercise should be performed at those paths were the pay-off is positive and the value from immediate exercise and one swing right remaining (equivalent to a Bermudan Option), is greater than not exercising and staying with two exercise rights. For those paths where condition (6.21) is satisfied, early exercise takes place. This means that for each corresponding path i, $C^2_{12}(i)$ is set equal to the payoff $Pay(P_{i,12})$ and the cash flows $C^2_{13}(i)$ and $C^2_{14}(i)$ (i.e. the 13th and 14th column of C^2) are replaced by $C^1_{13}(i)$ and $C^1_{14}(i)$ (i.e. the 13th and 14th column of C^1 in (6.20)), respectively, reflecting the fact that we have 1 up-swing left after having exercised 1 right at τ_{12}. After that, early exercise for τ_{12} is performed in C^1, i.e. at the previous iteration step. While C^3 and C^4 still remains unchanged in this step the other cash flow matrices in our example might look like as follows:

$$C^2 = \begin{pmatrix} \% & \% & \% & \cdots & \% & Pay_{12,1} & Pay_{13,1} & 0 \\ \% & \% & \% & \cdots & \% & 0 & Pay_{13,2} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,3} & Pay_{13,3} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,4} & 0 & Pay_{14,4} \\ \% & \% & \% & \cdots & \% & Pay_{12,5} & Pay_{13,5} & 0 \\ \% & \% & \% & \cdots & \% & 0 & Pay_{13,6} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,7} & Pay_{13,7} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,8} & 0 & Pay_{14,8} \end{pmatrix}$$

$$C^1 = \begin{pmatrix} \% & \% & \% & \cdots & \% & 0 & Pay_{13,1} & 0 \\ \% & \% & \% & \cdots & \% & 0 & Pay_{13,2} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,3} & 0 & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,4} & 0 & 0 \\ \% & \% & \% & \cdots & \% & 0 & Pay_{13,5} & 0 \\ \% & \% & \% & \cdots & \% & 0 & Pay_{13,6} & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,7} & 0 & 0 \\ \% & \% & \% & \cdots & \% & Pay_{12,8} & 0 & 0 \end{pmatrix} \quad (6.22)$$

For C^2 early exercise at τ_{12} was performed in paths 1, 3, 4, 5, 7 and 8. Note that in paths 3 and 7, and the cash flows at τ_{13} and τ_{14} had to be modified according to C^1 at the previous iteration step (STEP 2), as in (6.20). After performing the early exercise procedure for C^2 at τ_{12}, we can perform the early exercise for C^1, which is something that follows from equation (6.21). Note that at τ_{12} in C^1 early exercise occurs only in paths 3 ,4,7 and 8.

STEP 4:
In the same fashion early exercise tests have to be performed for C^3 at τ_{11} and C^4 at τ_{10}. Thus at τ_{10} early exercise takes place if the following conditions are satisfied for those paths where the immediate payoff is positive for C^4, C^3, C^2 and C^1 respectively:

$$\begin{aligned}
&Pay_{i,10} + Cont_{10}^3(i) > Cont_{10}^4(i)\\
&Pay_{i,10} + Cont_{10}^2(i) > Cont_{10}^3(i)\\
&Pay_{i,10} + Cont_{10}^1(i) > Cont_{10}^2(i)\\
&Pay_{i,10} > Cont_{10}^1(i)
\end{aligned} \tag{6.23}$$

Repeating the same procedure until τ_1, we end up with the final cash flow matrices C^1, C^2, C^3 and C^4. From these matrices we obtain the value of the corresponding Swing options by taking the sums of the discounted cash flows in every row, since they correspond to the cash flows using the LSM strategy, and then taking the average of all the row sums of the discounted cash flows to find the expected value of the swing option.

The pricing procedure described in this section is implemented in the next section to estimate the incremental premium required for swing options and identify how its value is driven by the different parameters that determine the stochastic behaviour of the electricity price.

6.6 Swing Options Analysis and Results

As discussed in the introductory section of the paper, we consider a contract with 14 swing dates, at daily intervals, and four swing rights. The option provides the holder with daily swing rights, where the notification should occur within a short time interval after the announcement of the following day's power price by Nord Pool. The incremental amount will then be delivered on the following day, since the spot refers to the day-ahead market. The first swing opportunity will be the first day after the contract is negotiated and the last swing opportunity will be for two weeks later. It will be assumed that the holder of the option has another forward contract with the obligation to take one unit of power per day, where the price is fixed to the forward price. For simplicity, the underlying base unit will be set to 1 MWh. The swing contract will give the holder the right to double the amount of electricity received four times during the options life within the two weeks of interest at the forward price. Thus the contract specifications are shown in Table 6.1; in the current analysis no penalties are assumed, as they will be examined later in the section on the effect of penalties.

Table 6.1: Parameter values of swing contract specification used in the sensitivity analysis

Swing Contract specifications	
Up-swing ammount (v^+)	1 MWh
Down-swing ammount (v^-)	0 MWh
Number of up-swings allowed (N^+)	4
Number of down-swings allowed (N^-)	0
Refraction period (Δt_R)	1 day
Exercise Price (K)	170.20 NOK/MWh

As discussed at the beginning of the chapter, our focus is on the incremental swing premium, which is defined as the difference in the price of the swing option against the strip of the most valuable European options, which also determine the lower boundary of the swing option[46], in percentage terms of the value of the strips. The strip of European call options consists of the four most expensive (valuable) options during the same period of delivery as that of the swing option i.e. the most valuable options out of the 14 different ones each one having maturities $\tau_1, \tau_2, ..., \tau_{14}$ respectively. Particularly in this section we are looking at the sensitivity of the option with respect to the parameters of the stochastic processes, X and Y. Thus each time in the simulation we vary one or two selective parameters and keep the remaining constant, as in the previous chapter in the sensitivity analysis of Asian options. The base case parameters are shown in Table 6.2.

Table 6.2: Parameter Values used in the spike model for the analysis of Swing Options

The table shows the value of the parameters of the spike model, used for the Monte Carlo estimation of Swing Option prices in the base case scenario.

Model Parameters	Parameters values
$X_0 = \left(\varepsilon - \frac{\lambda_X}{k_1}\right)$	5.137
σ_X	1.03
k_1	2.98
Y_0	0
l	5.57
μ_J	0.26
σ_J	0.35
k_2	287.61

The strike price is 170.20 NOK/MWh, which is the equilibrium spot price ($\log\left(\varepsilon - \frac{\lambda_x}{k_1}\right)$ = log (5.137)) in the risk neutral world for the models without seasonality. As discussed in the previous chapter, we do not consider the seasonal parameters and seasonality $f(t)$ since we do not want our results to be specific to any particular season, but rather we want to see how the different parameters from the

[46] Note that when $n=N$, and there are no penalties, the swing option is equivalent to N European Options.

stochastic factors affect the extra percentage premium paid for the swing option compared to a strip of European options.

We thus start the analysis with an ATM swing option, using the non-seasonal components of the spike model that is: $X_0 = 5.137$, $Y_0 = 0$. The LSM method described in the previous section is then implemented to find the value of the swing option. The Monte Carlo Simulations were run using 200,000 sample paths each time, and for each option price sensitivity simulation we used the same set of random variables, as described in the previous chapter.

6.6.1 Determinants of Additional Swing Premium

Consider first the sensitivity of the relative extra premium paid for the swing option with respect to the jump intensity l and mean jump size, μ_J. Figure 6.1 shows that the relative premium of the swing option value is increasing with the jump intensity, l and mean jump size, μ_J. The jump size effect is becoming much more pronounced at higher intensity levels. This is the opposite from what it has been shown for American calls when there is a GBM and their price is the same as European call options, as early exercise in this case is not the optimal strategy (Hull 2003). The reason is that when jumps are introduced in the model, as soon as a positive jump occurs the holder of the swing options has the advantage to exercise immediately rather than wait until the expiration date, when the jump may have died out due to mean reversion. Furthermore, jump events are rare and hence when the opportunity comes and the price reaches extreme price levels, immediate payoff will yield high profits, where as in the European options case the holder is restricted to exercise at maturity only. This is due to the fact that the probability of a jump occurring during a time interval is higher than the probability of the event happening at a specific point in time, as the cumulative probability is the sum of the independent probabilities of the jump happening at specific points in time. Using the same intuition we should also expect the extra swing premium for down-swing options to increase with intensity and as the mean jump size becomes more negative, since the holder of the option will benefit from the more probable downward movements, and has the advantage of exercising early rather than wait and miss the opportunities like the holder of European Put options.

Figure 6.1: Incremental Swing Option Premium Relative to Strips, with respect to changes in Mean Jump Size, μ_J, and jump intensity, l.

The graph displays a sensitivity analysis on the relative extra swing premium with respect to European call, for a Swing option with 4 up-swing rights and 14 exercise opportunities. The prices of Swing Options are estimated using the LSM method explained in this chapter, by changing the parameters of the Mean Jump size, μ_J, and Jump Intensity, l.

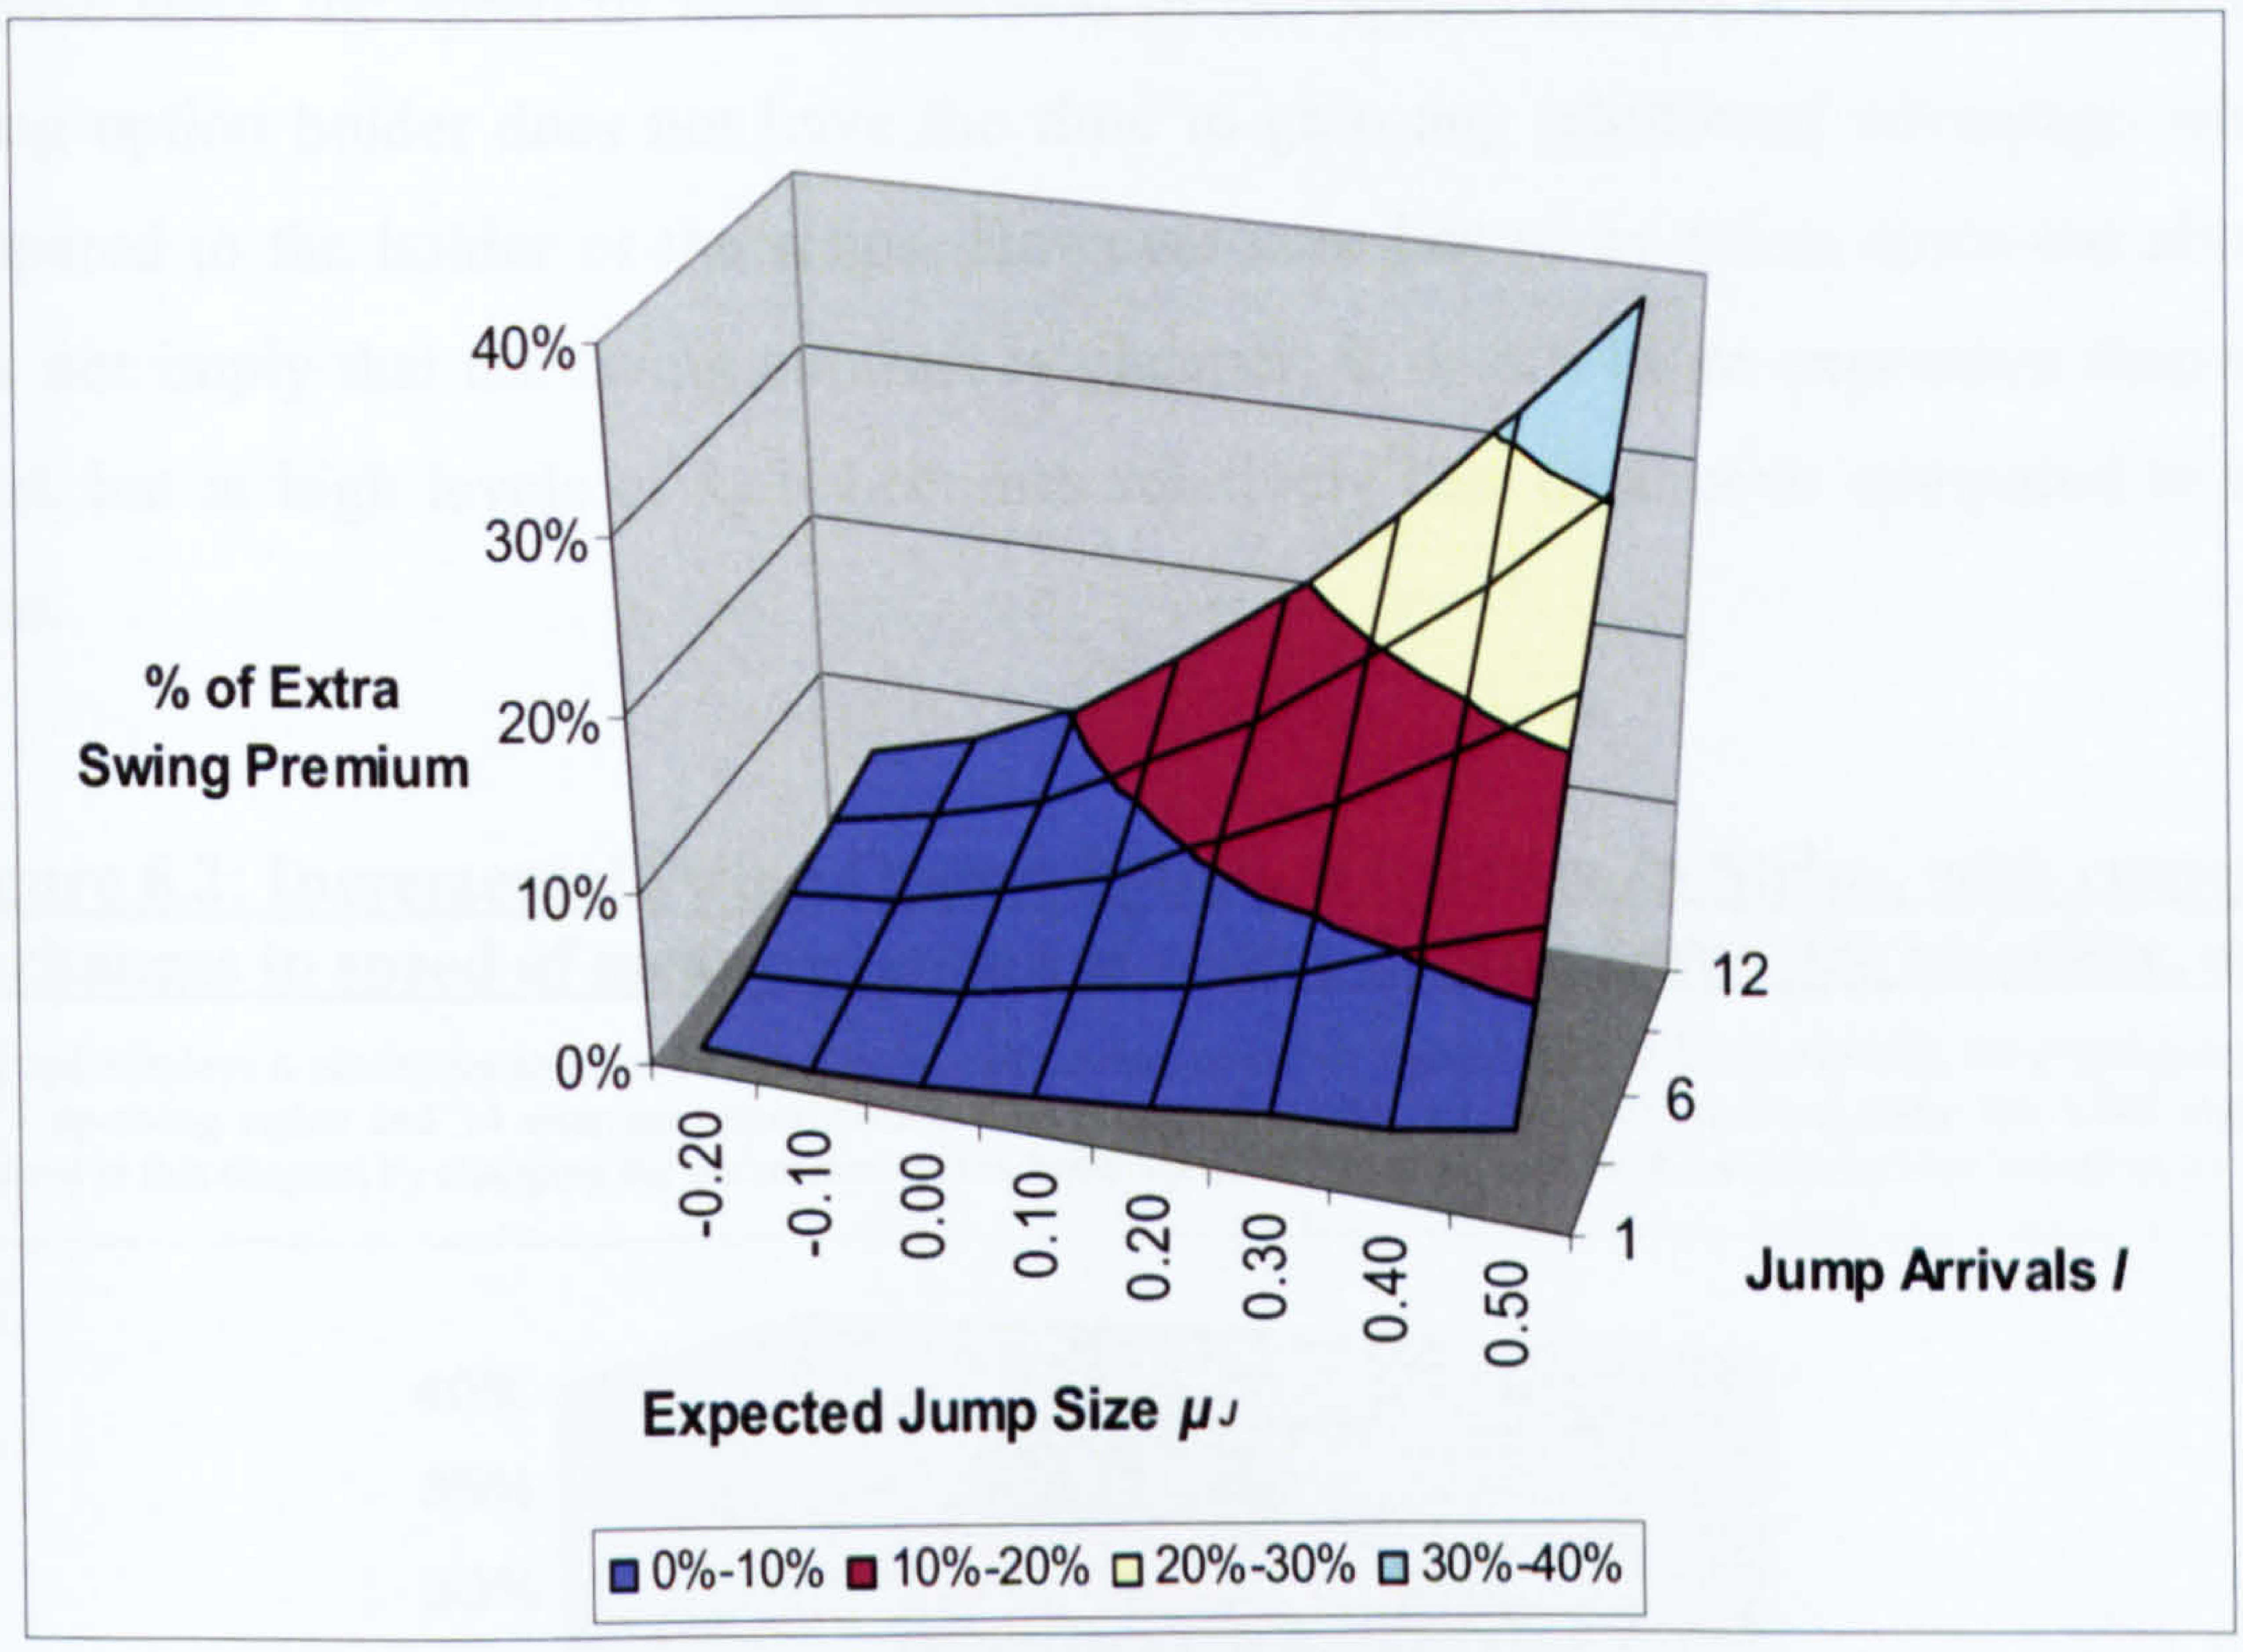

On the other hand Figure 6.2 looks at the sensitivity of the relative extra premium paid for the swing option with respect to the speed of mean reversion of spikes, k_2 and the jump size volatility, σ_J. The surface reveals that the relative extra premium is always increasing with respect to the jump size volatility. This is due to the fact that the more volatile the jumps are, the more extreme an event will be during the time interval and thus when an extreme event occurs the swing option holder will take the opportunity immediately and exercise one of his swing rights. This is also accentuated by the fact that $\mu_J>0$ hence increasing the likelihood of exercising up-swings. On the other hand the relation with the speed of mean reversion is parabolic. This is an interesting finding and can be intuitively explained as follows; first of all as we showed in the previous chapter, the higher the speed of mean reversion of a spike the lower the option price, because if a spike occurs the speed of mean reversion will pull the spike back to the equilibrium level very fast. This is an advantage for swing option holders, because when the spike occurs they can exercise immediately where as the European option holders will not benefit from the spike if it occurs before the option's maturity. On the other hand if k_2 is very high, spikes die out very fast therefore early exercise opportunities disappear quickly and the extra premium is reduced. In the current case, the level of spike mean reversion at which the swing option holder's advantage starts decreasing is approximately 100 (Half-life of 2.5 days). This is because since the speed of mean reversion of the spikes at this level is so fast, the Swing option holder does not have the time to gain any additional advantage when compared to the holder of the strips. However care has to be taken since the above does not imply that the swing contract is cheaper; it is still more expensive than the strips, but at high levels of k_2 it becomes relatively less expensive compared to the strips.

Figure 6.2: Incremental Swing Option Premium Relative to Strips, with respect to changes in speed of mean-reversion of spikes, k_2, and jump size volatility, σ_J.

The graph displays a sensitivity analysis on the relative extra swing premium with respect to European call, for a Swing option with 4 up-swing rights and 14 exercise opportunities. The prices of Swing Options are estimated using the LSM method explained in this chapter, by changing the parameters of the Spike speed of Mean Reversion, k_2, and Jump size volatility, σ_J.

