Preface

This thesis is the final thesis during the spring term at The Department of Industrial Economics and Technology Management, section for Managerial Economics and Operations Research.

A special thanks to our supervisor Stein-Erik Fleten for valuable comments and help on the theoretical and practical issues developed in this thesis. We also extend our gratitude to Professor Ånund Killingtveit at the Department of Hydraulic and Environmental Engineering help on Hydrological issues. Bjarne Schieldrop of former Enron Norway provided information about the market for precipitation derivatives and weather trader Lars Elmlund at AEP Energy took the trouble of providing price quotes on non-standard weather instruments. The Norwegian Meteorological Institute represented by Gudmund Dalsbø and The Norwegian Water and Energy Resources Directorate through Trine Fjeldstad gave us high-quality time series of precipitation and water discharge free of charge giving us firm data to back the analysis.

Throughout the process, we have gained valuable insight in modern financial theory and how it can be applied to value power and weather derivatives, as well as analyzing the investment problem in the context of a run-of-river power plant. Comments from the thesis supervisor Stein-Erik Fleten and our own ideas have generated a multitude of approaches to this problem, some of which are explored in this thesis, and others that due to space and time constraints remain at a sketching stage.

Trondheim, Norway, 27th of June, 2002

stud. tech. Anders Skogen

stud. tech. Narve Bjørdal

Summary

In this thesis, we present a survey of available models for spot- and forward prices for electricity, then the models are fitted to data obtained from the Scandinavian power Exchange, Nord Pool using non-linear regression analysis. A through discussion of the models is presented focusing on the relationship between the theoretical and empirical price trends, volatility structure and seasonal fluctations inherit in power prices.

Furthermore, using one of the frameworks originally suggested for power prices, we construct a model for the Enron Scandinavian Precipitation Index and derive analytical solutions for derivative instruments written on this thesis. Although a market for weather derivatives is not present in Norway at present, we obtain market quotes on precipitation options to estimate the market price of precipitation risk present due to the non-tradable nature of precipitation.

Based on the precipitation time series, we move on to consider whether precipitation derivatives can serve as a volumic hedge for a run-of-river power plant located at a specific site in Norway. The conclusion from this analysis is negative, and we consequently develop alternative methodologies for valuing such a power plant.

The first approach develops theoretical revenue futures by fitting a one-factor model to observed revenues of a hypothetical run-of-river power plant using price data and water discharge time series for the Gaula River. Through an arbitrage-free portfolio, we develop call options on plant revenues, since the time-varying discharge levels in the river and the presence of a minimum and maximum production threshold gives a payoff structure similar to a spread position engineered from plain vanilla options. The second approach combines two stochastic processes – one for water discharge assumed to follow zero-drift geometrical Brownian motion whereas prices follow a one-factor process fitted to long-term forward contracts.

Finally, we attack the problem of finding the early exercise boundary for the option to build a run-of-river power plant. Both a one-factor mean-reverting model for spot

prices as well as the one-factor GBM model fitted to the observed long-term price trend are analyzed using trinomial trees.

Table of Contents

Tables of figures and tables

1 Introduction

The final objective of this thesis is to value a run-of-river plant using a real option approach. Real option methodology is a central contribution to modern finance theory and incorporates factors such as for example the value of temporary project suspension, and the value of managerial flexibility. The framework also has the virtue of being independent of agent's risk preferences; consequently the appropriate discount rate is the risk-less rate. This introduces the concept of risk free valuation, and the possibility of using and developing financial derivatives directly in the analysis.

Much of the developed theory on continuous pricing using risk-neutral valuation is based on stochastics and dynamical systems. Whereas advanced models may yield more ”correct” results, there is always a trade-off between analytical tractability and the number of stochastic factors involved. If the model does not allow for analytical solutions or if the analytical solutions are too complex, simulation might be necessary. We limit our studies to analytically tractable models.

Central to the value of a run-of-river plant is the risk in power prices and instantaneous water discharge, since the plant has little or no storage capacity. Furthermore, a run-or-river plant has an asymmetric revenue structure because of upper and lower boundaries giving an upper bound of production and a minimum discharge required for the plant to operate at all. Between these bounds, discharge risk will be present most of the year. Due to this asymmetric revenue structure, we show that the plant can be valued as a spread position on the plant revenues.

Two of the main challenges are to gain a sound understanding of the stochastic processes assumed followed by power prices in the long and short run as well as the hydrological dynamics driving water discharge levels. To apply the real options framework, we need a way of removing both the price and the discharge risk, or obtain an estimate of the market price of these risks to develop a risk-natural framework. Provided that the correlation between a weather variable such as precipitation and instantaneous discharge is high, an alternative approach is to hedge the discharge risk using weather instruments, and move on to value the run-of-river plant as a portfolio of forward contracts and precipitation derivatives.

The right to exploit a site in a river for power production is normally valid within a time frame, and the real option approach allow us to model this as an American option on investing in a plant. Given a model for plant value and a model for power prices, we can find the early exercise boundary above which the option holder would benefit more from investing in a plant compared to holding the option to invest.

In sum, we need a good understanding of power price dynamics both in the short- and the long run as well as to investigate the relationship between precipitation and water discharge. Furthermore, the special characteristics of run-of-river power plants have to be taken into consideration.

2 Background

2.1 The Nordic energy market

Although the Norwegian energy market was only deregulated in 1991, a power exchange (market) has been present since the early seventies. After Norway's liberation of the market, Sweden followed in 1996, and a Nordic Energy market emerged. To organize the liberalized market, a common clearinghouse was needed, and the Nordic power exchange Nord Pool was created.

Because of its short existence relative to other markets, most energy markets would be categorized under emerging markets. Nord Pool is, however, one of the best-developed power exchanges in the world. Still, economists would argue that the volatility of the market is unnecessarily large, a sign of market immaturity. Energy, being a commodity, will always be more volatile than ordinary investment assets. Still, a closer integration of energy markets and other markets would increase market liquidity, as more arbitrageurs would even out temporal variations.

Still, the Nordic energy market is growing in size, as more foreign investors are entering the Nordic energy market. The secondary market, comprising energy derivatives, is also increasing. The gross turnover of the derivative markets increased form about 8 times the physical market in 2000 to about ten times in 2001. This is a sign that the energy market is still developing. In addition to Nord Pool, several independent brokers are present in the market, increasing the market turnover. Comparing the Nordic energy market to other energy markets, it is one of the best developed in the world.

Despite this growing market, the energy market can never be truly international, like the market for investment assets. Transmission grids limit the degree of energy export and import. Because of this, the physical market can never be truly global. The financial market, on the other hand, would in theory be open to anybody.

2.1.1 Market structure and Energy Products

Nord Pool offers several products. These can roughly be separated into three main categories, namely physical contracts, financial contracts and clearing. These will be given a brief explanation under.

2.1.1.1 The Physical Market

The only way to buy energy in the Nordic energy market is through a physical contract. Hence, the most basic product traded on the energy market, is spot electricity (Elspot). The Elspot is an agreement to deliver one kWh during a specific hour the following day for a settled amount of Norwegian Kroner (NOK). It could be argued that Elspot is actually a one-day forward contract for a one-hour delivery. We will, as is conventional, refer to it as the spot price.

Everything not traded spot, is considered being a derivative on the spot. For the practical purposes of this thesis, we will later let one-week forward contracts be a proxy the spot price. This is mainly due to convenience, as a higher resolution would not add anything significant to the analysis.

2.1.1.2 The Financial Market

The financial (derivative) market is divided into forwards and options markets. In this section, we will go through the most important characteristics of the two markets.

2.1.1.2.1 Forwards and futures

As in most developed markets, the energy business also has a secondary market. This market comprises forwards and future contracts, based on the underlying spot price. A forward is an obligation to buy stock at a predetermined price in the future, whereas a future essentially works the same way, but is settled on a daily basis, requiring a margin account[1]. At Nord Pool, as in most other markets, these contracts do not involve physical delivery. Contracts are settled in cash, based on the difference between the spot and the pre-settled delivery price.

[1] Forward and futures are identical for a constant risk free rate. See Cox, Ingersoll and Ross (1981) for details.

Energy forward contracts are traded in years and blocks. At Nord Pool, annual forward contracts start trading three years before delivery, and are traded until the start of the season preceding maturity. Then, annual contracts are split into seasonal forwards. These forward contracts divide the year into three separate seasons, called winter 1, summer and winter 2. Table 1 gives an overview of contracts and their duration.

Contract	Contract type	Duration	Hours of energy (Leap-year)
Year	Forward	Weeks 1-52	8760 (8784)
Winter1(V1)	Forward	Weeks 1-16	2879 (2903)
Summer (S0)	Forward	Weeks 17-40	3672
Winter2 (V2)	Forward	Weeks 41-52	2209
Block (Month)	Future	4-5 weeks	672
Week	Future	1 week	168
Day	Future	1 day	24

Table 2.1 Forward Contract structure at NordPool.

Source: www.nordpool.no

Seasonal forwards can be traded three years before actual delivery, until the beginning of the season before maturity. The contracts are then split into 3-6 blocks of four weeks that are tradable until the previous block matures, when the blocks are spilt into four weekly contracts traded until the day before delivery. These, and shorter contracts, are settled on a daily basis. Hence, they fall under the futures category.

Fleten and Lemming (2001) argued that the block structure complicated the process of finding prices for specific maturity times, or constructing a continuous term-structure curve. This is because the block structure only gives a partial picture of the prices. Furthermore, although the market for power contracts at Nord Pool is reasonably large and growing, long-term contracts are rather illiquid.

2.1.1.3 Contracts for difference

Due to capacity constraints on the transmission grids, Statnett has the possibility to declare a specific power grid zone as special when demand is high. Under these conditions, the system price quoted by Nord Pool would be different from the price prevailing in the actual area. In such situations, Nord Pool offers a contract for difference (CfD), available for seasons and years, as protection against exposure to these bottlenecks. CfD's are only available for a few big cities, and are fairly illiquid. Any hydropower plant will be exposed to risk of differences between system price and area price, if the system price is used for hedging. This exposure will not, however, be treated in this thesis.

2.1.1.4 Options

Belonging to the set of financial contracts, an option is a right, but not an obligation, to buy an underlying asset at a given (usually predetermined) price. This is the third "layer" in the electricity market. As for futures and forwards, electricity options are pure financial products settled in cash rather than delivery.

The amount of options available in the electricity market is not enormous. Contracts are mostly written on the longer-maturity forward contracts, and are mostly European in nature. There has been a market for Asian options, but the market demand seems to be low for these contracts at the time being. In fact, Nord Pool has stopped listing these products, and they are now traded over the counter (OTC).

During the recent years, a new market of various derivatives, modeled to fit other variables than price, has developed. Specifically, the market for weather derivatives has started growing. This is discussed in section 2.2.

2.1.1.5 Risk management and hedging in energy markets

Modigliani and Miller (1958) argued that a company's shareholders would always be able to hedge a risk more efficiently than the company itself. This might not apply to energy risk management. Due to the volatile nature of power markets, energy companies and industries heavily dependent on energy, could suffer big losses due to fluctuations in the market. Bjørkvoll et al. (2001) argued that due to market "imperfections", hedging at a company level might be profitable for shareholders. They further contended that due to economies of scale in power derivatives markets, firms are able to operate at lower cost than individual investors do. Traditionally, companies operating in the energy business do hedge their exposure.

2.1.1.6 Hedging strategies and separation problems

The energy market comprises several large actors that might see added value in hedging their interests. To be efficient, hedging should be done at a corporate level. Such hedging raises practical issues, however. Bjørkvoll et al (2000) claimed that production uncertainties make hedging very difficult and not necessarily separable from production planning. In effect, hedging at a corporate level might prove to be impossible in practice.

Secondly, it is crucial to distinguish between price hedges and volume hedges. A price hedge guarantees a price, but not a steady demand or supply. A volume hedge, on the other hand, is an instrument made to ensure market size, either on the supply or the demand side. Pure weather derivatives are volume hedges. Temperature mostly affects the magnitude of the demand, whereas precipitation affects the magnitude of supply. They both, of course, indirectly affect the price, but it is difficult to capture this relationship exactly.

In conclusion, the combination of lack of separabilty and the difficulties quantifying the exposure complicates energy risk management considerably. Nevertheless, the number of participants in the market is large and growing.

2.1.2 Transmission grids

In Norway, Statnett has controlled most transmission grids since the market liberalization in 1991, because the grids were considered a natural monopoly. Although a few minor interests still own some of the grid, this ownership is limited to the more local distribution grids, making the monopoly comparison liable for the main grids.

2.2 The weather market

Weather derivatives were originally developed to facilitate against non-catastrophic weather conditions that might negatively impact on a company's revenue stream. Two key features of weather derivatives are that they are written on a non-tradable underlying (some weather variable or weather index) and that they are used to hedge *volumetric* risk, while traditional derivatives hedge price or financial exposure.

2.2.1 Market development

The US remain the most developed weather market, with almost 2500 contracts traded in year 2000 amounting to a total volume of US$ 2.7 billion (Tigler and Butte, 2001). While numerous studies point out the sheer vastness of businesses whose revenues are affected by weather patterns, the market remains clearly dominated by energy companies. Temperature derivatives for ten major US cities are traded on the Chicago Mercantile Exchange. The European market is catching on, with the UK, Germany and France being the most active markets. Again, energy companies are the most active, while investment banks and insurance companies are gaining interest for the weather derivatives business. However, the market still remains illiquid and all contracts traded in Europe are arranged OTC[2].

2.2.2 Common underlying variables

The most common underlying variable in the weather market is temperature (98% of traded contacts); the rest of the market consists of precipitation (rain or snow), sunshine hours and wind speed. The criteria for a weather variable to serve as underlying can be listed in three main points:

- The variable needs measurable impact on enough agents to develop a market.
- The variable has to be accurately and objectively measurable.
- High-quality data and sufficiently long historical records are needed for agents to price contracts and study the properties of the relevant weather variable.

Historically, weather protection has been supplied by the insurance industry, but for several reasons, a weather derivatives market is a good supplement to traditional insurance contracts. The key differences are summarized in the following table.

	Insurance policies	**Weather Derivatives**
Usage	Protection against specific, low-probability catastrophic events	Protection against high-probability ”continuous” risks or natural fluctations of weather
Payment preconditions	Proof of financial losses due to the occurance of events as specified in the insurance contracts	Automatic payment based on the realization of some weather variable
Complexity	Usually high	Can be high for OTC-contracts, lower for exchange-traded contacts
Buyer’s counterpart	Insurance company	Arbitrary market participants (OTC) or clearing house (for exchange-traded contracts)

Table 2.2 Insurance policies versus weather derivatives

Source: Tigler and Butte (2001), Element Re (2001)

[2] Over The Counter. This refers to bilateral trading of contracts outside the ordinary exchange.

2.2.3 The rationale behind indices

Weather indices are preferred to individual measurements for several reasons. The most obvious is to make the instrument more liquid – in the Scandinavian power market, the power reservoirs are located in sparsely populated areas, and it is unlikely that there will be more than one actor interested in trading in the instrument for hedging purposes. Secondly, for business with geographically diverse activities, the outcome of some weather variable over a larger area might be more important than the outcome of the same variable in, say, a given city.

2.2.4 The Nordic weather market

The market for weather derivatives has yet to boost in Scandinavia. The few contracts traded are traded OTC. According to AEP Energy Norway, almost all current contracts are in practice building blocks in larger structured deals in which weather is one of more underlyings.

Enron Corp. developed a precipitation index to use as underlying for precipitation derivatives, but after Enron went bankrupt, contracts were no longer traded on this index. AEP Energy maintains the index and uses it as an internal reference for their structured deals. We will return to this index in chapters 7and 8 when we model and price precipitation derivatives.

2.2.4.1 Potential of the Nordic weather market

A commonly cited figure is that over 20% of the US economy is directly exposed to weather risk (see i.e. Geman, 2001), and AEP Energy claims that the corresponding number is at least as large in the Nordic market. The importance of hydropower could pave the way for a market for precipitation-based instruments, but a problem is that there is no natural counterpart for the generators. One could possibly expect that large *buyers* of power could act as contractual counterparts, but this seems not to be the case.

2.2.4.2 Reasons why the market has not caught on

The market lost its most central actor when Enron went bankrupt. Enron Nordic's operations were taken over by AEP Energy, but the market is not as big as it used to be, and was never really big in the first place. Moreover, the second biggest player, Aquilla Energy has withdrawn the bulk of its operations from Norway and is no longer a player in the Norwegian weather market.

AEP Energy points to the fact that the potential actors have been very slow to catch on, and experience shows that if the first weather deal struck ends with a loss, the actor tends to conduct no further deals.

As pointed out above, the market is one-sided due to a lack of counterparts for the energy businesses. This does not mean that there are no businesses with different exposure profile to weather phenomenon, they are simply not actors in the derivatives market. Smidt (2001) listed municipals, retailers, tourist operators, and breweries as businesses with huge exposure to weather, but they quite simply do not use derivatives to hedge or do not weather risk.

Cao and Wei (2001) claimed that large bid-ask-spreads caused problems in the US and European weather markets, and Williams (1999) pointed out that triple-digit bid-ask-spreads were not uncommon in the US market. Consequently, weather derivatives might be ineffective as a hedging mechanism. In Norway, the bid/ask spread is non-existing. It is either a bid or an ask.

In sum, there is no efficient market in Norway for weather instruments, and at the present time, not a market at all for pure weather instruments.

3 Theoretical framework

3.1 Project valuation approaches

There are various ways of estimating the value of a project. In the following sections, we will briefly touch upon the most common pricing methods.

3.1.1 Net Present Value

The net present value (NPV) approach was developed as a result of the risk-return framework developed by Markovitz (1953) and extended by, amongst others, Sharpe (1964). The approach comprises predicting future cash flows, and discounting at a risk-adjusted rate of return. The popularity of NPV is enormous due to its simplicity, and it is by far the most used pricing method today.

Whereas the NPV method is tractable due to its simplicity, it is inexact because of its static nature. Trigeorgis (2000) claimed that NPV analysis was far too static to capture the dynamic environment of a company. For example, NPV failed to capture the value of project flexibility, such as options to abandon and temporarily suspend projects. Trigeorgis (2000) further argued that traditional NPV was developed for passive portfolio management and that the value of active management would be better captured using other methods. Lastly, a fundamental problem with NPV is valuation of investment opportunities involving asymmetric payoff.

Still, NPV analysis could prove to be useful. Different projects have different characteristics, and require different valuation methods. Fleten (2000) argued that traditional NPV approaches would be adequate for pricing passive investments and for decision support for now-or-never projects.

3.1.2 Decision Tree Analysis

Decision Tree Analysis (DTA) is a project valuation approach similar to the NPV analysis. It involves developing several scenarios for the future, and assigning probabilities to these. The outcomes of the various scenarios are finally discounted at appropriate discount rates.

DTA is able to capture more flexibility than NPV analysis, because of the introduction of scenarios. It does, however, introduce other problems. Trigeorgis (2000) claimed that although the DTA approach was correct in principle, and gave a good overview over sequential investments, the approach had two major drawbacks.

Firstly, the appropriate discount rate has to be determined at every node. Say, for example, that at node *n*, there are two possible outcomes, *A* and *B*, with probability *p* and *1-p*. *A* is dependent on the energy market, and should be discounted at a rate reflecting the energy market risks. Outcome B, however, is risk-free. It should therefore be discounted at the risk-free rate. At node *n-1*, the combination of the two outcomes in *n* will need to be discounted at an appropriate combination of the two rates. On every single point in the tree, this problem is repeated, resulting in multiple discount rates, and increasing the complexity.

Secondly, the probabilities associated with the various scenarios would have to be set explicitly for each scenario. Although tedious, this problem might be solved using Monte Carlo methods. Finally, decision trees grow exponentially in size as the time horizon increases, making multistage problems difficult to handle.

3.1.3 Contingent Claims Analysis and Real Options

Contingent claims analysis is the "new" approach to project valuation. It was developed on the basis of the work of Cox, Ross and Rubinstein (1979). They developed a theory of replication of options in an arbitrage free portfolio, extending the framework of, amongst others, Black and Scholes (1973). The result of these advances, projects could be evaluated independent of individual risk preferences.

CCA uses a certainty equivalent approach, and utilizes the options framework instead of the static approaches discussed above. Hence, CCA methods avoid some of the problems introduced in the NPV and DTA approaches. The asymmetric payoff problem is taken care of in the same way as in DTA, by allowing scenarios, or a continuous specter of scenarios.

In a risk neutral world, all individuals are indifferent to risk. Hull (2000) argued that if the risk preference of the investors does not enter the equation, it could not affect the solution. Hence, utility theory is avoided, and individual preferences could be ignored. CCA is based on risk neutral valuation, and the value of a *real option* can essentially be found via a proper NPV analysis using risk-neutral valuation. The discount rate would be the risk-free rate r, and not the rate given by the CAPM or any other risk-adjusted discount-rate framework. This approach solves two problems. Firstly, preferences are eliminated. Secondly, only one discount rate is needed, which is a significant improvement compared to the DTA analysis above. At any given node in a decision tree, the discount rate is given as r.

3.2 Pricing of contingent claims

In this section, we are going to investigate how to price contingent claims using a risk-neutral framework. The section is largely based on Trigeorgis (2000), Dixit and Pindyck (1994) and Hull (2000). We start by investigating the statistical framework for modeling investment and commodity prices.

3.2.1 Stochastic processes

The modeling framework used is based on stochastic processes. A pure stochastic process is a process whose next step is independent of the previous steps. The process will still be dependent of the previous *state* of the process. If it is only dependent on the last state, the process possesses the Markov property.

The first building block is noise, based on the idea of Brownian motion, or Wiener processes. A pure Wiener process, as described in Hull (2000), is a process with no drift, given in equation (3.1).

$$dz = \varepsilon\sqrt{dt} \tag{3.1}$$

Here dz is a Wiener process, ε is a normally distributed value with mean zero and variance of one unit, and dt is a small time step. The only movement is “noise”, or random motion. In addition, a Wiener process’ noise ε follows the normal distribution, generally indicating that the likelihood of small movements is greater that that of larger magnitudes.