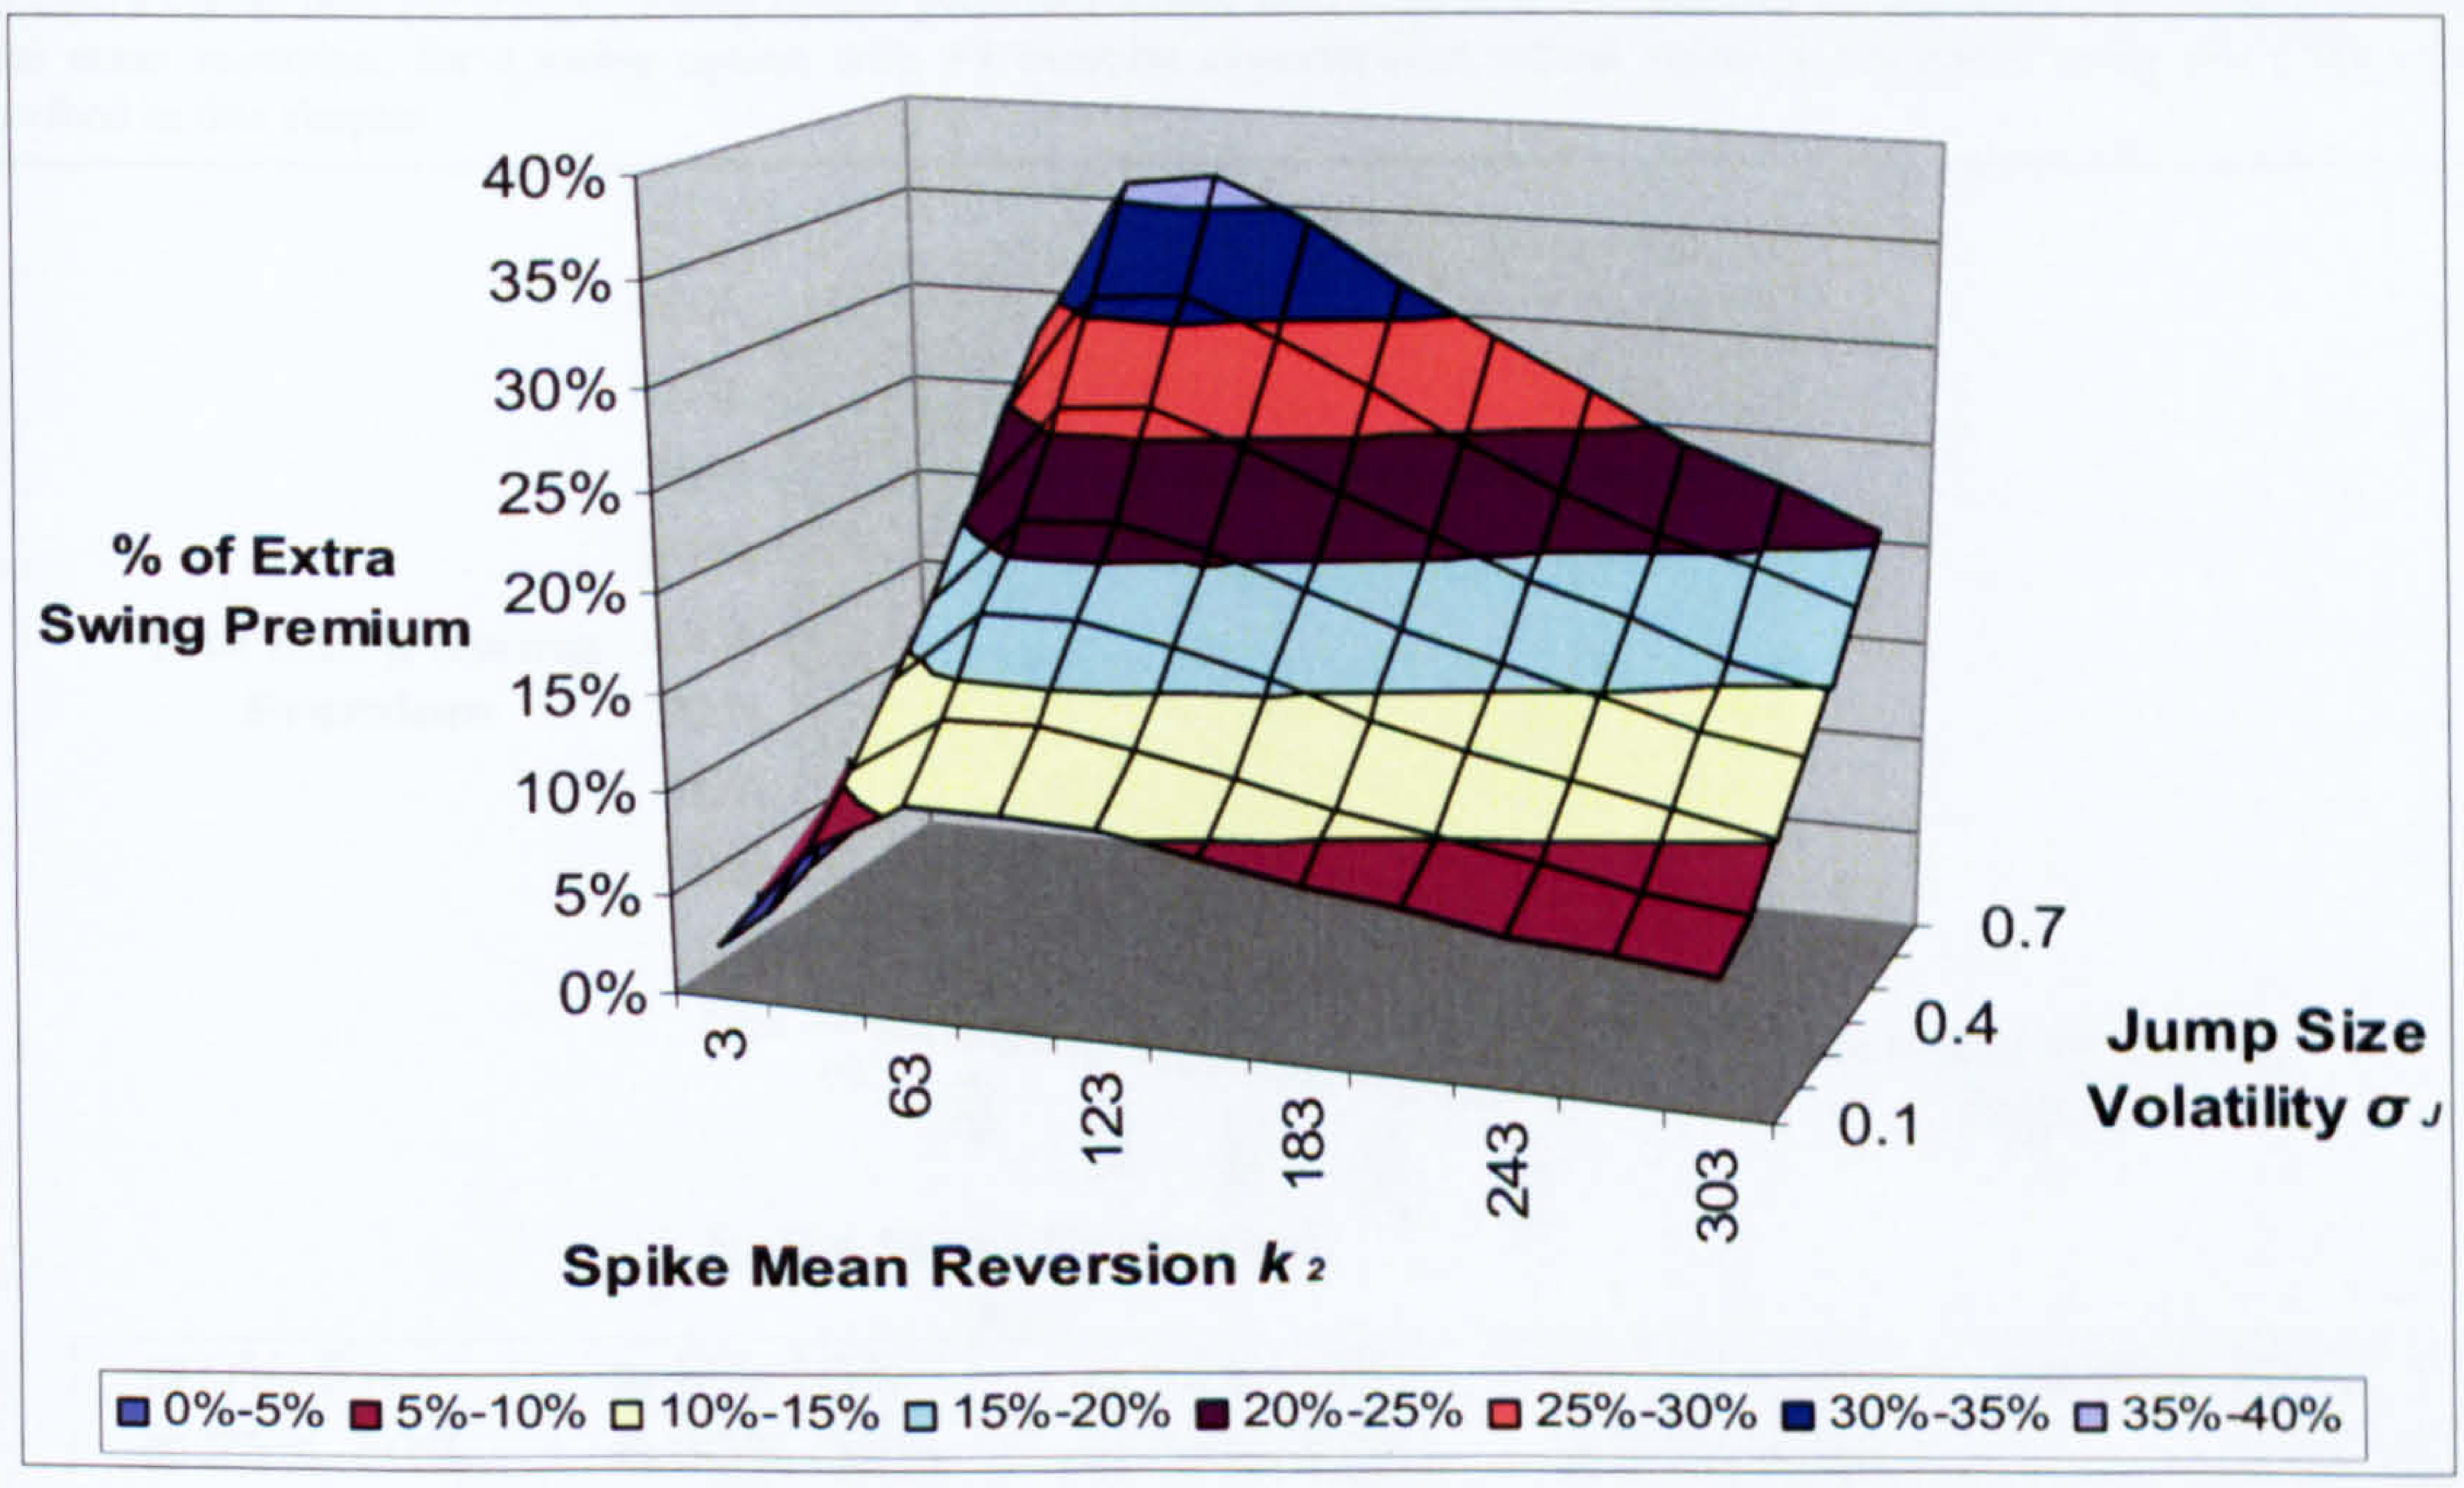

It is then interesting to see how this turning point, just discussed, actually changes with the number of swing opportunities? Figure 6.3 plots the surface for the extra swing premium with respect to different values of the spike speed of mean reversion and the number of up-swing rights. First of all from the graph we see that by increasing the number of swing rights the incremental swing premium decreases. This is expected since, as we will also explore in more detail later, as the number of swing rights approaches the number of swing opportunity dates τ_j, the option value approaches that of European strips. However, in the graph we see that the turning point where the relative extra swing premium starts decreasing, occurs at higher speeds of spike mean reversion as we decrease the number of swing rights for the same number of swing opportunities (14 days). This is also expected since if one has for example just one swing right (Bermudan option), by keeping all other parameter constant and increasing the spike speed of mean reversion, he will have a greater benefit in relative terms, than someone with more swing rights, as there will be more extreme events for him to exercise his unique right. Of course at this point, we should expect the same behaviour for down-swings if we had a negative mean jump size using the same intuition just described.

Figure 6.3: Incremental Swing Option Premium Relative to strips, with respect to changes in speed of mean-reversion of spikes, k_2 and number of exercise rights (up-swings).

The figure shows how the relative swing option premium varies with respect to the number of up-swing rights and the speed of spike mean reversion, for a swing option with 14 exercise opportunities whose value is estimated using the LSM method described in this chapter.

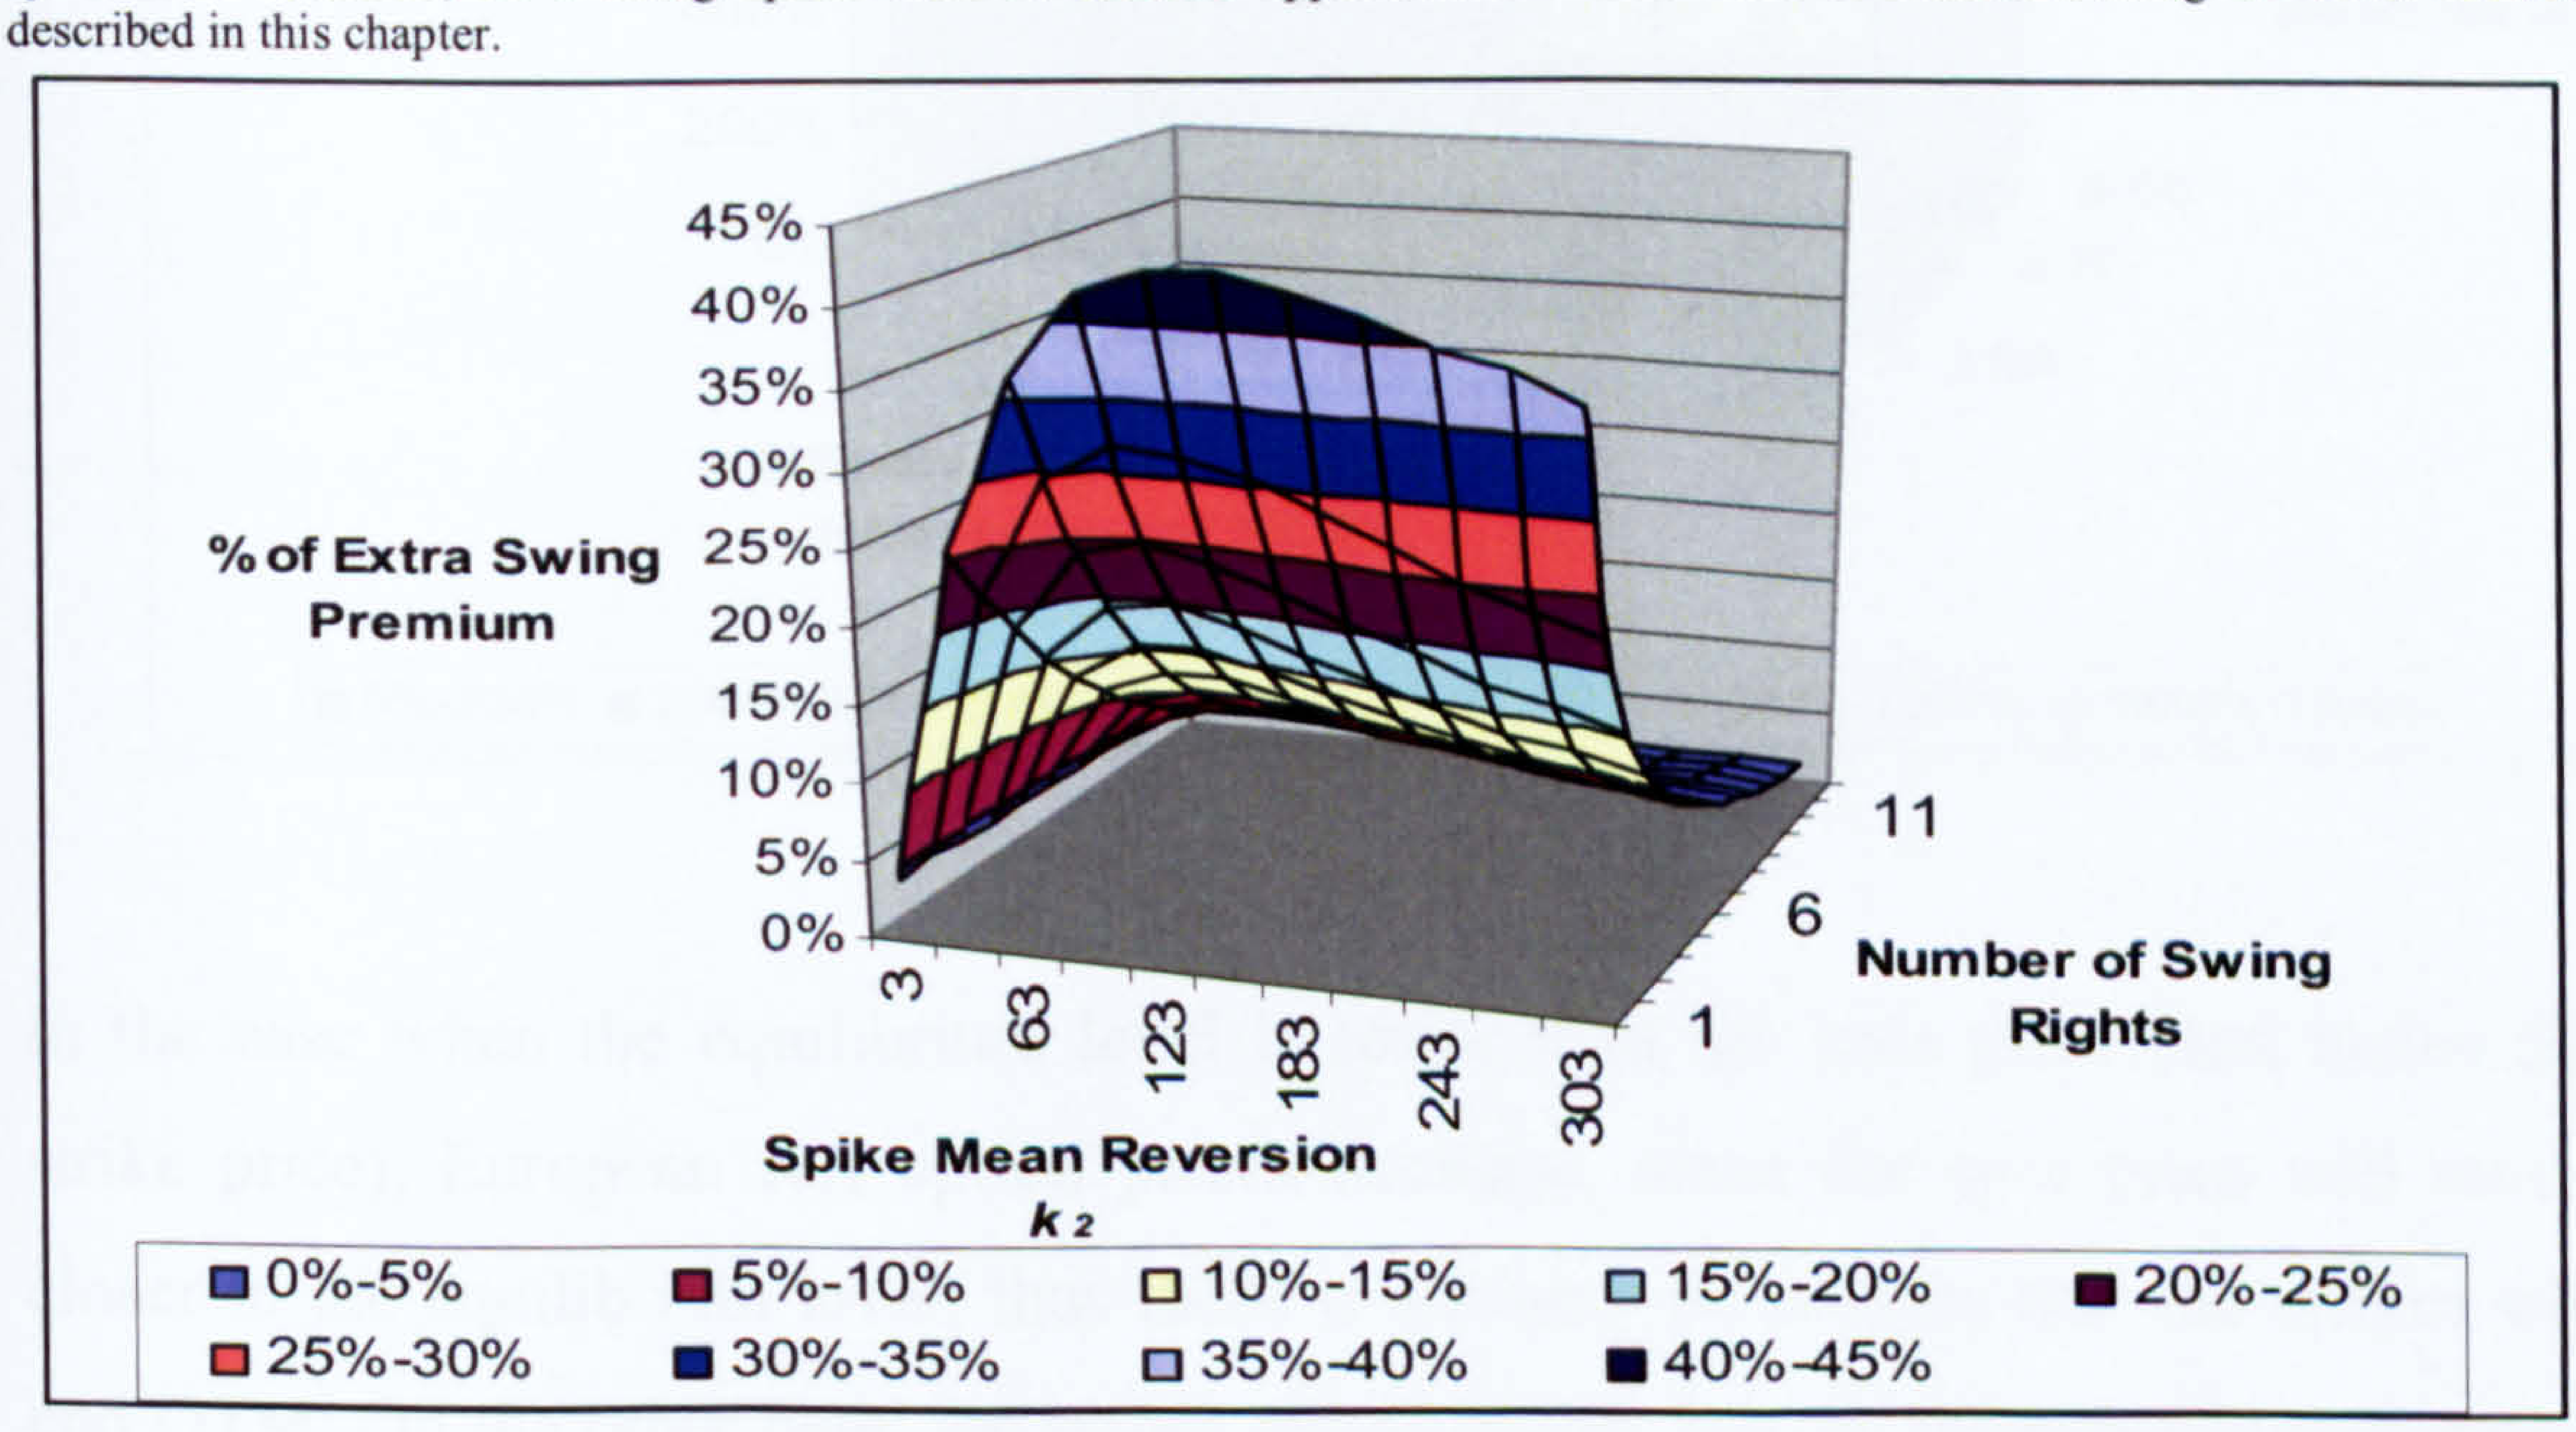

Figure 6.4 shows the sensitivity of the relative extra premium with respect to the diffusive component's equilibrium level and its speed of mean reversion. In this case, the extra relative risk premium is increasing with the speed of mean reversion as the equilibrium level decreases. This is happening because, for an At-The-Money option

as the one presented here, when the equilibrium level is higher than the current spot price and the speed of mean reversion is low, the swing option holder will benefit by waiting further until the spot price mean-reverts to the high equilibrium level, thus getting closer to maturity. Hence, the swing option value approaches that of the strips, since the most valuable European options are the ones with the longest time to maturity.

Figure 6.4: Incremental Swing Option Premium Relative to Strips, with respect to changes in diffusive speed of mean-reversion, k_1, and risk-neutral mean equilibrium level of X, $\left(\varepsilon - \frac{\lambda_X}{k_1}\right)$:

The graph displays a sensitivity analysis on the relative extra swing premium with respect to European call, for a Swing option with 4 up-swing rights and 14 exercise opportunities. The prices of Swing Options are estimated using the LSM method explained in this chapter, by changing the parameters of the speed of Mean Reversion, k_1, and the Risk Neutral Equilibrium Level to which the Diffusive component reverts to.

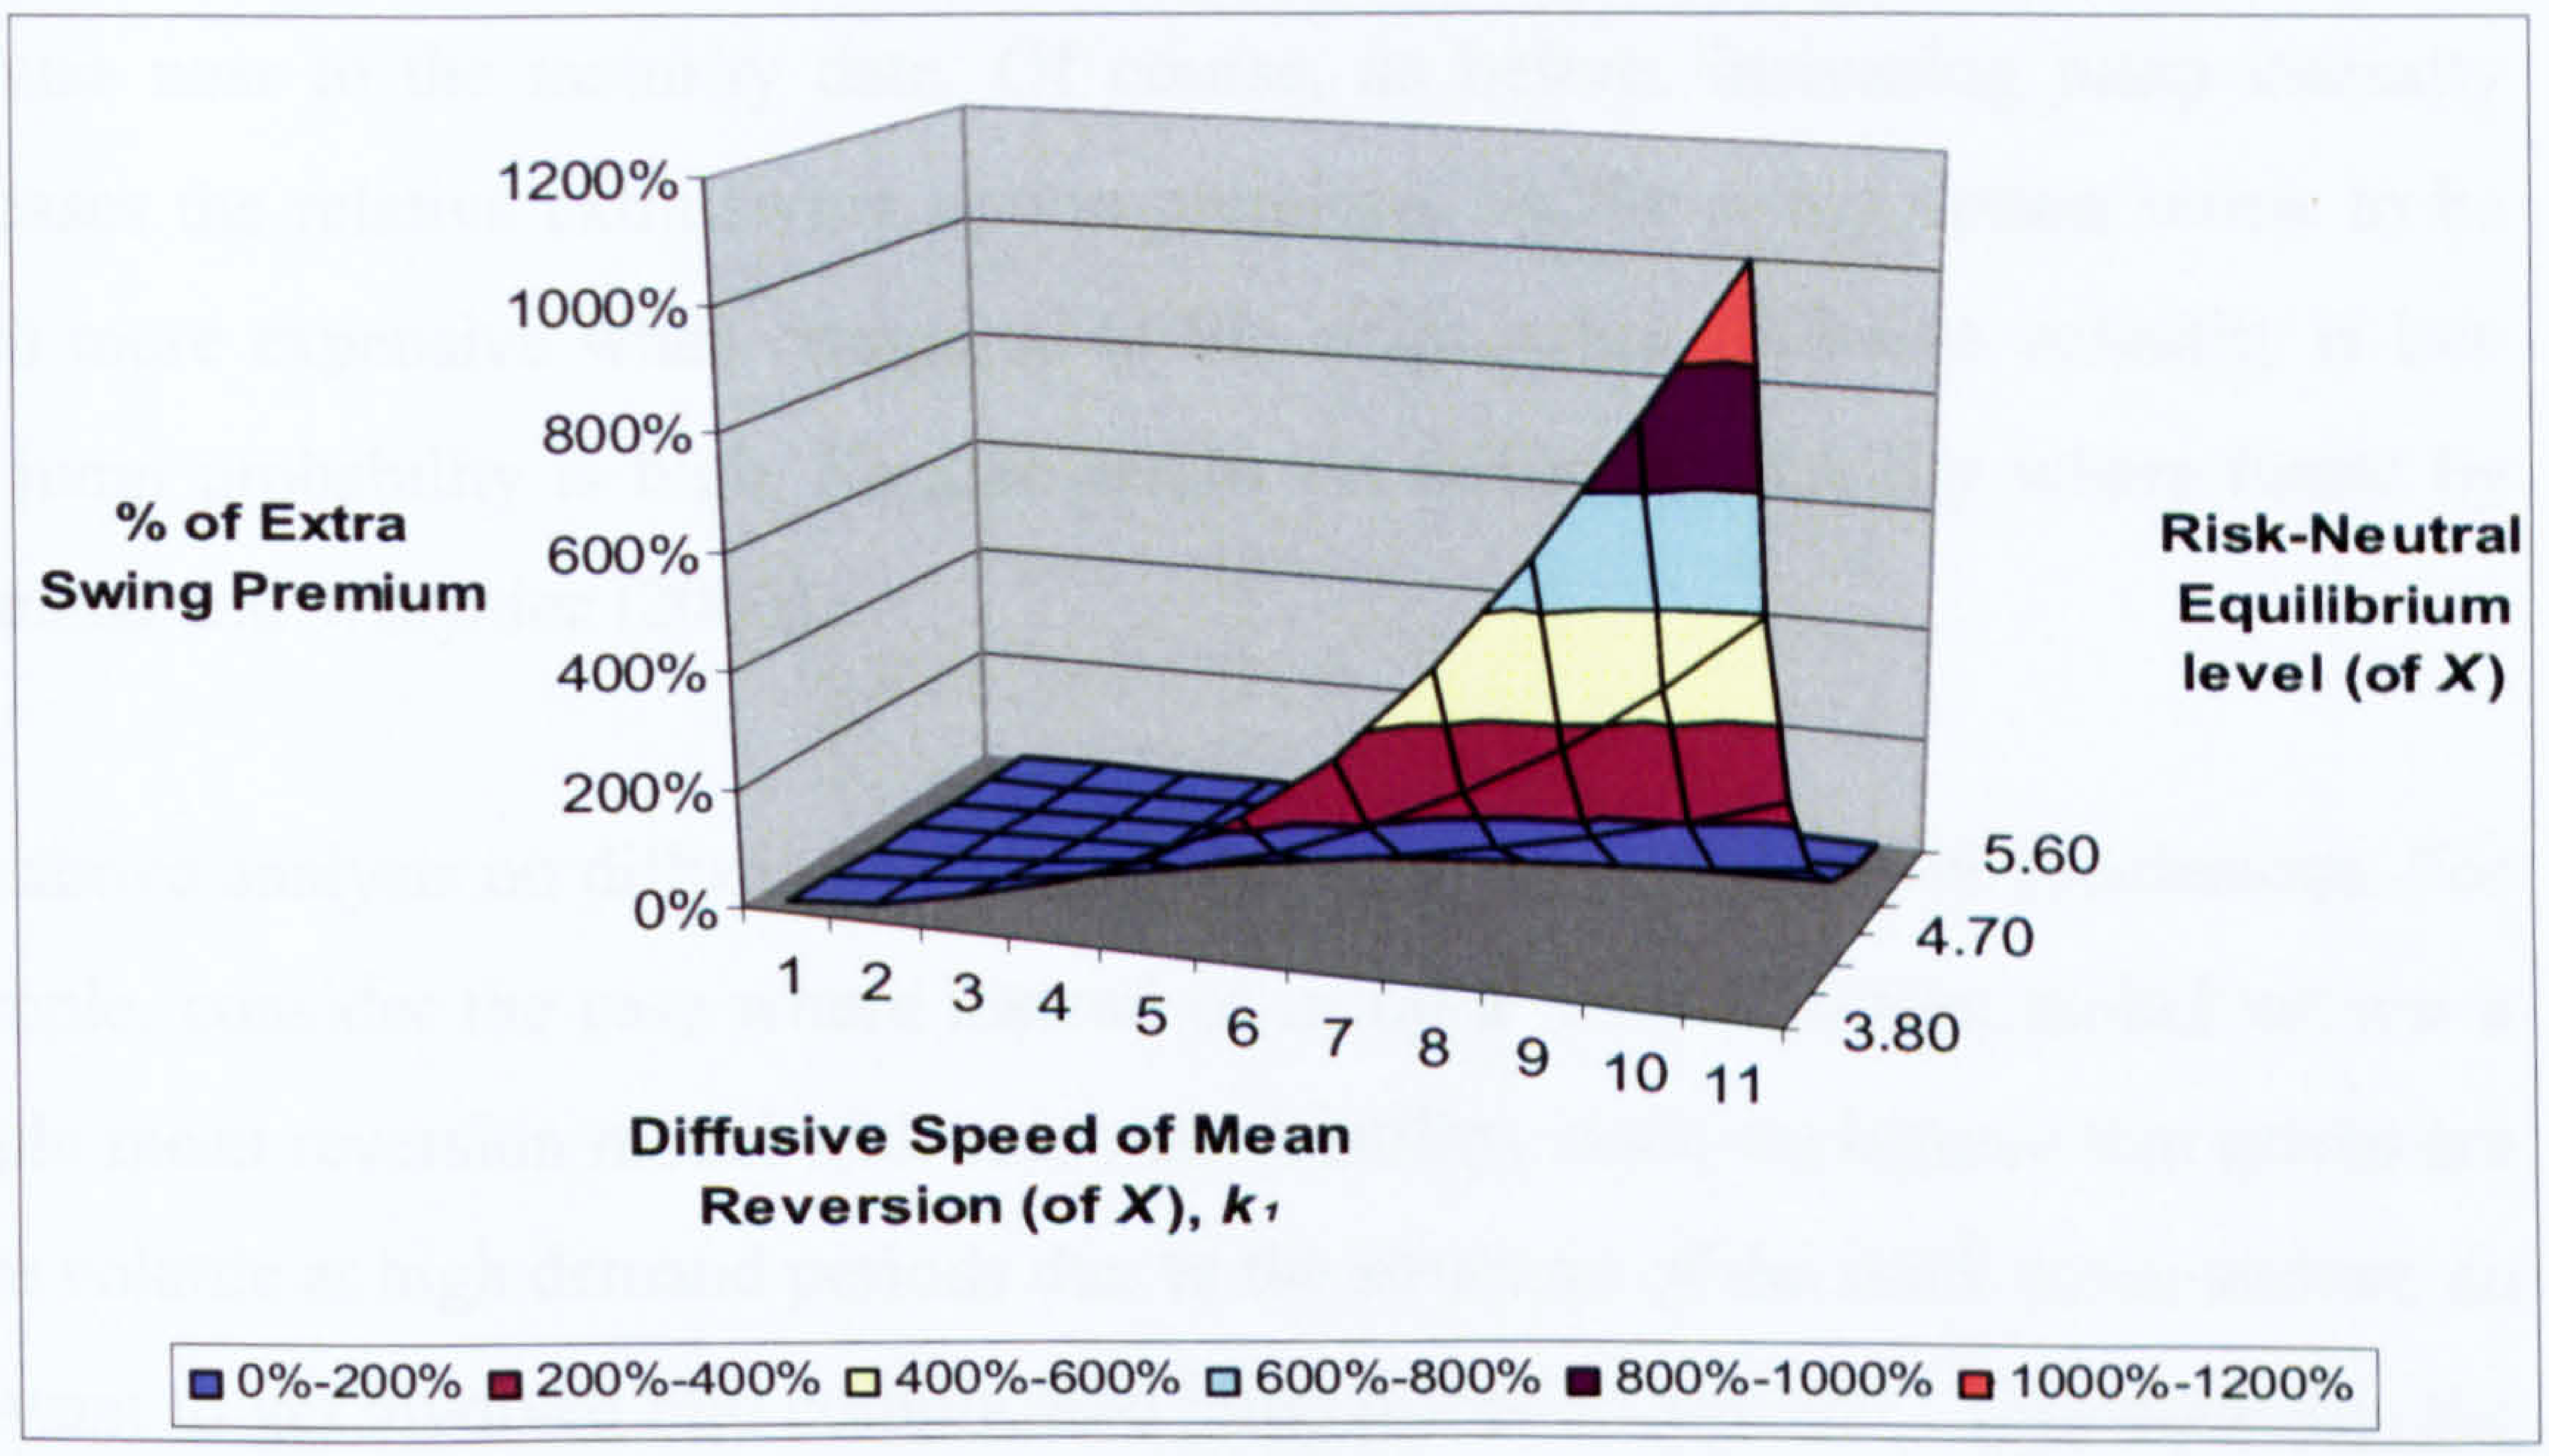

In the case when the equilibrium level is lower than the spot price (and hence the strike price), European call option prices decrease, since the spot price will revert closer to the equilibrium level, thus there is a higher probability that the option will end OTM. On the other hand the swing option holder has the opportunity to exercise early, and may exercise his swing rights early so as to maximise his profits, rather than wait until further in time when the spot price gets closer to the low equilibrium level. This additional benefit of course increases as the speed of mean reversion increases. Also note that the extra premium in percentage terms is quit high (maximum value being close to 1000%), due to the very low values of the European options, as explained in this paragraph. Figure 6.5 presents the incremental percentage premium with respect to changes in diffusive volatility and jump intensity, in order to study the effect of the diffusive and the jump volatilities in the total swing premium. Surprisingly, when increasing the diffusive volatility, the relative incremental swing premium decreases. This is due to the fact that the standard deviation of the increments of the Brownian motion dW is equal to $\sqrt{dt}$ which directly implies that there is a greater range of possible prices the further away we move into the future, and this range is scaled up proportionally by the diffusive volatility parameter, σ_X. Thus as time to maturity increases, the number of exercise opportunities increase as well, so it is not optimal to exercise early. Therefore, even though swing option values increase with the diffusive volatility, their prices converge to that of strips (since European Options prices increase with time to maturity), as it is now more optimal to exercise near to the maturity date. Of course, as before, increasing jump intensity increases the relative extra swing option premium. So the swing option seems to be much more expensive when compared to the strips, when diffusive volatility is low and jump probability is high. Similar results on diffusive volatility where found by Eydeland and Wolyniec (2003).

The above analysis on diffusive volatility provides some very useful conclusions. For example, consider the case where instead of using a jump diffusion model we use a simple mean reversion model with seasonal volatility, since we believe that prices are more volatile at high demand periods due to the structure of the stack curve and we do not want to get involved into complicated jump diffusion models[47]. Consequently, the swing option premium would be very much similar to that of the strips and thus no significant extra premium would be demanded; in this case one might also think that strips would provide an almost perfect hedge for swing options. On the other hand, using the spike model with seasonal jump parameters yields swing premia, which are relatively higher than those for strips, since jumps make it optimal to exercise early.

[47] A similar model with seasonal volatility was proposed by Kokebakker and Ollmar (2001) although they model the forward dynamics. Nevertheless the intuition of the results remains the same.

Figure 6.5: Incremental Swing Option Premium Relative to Strips, with respect to changes in diffusive volatility, σ_x, and jump intensity, *l*.

The graph displays a sensitivity analysis on the relative extra swing premium with respect to European call, for a Swing option with 4 up-swing rights and 14 exercise opportunities. The prices of Swing Options are estimated using the LSM method explained in this chapter, by changing the parameters of the diffusive volatility parameters, σ_X, and Jump Intensity, *l*.

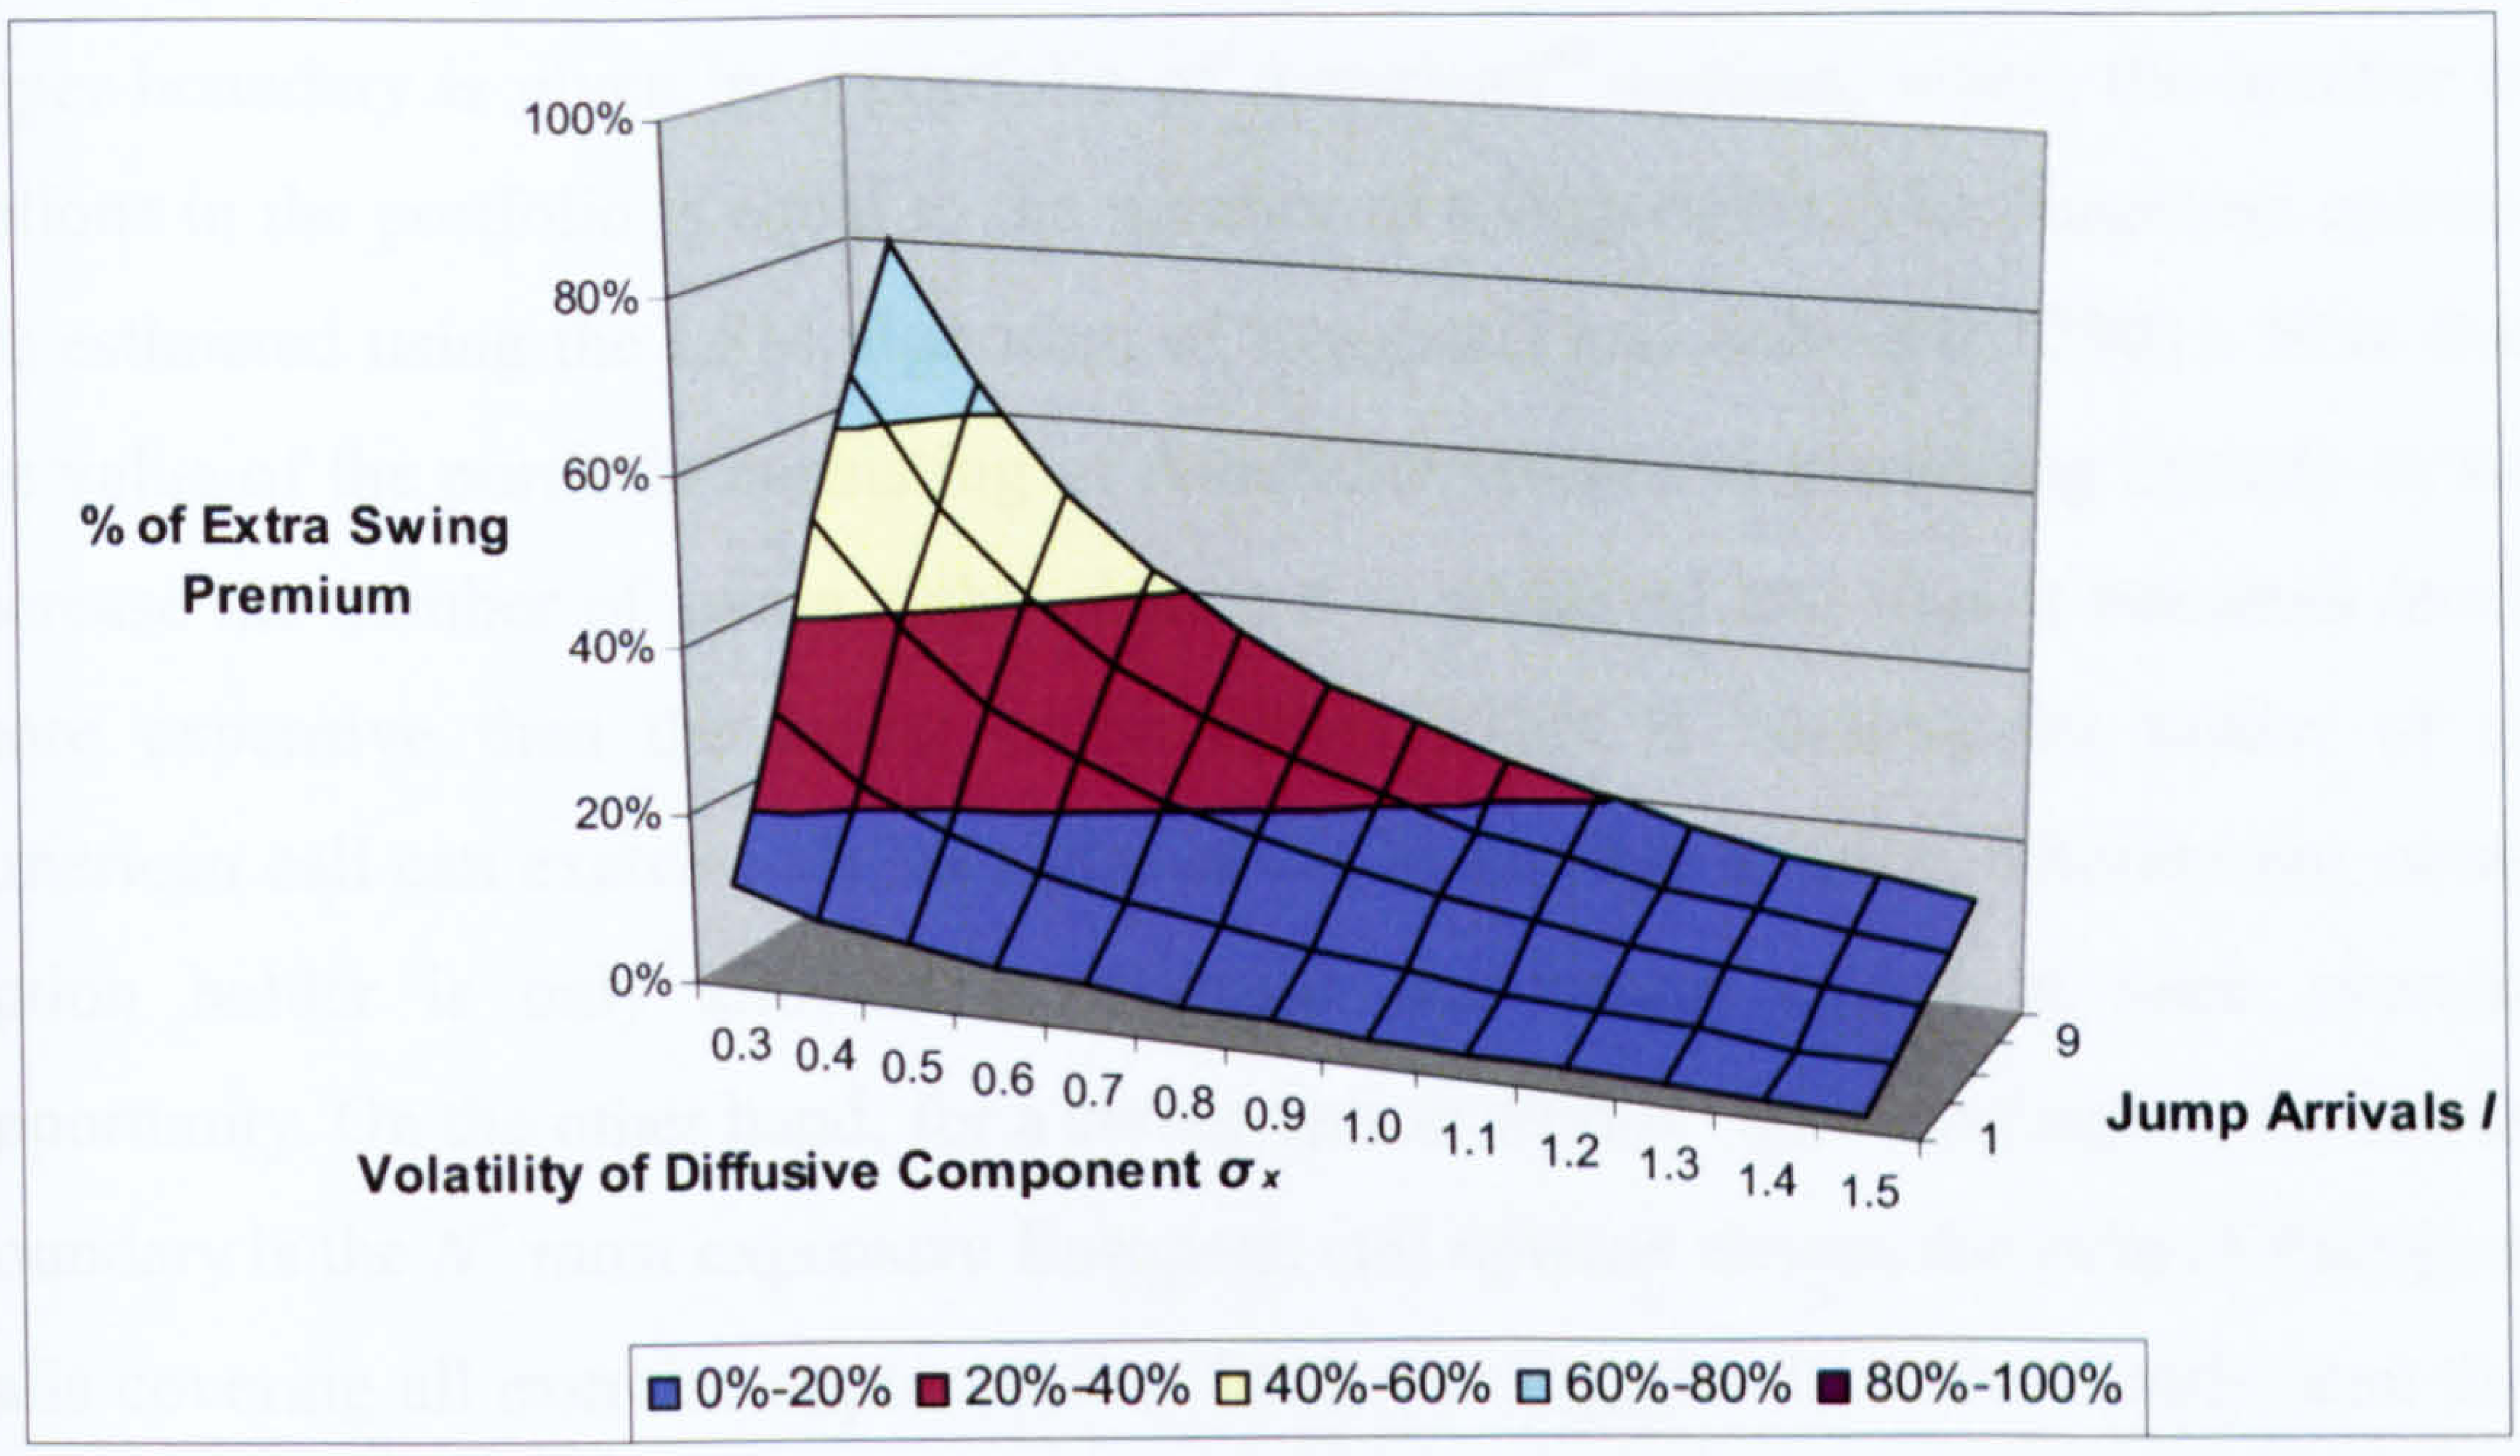

6.6.2 Swing Option Prices under the two Different Model

In this section we are going to explore the different swing option prices given by the MR and spike model respectively. We are going to make comparisons across different seasons, and we will also provide the upper and lower bounds of the swing options, which were discussed in the pricing section. The parameters used for the models are the same as the ones estimated in Chapter 4, Table 4.1.

Figure 6.6 and Figure 6.7 show how the value of the swing options increases with respect to the number of swing rights for the diffusive Mean Reverting model (MR) and Spike models, respectively. As expected, as the number of swing rights increases the value of the swing contract increases, however the marginal return in the value of the incremental swing premium is diminishing. This is due to the fact that as the number of swing rights increases, the extra benefit the swing holder receives is lower because at the same time interval there are limited exercise possibilities. Thus increasing the number of rights with the same possibilities for exercise in that interval increases the value but at a decreasing rate. A formal mathematical proof of this is shown by Meinshausen and Humbly (2004) using Doob's optimal decomposition theorem.