3.2.1.1 Geometric Brownian Motion

A pure Wiener process is not usually considered able to capture the movements of security prices, since the Wiener process has zero drift. Due to the time value of money, inflation, etc, a security would have an expected return, or a drift μ. Therefore, it has been normal to use a geometric Brownian motion to capture the movements of stock prices. Geometric Brownian Motion (GBM) is defined in equation (3.1).

$$dP = \mu P dt + bP dz \quad (3.2)$$

Here, dP is a small change in the security price, μ is the drift, dt is a small time step, b is the magnitude of the volatility[3], and dz is given in equation (3.1).

3.2.1.2 Arithmetic Brownian Motion

Unlike the more familiar GBM, arithmetic Brownian motion has the generic from

$$dP = \mu dt + b dz \quad (3.3)$$

An arithmetic Brownian motion is often referred to a generalized Wiener process. The difference between GBM and ABM lies in the noise term. GBM generates noise proportional to the stock price P, whereas ABM's noise is independent of the price P.

3.2.1.3 Application to asset pricing

Asset prices are normally assumed to follow GBM. It can be shown that if a stock price P follows GBM, then ln(P) follows a generalized Wiener process. Since a small increment dz is normally distributed, the Wiener process residuals should be normally distributed. In effect, the GBM residuals should be lognormally distributed. This is an important condition for the framework to be valid.

3.2.1.4 Itô processes

An Ito process is generally speaking the same as a Wiener process, only that the drift μ and the standard deviation b are replaced with functions μ(P,t) and b(P,t). Hence, we can restate equation (x) as

$$dP = \mu(P,t)dt + b(P,t)dz \quad (3.4)$$

[3] By volatility , we mean a standardized measure of the standard deviation.

The difference is that a more general drift and standard deviation can be incorporated. For example, the Itô process could allow seasonal variations. This property is particularly useful when modeling commodity prices with seasonal trends.

3.2.1.5 Mean reversion processes

A mean reversion process is a combination of a deterministic and stochastic process. Whereas GBM processes are characterized by the Markov property (all previous states are insignificant), the mean reversion process is dependent on the deviation from a long-term mean. A mean reversion process can be given as

$$dP = \kappa(m - P)dt + \sigma dz \tag{3.5}$$

In equation (3.5), κ is a mean-reversion factor, m is the long-term mean, P is the current state (price), σ is the volatility, and dt and dz increments in time and the Brownian motion.

3.2.1.6 The choice of process

The choice of process depends on the behavior of the price time series. The choice of an adequate model is crucial, considering that the process forms the basis for the pricing of an asset. Therefore, it is often necessary to try several models before deciding which is most suited for the task. Furthermore, simplicity is an important issue, as a more advanced model might be notoriously difficult to communicate and work with, without adding significant improvements compared to an easier model.

3.2.2 Creating a simple arbitrage-free portfolio

Consider the case of a tradable and storable underlying asset. In the discrete case, a portfolio consists of N shares of one asset, financed by a loan B at the risk free rate r. For simplicity, we assume there are two possible outcomes, C^+ and C^-, with outcome probability p and 1-p respectively. The value of the portfolio P at time 0 and 1 can be visualized through a tree:

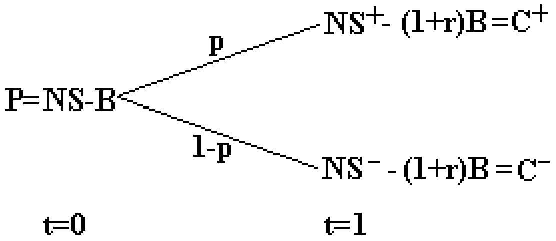

Figure 3.1: The payoff of an asset S

We now design a portfolio whose payoff at time t=1 does not depend on the path chosen. In other words, we eliminate uncertainty. Setting the values at the two states, C^+ and C^-, equal, we can find an N that makes the outcomes identical, regardless of uncertainty. After rearranging, we get

$$N = \frac{C^+ - C^-}{S^+ - S^-} \tag{3.6}$$

At this ratio, the payoff is certain, and should be discounted at the risk free rate r. Following Trigeorgis (2000), we hence define the risk-neutral probability as

$$p = \frac{(1+r) - S^-}{S^+ - S^-} \tag{3.7}$$

This is the probability that would prevail in a risk-free world. Furthermore, the expected return on an option must equal the risk-free rate in a risk-free world, r:

$$\frac{pC^+ + (1-p)C^-}{C} - 1 = r \tag{3.8}$$

This exercise can be repeated for the continuous case, and has one important message. If the risk-free return of the option is not equal to the risk-free rate, then arbitrage opportunities would exist. In other words, the relationship in equation (3.8) must hold.

3.2.2.1 The market price of risk

Following Hull (2000), in absence of arbitrage, two derivatives, f_1 and f_2, both dependent on the underlying asset θ, must be priced internally consistent, so that a risk-less portfolio comprising the two would provide the risk-free return. This can be illustrated by the two derivative processes below, both dependent only of θ and dt:

$$\frac{df_1}{f_1} = \mu_1 dt + \sigma_1 dz \tag{3.9}$$

$$\frac{df_2}{f_2} = \mu_2 dt + \sigma_2 dz \tag{3.10}$$

All the risk in the above securities lies in dz, which is equal for both securities. Hence, eliminating dz would create a risk free portfolio. We construct this portfolio by the use of $f_2\sigma_2$ portions of the first derivative, and $-f_1\sigma_1$ of the second. This portfolio will have a certain payoff. The value Π of the portfolio is then given in equation (3.11).

$$\Pi = (\sigma_2 f_2) f_1 - (\sigma_1 f_1) f_2 \tag{3.11}$$

A small change in Π can be written as

$$d\Pi = (\sigma_2 f_2)df_1 - (\sigma_1 f_1)df_2 \quad (3.12)$$

Since the portfolio is risk-less, it should earn the risk-free rate of return *r*. The payoff of the portfolio over a small time step *dt* is stated in equation (3.13):

$$d\Pi = r\Pi dt \quad (3.13)$$

Substitution into this equation from (3.11) and (3.12) yields

$$\frac{(\mu_1 - r)}{\sigma_1} = \frac{(\mu_2 - r)}{\sigma_2} \quad (3.14)$$

This term is commonly referred to as the market price of risk, and labeled λ. Dropping the indices, we can restate the market price of risk as

$$\frac{(\mu - r)}{\sigma} = \lambda \quad (3.15)$$

The market price of risk is hence a price of the volatility. The product of the volatility and the market price of risk constitute the difference between the value of a security's expected value and its certainty equivalent, the futures/forward contract. This difference is seen as a compensation for bearing the risk of trading spot. Pirrong (2000) pointed out that there are large differences between these two values in the PJM[4] market, even for one-day forward contracts. This indicates the presence of a market risk premium. Rewriting equation (3.15), a risk-neutral measure of the return of a risky asset is given in equation (3.16).

$$r = \mu - \lambda\sigma \quad (3.16)$$

3.2.2.2 Ito's lemma

Hull (2000) contended that any derivative is a function of its underlying and time. It is difficult in practice, however, to derive the explicit relationship between the derivative and its underlying. Ito's lemma provides us with the necessary tools to perform the transformation of the stochastic process for the underlying processes to the new derivative process.

The method Ito suggested involves simplifying the derivative expression through a Taylor expansion. Ito claimed that using the first two Taylor expansions of the function with respect to the underlying, and the first Taylor expansion of the function

[4] Pennsylvania, New Jersey and Maryland

with respect to time would provide the necessary approximation. In other words, a change in a function G of an underlying stock u will be approximately equal to

$$dG = \frac{\partial G}{\partial u}du + \frac{\partial G}{dt}dt + \frac{1}{2}\frac{\partial^2 G}{\partial u^2}(du)^2 \tag{3.17}$$

For a stock following the dynamics given in equation (3.18), the process for a derivative G will follow the relationship given in (3.19).

$$du = a(x,t)dt + b(x,t)dz \tag{3.18}$$

$$dG = \left(\frac{\partial G}{\partial u}a + \frac{\partial G}{dt} + \frac{1}{2}\frac{\partial^2 G}{\partial u^2}b^2\right)dt + \frac{\partial G}{\partial u}bdz \tag{3.19}$$

3.2.2.3 Creating a continuous instantaneous arbitrage-free portfolio

After finding the process followed by G, we construct a risk-free portfolio of the underlying and the derivative, by going short in the derivative and long ∂G/ ∂u units of in the underlying. This operation is funded[5] by a bank loan at the risk free rate. The portfolio value Π is then expressed in equation (3.20).

$$\Pi = -G + \frac{dG}{du}u \tag{3.20}$$

Correspondingly, a change in the portfolio value, ΔΠ, is given in equation (3.21).

$$d\Pi = \left(-\frac{\partial G}{\partial t} - \frac{1}{2}\frac{\partial^2 G}{\partial u^2}b^2u^2\right)dt \tag{3.21}$$

This portfolio is instantaneously risk-less. Therefore, it should earn the risk-less rate of return:

$$d\Pi = r\Pi dt = r\left(-G + \frac{\partial G}{\partial u}u\right)dt \tag{3.22}$$

Rearranging, we obtain the differential equation satisfied by G:

$$\left(\frac{\partial G}{\partial t} + ru\frac{\partial G}{du} + \frac{1}{2}\frac{\partial^2 G}{\partial u^2}b^2u^2\right) = rG \tag{3.23}$$

Similarly, it can be proven that for a stock with convenience yield c, the process in equation (3.23) changes to equation (3.24).

---

[5] This is only needed if the underlying requires investment at time 0. If the underlying is a forward contract, no initial investment is needed.

$$\left(\frac{\partial G}{\partial t}+(r-c)u\frac{\partial G}{du}+\frac{1}{2}\frac{\partial^2 G}{\partial u^2}b^2u^2\right)=rG \tag{3.24}$$

Finally, in a similar way, it can be shown that if a forward contract is the underlying, the differential equation of the derivative G is given as

$$\left(\frac{\partial G}{\partial t}+\frac{1}{2}\frac{\partial^2 G}{\partial F^2}b^2F^2\right)=rG \tag{3.25}$$

Here, F represents a forward contract. Because the forward contract has a fixed contract price at a future time T, the drift term of *(r-c)* is not included in the differential equation.

3.2.2.4 Forward risk neutrality and equivalent martingale measures

When the underlying is not tradable, we cannot use a derivative and its underlying to find the risk-neutral probability measure, as given in equation (3.7). Instead, we need two derivatives. Using the approach suggested by Hull (2000), we define f_1 and f_2 to be two derivatives dependent on a single source of uncertainty. Let the relative price of f_1 with respect to f_2 be expressed as θ so that the relationship given in (3.26) holds.

$$\theta=\frac{f_1}{f_2} \tag{3.26}$$

A martingale is a stochastic process with zero expected drift. The *equivalent martingale measure* result shows that, in the absence of arbitrage opportunities, θ is a martingale for some choice of the market price of risk. Furthermore, if the volatility of f_2 is the market price of risk, then $f1/f2$ is a martingale for all security prices f_i. A complete proof of this result can be found in Hull (2000).

This result means that to appropriately value energy derivatives, we could use the forward curve as the underlying asset, since this curve represents the risk-neutral expected future spot price. Since the income from a predetermined forward contract is fixed, there is no uncertainty in the future value. Consequently, valuation is risk-free, and should be discounted at the risk-free rate.

3.2.2.5 Solving the differential equation

The differential equation given in equation (3.24) has an unlimited set of solutions. The appropriate solution is conditioned on the boundary conditions of the derivative.

For forward contracts, the boundary condition would be $F_0=E(S_0,T)$, making the portfolio instantaneously risk-less. For options, the boundary condition is given as G=max(S-X,0) for call options, and G=min(S-X,0) for put options.

3.3 Options

In financial theory, an option is the right, but not the obligation to purchase spot at a predetermined point in time. In practice, however, an option is a pure financial tool, that will be settled in cash if the option expires in the money. The settlement will be the difference between the actual price and a predetermined strike. If the option expires out of the money, whatever the amount, the settlement is zero. Hence, an option has a non-linear payoff structure. The tractability of options is mainly due to the asymmetric payoff, giving it a unique position in risk management.

3.3.1 Purpose of options in this thesis

Within the framework presented in this thesis, options are used in several contexts. A power plant with a lower capacity limit will have an asymmetric revenue structure similar to that of options. Hence, the production could be viewed as an option. Weather contingent claims are also, to a large extent, based on options. Finally, American options will be used to evaluate the value of delaying the investment decision, and find an optimal investment price.

3.3.2 Framework

Options come in a variety of forms. To price an option, we need the following parameters:

- Risk-neutral expected forward price $F_0(T)$. As discussed in section 3.2.2.4, this is equivalent to the forward price.
- Strike price X. This is the predetermined price above or below which the option generates payoff.
- Risk-free rate r.
- Volatility σ. This is needed to find the accumulated variance.
- Time to maturity T.

The following table explains the most usual option types:

Options and payoffs

Option style	Type	Exercise	Payoff
European	Call	At expiry	Max(S_T-X,0) at maturity
European	Put	At expiry	Max(X- S_T,0) at maturity
American	Call	Within expiry	Max(S_T-X,0) at time of exercise
American	Put	Within expiry	Max(X- S_T,0) at time of exercise
Asian European	Put/Call	At expiry	Some average
Asian American	Put/Call	Within expiry	Some average

Table 3.1: Option styles

In this thesis, European options are prices using the framework developed by Black and Scholes (1973). This framework is based on a stock following GBM, and a necessary condition for the framework to work, is that the price shocks are lognormally distributed.

For American options, we use Clewlow and Strickland (1998) and Hull(2000). When needed, the original framework is altered. This is especially necessary when pricing options based on mean-reversion.

3.3.3 Pricing of European Options

Options are priced using the risk neutrality framework. The differential equation satisfied by all derivatives can be written as

$$\frac{\partial c}{\partial t}+rP\frac{\partial c}{\partial P}+\frac{1}{2}\sigma^2P^2\frac{\partial^2 c}{\partial P^2}=rc \tag{3.27}$$

Here, c is the call option, P the spot price, r the risk free rate and σ the volatility. The solution of this differential equation, using the boundary condition given in equation (3.28) yields the option equation for a call option.

$$c=\max(P_T-X,0) \tag{3.28}$$

3.3.3.1 Variations of the differential equation

Equation (3.27) is a standard differential equation followed by any derivative. The equation might change slightly when a deterministic seasonal function and convenience yield enters the equation. Of course, the solution of the equation is also dependent on the boundary condition.

3.3.3.2 European Options based on variables following GBM

A European option is an option with the only exercise possibility at maturity. Black and Scholes (1973) showed that a European call option with strike *X*, forward price *F*, volatility σ and time to maturity *T* could be expressed as follows:

$$c_{0,T} = e^{-rT}\left(F_{0,T}N(d_1) - XN(d_2)\right) \tag{3.29}$$

$$d_1 = d_2 + \sigma\sqrt{T} = \frac{\ln\left(\frac{F_{0,T}}{X}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} \tag{3.30}$$

Here, the parameters are the same as in equation (3.27), except for the forward price at time *0*, maturing at time *T*, given as $F_{0,T}$. $N(d_i)$ is the cumulative normal distribution function, and d_i is given in equation (3.30).

3.3.3.3 Options based on mean-reversion dynamics

The option price formula in the previous section is subject to slight changes when implemented. Looking at equation (3.30), mean reversion models have the same payoff structure, but not the same volatility. This is because the volatility in a mean-reversion is in itself mean-reverting.

3.3.3.3.1 One-factor volatility

Using a one-factor model with no long-term price trend (and hence no uncertainty in the long run) the volatility is given as:

$$\sigma_F = \sigma e^{-\kappa T} \tag{3.31}$$

The Black/Scholes equation requires the cumulative variance. For this model, the cumulative variance can be given as

$$w = \int_0^T \sigma_F{}^2 dt = \int_0^T \left(\sigma e^{-\kappa T}\right)^2 dt = \frac{\sigma^2}{2\kappa}\left(1 - e^{-2\kappa T}\right) \tag{3.32}$$

Using this volatility, equation (3.30) can be replaced by equation 7.

$$d_1 = d_2 + w = \frac{\ln\left(\frac{F_{0,T}}{X}\right) + \frac{1}{2}w}{\sqrt{w}} \tag{3.33}$$

Here, w is the general cumulative volatility. Note that the volatility given in (3.31) approaches zero. Thus, the cumulative volatility w approaches a limit given as $\sigma^2/2\kappa$

from equation (3.32). This limits the tractability of the problem to short-run problems, or long run problems without equilibrium uncertainty.

3.3.3.3.2 Volatility

The cumulative volatility function of the simple mean-reversion model is given as

$$w = \frac{\sigma^2}{2\kappa}\left(1 - e^{-2\kappa T}\right) \tag{3.34}$$

This volatility approaches a limit, and is most appropriate when there is no uncertainty in the long-run equilibrium.

In order to incorporate long-term realism into the option prices, we could assume that a mean-reversion model long-term trend follows GBM, and short term shocks follows mean-reversion. Defining a short run volatility function $f(\sigma_S)$ and a long-run equilibrium volatility σ_L equal to the long-run implied volatility or forward volatility, the cumulative volatility w is given as

$$w = f\left(\sigma_S\right)^2 + \sigma_L^2 + 2\rho_{SL}\sigma_L f\left(\sigma_S\right) \tag{3.35}$$

Using these results, it can be shown that in the long run, prices using a two-factor model approaches prices using a one-factor GBM model with accumulated variance equal to $\sigma_L^2 T$ plus a constant. Schwartz (1998) utilizes this relationship when reducing a two-factor model to a one-factor model for long maturity contracts.

3.3.4 Pricing of American Options

An American option can be exercised at any point within a predetermined period of time. Due to the nature of the American option, there has not yet been developed an analytical pricing tool for this type of option. Therefore, these options have to be valued through the use of a tree structure, or alternatively, through Monte Carlo simulation. This thesis will consider trees for option valuation, and the following section explores a tree-fitting approach for a GBM process. Then, a tree is fitted to a mean-reversion process. Finally, the option is priced, and a brief discussion follows.

3.3.4.1 Tree-building using a GBM process

This section illustrates the procedure of pricing American options through a tree structure. The precision of the solution will be dependent on both time steps and branching complexity. This section is concerned with modeling the options using a trinomial tree, and will use an approach obtained from Clewlow and Strickland (1998) and Hull (2000). The section starts by investigating initial requirements and continues by building a tree for a GBM process.

3.3.4.1.1 Initial requirements

Tree building procedures are discrete in nature, and American option trees are not radically different from other tree structures. Consequently, the building of American style options involves the creation of discrete time steps. We furthermore need the following information to build the tree:

- Transition probabilities
- Risk-free discount rate *r*. This is obtained in section 4.5.
- Risk-neutral drift $\alpha = r - c$ for commodities.
- Up- and down steps defined.
- Volatility σ.
- Initial price P_{0}.

3.3.4.1.2 Transition probabilities

Transition probabilities are discrete probabilities of price movement in a tree. For a trinomial tree, the natural choices of outcomes are moving up, staying in the middle or moving down. For an asset following GBM, the transition probabilities are equal for all nodes. Hull (2000) claimed that the following transition probabilities for a non-dividend paying stock yielded good results:

$$P(up) = p_U = \sqrt{\frac{\Delta t}{12\sigma^2}}\left(r - \frac{1}{2}\sigma^2\right) + \frac{1}{6} \tag{3.36}$$

$$P(middle) = p_M = \frac{2}{3} \tag{3.37}$$

$$P(down) = p_D = -\sqrt{\frac{\Delta t}{12\sigma^2}}\left(r - \frac{1}{2}\sigma^2\right) + \frac{1}{6} \tag{3.38}$$

Of course, for these probabilities to be correct, prices must be lognormally distributed and follow GBM. For a stock paying continuous dividend yield, or having convenience yield *c*, *r* is replaced by *(r-c)*. Notice that the price drift is inherent in the transition probabilities.

3.3.4.2 Upwards and downward moves

In section 3.2.2, we investigated the up- and down states of an asset with two possible outcomes. A trinomial tree is basically a discrete approximation of an asset having continuously lognormally distributed payoff. Hull (2000), following Cox, Ross and Rubinstein (1979), suggested three suitable outcomes from the lognormal distribution as up *u*, down *d* together with the expected value:

$$u = e^{\sigma\sqrt{3\Delta t}} \tag{3.39}$$

$$d = \frac{1}{u} = e^{-\sigma\sqrt{3\Delta t}} \tag{3.40}$$

Since the drift is inherent in the probability measures, the total tree can still be assumed to have an upward drift. Figure 3.2 illustrates the trinomial tree and its corresponding values of p and d:

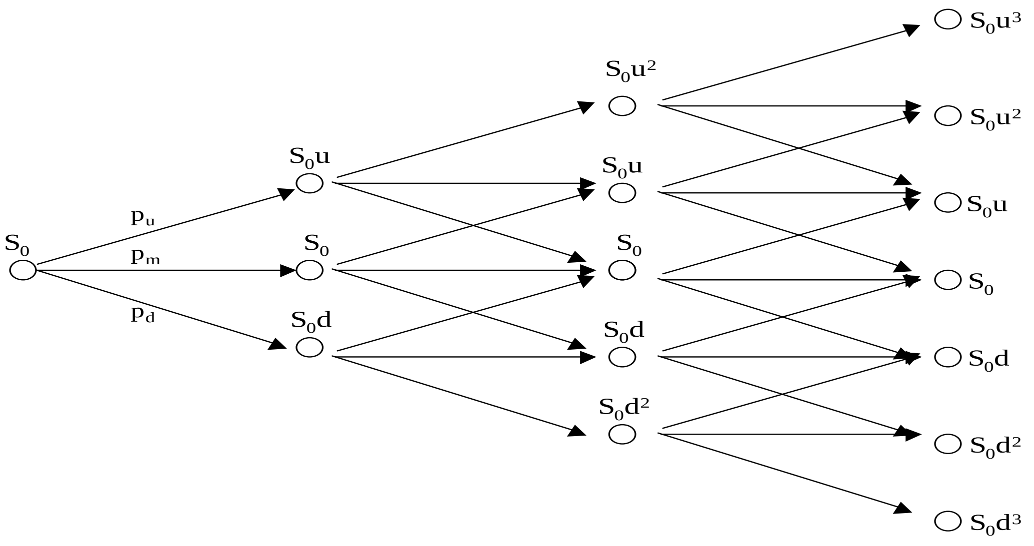

Figure 3.2: Trinomial lattice

After finalizing this part, the tree needs to be initialized and the option priced. This is done in section 3.3.4.4.