In each graph we also show the upper and lower boundaries of the swing options. The upper boundary is given by a portfolio of American[48] options, where the number of options in the portfolio is equal to the number of swing rights. The American options are estimated using the LSM algorithm of Longstaff and Schwartz (2001). Note that the value of the portfolio consisting of American options is increasing linearly as we increase the number of swing rights, (since it is additive) and thus it becomes much more expensive than the actual swing option. This is because the holder of an American call can exercise all his rights at the same time at once, whereas the swing option holder is only allowed to exercise one swing option at each exercise opportunity. On the other hand, for a swing option with N^+ up-swing rights, the lowest boundary is the N^+ most expensive European call options among the strip of European calls covering all exercise opportunities. From the Figures it can be clearly seen that the value of the strips approaches the value of the swing option and is the same when the number of swings is equal to the number of exercise rights. As explained earlier, this mainly occurs because in this case there are no more early exercise opportunities and thus the swing option becomes equivalent to the sum of European options, which is equivalent to a Cap with daily exercise possibilities.

One clear difference between the values for the swing options produced by the Mean-Reverting and the spike model is the fact that the swing option prices produced by the Mean-reverting model are much closer to their lowest boundary than the prices produced by the spike model. This is due to the fact that in the spike model we have jumps which, as we saw from the previous sensitivity analysis, give incentive for early exercise. On the other hand the Mean-Reverting model has a greater diffusive volatility, and the analysis in Figure 6.5 showed that this decreases the possibility of early exercise thus the value of the swing approaches the value of the strip of European options.

[48] Here the term American option is actually more accurate to be referred as Bermudan options, since the rights can be exercised at specific points in time (i.e. daily) rather than continuously.

Figure 6.6: Swing (up-Swings) Option Values (NOK/MWh) and their boundaries with respect to the number of swing rights for the Mean Reverting (MR) Model

The figure shows the value of Swing options with respect to the number of up-swing rights, and its upper and lower boundaries, which consist of American and European options respectively, using the Mean-Reverting model described in Chapter 4 without taking into account deterministic seasonality, *f(t)*.

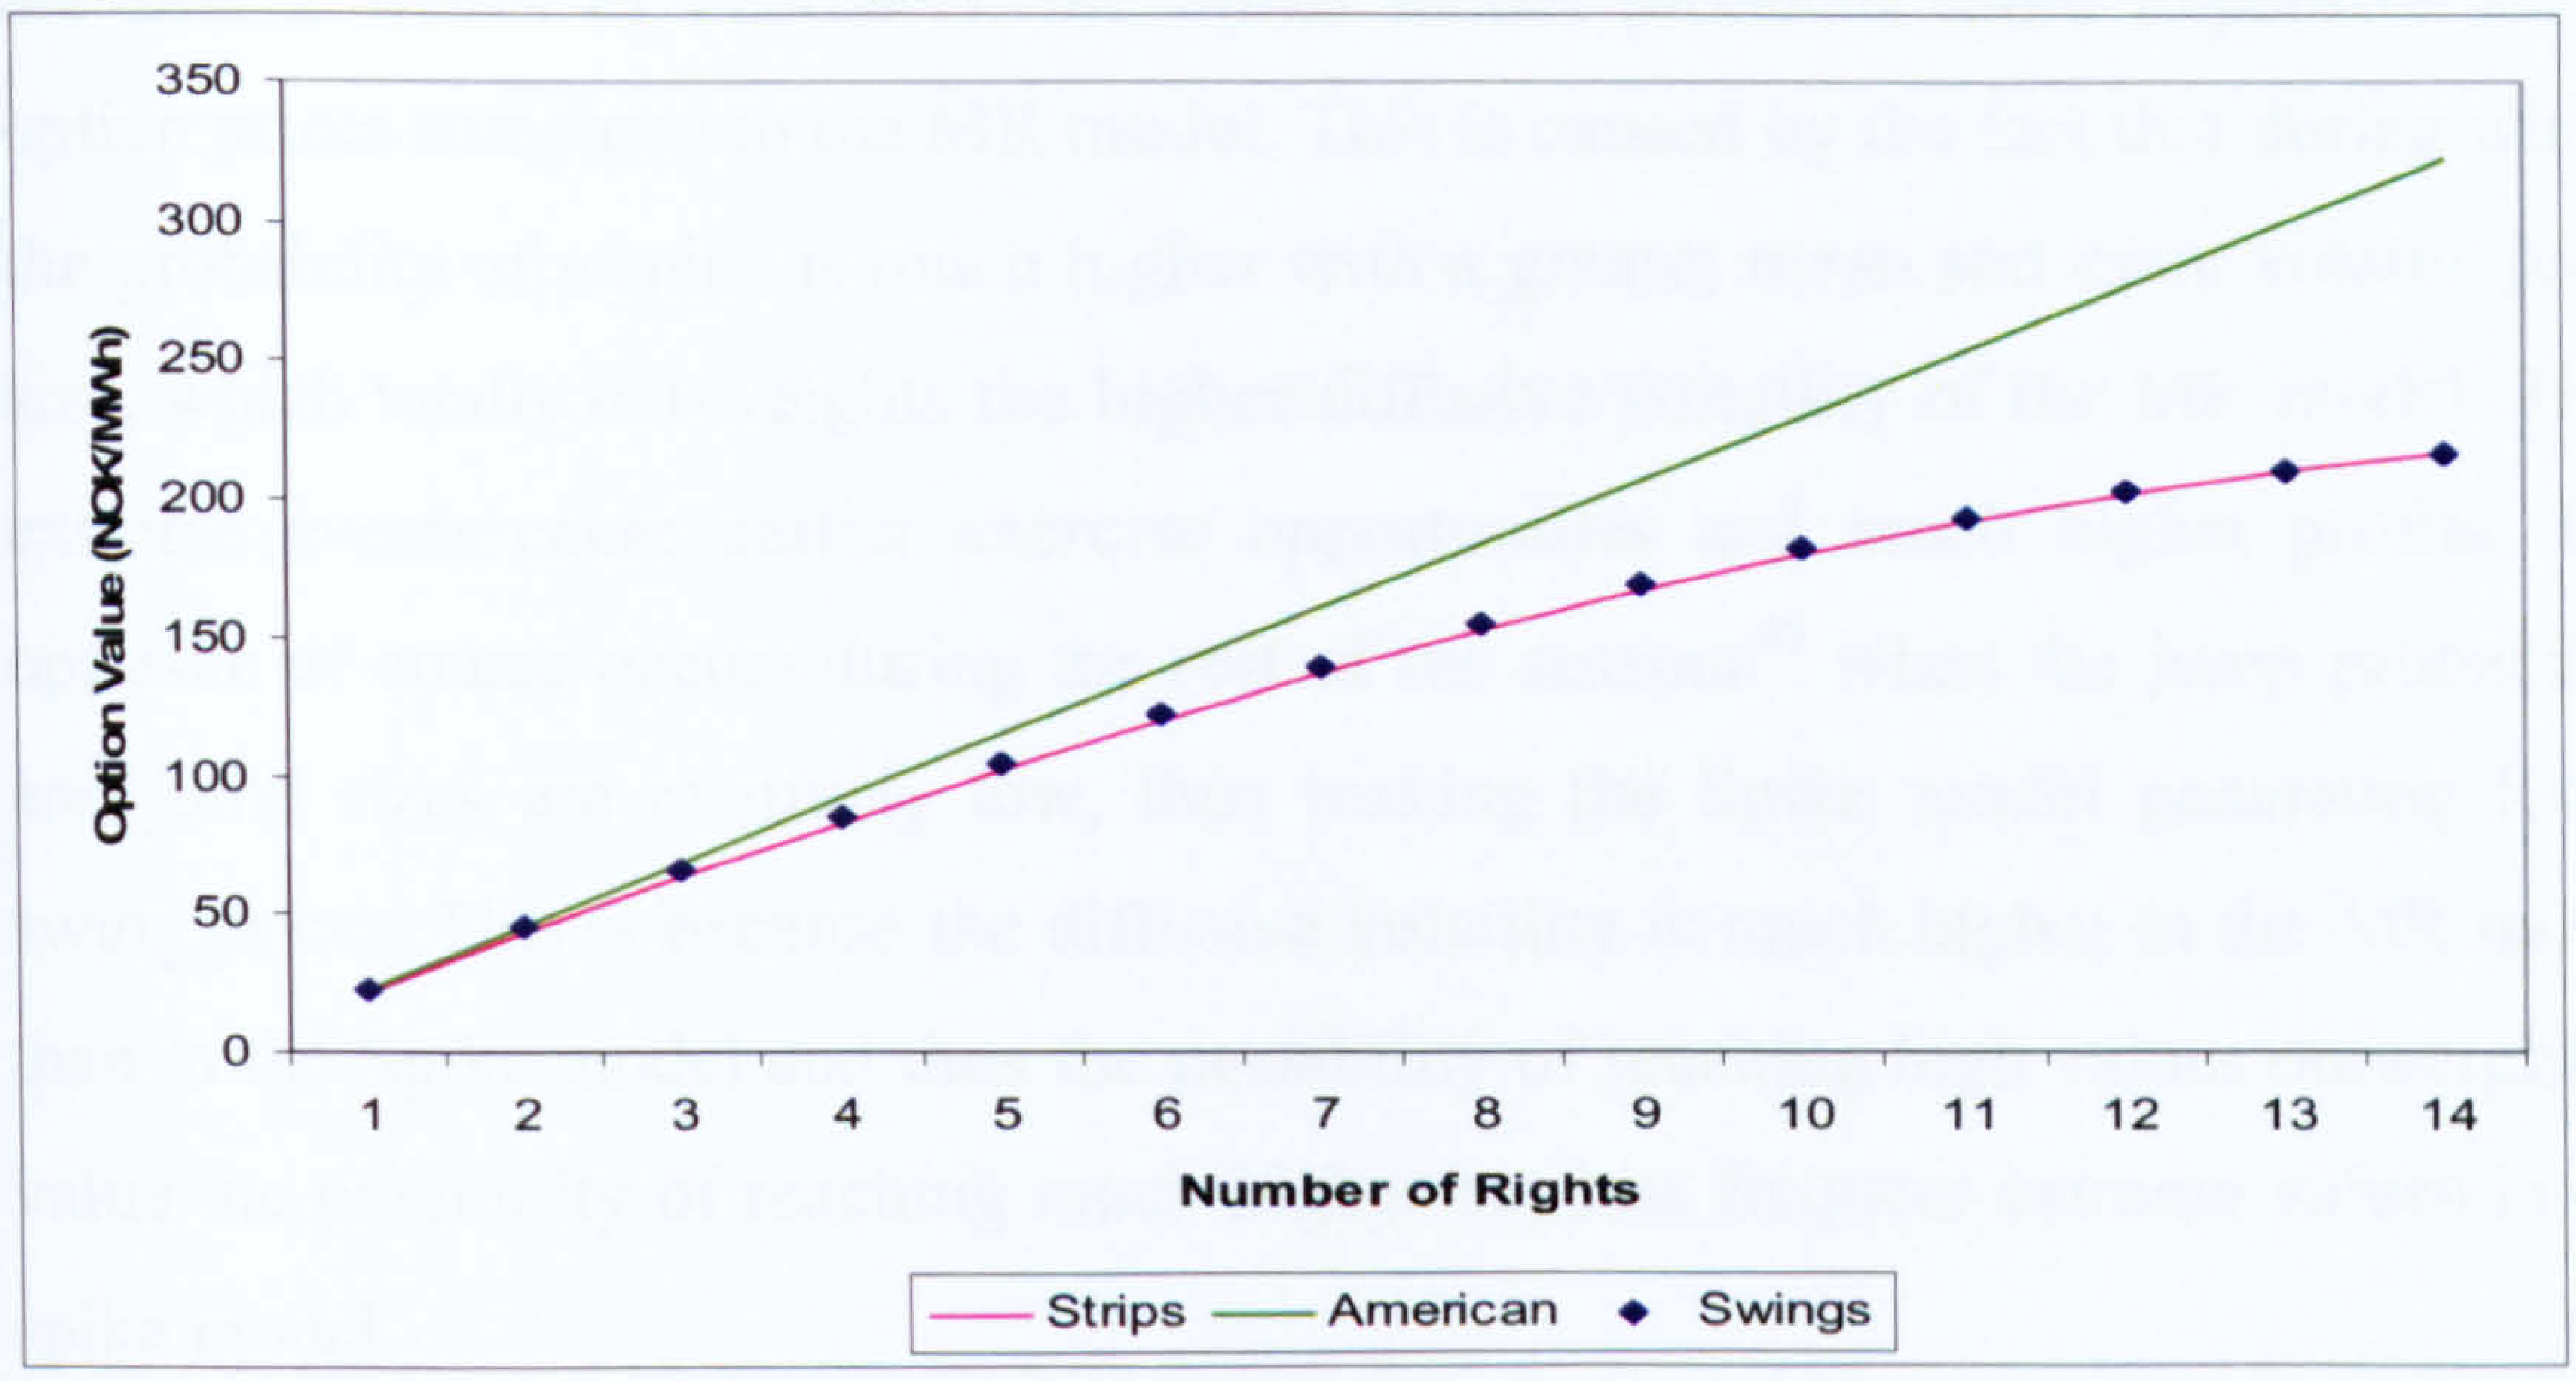

Figure 6.7: Swing (up-Swings) Option Values (NOK/MWh) and their boundaries with respect to the number of swing rights, for the General Spike Model

The figure shows the value of Swing options with respect to the number of up-swing rights, and its upper and lower boundaries, which consist of American and European options respectively, using the Spike model described in Chapter 4 without taking into account deterministic seasonality, *f(t)*, and seasonality in the jump parameters.

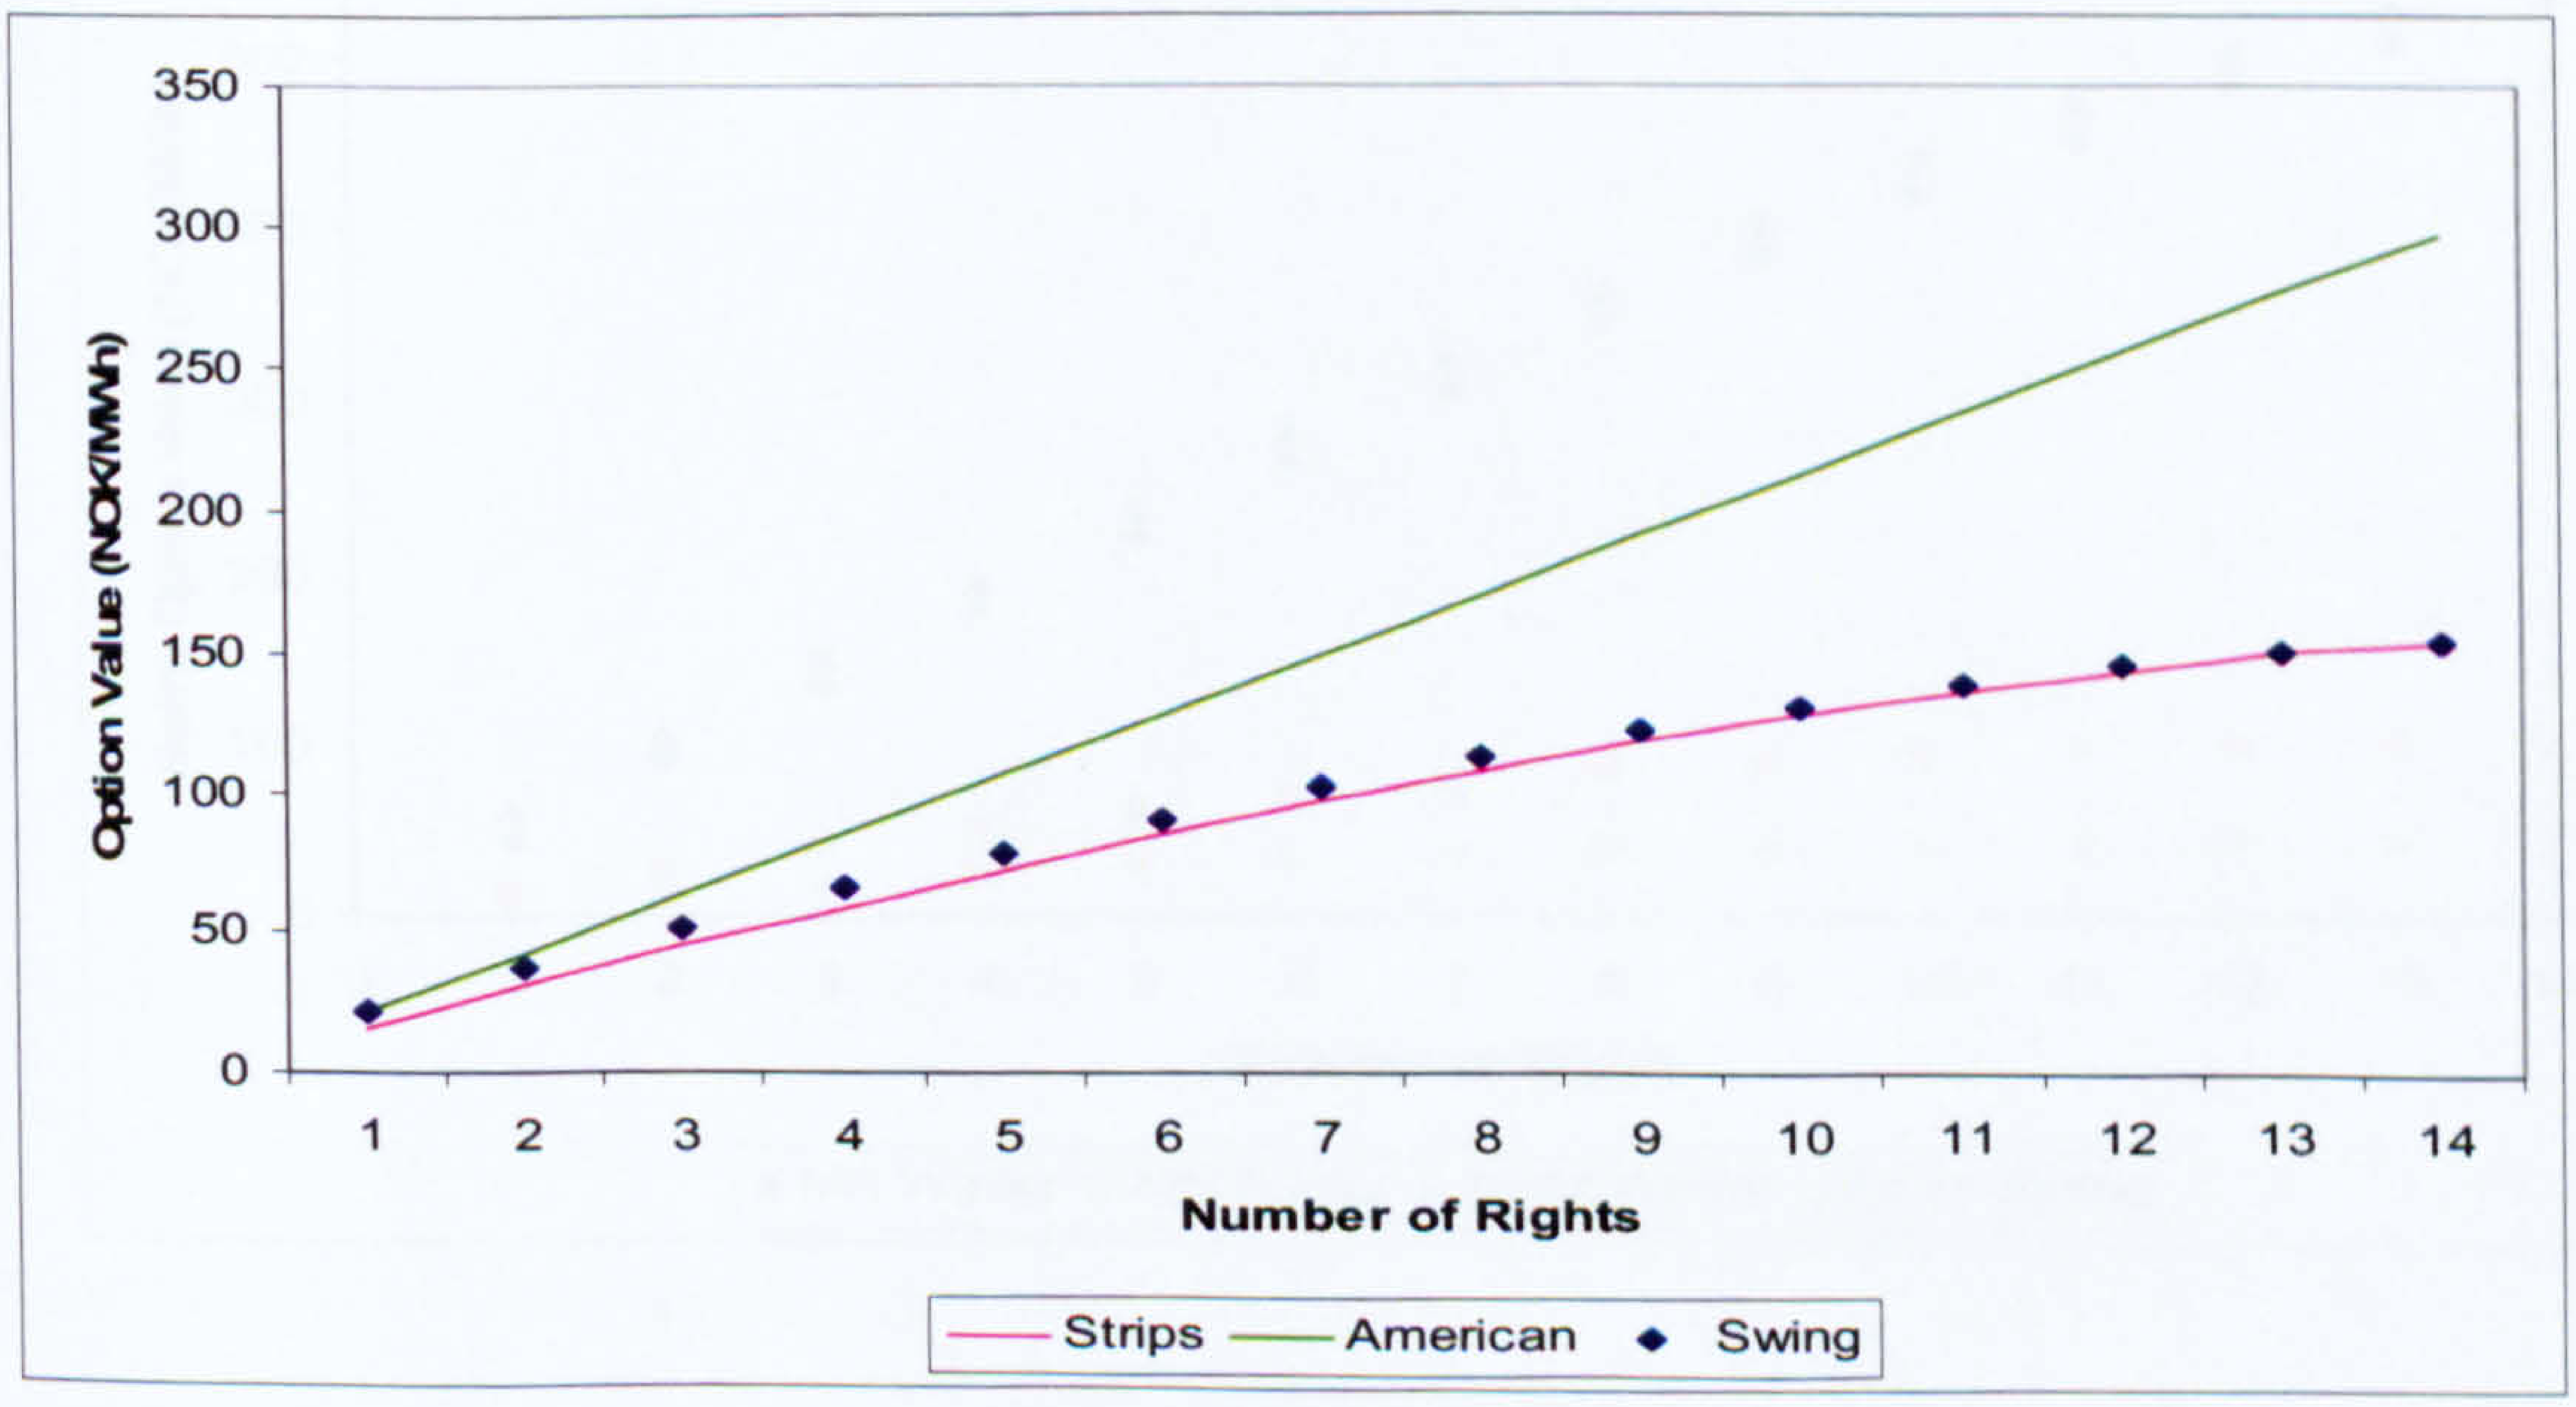

Next we make a comparison between the two basic modes (spike and MR) using deterministic seasonality (*f(t)*) and the seasonal parameters of our models (i.e. the seasonal jump intensity volatility and mean jump size). More specifically in Figure 6.8 we see that during winter time (in this case we estimated swing option prices for the first 2 weeks of February), the Spike model produces more expensive swing option prices compared to the MR model. This is caused by the fact that during winter the probability of a spike is much higher with a greater mean and more volatile jump size, which totally out-weights the higher diffusive volatility of the MR model. Thus extreme events cause earlier exercise opportunities and much higher profits. The opposite of course occurs during the rest of the seasons[49] when the jump probability and jump sizes are relatively low, thus making the Spike model generating lower swing prices. This is because the diffusive volatility is much higher in the MR model than in the Spike model and thus the possibility of reaching high values outweighs in value the possibility of reaching much higher but less frequent extreme values in the spike model.