3.3.4.3 Tree-building using mean-reversion

Although our analysis thus far has assumed a stock following GBM, the procedure can be adjusted to fit alternative stochastic processes. This section outlines the changes necessary to build a tree for assets following mean reversion dynamics. We start by briefly introducing an appropriate mean-reversion model, and continue by redefining branching, and defining the steps necessary to build the tree.

3.3.4.3.1 A mean-reversion process

Hull and White (1990) extended an interest rate model introduced by Vasicek (1977), of the following form:

$$dR = \kappa\left[\frac{\theta(t)}{\kappa} - R\right]dt + \sigma dz \tag{3.41}$$

In this model, R represents an interest rate, but can easily be adopted as the asset price, whereas κ is the mean-reversion parameter as it is in Lucia and Schwartz (2001). The model can be shown to be equal to our one-factor mean reversion model presented in section 4.4.1, with a constant drift $(r\text{-}c)R_t$ replacing $\theta(t)$. A function $\theta(t)$ consistent with this model is then given as

$$\theta(t) = \frac{dF_0(t)}{dt} + \kappa F_0(t) + \frac{\sigma^2}{2\kappa}\left(1 - e^{-2\kappa t}\right) \tag{3.42}$$

In this equation, $F_0(t)$ represents the present value of a forward maturing at time t. For a more thorough discussion on mean-reversion models, see section 4.4.

3.3.4.3.2 Branching patterns

For mean-reversion processes, it is necessary to alternate the branching pattern slightly. Figure 3.3 illustrates the three possible branching structures of a mean-reversion tree.

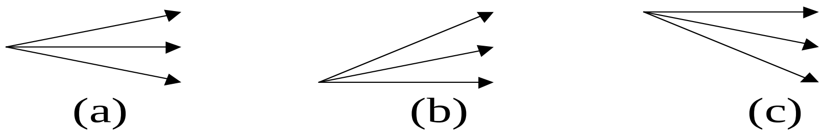

Figure 3.3: Alternative branching for a mean-reversion process

3.3.4.3.3 Generating the mean-reverting noise term

The tree building requires two steps. The first step comprises building a tree for the mean-reverting noise term, and the second step converts this tree to the term structure of the mean-reversion process. As seen in section 4.4.1, a mean-reversion process has a stochastic parameter following the dynamics given in equation (3.43):

$$dR^* = -\kappa R^* dt + \sigma dz \tag{3.43}$$

The discrete version of the process will assume a constant time step Δt, and the spacing between the nodes in the tree are taken as

$$\Delta R = \sigma\sqrt{3\Delta t} \tag{3.44}$$

Hull (2000) then defined[6] *(i,j)* as node coordinates, so that $t=i\Delta t$ and $R^*=j\Delta R$. A tree branching procedure with $\kappa>0$ will then follow branching method *(a)* from figure 3.3 until a barrier is reached when j is sufficiently large. Hull (2000) defined this barrier to be

$$j_{MAX} = -j_{MIN} = \left\lceil \frac{0.184}{\kappa \Delta t} \right\rceil \tag{3.45}$$

Beyond this barrier, the tree does not grow wider, and alternative branching procedure *(b)* or *(c)* from figure 3.3 are used as lower or upper branch respectively.

3.3.4.3.4 Risk-neutral probabilities

Using the branching above, three sets of risk-neutral probabilities are necessary. For branching method *(a)*, the following set was suggested by Hull (2000):

$$\begin{aligned} p_U &= \frac{1}{6} + \frac{\kappa^2 j^2 \Delta t^2 - \kappa j \Delta t}{2} \\ p_M &= \frac{2}{3} - \kappa^2 j^2 \Delta t^2 \\ p_L &= \frac{1}{6} + \frac{\kappa^2 j^2 \Delta t^2 + \kappa j \Delta t}{2} \end{aligned} \tag{3.46}$$

Similarly, for branching method *(b)*, the probabilities are given in equation (3.47):

$$\begin{aligned} p_U &= \frac{1}{6} + \frac{\kappa^2 j^2 \Delta t^2 + \kappa j \Delta t}{2} \\ p_M &= -\frac{1}{3} - \kappa^2 j^2 \Delta t^2 - 2\kappa j \Delta t \\ p_L &= \frac{7}{6} + \frac{\kappa^2 j^2 \Delta t^2 + 3\kappa j \Delta t}{2} \end{aligned} \tag{3.47}$$

Finally, the probabilities in a top node are given in equation (3.48):

$$\begin{aligned} p_U &= \frac{7}{6} + \frac{\kappa^2 j^2 \Delta t^2 - 3\kappa j \Delta t}{2} \\ p_M &= -\frac{1}{3} - \kappa^2 j^2 \Delta t^2 + 2\kappa j \Delta t \\ p_L &= \frac{7}{6} + \frac{\kappa^2 j^2 \Delta t^2 - \kappa j \Delta t}{2} \end{aligned} \tag{3.48}$$

[6] Observe that *i* is a positive integer indicating the depth of the tree, whereas *j* indicates the distance from the expected return in units of ΔR. At any depth i, $2j+1$ is the number of nodes at the same level.

3.3.4.3.5 Converting the tree

Given the process tree for R^* of the previous section, the task now becomes transforming this into a tree for the price process *R*. Define

$$\alpha(t) = R(t) - R^*(t) \tag{3.49}$$

Now, knowing the process of *R* from equation (3.41) and R^* from equation (3.43), it follows that

$$d\alpha = [\theta(t) - \kappa\alpha(t)]dt \tag{3.50}$$

This means that the solution for $\alpha(t)$ is given as

$$\alpha(t) = F_0(t) + \frac{\sigma^2}{2\kappa}\left(1 - e^{-\kappa t}\right)^2 \tag{3.51}$$

For small values of k, this converges to $\alpha(t)=F_0(t)+$**Error!**$\sigma^2 t^2$.

3.3.4.4 Finalizing the tree and obtaining a call option value

The construction of the actual tree is then done by starting at node 0, multiplying S_0 by an appropriate u or d so that this price branches out to create three values at the next level of the tree. The procedure is continued for the whole tree. Then, the value of the call option of the termination node is found as

$$c_{N,j}(T) = \max(S_0 u^j d^{N-j} - X, 0) \tag{3.52}$$

Here, $c_{N,j}(T)$ indicates the value of the call option at time *T*, the final termination date. *N* indicates the number of subintervals, measured as $T/\Delta t$, and j is the number of up movements at the current node. The value of an American option at earlier time *t* in the tree is given as the maximum of early exercise and waiting:

$$c_{N,j}(t) = \max\left\{S_0 u^j d^{N-j} - X, E_t\left[c_{N,j}(t+\Delta t)\right]\right\} \tag{3.53}$$

where Δt is the time step defined in the tree, S_t is the price in node *t*, and $E_t(\Theta)$ is defined below:

$$E_t\left[c_{N,j}(t+\Delta t)\right] = e^{r\Delta t}\left(p_U S_t u^{j+1} d^{N-j} + p_M S_t u^j d^{N-j} + p_D S_t u^j d^{N-j+1}\right) \tag{3.54}$$

After finding the value of the option in all the nodes by backtracking, the final value of the call option at time zero is found in the first node of the tree.

3.3.4.5 Discussion

This is a generic presentation of the approximation of an American call option using a trinomial tree. T solution is dependent on both the time resolution and branching rules considered. The approach presented in this chapter has utilized recommendations from Hull (2000), and is by no means perfect. The method nevertheless gives a very exact approximation for the option price if the tree is sufficiently large.

3.4 Optimal investment timing

This section is concerned with the theory of optimal investment timing. It is to a large extent based on Dixit and Pindyck (1994). Optimal investment timing is concerned with finding the maximum value of an investment by postponing it for the purpose of increasing the overall value. In practice, the approach returns a price above which investment is optimal.

McDonald and Siegel (1986) considered the following problem for an investment: At what point is it optimal to pay a sunk cost I in return for a project whose value is V, given that V evolves according to GBM:

$$dV = \alpha V dt + \sigma V dz \qquad (3.55)$$

Here, α represent the drift in the investment value, σ its volatility and dz a standard Brownian motion increment. Dixit and Pindyck (1994) commented that this implied that the current project value is known, whereas the future value is lognormally distributed with linear annual growth. The present value of the project payoff initiated at any future time *T* is given in equation (3.56):

$$F(V) = \max E\left((V_T - I)e^{-\rho T}\right) \qquad (3.56)$$

Here, *E* is the expectation operator, V_T the value of the revenues at time *T*, *I* the investment, and ρ the discount rate, in our case replaceable by *r*, the risk-free rate. V_T is assumed to follow GBM, and has expected value at time *T* given as

$$E(V_T) = V_0 e^{\alpha T} \qquad (3.57)$$

3.4.1 The deterministic case

The maximization of F(V) for a deterministic growth rate α, is trivial. This approach returns an optimal investment time T^*. Dixit and Pindyck (1994) showed that this can be done by setting the derivative of equation (3.56) equal to zero, and solving for *T*. This yields the following first-order condition:

$$\frac{dF(V)}{dT} = -(\rho - \alpha)Ve^{-(\rho-\alpha)T} + \rho I e^{-\rho T} \qquad (3.58)$$

Provided[7] $\rho > \alpha$, solving for *T* yields the optimal waiting time:

$$T^* = \max\left\{\frac{1}{\alpha}\ln\left(\frac{\rho I}{(\rho-\alpha)V}\right), 0\right\} \tag{3.59}$$

Looking at equation (3.59), it is evident that if *V* is close to *I*, $T^*>0$, and investment is suspended until T^*.

3.4.2 The stochastic case and the contingent claims approach

For the stochastic case, the previous approach is not possible, due to the fact that a definite point in time where $(V_T\text{-}I)$ reaches the value of optimal investment is not possible to determine. Instead, the decision rule changes into finding a minimum value V^*, above which investment is optimal. The investment is then initiated once $V>V^*$. The solution for the optimal value V^*, can be obtained either by dynamic programming, or by contingent claims analysis (CCA). This thesis is concerned with the latter approach.

3.4.2.1 A necessary precondition

Dixit and Pindyck (1994) described one necessary precondition for the CCA approach to function satisfactory. The exiting assets in the economy must span the stochastic changes in V. This means that it is essential that existing assets of the economy can replicate asset *V*'s payoff. As a result, there exists an asset x whose payoff is perfectly correlated with V. This is to capture the nondiversifiable risk of V.

3.4.2.2 The value of the project

Assuming constant investment costs, the isolated decision to invest in a project *F* at a time *T* in the future, has a present value in the form given in equation (3.56). This value is the present value of initiating the project at time *T*, receiving the benefits accrued from future contingent income, at the (sunk) investment cost *I*. McDonald and Siegel (1986) claimed that this could be seen as an exchange of *I* for the contingent asset *F(V)*. The decision to exchange is viewed as an irreversible decision.