Figure 6.8: Swing (up-Swings) Option Values (NOK/MWh) in Winter and Spring using the original Spike and MR model with seasonality and seasonal Parameters

The figure compares the value of the Swing options with different number of up-swing rights, produced by the Spike and the Mean-Reverting Models, during winter and spring.

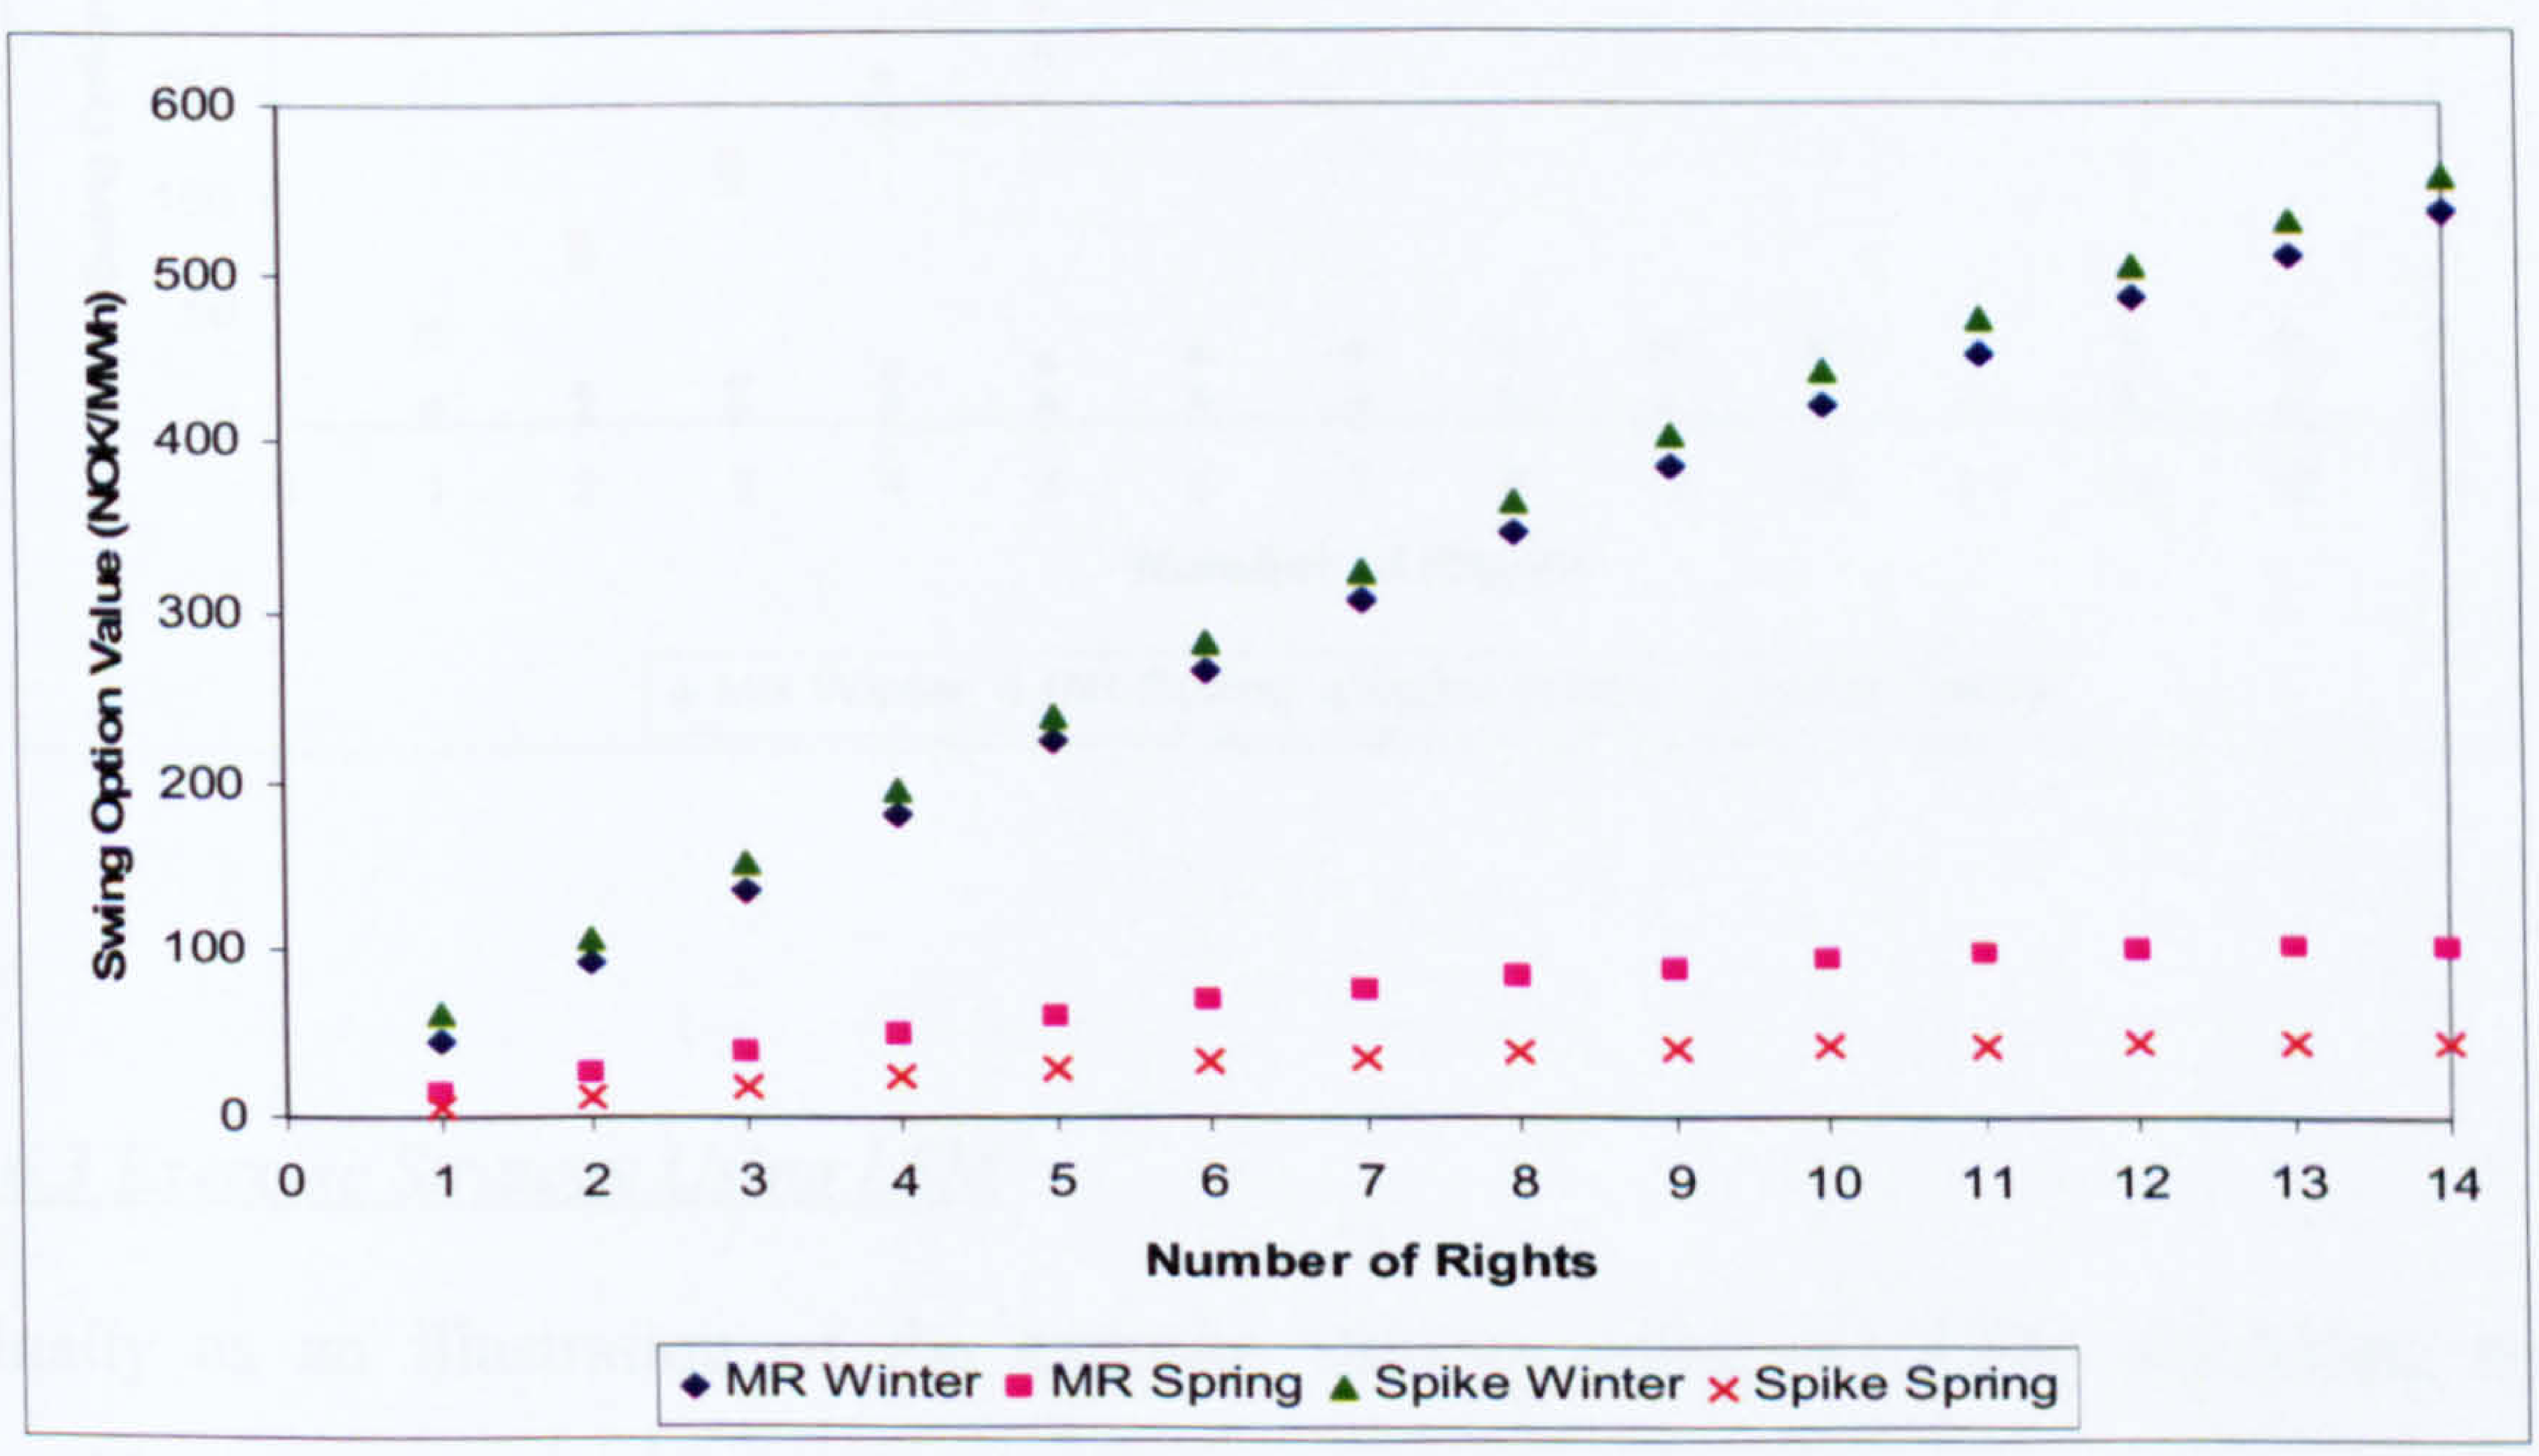

[49] Note that the swing options presented in the graph correspond to the first two weeks of May, but the comparative results are the same for the other months, excluding the months corresponding to winter.

Regarding down-swing options, the methodology in the Monte Carlo method is exactly the same, only this time the pay-off at each exercise opportunity is the same as that of a put option, *max(K-P,0)*. Figure 6.9 shows the swing option values, by varying the number of down-swings, for the two different models in winter and spring. Not surprisingly, the Mean-Reverting model always produces more expensive values than the spike model in both seasons. This occurs due to the fact that spikes have a positive mean jump size, and thus any jump that occurs would be upward biased limiting any possibility of the spot price ending below the strike price. Thus the more positive the jump size, the lower the value of this put like option.

Figure 6.9: Swing (down-Swings) Option Values (NOK/MWh) in Winter and Spring using the original Spike and MR model with seasonality and seasonal Parameters

The figure compares the value of the Swing options with different number of down-swing rights, produced by the Spike and the Mean-Reverting Models, during winter and spring.

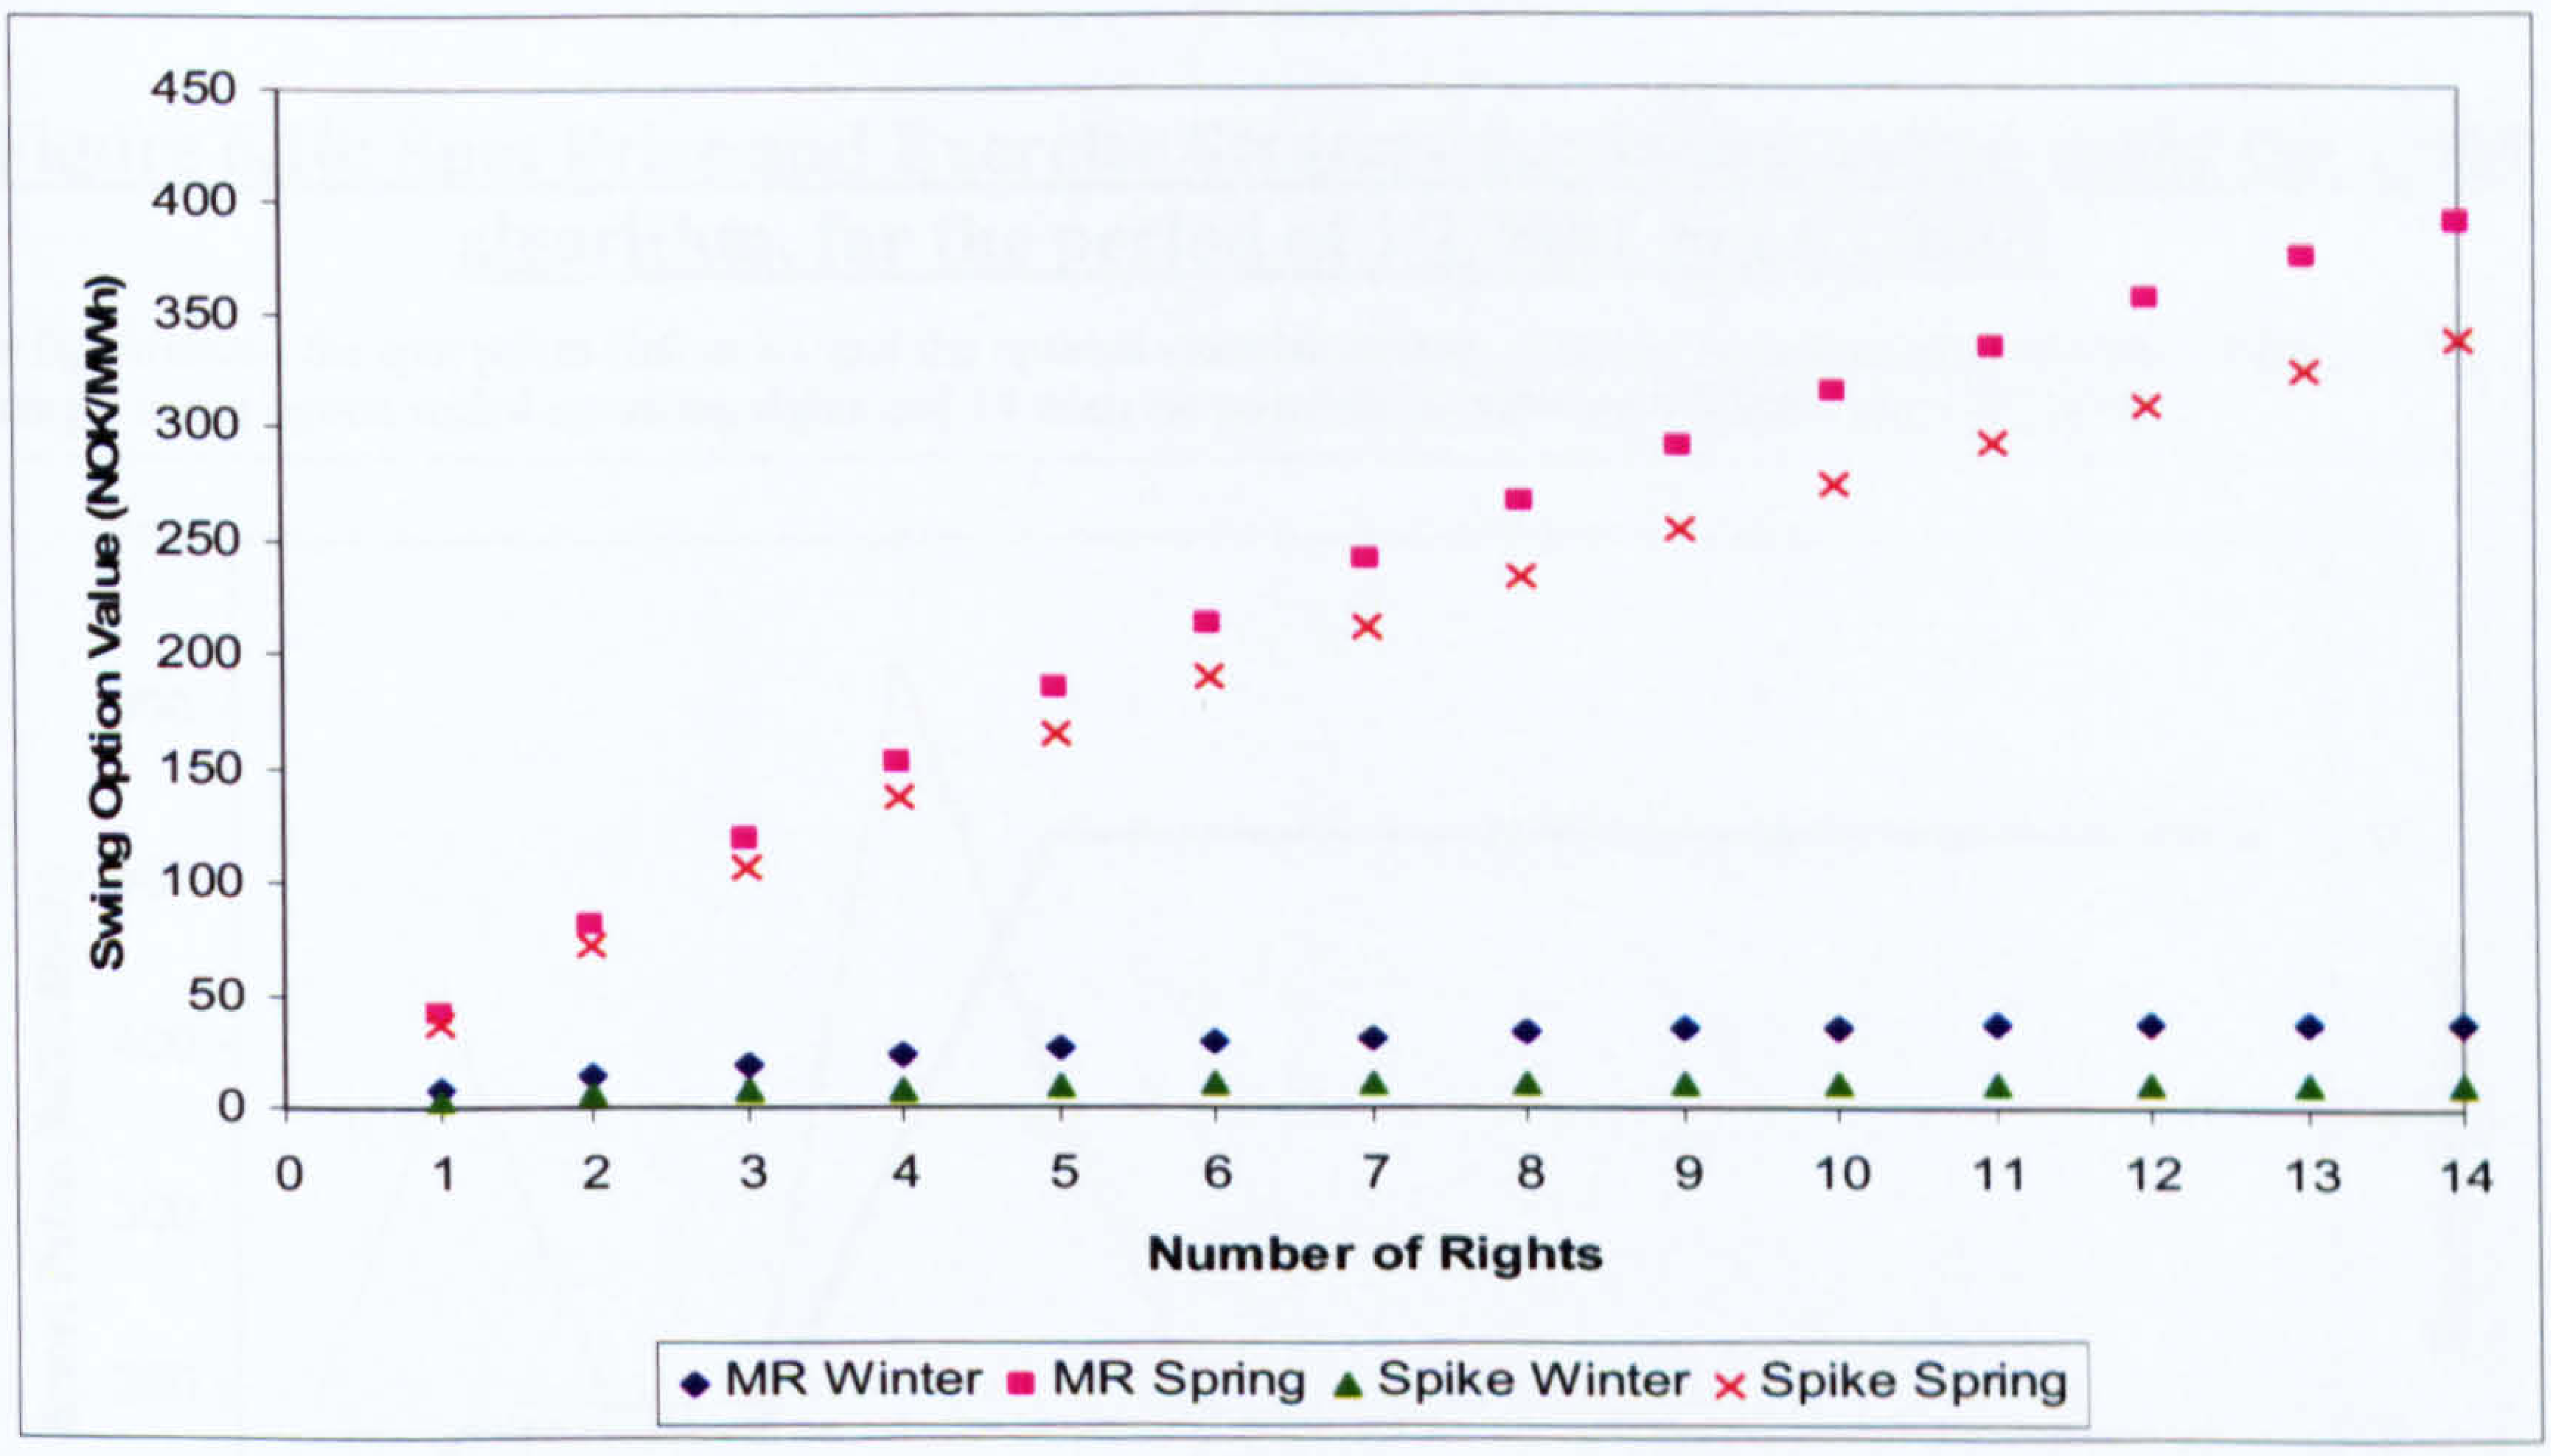

6.6.3 Exercise Strategy Using LSM

Finally as an illustration of the exercise strategy using the LSM algorithm, we consider again our swing option with 4 upswings rights. We present this strategy using actual data, for a spiky period such as February 2001. Thus, our swing option has 14 exercise date opportunities starting from February 1st 2001 and the final opportunity is on February 14th 2001. The contract is supposed to be agreed on the

31st of January 2001, with an exercise price equal to the spot price on *that date* (207.84), so the option is ATM. Figure 6.10, shows the spot price and the exercise strategy yielded by the LSM algorithm. Both the MR and Spike models yield the same strategy. From the figures it can be clearly seen that only 3 rights are exercised out of a possible maximum of 4. This is mainly due to the fact that the seasonal function *f(t)* starts to decrease at that period as February is the last month of the winter period and we slowly enter spring time. Thus spot prices decrease, making the option to have limited exercise possibilities. However, there are 3 particular dates when spikes occur which take place on exercise date opportunities, 2, 5 and 6. These are the dates, which the LSM algorithm chooses to exercise the option. Thus this analysis gives indication that spikes and seasonality play a crucial part in the exercise strategy as well as the pricing of swing options, and should be taken into consideration when pricing and trading these instruments.

Figure 6.10: Spot Price and Exercise Strategy for Swing option using the LSM algorithm, for the period of 1/2/2001 to 14/2/2001

The figure shows the spot prices (left axis), and the optimal exercise strategy (right axis) produced by the LSM algorithm for pricing a swing option with 4 up-swing rights and 14 exercise possibilities between 1/2/2001 and 14/2/2001.

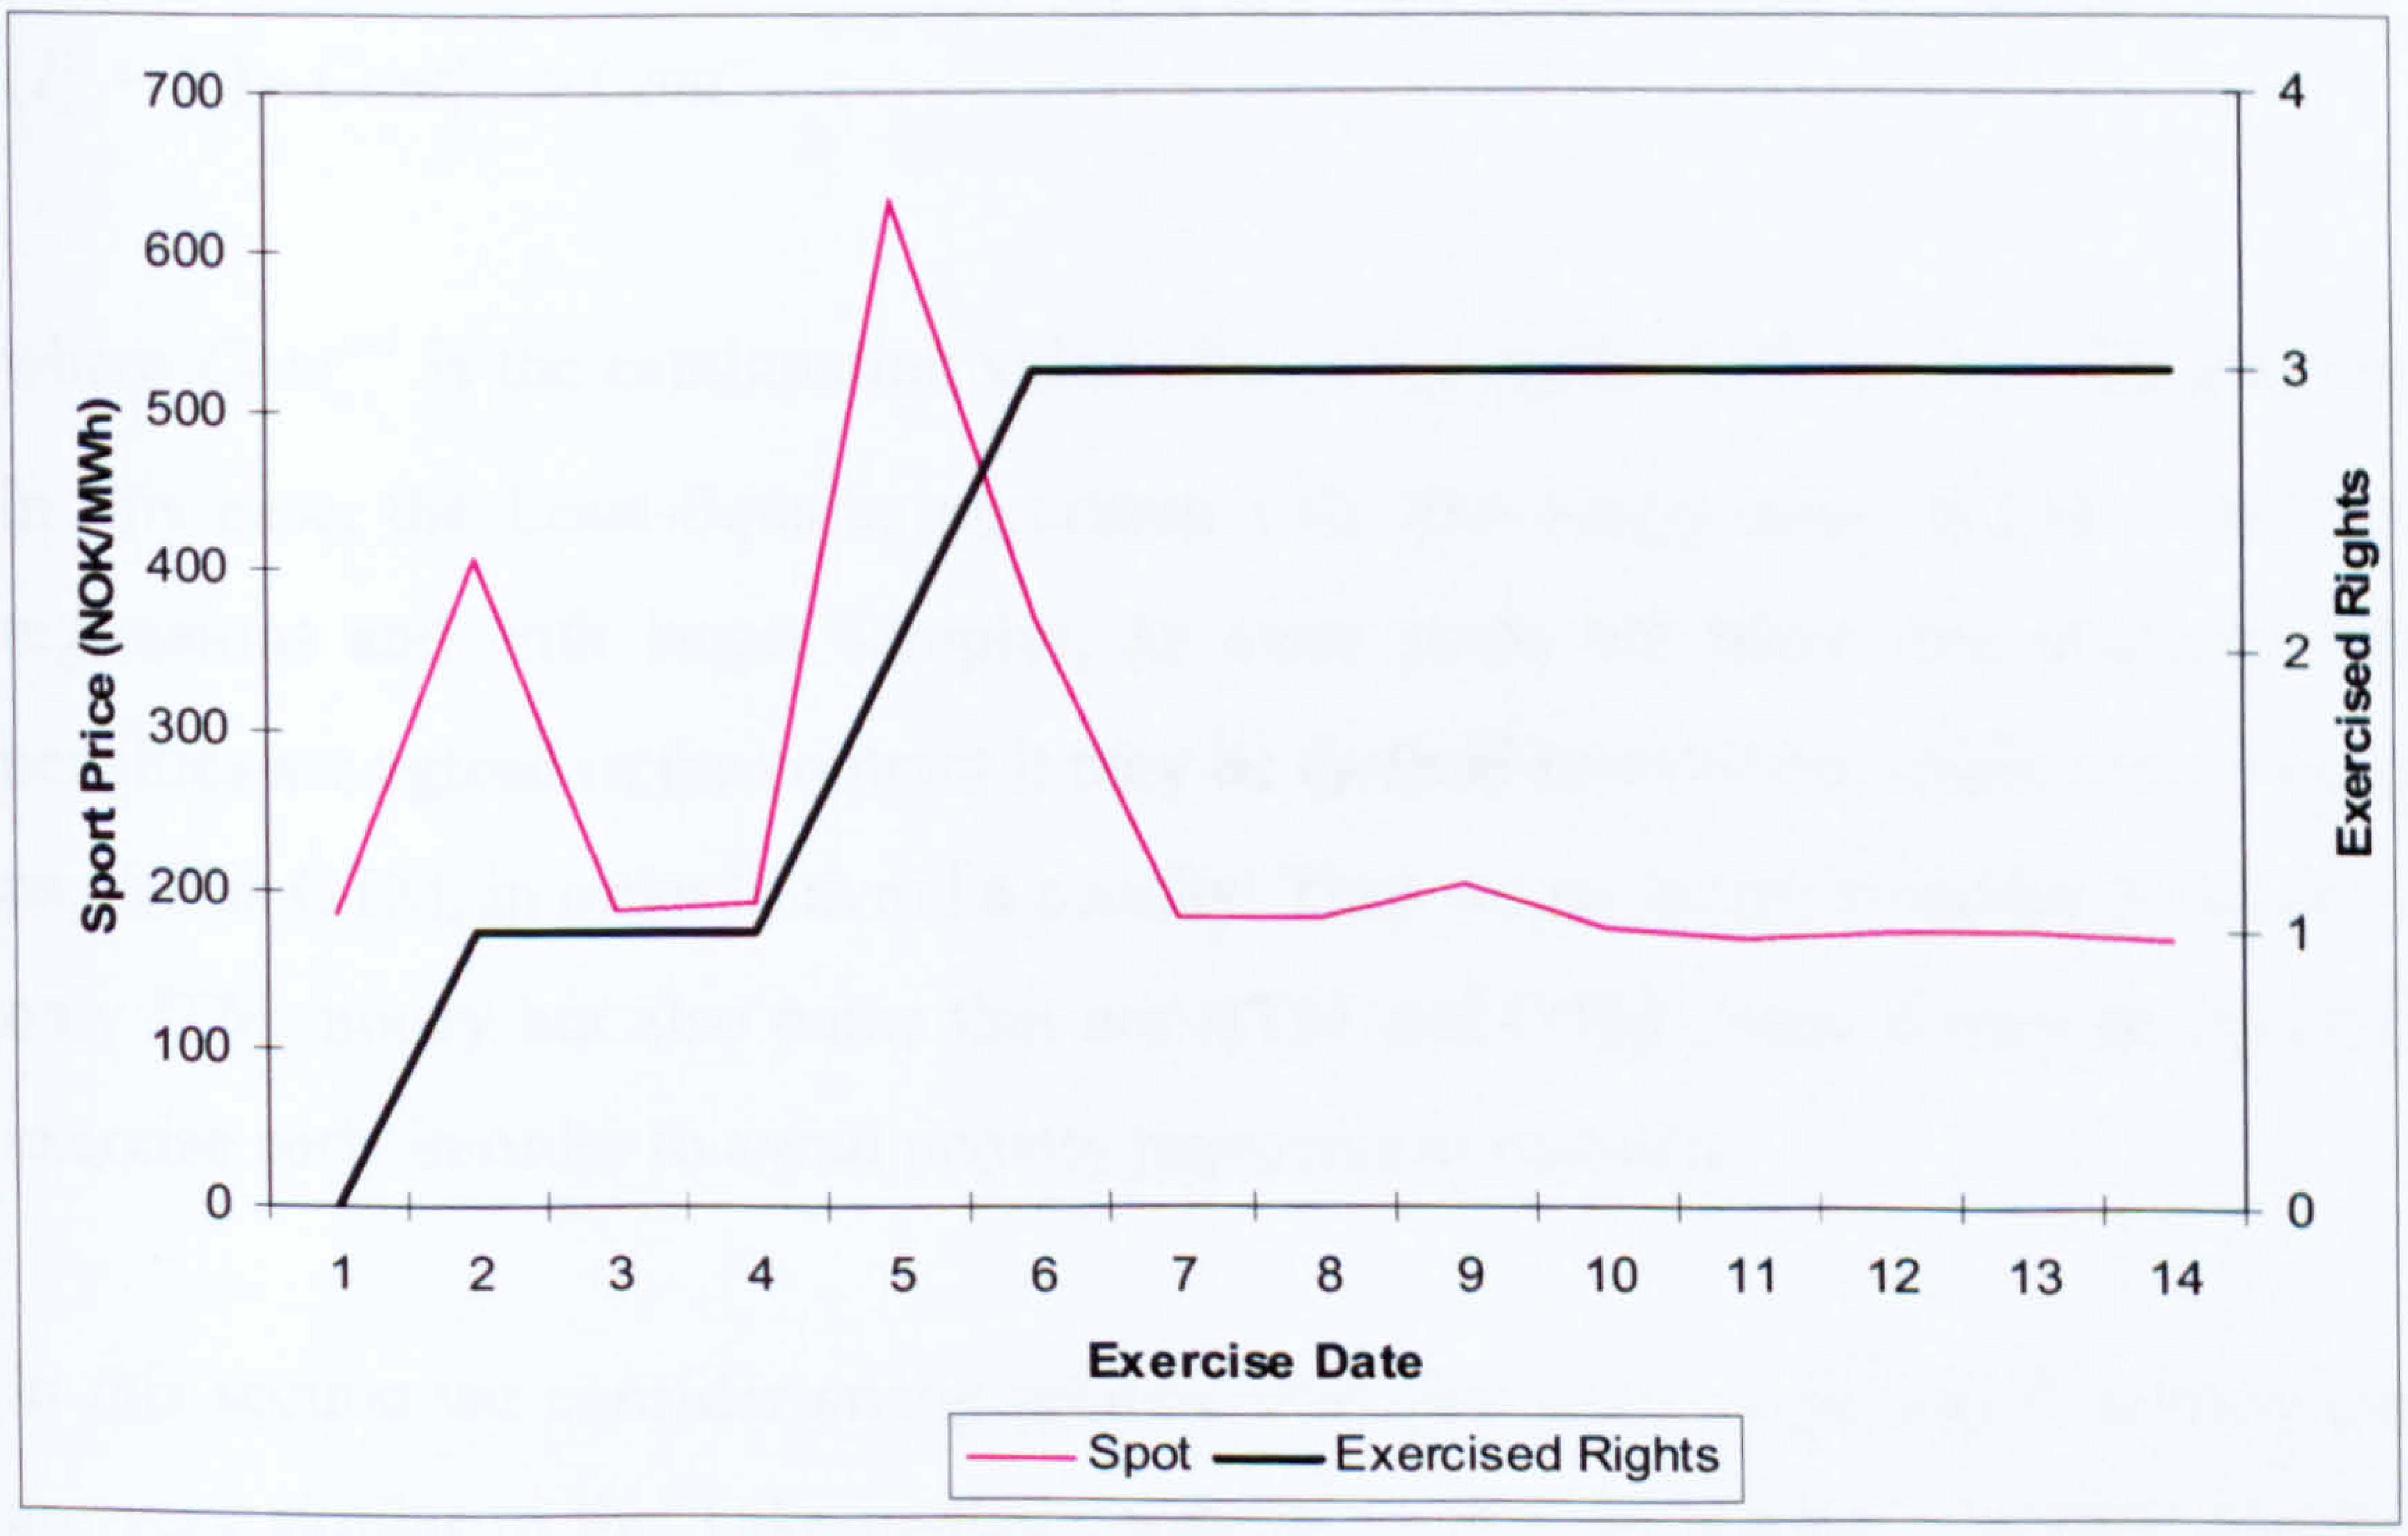

6.6.4 The Effect of Penalties

When valuing swing options with penalties, the valuation problem using the Monte Carlo simulation is slightly different this time. This is caused by the fact that eventhough, e.g. for up-swings early exercise might be possible only if the spot price is above the strike price, in the presence of penalties early exercise might be possible at lower prices than the exercise price, in order to avoid any penalties payable at the maturity date. Thus, if we use the same example as before with 4 up-swings, with a penalty function, ϕ^u the pay-off for 4 swings left unexercised at maturity $T_2 = \tau_{14}$, is $Pay(P) = max\ (P-K-\phi^3,\ -\phi^4)$. At earlier dates, however we have to consider all spot prices generated by the Monte Carlo simulation and thus the least squares is performed for all possible pay-offs, since it might be possible to exercise earlier and avoid possible penalties. Thus, let C^u denote the cash flow matrix for u up-swings exercised, hence each time we have to step forward from $u=0$ to $u=4$. Therefore, early exercise is performed when the following condition is satisfied for all spot price paths:

$$(P_{t,i} - K) + Cont_t^{u+1} > Cont_t^u,$$

where $Cont_t^{u+1}$ is the continuation value of a swing option with $u+1$ swings exercised. In this case, the Least-Squares algorithm will take longer since it has to run more regressions and with larger samples, as more paths are taken into account. When penalties are agreed in the contract it may be optimal to exercise a right even when the pay-off is OTM, in order to avoid a penalty. Thus we no longer consider paths that are only ITM money but also paths that are ATM and OTM, since it may be optimal to exercise early in order to avoid penalty payments at maturity.

In this section we consider swing options with two types of penalty functions which are very similar to the Take-or-Pay contracts in the gas market. Let *Pen* denote the penalty function as a percentage of either the spot or the exercise price. We then consider the following two scenarios:

$$Scenario\ (1):\ \phi(u)=\begin{cases} Pen\left(u^{+}-u_{\tau_n}\right)K & if\ u^{+}>u_{\tau_n} \\ 0 & otherwise \end{cases}$$

$$Scenario\ (2):\ \phi(u)=\begin{cases} Pen\left(u^{+}-u_{\tau_n}\right)P_{T_2} & if\ u^{+}>u_{\tau_n} \\ 0 & otherwise \end{cases}$$

where u^{+} is the maximum number of swing options that can be exercised until the maturity of the option, and u_{τ_n} is the number of exercise rights that have been expected during the life of the contract[50]. Thus the penalty which is paid at maturity is proportional to the number of rights that have been exercised during the lifetime of the swing option. Note also that under the two alternative scenarios, the penalty is either a proportion of the strike price, K, or of the actual spot price at maturity, P_{T_2}. Since an exercise right allows the holder to receive an extra MWh volume of electricity, the maximum volume that the holder can buy from the swing optionality is $V^{+}=4$ and if all the rights are thus exercised given the constraint that the holder cannot exercise more than the maximum u^{+}, the penalty is zero. Hence, in this case the Penalty function is linked with the number of rights, however it could also be linked to the delivered volume as discussed initially in the pricing section of the chapter.

Note that for the pricing of swing options with variable daily incremental swing volumes (i.e. $v^{+}(\tau_j)$ and $v^{-}(\tau_j)$), the LSM algorithm is similar to the one proposed in section 6.5, since whenever it is optimal to exercise early, the holder should purchase (sell) the maximum incremental volume. However, in cases when penalties are imposed, the algorithm has to change, since this time the generalised cash-flow matrix of our algorithm must have four dimensions instead of three, since the optimal exercise decision can be computed by searching over the range of possible purchase (sell) volumes ($v^{-}_{\max}$ to $v^{+}_{\max}$) for the volume which maximises the sum of discounted expectation and the value of the current purchased (sold), as the discounted expectation depends on the current volume purchase (sell) decision, since this affects the volume remaining to be purchased (sold) in the future in order to avoid the penalty payment, as shown by Clewlow and Strickland (2000).

[50] Note that each exercise right gives the owner the option to buy 1MWh of electricity, as in the previous analysis.

Figure 6.11: Swing (up-Swings) Option Value with respect to number of swings and penalty (% of Exercise Price, *K*).

The figure shows the value of swing option prices with respect to the number of up-swing rights and the penalty function as a proportion of the Strike price, under *scenario (1)*.