3.4.2.3 The value of waiting

The value of waiting to invest can be seen as an option. This option is the right, but not the obligation, to invest in project *F* at one point in the future. Let us assume this

---

[7] If $\alpha>\rho$, it is always optimal to wait, and the value of the project increases indefinitely.

right can be exercised at any given point in a predetermined time period $\{T_{MIN}...T_{MAX}\}$. For the sake of simplicity, we assume T_{MIN} is zero. At any given time T, the value of waiting W_T is given as

$$W_T = \max[F_{T+\Delta T}, W_{T+\Delta T}]e^{-\rho\Delta T} \tag{3.60}$$

Equation (3.60) might require some explanation. The value of investing in the next time period is $F_{T+\Delta T}$, and the value of waiting is given as $W_{T+\Delta T}$. This value of waiting has structure similar to an American option, as discussed in section 3.3.4. The value of the option at time T is given as the discounted maximum of the two future values. Finally, the strike of the option to build is I.

3.4.2.4 Optimal exercise value

Following the above structure, the optimal exercise value V^* is found as the project value above which the value of investment exceeds the value of waiting. Specifically, V^* is the value of F, so that

$$F_T > W_T \tag{3.61}$$

Now, why not wait until prices increase even more? This can be explained by the dynamic relationship between growth in prices and discount rates. If exercise is delayed too long, the project might be worth less. This completes the option analogy.

3.4.2.5 Interpretation of the model

Recapping equation (3.59), this process tries to determine an optimal investment time, or in the presence of uncertainty, the price above which exercising a right to build a power plant yields the maximum payoff. The option structure uses the investment cost I as the strike and V_t as the project income. The value of an American option on the right to delay is defined as the maximum of exercising the option in the next period, or delaying the investment. The boundary value V^* above which exercise is optimal, is then found by comparing the value of waiting to the value of exercising for each node.

As mentioned, the strike of the project is its investment cost I. If I is increasing, the problem changes slightly. This was the original idea, put forward by McDonald and Siegel (1986), as they considered drift in both V_t and I_t. These is easily coped with by implementing minor changes to the tree.

4 Commodities and Energy models

4.1 Commodity characteristics

Energy is a commodity, or a consumption asset, as opposed to an investment asset, representing ownership only. Commodities differ from investment assets in several aspects. In this section we will highlight the main distinctions between a commodity and an investment asset. The differences can be categorized as follows:

1) Delivery
2) Convenience yield and storage costs
3) Volatility and price jumps
4) Short-run behavior

Finally, we will discuss the long-term properties for commodities, and some special features of energy prices.

4.1.1 Delivery

Commodities are based on actual delivery. The transfer of a commodity is hence troublesome, as great resources are needed. For example, cotton and wheat requires a vehicle, oil requires pipes, and energy requires a power grid. Investment assets are essentially paper assets, and delivery costs are virtually zero. Temporary delivery problems might hence affect the commodity price.

4.1.2 Convenience yields and storage costs

The relationship given in (3.15) indicates a non-negative relationship between the market price of risk and a positive return, since a risky asset would have higher required return than a risk free. This is true for an investment asset, but not, in general, for a consumption asset, due to convenience yield and storage costs.

4.1.2.1 Convenience yield

Hull (2000) defined convenience yields as the benefits obtained from owning an asset, but not obtained by holding a futures contract. These benefits include the ability to profit from local shortages, or the ability to keep production running. This effect has been investigated by, amongst others, Telser (1958), who investigated futures prices of cotton and wheat.

Gjolberg and Johnsen (2002) added that if inventories are full, and shortages are extremely unlikely, convenience yield would approach zero. They furthermore claimed that convenience yield for energy was highly convex in inventories, hence preventing stock-outs and acute power shortages.

4.1.2.2 Storage costs

For most commodities, ownership involves storage costs, as commodities involve physical delivery. For some commodities, such as steel, these costs might be substantial. Everything else equal, storage costs will shift the required return of the commodity above the risk-free rate.

4.1.2.2.1 Energy storability

For normal traders, energy is not storable. This has one major implication for pricing. Setting up an arbitrage portfolio using spot prices can not be created. Following this line of reasoning it is tempting to contend that if the arbitrage argument cannot be used for pricing, it cannot be used to make risk-less excess returns either. This is only partially correct. Gjolberg and Johnsen (2002) reported that arbitrage opportunities might occur for energy producers with water reservoirs, because a reservoir might be considered a type of storage available to suppliers only. Provided that production can be quickly initiated or increased, this could be considered an arbitrage opportunity. For ordinary investors, however, energy can not be stored.

The non-storability amplifies the price jumps, and hence the volatility, because inventories can not be used to “smooth away” temporary shortages.

4.1.2.3 Implications for pricing

In a risk neutral world, the growth of a commodity should include convenience yield c and storage cost u. The relationship in equation (3.16) should therefore be changed to the following risk-adjusted expected growth:

$$r-(c-u)=\mu-\lambda\sigma \tag{4.1}$$

Here, μ is the expected growth, λ the market price of risk, σ the volatility and r the risk free rate of return. A positive convenience yield thus lowers the required return on the asset, whereas the storage cost increases the required return. Finally, Gjolberg and Johnsen (2002) argued that the nature of hydropower producers implied that marginal storage cost *u* would be zero for electricity producers.

4.1.3 Volatility and price jumps

Due to the risk of temporary shortages, commodity prices are far more volatile than ordinary asset prices. Extreme situations of price shifts are normally referred to as jumps. These are results of either temporary supply shortages or demand shocks. For example, a sudden temperature decrease increases the demand for heating. As producers engage in arbitrage to benefit from these shortages, the prices are forced down again, leading to a mean-reverting behavior. These temporary irregularities finally lead to an advantage to holding spot over holding a forward contract, as discussed in the section above.

4.1.3.1 Energy volatility

In the Nordic energy market, hourly price jumps of 100% are known to happen. Lucia and Schwartz (2001) studied the Nordic power market, and found that the annual volatility of the daily log-prices was 189%. This is substantial, compared to other markets. The effect is mainly due to non-storability and bottlenecks.

4.1.3.2 Long-run volatility and futures prices

The volatility of long-run energy contracts is known to be much less than the volatility of short time contracts. As contracts are maturing and blocks are split up, prices become increasingly volatile. Hence, a good volatility function should decrease as a function of maturity. The low volatility of long-run prices has several reasons. Firstly, the uncertainty in long-term equilibrium should be reasonably unaffected by temporary shortages. Secondly, due to the illiquid nature of long-maturity contracts, the resulting volatility decreases. This can be viewed in light of non-trading effects, as investigated by for example Lo and MacKinlay (1999).

4.1.4 Short-run behavior

Schwartz (1997), amongst others, analyzed price behavior in commodity markets. His research indicated, as opposed to security markets, that prices did not seem to follow geometric Brownian motion. The effects, mainly attributed to the preceding three subsections, manifested themselves in autocorrelation patterns in commodity prices, and a tendency of prices reverting to a long-term equilibrium. Schwartz (1997) commented that when prices move out of an equilibrium setting (because of the previously discussed jumps), producers would try to profit from these conditions by adjusting supply, hence restoring equilibrium[8]. This pull-to-equilibrium pattern seems to be typical for commodities, but not for investment assets.

Knittel and Roberts (2000) investigated US energy prices, and commented that the prices in the Californian energy market did not pass unit root[9] tests for hourly price data. This should not be surprising. Autocorrelation is a basic characteristic of any mean-reverting model, due to the existence of a pull towards a long-term equilibrium. As already discussed in section 3.2.1.5, consecutive prices in a mean-reversion model are not as arbitrary as prices following a pure Markov process such as GBM.

4.1.4.1 Properties of energy price distributions

Because of the large volatility, and the likelihood of extreme price jumps, the distribution of the random component of energy prices is very fat tailed. Knittel and Roberts (2001) argued that the forecasting performance of energy models were unacceptable for any practical purposes, such as pricing of energy based financial products. They looked at heavy tailed distributions like students t and Levy processes as alternatives to the normal distribution.

Knittel and Roberts (2001) also reported an inverse “leverage effect”. This effect describes an asymmetric response in volatility to positive and negative price shocks. This asymmetry is mainly due to the impossibility of negative prices. Large negative price jumps are simply not possible.

Eydeland and Geman (1998) emphasized the need for a stochastic volatility model in order to capture the fat tails and spikes displayed by the distribution of realized power

---

[8] Fleten (2000) argued that the downward pressure is due to high-cost manufacturers entering the market, and consumers changing to cheaper sources of energy.

[9] A unit root test shows whether the difference between data at points t and t-1 has a trend. A failed test typically indicates possible autocorrelation in the data set, and that the time series is not stationary. Most autocorrelation tests should, however, be judged with a healthy amount of skepticism, as results in many cases are far from reliable.

prices. In conclusion, a stochastic volatility model might be crucial for making a good energy price model.

4.1.5 Equilibrium prices

As has been noted earlier, normal stocks tend to follow Geometric Brownian Motion (GBM), whereas commodity and energy prices are largely mean reverting. The problem with the mean reversion is that the actual mean is not observable, due to the reasons mentioned above. Nevertheless, using time series analysis, it is possible to generate a price mean to be used for modeling purposes. Whereas mean reversion models and autocorrelation can teach us a lot about short-time price behavior, it will not add anything significant to a model in the long run. This is because the long-run equilibrium is driven by the unobservable long-term mean.

Several authors have discussed the long-run characteristics of energy prices. For example, Pilipovic (1997) modeled long-run energy prices as GBM. Her research indicated the presence of a long-run trend in energy prices.

4.1.5.1 Seasonal prices

Energy prices are subject to large seasonal variations. As can be seen from figure 4.1, the one-week energy forward price is subject to large weekly variations. In the summer, the spot price is a lot less than in the winter. The Nordic energy market is different than most markets when it comes to this seasonal pattern. Because the demand for cooling is low, and heating is high, the peaks are in the winter season, instead of in the summer, like in, for example, California.

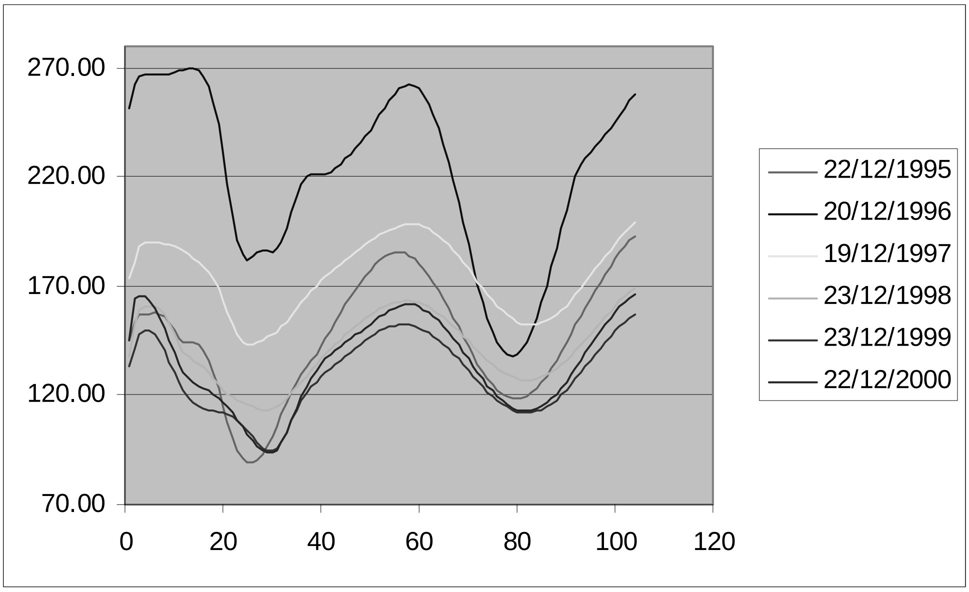

Figure 4.1: Seasonal pattern in energy prices

4.1.5.2 Short-term patterns

Johnsen (1999) and Lucia and Schwartz (2001), amongst others, commented that there are noticeable intra-week and intra-day patterns in the Nordic energy market. Figure 3.1 is an excerpt of daily price patterns from Nord Pool's spot (one-day forward). The graph shows that prices are peaking around 9 a.m. in weekdays, and a little later in the weekends. The general demand for electricity is also less in the weekends than in weekdays.

Intra-day and intra-week patterns will not have a large effect in our analysis, as this thesis will use one-week forward prices as basis for our decisions. If priced correctly, the one-week forward contract should reflect the average price during a week's production. For a run-of-river power plant that produces 24 hours a day and seven days a week, this should not have large impacts.

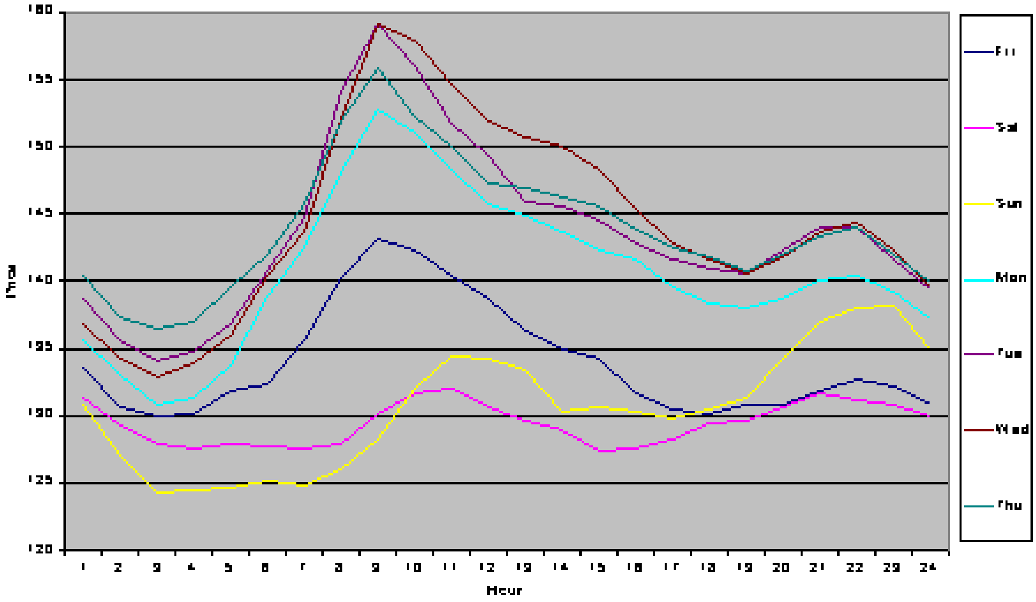

Figure 4.2: Intra-day prices for a random week. Source: www.nordpool.no

4.1.5.3 Market efficiency in electricity markets

Whereas normal capital markets are considered to be informationally efficient, Gjolberg and Johnsen (2002) commented that this might not be true for the Nordic energy market, and that arbitrage opportunities for producers might exist in combining future positions with storage. The market is still immature, at worst inefficient, although being one of the best-developed electricity markets in the world.

Market inefficiency is typical for a market not yet fully developed. In stock markets, it has been argued that brokers might be entitled to extraordinary returns because the brokerage right is very expensive. A similar argument can be used with power plants. For example, one might argue that the large investment costs of a power plant entitles the owner to extra-ordinary income. In that case, flexibility should be embedded in the sales price of traded power plants.

4.2 Commodity price models

In this section, we will present some commodity price models. These models are modeled for general commodities, or tailored to energy spot prices. They are, however, transferable to other variables. In this thesis, we will model precipitation, prices, revenues and river flow, so more than one pricing model is needed.

4.2.1 General commodity price models

Developed from the interest rate framework presented by amongst others, Vasicek (1977), commodity price models are often based on mean reversion. A mean reversion model is considered more able to capture the commodity characteristics outlined in section 4, than GBM or ABM. This only provides half the picture, however, as mean-reversion is a short-term property. Long-run commodity prices might have patterns closer to GBM. Nevertheless, most commodity and energy price frameworks utilize mean-reversion dynamics. The mean-reversion process has the property of reverting to an underlying function with time, and deals at least partially with local price jumps, by smoothing the differences down over time.

4.3 The number of factors

When working with stochastic models, we usually consider one, two or perhaps three factors in our analysis. The number of factors refers to the stochastic components in the model, and it is usually fair to assume that the addition of an extra factor increases the model complexity.

The tractability of the model depends on its purpose and the resources involved in maintaining it. Lucia and Schwartz (2001) commented that one-factor models were analytically very tractable. The simplest models might have an added advantage of being easy to communicate. Their explanatory powers, however, might be significantly lower compared to more advanced systems. The art of modeling is the art of selecting the appropriate tradeoff between simplicity and accuracy.

Pilipovic (1997) commented that the recent deregulation of the energy markets was bound to cause changes in the way the prices act. Hence, the price models working today might not work tomorrow.

4.3.1 Analytical modeling or simulation

The number of price models available is large and increasing. There is a basic tradeoff between deriving analytical solutions and obtaining prices through simulation. The models in this thesis are limited to those whose solution is developed analytically. We are aware of the fact that many good models may be based on simulation. Nevertheless, we will, due to time and spatial considerations, base our analysis on analytically tractable models.

4.4 Common spot price models

A number of authors have investigated commodity and energy spot price behavior. Gibson and Schwartz (1990) analyzed oil prices using a two-factor model where the long-run price followed geometric Brownian motion and the convenience yield followed mean-reversion dynamics. Schwartz (1997) tested this model against a one-factor pure mean-reversion model, and a three-factor model where the long-run price followed GBM, and both interest rate and convenience yield were assumed to possess mean-reversion characteristics. Although the models are based on various commodities, they seem to be more or less transferable to energy pricing, and offer analytical solutions of the SDE[10]. Gibson and Schwartz' convenience yield model is, however, perhaps not intuitive enough to catch the masses.

Lucia and Schwartz (2001) investigated various one-and two-factor models and their behavior in a pure electricity perspective. Pilipovic (1997) also developed a model, named the long-run/short run model, where the short-run deviations in price mean-revert to a stochastic long-term mean modeled as GBM.

[10] Stochastic Differential Equation.

Lucia and Schwartz (2001) discussed several models based on mean reversion. The following sections are based on their framework for modeling energy prices. Their models are largely based on Schwartz (1997). We start with the simple case of a single factor model in the price P_t.

4.4.1 One-factor mean reversion model of P_t

Let P_t be the spot price at time t, $f(t)$ a deterministic seasonal function of the price, and X_t the stochastic part so that prices follow the model given in equation (4.2).

$$P_t = f(t) + X_t \tag{4.2}$$

Furthermore, assume X_t follows a stochastic process of the form

$$dX_t = -\kappa X_t dt + \sigma dz \tag{4.3}$$

where $\kappa>0$, $X_0 = x_0$, and dz represents an increment in standard Brownian motion, explained in section 3.2.1. Then, letting $X_t = P_t - f(t)$, and applying Ito's lemma, we can rearrange the equations to produce a mean reversion process in P_t:

$$dP_t = \kappa\left(a(t) - P_t\right)dt + \sigma dz \tag{4.4}$$

where a(t) is a deterministic function of t, given as

$$a(t) = \frac{1}{\kappa}\frac{df}{dt}(t) + f(t) \tag{4.5}$$

This process is often referred to as an Ornstein–Uhlenbeck process. Solving the above one-factor model for P_t yields the following solution of the PDE:

$$P_t = f(t) + X_0 e^{-\kappa t} + \sigma \int_0^t e^{\kappa(s-t)} dz(s) \tag{4.6}$$

This model has several basic shortcomings. Firstly, it assumes the processes are risk-neutral, a fact that might be far from true. Secondly, Lucia and Schwartz (2001) commented that a great Brownian motion increment might result in negative prices, and would hence lack realism. Finally, as the model stands here, its forward variance approaches zero in the long run, as the long run mean is deterministic. This is a quite unrealistic assumption. Finally, Pilipovic (1997) argued that the one-factor spot price model performed unsatisfactory with respect to the fat-tailed distribution of returns.

4.4.2 One-factor mean reversion model of $\ln(P_t)$

For log-prices, this relationship becomes more difficult. Let $\ln(P_t)$ follow the dynamics given in equations (4.7) and (4.8).

$$\ln(P_t) = f(t) + Y_t \tag{4.7}$$

$$dY_t = -\kappa Y_t dt + \sigma dz \tag{4.8}$$

After combining the two, we get the following mean-reverting process

$$\mathrm{dP}_t = \kappa(b(t) - \ln P_t) P_t dt + \sigma P_t dZ \tag{4.9}$$

$$\mathrm{b(t)} = \frac{1}{\kappa}\left(\frac{\sigma^2}{2} + \frac{df}{dt}(t)\right) + f(t) \tag{4.10}$$

The derivation of equation (4.7) throughout (4.10) is given in appendix A. From the above equations, it is relatively easy to see that a mean-reverting process reverts to a long-term mean of *f(t)*. It should also be noted that the noise term in equation (4.9) is geometric in nature, and will grow for increasing values of P_t. The process for P_t, given in equation (4.4) is, in contrast, following a generalized Wiener process.

The one-factor model in $\ln(P_t)$ solves one of the problems presented in the first model. The Brownian motion increments in the log model can not result in negative prices. If the random shocks of $\ln(P_t)$ are normally distributed, the shocks of the true process will be log-normally distributed, and always positive. But the long-term volatility problem will still prevail.

Secondly, Pilipovic (1997) argued that whereas the log of price mean-reverting model seemed to perform satisfactory in capturing distribution width, it did not solve the fat-tails problem. Despite this, the log-price model still provides useful improvements.

4.4.3 One-factor volatility

A volatility model shown to be consistent with both the above one-factor models is given in equation (4.11). We will refer to this model as model A.

$$\sigma^A(t,T) = \sigma e^{-\kappa(T-t)} \tag{4.11}$$

In this equation, σ is the standard deviation, and κ is the mean-reversion factor.

Clewlow and Strickland (2000) discussed this type of volatility models, based on the analysis in Schwartz (1997). The model seemed to have great advantages in being simple and easy to interpret. It performed reasonably well in the short run, but as maturity increased, volatility approached zero. Consequently volatility in long maturity forward contracts would be understated, and for long-term planning purposes, the model would lose its tractability. This is a common problem with one-factor models. To increase the volatility to fit longer maturity contracts, an empirical model like the one presented in the next section could be used.