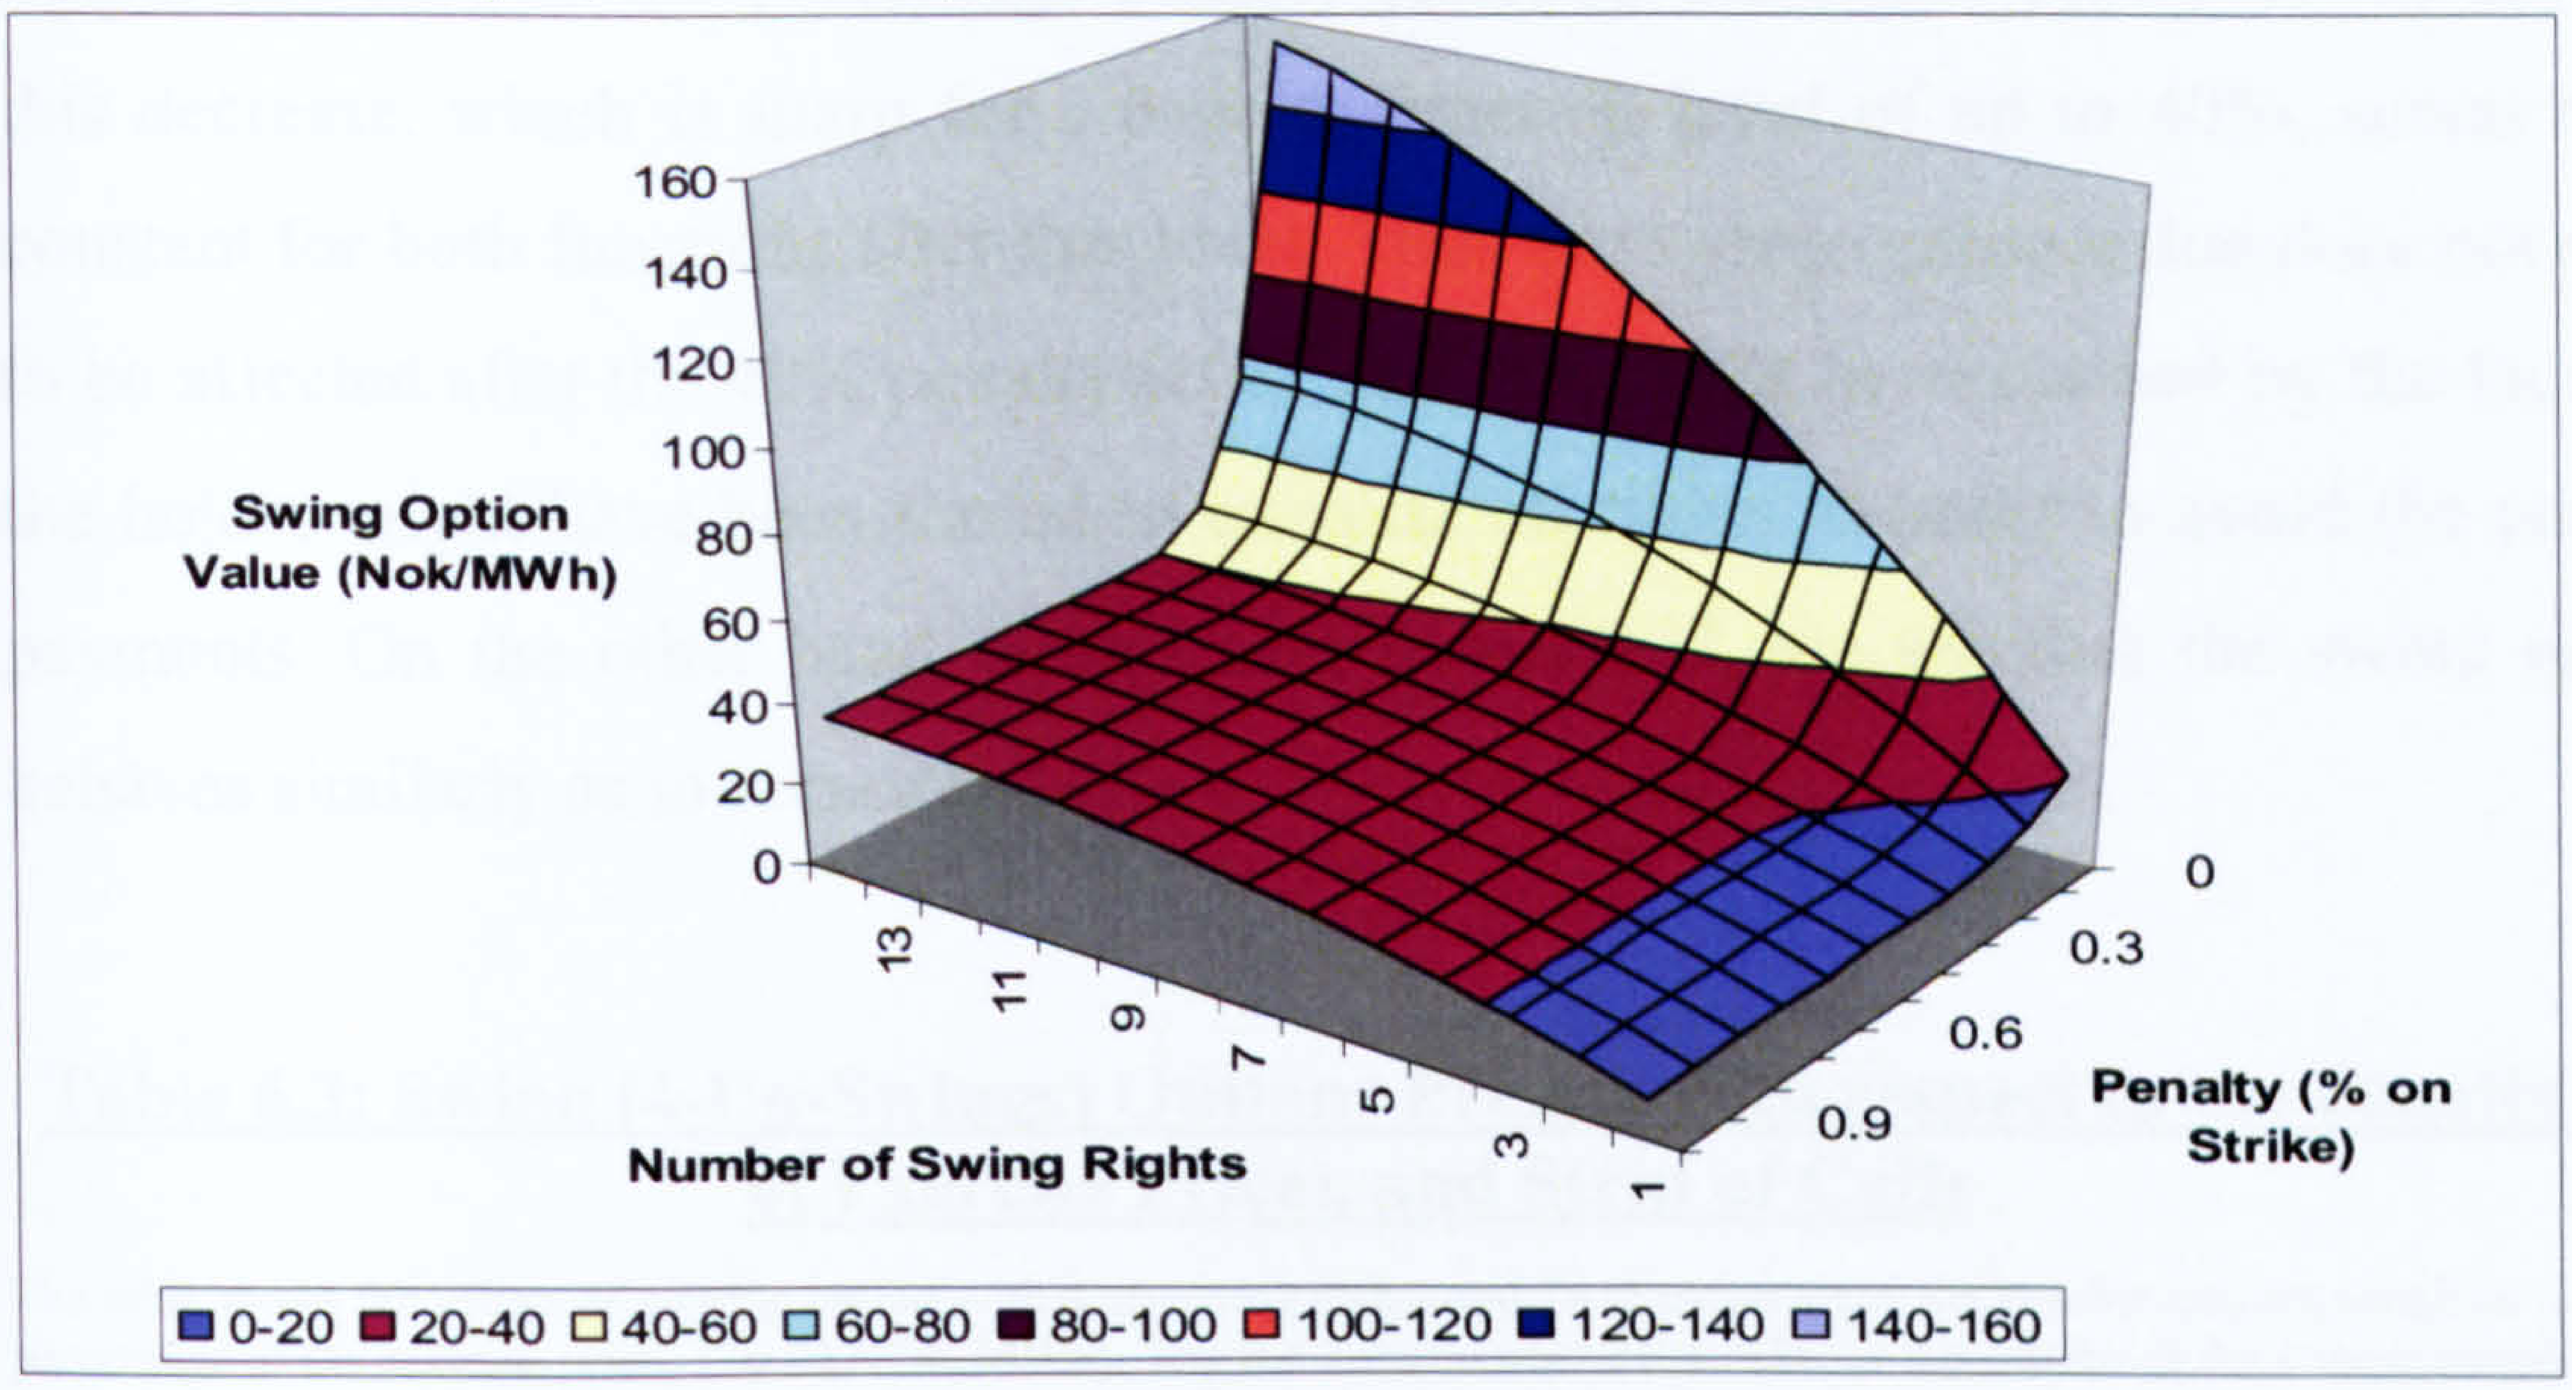

Figure 6.12: Swing (up-Swings) Option Value with respect to number of swings and penalty (% of spot price at Maturity, P_T)

The figure shows the value of swing option prices with respect to the number of up-swing rights and the penalty function as a proportion of the spot price at maturity, under *scenario (2)*.

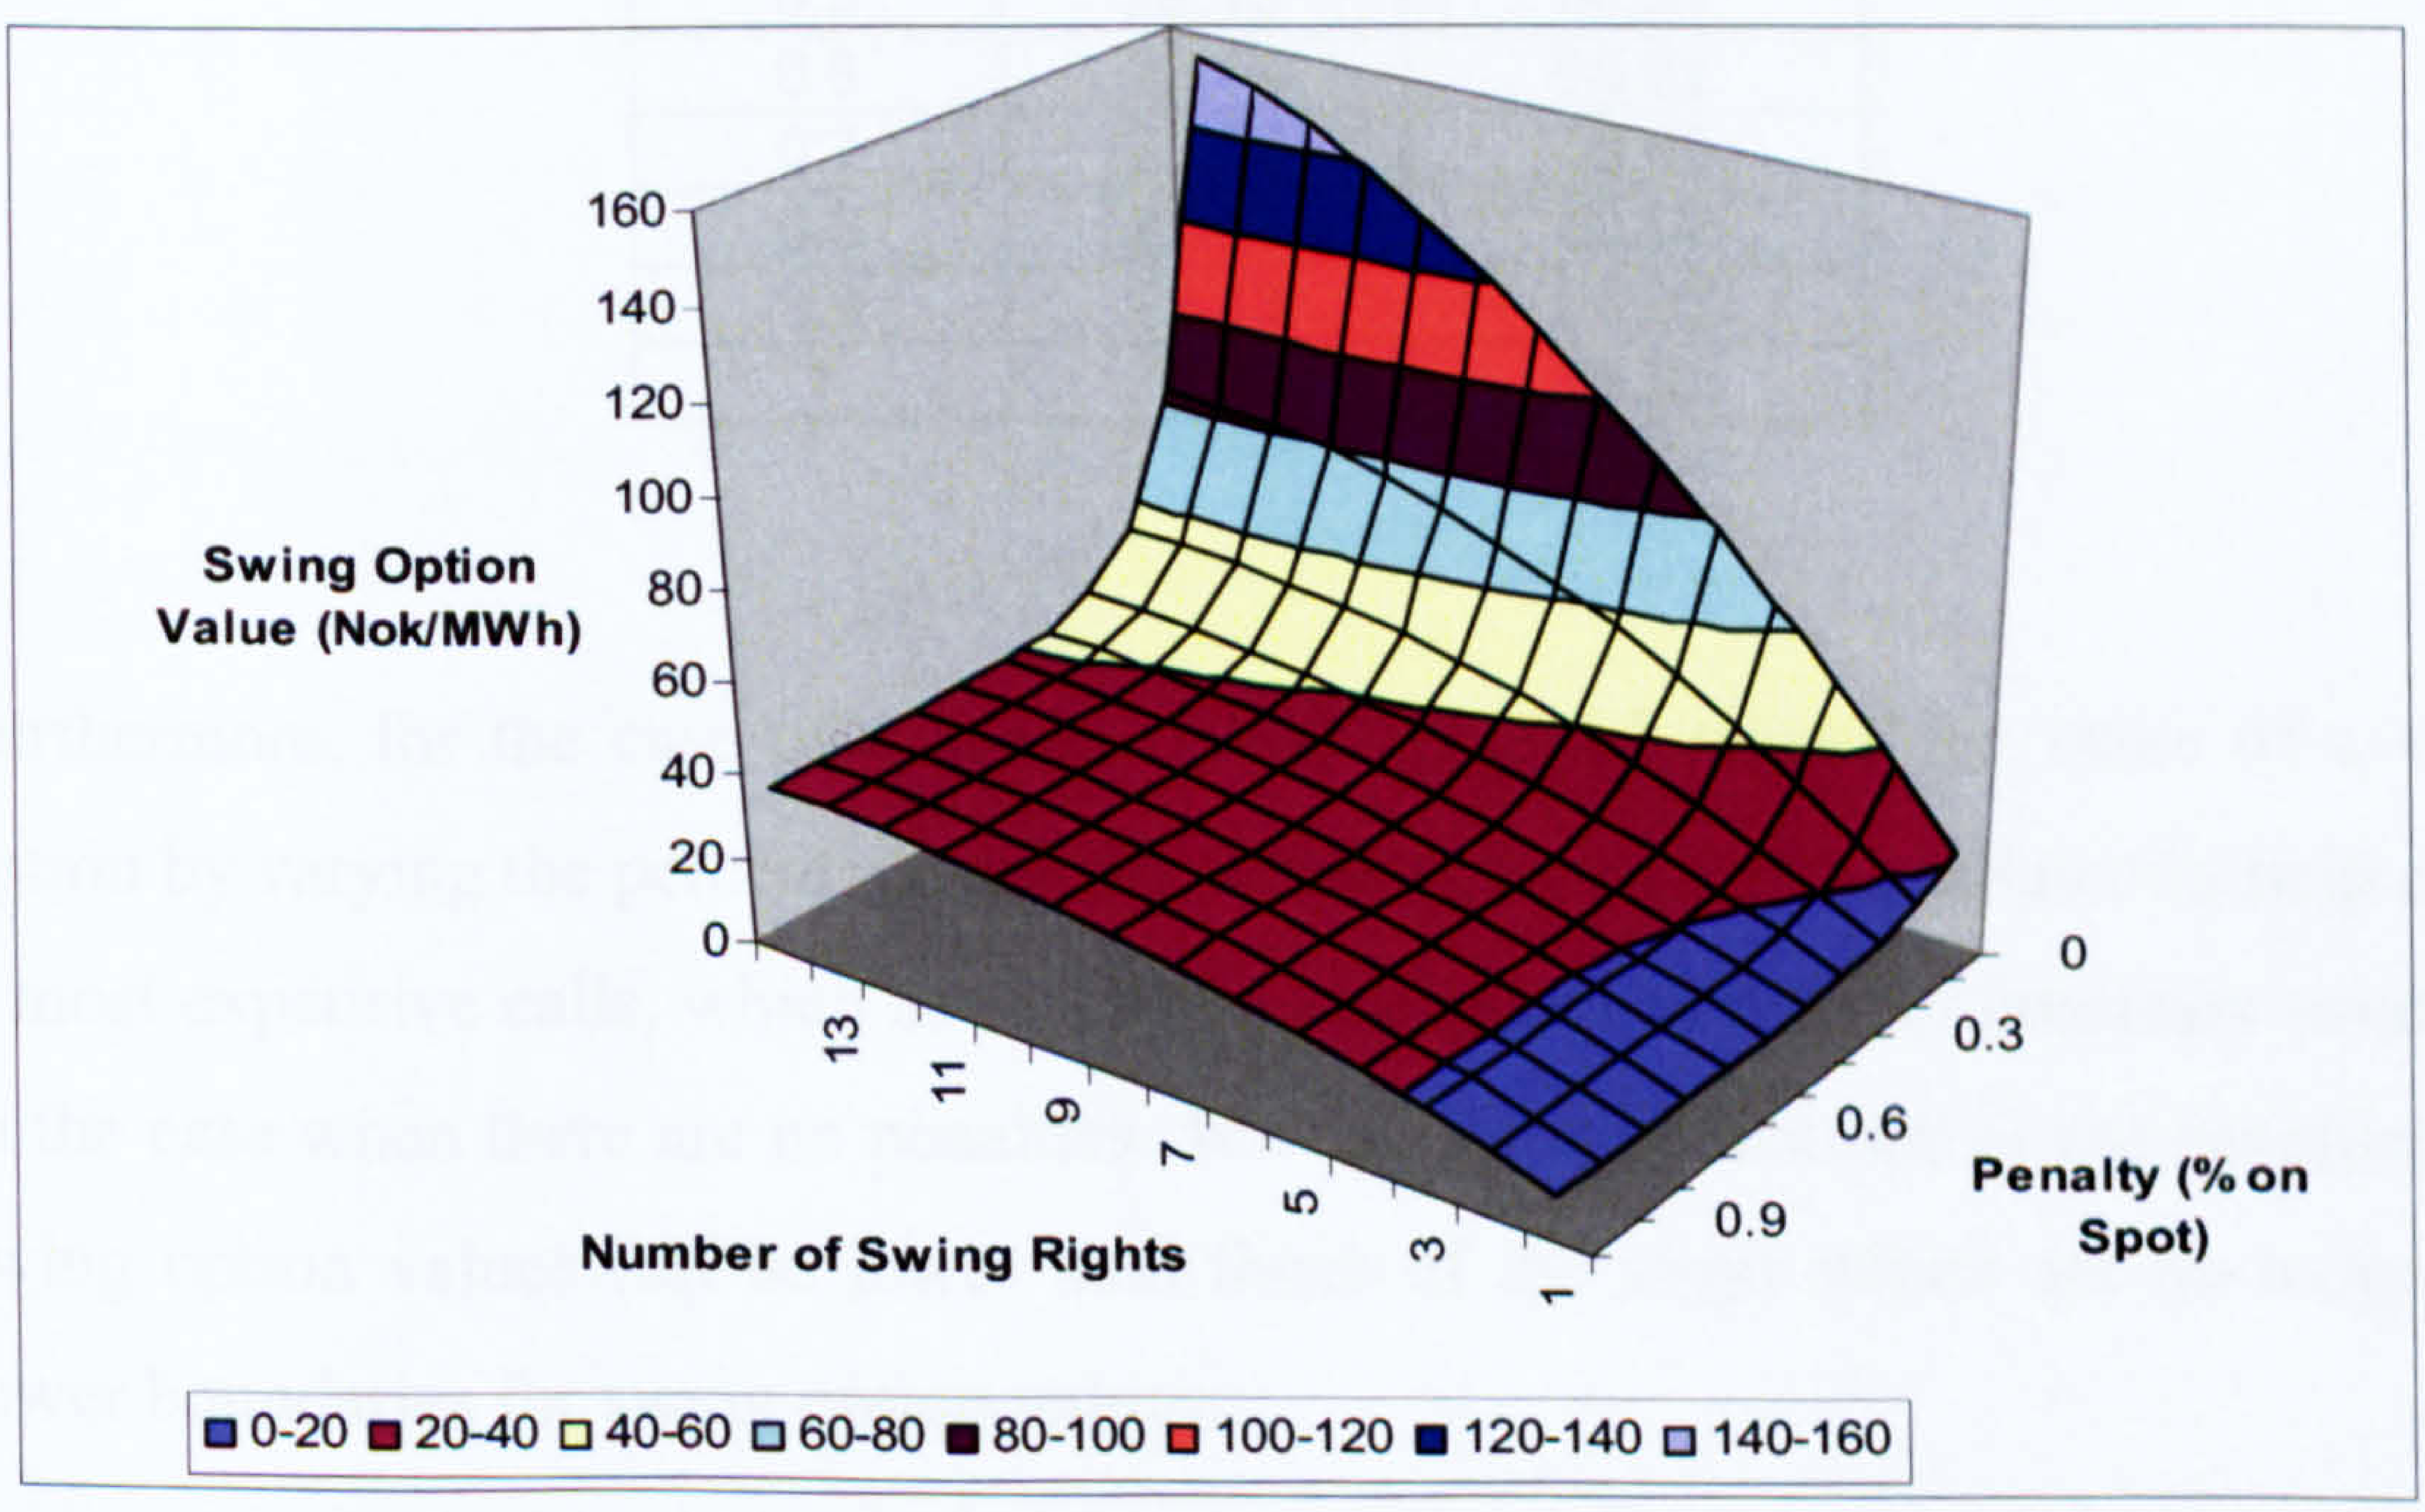

The results under the two scenarios are shown in Figure 6.11 and Figure 6.12, respectively. As expected the swing option values (with varying number of swing rights) decrease as the percentage penalty function increases. However we see that this decrease, which is sharp for a penalty function level of up to 40%, seems to be constant for both functions after this level. Thus the swing option value does not seem to be affected after the 40% penalty level, and this might be explained by the fact that the holder might have been forced to exercise all rights in order to avoid the penalty payments. On the other hand when using *scenario 2*, we see that the swing option behaves similarly as in *scenario 1*.

Table 6.3: Swing (4-Up-Swings) Options Premia with respect to the Penalty (% of Exercise Price), and Strip of Calls

The table shows the values of a swing option with 4 up-swing rights and 14 exercise opportunities for varying penalties (as a percentage of the exercise price, 170.20 NOK.MWh). The last column shows the value of a portfolio of the 4 most valuable European call options within the same period.

Penalty (% on Strike)	Swing Option Value	Strip of Calls
0	65.40	58.92
0.1	39.41	58.92
0.2	27.16	58.92
0.3	22.66	58.92
0.4	22.14	58.92
0.5	22.13	58.92
0.6	22.13	58.92
0.7	22.13	58.92
0.8	22.13	58.92
0.9	22.13	58.92
1	22.13	58.92

Furthermore, for the case of 4 swing rights, Table 6.3 shows the value of a swing option by varying the percentage penalty function, as well as the value of a strip of the 4 most expensive calls, which as explained earlier, act as the low boundary condition in the case when there are no penalties. We can see now that due to the penalties, the swing option values can be lower than those of the strips which are no longer the lower boundaries for swing option values.

6.7 Conclusions

In this chapter of the thesis we examined swing options, which are a very effective hedging tool in the energy markets, in terms of price, demand and supply risk. We explained their pricing methodology using Monte Carlo Simulations which is an extended implementation of the Least-Squares algorithm by Longstaff and Schwartz (2001). In addition we also showed the lowest and highest boundary conditions of swing option values in terms of strips of European options and Bermudan option; however we also showed that these work, only if there are no penalties in the contract.

Furthermore, we performed a sensitivity analysis on the extra premium for up-swing option relative to the premium paid for strips (i.e. the lowest boundary conditions), in order to identify the value drivers of the swing option premium. The analysis gave very intuitive results; increasing the jump intensity and jump sizes increases the extra swing option premia since the infrequent jumps when they occur give a very advantageous profit for early exercise, something that the holders of European options are not able to exploit. On the other hand, by increasing the speed of mean reversion of the jumps there seems to be an increase in the relative extra swing premium, but up to a level of approximately 100. From that level onwards, jumps die out at a much faster rate and hence early exercise seems to be not so much in the benefit of the swing option holder maybe because, he may not be able to act fast enough to exploit all the early exercise possibilities. In terms of the diffusive mean reversion and equilibrium level, the swing option holder has a comparative advantage the fastest the speed of mean reversion and the lowest the equilibrium level, since any deviation from the equilibrium level can be exploited immediately using the early exercise feature. Finally, the higher the diffusive volatility the lower the extra swing option premium, since higher diffusive volatility translates into more possible spot prices in the future and thus early exercise is not optimal.

Furthermore, we examined the swing option values for both the Mean-Reverting and the Spike model and, as expected, the higher possibility of extreme price movements over winter makes the spike model to yield more expensive up-swing option prices whereas the opposite is true for the rest of the seasons when the probability of a spike is low and their expected jump size is even smaller. On the other hand, the Mean-Reverting model always yields higher swing option values, when down-swing are examined, since this time the upward spikes do not benefit the put structure of the option. Finally, we also studied how penalties affect the value of swing options and we found that they decrease very fast up to a certain level. Furthermore, when penalties are imposed European strips stop acting as the lower boundaries of swing options since the holder of the option may be forced to exercise un-economical rights in order to avoid the penalty payment at maturity.

7. CONCLUSION OF THE THESIS AND EXTENSIONS

7.1 Summary of the Findings

The aim of the thesis is to investigate the most significant risks in the electricity spot market, and more specifically the Scandinavian market, and provide models that are able to capture them in a simplistic but accurate way for the purposes of derivatives pricing.

Chapter 1 explains the economic fundamentals of the electricity market and the main drivers of the spot price. More specifically, electricity is defined as a commodity that cannot be stored and cannot be transported in large amounts across different areas due to transmission constrains. Furthermore, demand and supply have to be in equilibrium at all times and, given the fact that electricity is a necessity good, electricity demand is inelastic. Also, using the marginal cost of production of a generator, which is based on the fuel used to generate electricity, one can derive the convex shape of the supply stack and thus explain the inverse leverage effect in electricity spot prices. Moreover, the instantaneous balance of supply and demand in electricity makes the electricity spot prices to follow a seasonal path through out the year due to weather dependence. In case of outages of major generating plants or transmission constraints, spot prices jump to extreme price levels for a very short period of time, and then revert back to their equilibrium level as soon as the problem is fixed. We also explained that electricity prices follow a mean-reverting pattern to an equilibrium level, which is mainly determined by the cost of production and the level of demand. Using these fundamentals we explored the different pricing theories in the construction of the forward prices in commodities. The two main theories rely on the Theory of storage introduced by Kaldor (1939) and Working (1933) and the Theory of risk premium. However, since electricity cannot be stored in an economic and feasible way yet, market participants cannot delta hedge the risks when trading in the derivatives market and thus they have to ask for a premium for every risk they face. Therefore the theory of risk premium seems the most appropriate in the electricity market. Under these circumstances, market participants can use forwards and futures which are the most liquid derivatives, and thus delta hedge other contingent claims in the power market. Following that, we arrive at the conclusion that the most appropriate model in the electricity market, is the one that is able to capture the most significant path and distributional characteristics of the spot market, and at the same find the best Equivalent Martingale Measure that is able to fit the forward curve at every observation date, and thus price the most significant risks in the market.

In Chapter 2 we go through the literature review in terms of modelling and pricing derivatives in the electricity market. We divide the literature review into four parts; the first way of modelling is using the reduced form models that describe the stochastic dynamics of the spot market in a parsimonious way using mean reversion, stochastic volatility and jumps. However there is limited literature in terms of describing the spot dynamics accurately and at the same time finding the best measure that is able to fit the forward market which can be used for hedging. Moreover, the spiky nature of the jumps is not captured by these models and the empirical evidence and applications to any market is limited. The second approach is the fundamental equilibrium approach where, supply and demand relations are modelled and the power process is determined via the solution of a certain optimisation problem. However these models are mainly used for policy making rather than derivatives pricing, since they fail to capture the price dynamics in a robust quantitative manner. The third way of modelling, is via the stochastic behaviour of the whole forward curve, however this way of modelling has not yet been successful since empirical evidence has shown that more than 10 factors are needed to capture 94% of the variation in the power forward term structure. Furthermore, since the forward contracts entail a delivery period, and due to the spiky nature of the jumps in the spot, forward prices do not jump. Finally we have the Hybrid models, where the stochastic behaviour of the fundamental drivers is modelled, and then using the fundamental equilibrium approach, the dynamics of the power prices are produced. These models seem to be the most appealing, since every source of risk is priced and hedged using the corresponding derivatives market. However they have a major drawback, as they rely on detailed data that are not even available to most market participants. Finally we explain our contribution to the literature and justify our approach in comparison with the other approaches.

In Chapter 3 we explain the different types of market structures that generally exist in electricity, and then gave a short summary of the Scandinavian electricity market and more specifically Nord Pool, in terms of its history and its organisational bodies. Then we provide empirical evidence of the stylised facts of electricity spot prices and the shape of the forward curve in the Nord Pool market. Our main findings are that electricity spot prices are highly volatile, and the major contribution to that are the short-lived spikes, caused mainly by outages or transmission failures. We also evidence a distinct pattern between different seasons and days of the year (working and non-working days) mainly driven by electricity demand that in turn is affected by human activity. This seasonal pattern is also taken into account when valuing derivatives on the underlying electricity price. Our analysis also shows the spot price in Scandinavia is mean-reverting, and that the equilibrium level is strongly dependent on the level of water in the reservoirs, due to the high share of hydro-generation in the market.

In Chapter 4 all of the above stylised facts are incorporated in reduced form models, that are tractable, easy to estimate, and can provide closed-form solution for derivative prices. First we introduce the spike model, which accounts for the different speeds of mean reversion between normal and spiky shocks in the market. Our proposed model improves significantly the fit between theoretical and observed forward prices as the spikes in the model are discounted at a very fast rate due to their high speed of mean reversion. Consequently, spikes in the spot market do not spill-over to the forward market, a finding that is consistent with the actual patterns observed in the market. Furthermore, our proposed model also accounts for seasonality in the risk premium, which reflects the increasing need to hedge against spikes when their probability of occurrence is higher during the colder winter months. Moreover, we extend the spike model to a Regime Switching Spike model that incorporates two separate regimes to distinguish between periods of high and low water levels in the reservoirs, which translate to two different equilibrium levels in the spot prices. Another extension of the spike model is the Three-Factor spike model which allows the equilibrium level not to be dependent just on the level of water in the reservoirs but by a general factor which captures risk stemming from long-run changes in power supply and demand. The main disadvantage of the latter model is that all three stochastic variables are unobservable, and two of them have to be estimated using the Kalman Filter, therefore we have to use long-maturity illiquid forward contracts. The performance of our proposed models is compared to that of other models proposed in the literature such as a simple mean-reverting (MR) and a mean-reverting jump diffusion model (MRJD). The models are tested in term of fitting the observed term structure, as well as in terms of capturing the trajectorial and distributional properties (where closed-form solutions are derived for each moment) of the spot prices. In particular both the Spike and Regime Switching Spike models provide more realistic simulated price paths and a better fit to the forward term structure than the MR and MRJD models. Overall, it seems that the proposed models' trade-off between complexity and simplicity is well behaved, as they account for the most significant risks in the highly volatile power market, in a parsimonious and intuitive way.

Chapter 5 looks at the implications of the models proposed in Chapter 4, in terms of options pricing. Option pricing formulas are derived under each model; in particular we find that under the simple MR model the closed-form solution is very similar to that of Black and Scholes (1973) where the underlying is the discounted deseasonalised forward price and the volatility is replaced with that of the forward's. On the other hand, when introducing jumps in the model like in the spike model, only semi-closed form solutions can be derived. Finally, the pricing formulas showed that the moneyness of the option in the electricity market is affected by both the exercise price and the deterministic factor. The analysis on the model-implied volatilities showed that even a mean-reverting process displays volatility skews but mainly for ITM options. However the presence of jumps, displays volatility skews for OTM options depending on the sign of the mean jump size. We also show that mean reversion reduces the volatility smile as time to maturity increases. Furthermore, the jump size volatility and jump intensity are the main factors affecting the curvature of the volatility smile. Moreover, we provide evidence that for a desired jumpiness level, the jump size volatility affects the volatility smile more than the jump intensity. Comparing the different proposed models of Chapter 4 we see that the mean-reverting model seems to converge to an option price very fast as time to maturity increases, and from a point onwards the discounting effect of the interest rate starts playing a more important role. We also explored the effect of the water uncertainty in the prices of European options; More specifically, if we are in a state where there is enough water in the reservoirs, we see that call option prices increase with time to maturity consistently and in this case, the premium paid is mainly driven by the increase of the probability that we may end up in a state where the reservoir level is low. The opposite effect happens when we start from a state where the reservoir levels are low, and thus we have a downward sloping option price term structure. The same analysis was performed using the three-factor spike model, however in this case the long-run factor, which displays a GBM behaviour for long maturity options, increases the longer the time to maturity. Finally we performed a sensitivity analysis, using the spike model, on Asian call options. Since no closed-form solution is available and also the model displays jumps, the most appropriate method is the Monte Carlo simulation. The sensitivity analysis shows that Asian call option prices increase with jump intensity and that the jump size starts playing a more significant role at high intensity levels since the jumps are no longer averaged out. The effect of the speed of mean reversion depends on the equilibrium level; if the risk neutral equilibrium level is above the spot, increasing the speed of mean reversion increases the value of the option, but if the risk-neutral equilibrium level is below the spot, increasing the speed of mean reversion decreases the value of the Asian option. On the other hand, the spike speed of mean reversion plays a vital role, as it may give unrealistic Asian option prices if its coefficient is too low due to the fact that any jump that may occur in the market will last for a long time and in turn affect the price of the Asian option as it is no longer averaged out.

Finally in Chapter 6 of the thesis we examine swing options, which are effective hedging tools in the energy markets, in terms of price, demand and supply risks. We explained the pricing framework of swing options in terms of stochastic optimal control theory. Their pricing is implemented via the use of and an extended Least Squares Monte Carlo Simulations algorithm by Longstaff and Schwartz (2001) and Dörr (2003). We define the lowest and highest boundary conditions in terms of European strips and Bermudan option, which hold when there are no penalties in the contract. Furthermore, we perform a sensitivity analysis on the incremental swing premium relative to European strips, for up-swing contracts. The analysis gave very intuitive results; first, increasing the jump intensity and jump sizes increase the incremental swing option premia since the infrequent jumps, when they occur, give a very advantageous profit for early exercise, something that the holder of European options is not able to exploit. On the other hand, increasing the speed of mean reversion of the jumps seems to increase the incremental swing premium, but up to a certain level at which the opportunity of early exercise cannot be taken since the holder cannot act as fast as the spike may last. In terms of the diffusive mean reversion and the equilibrium level, the swing option holder has a comparative advantage over the holder of European strips as the speed of mean reversion increases and the equilibrium level decreases, since the swing option holder has the opportunity to exploit any opportunities that may arise from any deviations of the spot price from the low equilibrium level. Finally, the higher the diffusive volatility the lower the incremental swing premium, since any increase in the diffusive volatility translates to a greater range of spot prices in the future and thus early exercise is not optimal. Last, we examined the swing option values for both the mean-reverting and the spike models; first, we find that for up-swing options the high possibility of big price movements over the winter in the spike model overweighs the bigger diffusive volatility in the mean-reverting model and thus the swing option price is higher; of course the opposite happens for the rest of the seasons since the probability of spikes is very small. On the other hand, the Mean-Reverting model always gives higher down-swing option values, since this time the upward spikes do not benefit the put payoff structure of the option.

7.2 Suggestions For Further Research

Through the development of the proposed models and the subsequent empirical analysis we have identified some satisfactory results. However there are a number of potential extensions which can to some degree complement the study and consequently shed some light on the issues not covered in this thesis.

First of all, a more accurate representation of the spike model will be to include stochastic volatility in the normal factor along the lines of Heston (1993). The inclusion of this factor will induce more complexity in the model in terms of finding closed-form solutions for derivatives pricing, running Monte Carlo simulations (due to the squared root process) as well as calibration, however it may improve the fit to the forward term structure and may also capture the moments of the price distribution more accurately.

Another extension is the test of the models to option prices, when the market becomes liquid, which will be very important for market participants. On the other hand, a more practical extension is to test different strategies for hedging derivatives and more particular for hedging swing options. As we showed in chapter 5, hedging of options becomes very difficult when a model contains jump since the hedging parameters depend on the size of the jump, however one can develop hedging strategies based on a variance minimisation approach as in Kennedy (2006). This will be very important for market participants and especially investment banks, when they sell derivatives in the power market and they need a hedging strategy in order to replicate the product sold.

Last but not least, one may use the proposed models and apply them for the valuation of spread options, which are particularly important for the hedging of power plants that use fossil fuels for power generations (one example is the spark spread which can be used for the hedging of Gas fired plants). This is also important for market participants as they need to hedge for periods when the costs of running the generators are above the price of electricity, and thus they may end up without a profit for that period.

REFERENCES

Alizadeh, A. H. and Nomikos, N. K. (2002): "Efficiency of Electricity Futures Prices: Evidence from the Nord Pool Exchange", Working Paper, Cass Business School.

Alvarado, F. L., and Rajaraman, R.: "Understanding Price Volatility in Electricity Prices", Proceedings of the 33rd Hawaii International Conference on System Sciences – 2000.

Andersen, K., (2000), "A Simple approach to the Pricing of Bermudan Swaptions in the Multi-Factor Libor Market Model", *Journal of Computational Finance*, 3, 5-32.

Audet, N., Heiskanen, P., Keppo, J. and Vehviläinen, *I. (2003): "Modelling* of electricity forward curve dynamics", working Paper.

Babbs, S.H. and Nowman, K.B. (1999), "Kalman filtering of generalized Vasicek term structure models", *Journal of Financial and Quantitative Analysis,* Vol. 34, 115-130.

Ball, C. A. and Torous W.N., 1983, "A simplified Jump Process for Common Stock Returns", *Journal of Financial and Quantitative Analysis*, Vol. 18(1), 53-65.

Ball, C. A. and Torous W.N., 1985, "On Jumps in Common Stock Prices and their Impact on Call Option Pricing", *Journal of Finance*, Vol. 40, 155-173.

Barlow, M. T., (2002): "A Diffusion Model for electricity Prices", *Mathematical Finance*, 12, No.4, 287-298

Barlow, M., Yuri, G. and Lai, M., (2004): "Calibration of Multifactor Models in Electricity Markets", *International Journal of Theoretical and Applied Finance*, vol. 7, 101-120.

Barbieri, A., and Garman, M.B., (1996), "Putting a Price on Swings", Energy and Power Risk Management, 1 (6).

Barone-Adesi, G. and Gigli, A. (2002), "Electricity Derivatives", Working paper

Bastes, D.S., (1997), "The Skewness Premium: Option Pricing under Asymmetric Information", *Advances in Futures and Options Research*, 9, 51-82.

Benth, F.E and Koekebakker, S., (2005), "Stochastic Modelling of Electricity Contracts", Working Paper.

Benth, F.E., Kallsen, J. and Meyer-Brandis, J., 2005, "A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modelling and derivatives pricing", *Applied Mathematical Finance*, forthcoming.

Bessembinder, H. and Lemmon, M.L. (2002), "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets", *Journal of Finance* 57, 1347-82.

Bhanot, K. (2000), "Behavior of Power Prices: Implications for the Valuation and Hedging of Financial", *Journal of Risk,* 2 , 43-62.

Bjerksund, P., Rasmussen, H. and Stensland, G., (2000): "Valuation and risk management in the Nordic electricity market", Working paper, Institute of Finance and Management Science, Norwegian School of Economics and Business Administration.

Black, F. and Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities", *Journal of Political Economy*, 673-54.

Bodily, S. and Del Buono, M. (2002):"Risk and reward at the speed of light: a new electricity price model", EPRM

Borovkova, S., Permana, F. J., 2004, "Modelling electricity prices by the potential jump-diffusion", *Proceedings of the Stochastic Finance 2004 Conference.*

Branger, N., (2004), "An Option Pricing Anatomy", Working Paper

Brennan, M. (1991), "The Price of Convenience and the Valuation of Commodity Contingent Claims", in D. Lund and B. Oksendal (eds.), *Stochastic Models and Option Models*, Elsevier Science Publishers.

Burger, M., Klar, B., Muller, A., Schindlmayr, G., 2004, "A spot market model for pricing derivatives in electricity markets", *Quantitative Finance* 4, 109-122.

Carr, P., and Madan, D., (2001), "Optimal Positioning in Derivatives Pricing", Quantitative Finance, 1, 19-37

Cartea, A. and Figueroa, M. G., 2005, "Pricing in electricity markets: A mean reverting jump diffusion model with seasonality", *Applied Mathematical Finance*, forthcoming.

Chacko, G. and Das, S., (2002), "Pricing Interest Rate Derivatives: A General Approach", *Review of Financial Studies*, 15(1), 195-241.

Clement, E., Lamberton, D., and Protter, P., (2002), "An Analysis of a Least Squares Regression method for American Option Pricing", *Finance and Stochastics* 6, 449–472.

Clewlow L., Strickland C., and Kaminski V. (2001): "Extending Mean Reversion Jump diffusion", EPRM

Clewlow L., Strickland C., and Kaminski V. (2000): "Jumping the gaps", EPRM

Clewlow, L. and Strickland C., (1999a): "A multifactor model for energy derivatives", Quantitative Finance Research Group, Working paper, University of Technology, (Sydney).

Clewlow, L. and Strickland C., (1999b): "Valuing energy options in a one factor model fitted to forward prices", Quantitative Finance Research Group, Working paper, University of Technology (Sydney).

Clewlow, L., and Strickland, C. (2000), *Energy derivatives: Pricing and Risk Management*, Lacima Publications.

Clewlow, L., and Strickland, C. (2001), "Risk Analysis of Swing Contracts", Energy and Power Risk Management.

Cortazar, G. and Schwartz E. S., (1994),: "The valuation of commodity contingent claims", *The Journal of Derivatives*, 27-39.

Cortazar, G. and Schwartz, E. (2003),: "Implementing a stochastic model for oil futures prices", *Energy Economics,* vol. 25, 215-238

Cortazar, G., Naranjo, L., and Schwartz, E. (2003), "Term Strucutre Estimation in Lpow Frequency Transaction Markets: A Kalman Filter with Incomplete Panel Data", Working Paper.

Cox, Ingersoll, and Ross (1985), "An intertemporal general equilibrium model of asset prices", *Econometrica,* vol. 53, 363-384.

Cox, J. C., Ingersoll, J. E. and Ross S. A., (1981): "The relation between forward prices and futures prices", *Journal of Financial Economics* 9, 321- 346.

Curran, M., (1992), "Beyond Average Intelligence", Riskm 5, 10, 60.

Dai, Q. and Singleton, K.J. (2000), "Specification analysis of affine term structure models" *Journal of Finance,* Vol. 55, N° 5, 1943-1978.

Das, S.D., (2002),"The Surprice Element: Jumps in Interest Rates", *Journal of Econometrics,* 106, 27-65.

De Jong, F. (2000), "Time-series and cross-section information in affine term structure models" *Journal of Business & Economics Statistics,* Vol. 18, N° 3, 300-314.

De Jong, F. and Santa-Clara, P. (1999), "The dynamics of the forward interest rate curve: A formulation with state variables", *Journal of Financial and Quantitative Analysis,* Vol. 34, 131-157.

Deng, S.J., (2001), "Stochastic models of energy commodity prices and their applications: mean-reversion with jumps and spikes", UCEI, PWP-073.

Deng, S.J., and Oren, S.S., (2006), „Electricity Derivatives and Risk Management", *Energy,* 31, 940-953.

Diebold, F. X. and Mariano, R. S (1995), "Comparing Predictive Accuracy", *Journal of Business & Economic Statistics,* Vol. 13 Issue 3, p253-263.

Dörr, U., (2003), "Valuation of Swing Options and Examination of Exercise Strategies by Monte Carlo Techniques", University of Oxford Thesis.

Dixit, A.K. and Pyndyck, R.S. (1994) "Investment Under Uncertainty", Princeton University Press

Du Y. (2003): "Exploring Option Pricing with mean-reversion jump diffusion", EPRM

Duffie D.(1999), "Volatility in Energy Prices", in *Managing Energy Price Risk* (second edition), RISK books, London.

Duffie, D., Pan, J. and Singleton, K. (2000) "Transform analysis and asset pricing for affine jump-diffusions", *Econometrica,* vol. 68(6), 1343-1376.

Elliott, R.J. and Sick, G.A., (2003), "Modelling Electricity Price Risk", Working Paper

Escribano, A., Peña J.I. and Villaplana P., (2002) "Modeling electricity prices: international evidence", Universidad Carlos III, working paper 02-27.

Ethier, S., and Kurtz, T., 1986, "Markov Processes, Characterization and Convergence", New York: John Willey & Sons

Eydeland, A. and Geman H., (1998), "Pricing power derivatives", *Risk*, October, 71-73.

Eydeland, A., and Wolyniec, K. (2002), "Energy and Power Risk management: new developments in modelling, pricing and hedging" (Wiley, Chicago).

Geman, H. and Roncoroni, A., 2006, "Understanding the Fine Strucutre of Electricity Prices", *Journal of Business*, Vol. 79, No. 3.

Geman, H. and Vasicek, O., 2001, "Forwards and Futures contracts on non-storable commodities: the case of electricity", *Risk*, August.

Geman, H., 2005, "Commodities and commodity derivatives: Modelling and Pricing for Agricalturals, Metal and Energy" (Wiley Finance).

Geman. H., Karoui N. E., and Rochet J.C. (1995): "Changes of Numeraire, changes of Probanility Measure, and Option Pricing," *Journal of Applied Probability trust*, 32, 443-458.

Gendra, T.W., (1991), "Optional Forward Contracts for Electric Power Service Contracts", PhD Thesis, University of California, Berkley

Gendra, T.W., (1992), "Markets and Pricing for Interruptible Electric Power Transactions", IEEE Trans Pwer Syst.

Gibson, R. and Schwartz, E.S., (1990): "Stochastic Convenience Yield and the Pricining of Oil Contingent Claism", *Journal of Finance*, 30, 958-976.

Gibson, R., and Schwartz, E. (1990), "Stochastic Convenience Yield and the Pricing of Oil Contingent Claims", *Journal of Finance*, 45 (July), 959-976.

Glasserman, P., (2004), "Monte Carlo Methods in Financial Engineering (Stochastic Modelling and Applied Probability)", Springer.

Glasserman, P., Boyle, M., and Broadie, M., (1997), "Monte Carlo Methods for Security Pricing", *Journal of Economic Dynamics and Control*, 21, 1267-1321.

Glasserman, P., and Broadie, M., (1997), "Pricing American Style Securities Using Simulation", *Journal of Economic Dynamics and Control*, 21, 1323-1352.

Glasserman, P., and Broadie, M., (1997), "Them Valuation of American Options on Multiple Assets", *Mathematical Finance*, 7, 241-286.

Greene, W.K., 1993, *Econometric analysis*, 2nd ed., (Prentice-Hall, NJ).

Hamilton (1994), "Time Series", Princeton University Press.

Harrison, J.M., and Kreps, D.M. (1979), "Martingales and arbitrage in multiperiod securities markets", Journal of Economic Theory, 20, 381-408.

Harvey, A.C, (1989), *Forecasting Stuctural time series models and the Kalman Filter*, (Cambridge University Press).

Heath, D., Jarrow R. and Morton A., (1992): "Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation", *Econometrica*, 77-105.

Heston, S. (1993). "A Closed-Form Solution of Options with Stochastic Volatility with Applications to Bond and Currency Options", *Review of Financial Studies*, 6, 327-343.

Hjalmarsson, E., (2003), "Does the Black-Scholes Formula Work for Electricity Market? A Non-parametric Approach", Working Paper

Huisman, R. And Jong, C., (2002), "Option Formulas for Mean-Reverting Power Prices with Spikes", Working Paper

Huisman, R. And Manhieu, R., (2003), "Regime Jumps in Electricity Prices", *Energy Economics*, 25, 425-434.

Hull J. and White A., (1993), "One-factor interest-rate models and the valuation of interest-rate derivative securities", *Journal of Financial and Quantitative Analysis*, 28, 235-254.

Hull, J. and White, A., (1990), "Pricing interest-rate derivative securities", *Review of Financial Studies*, 3, 572- 592.

Hull, J., "Options, Futures and other Derivatives", 4'th ed., Prentice-Hall, 1999.

Jaillet, P., Ronn, E., and Tompaidis, S., (2003), "Valuation of Commodity Based Swing Options", *Management Science*, 50, 909-921.

Ibañez, A., and Zapatero, F., (2004), "Valuation of Amerucan Options through Computation of the Optimal Exercise Frontier", *Journal of Financial and Quantitative Analysis*, 39, 253-275

Jaskow, (1985), "Vertical Integration and Long-Term Contracts: The Case of Coal-Burning Electric Generating Plants", *Journal of Law, economics and Organisation*, 1, 33-80.

Jaskow, (1987), "Contract Duration and Relation-Specific Investments: Empirical Evidence form Coal Markets", *American Economic Review*, 77, 168-185.

Johnson, B., and Barz. G.: "Selecting Stochastic Processes for Modelling Electricity Prices" in R. Jameson, ed., Energy Modelling and the Management of Uncertainty. London: Risk Publication, 1999.

Kaminski, V., (1997), "The Challenge of pricing and risk managing electricity derivatives", Ch. 10 in: *The US power market*, 149-171, (Risk Publications, London).

Kelleharhals B. Phillip (2002), "Pricing Electricity Forwards under Stochastic Volatility", Working Paper

Kemna, A.G,Z and Vorst, A.C.F., (1990), "A Pricing Method for Options Based on Average Asset Values", *Journal of Banking and Finance*, 14, 113-129.

Kennedy, J. S., Forsyth, P. A., and Vetzal, K. R., (2006), "Dynamic Hedging under Jump Diffusion and Transaction Costs", Writing Paper, University of Waterloo.

Keppo, J., (2004), "Pricing Electricity Swing Options", *Journal of Derivatives*, 11, 26-43.

Knittel, C.R. and Roberts M., (2001) "An empirical examination of deregulated electricity prices", UCEI, PWP-087.

Lari Lavassani, A., Simchi, M., and Ware, A., (2000), "A Deiscrete Valuation of Swing Options", Canadian Applied Mathematics, 9, 35-74.

Lien, G., and Koekebakker, S., (2004) "Volatility and Price Jumps in Agricultural Futures Prices - Evidence from Wheat Options," American Journal of Agricultural Economics, 86(4), 1018-1031

Longstaff, F., and Schwartz, E., (2001), "Pricing American Options by Simulation: A Simple Least Square Approach", *Review of Financial Studies*, 14, 113-147.

Longstaff, F., and Wang, A., 2004, "Electricity Forward Prices: a high-frequency empirical analysis", *Journal of Finance*, vol. 59(4), 1877-1900

Lucia, J. and Schwartz, E., (2001) "Electricity prices and power derivatives. Evidence from Nordic Power Exchange", *Review of Derivatives Research*, vol. 5 (1), 5-50.

Meinshausen, N. and Hambly. B. M., (2004), "Monte Carlo Methods for the Valuation of Multiple-Exercise Options", *Mathematical Finance*, 14, 557-583.

Musiela, M., and Rutkowski, M., (1998), *"Martingale Methods in Financial Modelling"*, Springer.

Managing Energy Price Risk, (A Collection of different essays from academics and Market Practitioners), Risk Books (1999).

Manoliu. M., and Tompaidis. S., (2002): "Energy futures prices: term structure models with Kalman Filter estimation", *Applied Mathematical Finance*, 9, 21-43

Merton R, (1976), "Option pricing when underlying stock returns are discontinuous," *Journal of Financial Economics* 3, 125-144.

Merton, R. (1973), *"The Theory of Rational Option Pricing"*, Bell Journal of Economics and Management Science, 4 (Spring), pp. 141-83.

Merton, R., (1973), "An Intertemporal Capital Asset Pricing Model", *Econometrica*, 41, 867-888

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 & Quantitative Analysis*, vol 33, 33-59.

Nord Pool, *Annual Report 2003*, (Nord Pool ASA).

Nord Pool, *Options*, (April 2003, Nord Pool ASA).

Nord Pool, Security calculation and settlement of financial power contracts, (Sept., Nord Pool ASA)

Nord Pool, *Trade at Nord Pool's Financian Market*, (April 2004, Nord Pool ASA).

Nord Pool, *Trade at the Nordic spot market*, (April 2004, Nord Pool ASA).

Pablo Villaplana (2003), "Pricing Power Derivatives: A Two Factor Jump-Diffusion approach, Working Paper, Universitat Pompeu Fabra.

Pilipovic, D., (1998), *Energy risk: Valuing and managing energy derivatives*, (McGraw-Hill, New York).

Pilipovic, D., and Wengler, J., (1998), "Getting into the Swing", Energy and Power Risk Management, 2.

Pindyck, R., (1999), "The long-run evolution of energy commodity prices", *Energy Journal*, April, IAEE.

Pirrong, C. and Jermakyan, M. (2000). "The Price of Power: the Valuation of Power and Weather Derivatives", working Paper, Olin School of Business, Washington University.

Rebonato, R., (2004), "Volatility and Correlation: the Perfect Hedger and The Fox", Wiley.

Routledge, B.R., Seppi, D. J., and Spatt, C.S. (2000), "Equilibrium forward curves for commodities", *Journal of Finance* 55, 1297-1338.

S. Koekebakker, Ollmar, F. (2001), "Forward curve dynamics in the Nordic electricity market", Preprint, Norwegian School of Economics and Business Administration

Schwartz, E. and Smith, J.E., (2000), "Short-term variations and long-term dynamics in commodity prices", *Management Science*, 46 (7), 893-911.

Schwartz, E., (1997), "The stochastic behaviour of commodity prices: implications for valuations and hedging", *Journal of Finance*, vol. 52(3), 923-973.

Simonsen, I., Weron, R. and Mo, B., 2004, "Structure and stylised facts of a deregulated power market", preprint submitted to *Elsevier Science.*

Skantze, P. and Ilic M. (2000): "The Joint Dynamics of Electricity Spot and Forward Markets: Implications on Formulating Dynamic Hedging Strategies", Energy Laboratory Massachusetts Institute of Technology.

Sørensen, C., (2002), "Modelling seasonality in agricultural commodity futures", *Journal of Futures Markets,* Vol. 22, 393-426.

Tomposon, A.C., (1995), "Valuation of Path-Dependent Contingent Claims with Multiple Exercise Decisions over Time: the case of Take-or-Pay", *Journal of Financial and Quantitative Analysis*, 30, 271-293

Turnbill, 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.

Vasicek, O. (1977),"An Equilibrium Characterization of the term structure," *Journal of Finance* 5, 177-188.

Vehviläinen, I., (2002), "Basics of Electricity Derivative Pricing in competitive Markets", Applied Mathematica Finance, 9, 45-60.

Weron, R., 2005, "Market price of risk implied by Asian-style electricity-options", *Energy Economics*, forthcoming.

White, H., 1980, "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct tests for Heteroskedasticity", *Econometrica*, 48:817 – 838.