4.4.3.1 An empirical volatility model

Bjerksund, Rasmussen and Stensland (2000) suggested an approach to keep long-term volatility at a sensible level. Instead of the volatility given in equation (4.11) they introduced an empirical one-factor volatility model of the type given in equation (4.12). We will refer to this model as model B:

$$\sigma^{B}(t,T) = \frac{a}{(T-t+b)} + c \tag{4.12}$$

This model converges to a long-run volatility *c*. The parameters *a*, *b* and *c* are constants given by a nonlinear regression of the volatility of observed forward prices. Later, Koekebakker and Ollmar (2001) investigated this model and others by decomposing the variance through principal component analysis. This model will provide a more reasonable estimate for the long-term volatility.

4.4.4 Risk neutralization of the mean-reverting processes

To correctly price expected future spot prices, the spot price process has to be risk-neutral. By subtracting the risk from the expected return function, we end up with the risk-neutral function. From section 3.2.2.1, we know that if an asset follows a Wiener process with return μ and a standard deviation of σ, then the risk-neutral return is given as

$$\mu_{neutral} = \mu - \lambda\sigma \tag{4.13}$$

The mean-reverting processes can be neutralized in a similar way. Let equation (4.8) be replaced by the following equation:

$$dX_t = \kappa(\alpha^* - X_t)dt + \sigma dz^* \tag{4.14}$$

where

$$\alpha^* = -\frac{\lambda\sigma}{\kappa} \tag{4.15}$$

Investigating this equation, the only difference from equations (4.9) and (4.10), is that the drift is reduced by a factor of $\lambda\sigma$. In the above equations, dz^* represents a standard increment in the risk-neutral Brownian motion process under the risk-neutral probability measure and λ is the market price of risk, assumed to be constant. After rearranging, following the same steps as before, Lucia and Schwartz (2001) derived the following solution for the SDE:

$$P_t = f(t) + Y_0 e^{-\kappa t} + \alpha^*(1 - e^{-\kappa t}) + \sigma \int_0^t e^{\kappa(s-t)} dz^*(s) \tag{4.16}$$

The above equation is equal to equation (4.10), except for the risk factor $\alpha^*(1-e^{-\kappa t})$ and the risk-adjusted Brownian motion *dz*(s)*. Consequently, the price of a forward contract maturing at time T should be equal to the risk-neutral expected spot price at the same point in time.

$$F_0(P_0, T) = E(P_t) = f(T) + (P_0 - f(0))e^{-\kappa T} + \alpha^*(1 - e^{-\kappa T}) \tag{4.17}$$

As mentioned in section 4.4.1, one of the main disadvantages of the model above is that a large negative Brownian motion increment might result in negative prices. This problem can be overcome by changing to a log-price process, i.e. defining the process for $ln(P_t)$ instead of P_t as above. A risk neutral log price process has an analytical solution given as

$$\ln P_t = f(t) + Y_0 e^{-\kappa t} + \alpha^*(1 - e^{-\kappa t}) + \sigma \int_0^t e^{\kappa(s-t)} dz^*(s) \tag{4.18}$$

The expected value of the price, $E(P_t)$ is given as

$$E_t(P_T) = \exp\left(E_t(\ln(P_T) + \frac{1}{2} Var_t[\ln(P_t)]\right) \tag{4.19}$$

The extra term including the variance of $ln(P_t)$ is a consequence of Ito's lemma. Following this, the forward price is given as the expected future spot price under the risk-neutral probability measure.

$$F(P_0, T) = \exp(f(T) + (\ln P_0 - f(0))e^{-\kappa t} + \alpha^*(1 - e^{-\kappa t}) + \frac{\sigma^2}{4\kappa}(1 - e^{-2\kappa t})) \tag{4.20}$$

Clewlow and Strickland (2000) discussed the adequacy of the one-factor model presented by Schwartz (1997). They concluded that it behaved reasonably well in the short run, but equally bad in the long run, as the volatility function approaches zero for long maturity contracts.

4.4.5 Performance of one-factor models

How do the models behave in the market? Clewlow and Strickland (2000) discussed the adequacy of the one-factor model presented by Schwartz (1997). They concluded that it behaved reasonably well in the short run, but lacked realism in the long run, as the volatility function approaches zero for long maturity contracts. The forward volatility is given as

$$\sigma_F(t,T) = \sigma_{spot} e^{-\kappa(T-t)} \tag{4.21}$$

The problem is that as the maturity T increases, volatility converges to zero. As we will see later, Bjerksund, Rasmussen and Stensland (2000) avoided this problem by changing the volatility function to one converging to a non-zero value.

Secondly, if a one-factor is used, it is normal to assume zero drift. This implicitly means no uncertainty in the long run, making price predictions deterministic of looking sufficiently into the future. For variables with non-zero time value, this assumption is not adequate. We hence need a more sophisticated model.

Finally, the seasonal function of the one-factor models is very different. The model for P_t assumes constant seasonal variations, whereas the log-model has seasonal variations growing with P_t. For later models with drifts, this difference is not unimportant.

4.4.6 The Lucia and Schwartz 2-factor model

Lucia and Schwartz (2001) suggested a model composed of two factors, short run deviations χ_t and long-run equilibrium ξ_t where χ_t follows a mean-reverting process given in equation (4.23) and ξ_t follows a Wiener process given by equation (4.22). For brevity, the processes presented are assumed already risk-neutralized.

$$d\xi_t = (\mu_\xi - \lambda\sigma_\xi)dt + \sigma_\xi dz_\xi \tag{4.22}$$

$$d\chi_t = \kappa(\alpha - \chi_t)dt + \sigma_\chi dz_\chi \tag{4.23}$$

In the equations above, κ represents the mean-reversion parameter, μ_ξ the mean risk-neutral drift, dt the standard time increment, and σ_χ and σ_ξ are the volatility parameters of the processes. The Brownian motion increments dz_χ and dz_ξ are correlated as shown in equation (4.24):

$$dz_\chi dz_\xi = \rho_{\chi\xi} dt \tag{4.24}$$

This model can be applied to several commodities. As we have previously seen, for energy applications, the price seems to be best modeled using some seasonal factor *f(t)*. Lucia and Schwartz (2001) incorporated a deterministic seasonal function into the Schwartz and Smith (2000) two-factor model. This resulted in the following model:

$$P = f(t) + X_t + \varepsilon_t \tag{4.25}$$

In the above model, *f(t)* is a deterministic function, X_t follows mean reversion, and ε_t follows a generalized Wiener process.

4.4.6.1 Analytical solution for forward prices

Lucia and Schwartz (2001) presented the expected forward price for the two-factor model as

$$F_0(P_0,T) = f(t) + \varepsilon_0 + \left(P_0 - f(0)\right)e^{-\kappa T} + \alpha^*(1-e^{-\kappa T}) + \mu^* T \tag{4.26}$$

The model parameters P_0, *f(T)* and κ are the same as in the one-factor model. In addition, ε_0 is the long-term mean, μ^* represents the risk-adjusted long-term drift, and α^* is the risk-adjustment in the mean reversion, given in equation (4.27). Finally, σ_S is the short-run volatility.

$$\alpha^* = \frac{\lambda \sigma_S}{\kappa} \tag{4.27}$$

4.4.6.2 Short- and Long-run risk premiums

Schwartz and Smith (2000) argued that using two-factor models of this type, the short-term risk premium would be driven by the mean-reversion risk premium. The value of the mean reversion risk approaches zero for long maturity contracts, however, and will after approximately three years be very close to the long term risk premium. Hence, for a long-term model, we might choose to ignore mean reversion risk. This approach is similar to that of Schwartz (1998). This model will be discussed shortly.

4.4.7 Schwartz' one factor approximation

Schwartz (1998) suggested a different approach. A two-factor model following long-term GBM, and short-term mean-reversion can be expressed as follows:

$$dP = (r-\delta)Pdt + \sigma_L P dz_L \tag{4.28}$$

$$d\delta = \kappa(\hat{\alpha} - \delta)Sdt + \sigma_S dz_S \tag{4.29}$$

$$dz_L dz_S = \rho dt \tag{4.30}$$

This model has an analytical solution, but parameters are not easy to estimate. The model can, however, be approximated. Schwartz claimed that in the long run, this two-factor model could be simplified to follow dynamics of a single factor model following GBM:

$$dZ = (\mu - \lambda\sigma)Zdt + \sigma Zdz \tag{4.31}$$

Here, *Z* represents a shadow spot price, found by discounting long-term forward contracts at the drift rate. Using this approach, short-term deviations from the mean do not affect the shadow price *Z*. Continuing along these lines, he claimed that the risk-neutral drift should be fitted to forward prices. Hence, the drift in equation (4.31) can be substituted by

$$dZ = \frac{1}{F}\frac{\partial F}{\partial T} Zdt + \sigma Zdz = (r-c)Zdt + \sigma Zdz \tag{4.32}$$

Schwartz commented that this model seemed to be working satisfactory when time to maturity was more than three years. The final issue of this approach, is volatility, which is covered in section 4.4.7.2.

4.4.7.1 Seasonal variations on the Schwartz 1-factor approximation

The inclusion of a seasonal factor using Schwartz' 1-factor is trivial. Let

$$P = f(t) + Z \tag{4.33}$$

$$dZ = (r-c)Zdt + \sigma Zdz \tag{4.34}$$

Combining the two yields

$$dP = f'(t) + (r-c)(P - f(t))dt + \sigma(\mathrm{P - f(t)})dz \tag{4.35}$$

A solution of this differential equation is then given as follows:

$$P_t = f(t) + (P_t - f(t))e^{(r-c)t} = f(t) + \varepsilon_0 e^{(r-c)t} + \int_0^t \sigma e^{(r-c)t} dz \tag{4.36}$$

This model assumes deterministic seasonal variations, and is very similar to the mean-reversion model of section 4.4.1. In the model, ε_0 is the shadow spot price explained in the model above. This grows geometrically using the risk-free rate r-c. On top of this, a seasonal function *f(t)* gravitates around the long-term mean.

4.4.7.2 Two-factor volatility and Schwartz' approximation

Schwartz (1998) suggested using the forward volatility as an estimate for the variations in price. The volatility can be divided into long-run and short-run volatility. Long-run volatility can be seen as macroeconomic changes, and uncertainties in the long-term risk-free rate, etc, whereas the short-run volatility is more or less taken as temporary shortages, demand deviations, etc.

Now, for the two-factor model, volatility comprises a long-run part σ_{LR} and a short-run part σ_{SR} correlated with correlation coefficient ρ_{LS}. The forward variance is then given in equation (4.37).

$$\sigma_F^2 = \sigma_{LR}^2 + \sigma_{SR}^2 \frac{1-e^{-\kappa T}}{\kappa} - 2\rho_{LS}\sigma_{LR}\sigma_{SR}\frac{1-e^{-\kappa T}}{\kappa} \tag{4.37}$$

This model has several unobservable parameters, and the determination of the parameters is beyond the scope of this thesis.

4.4.8 Joining two dependent stochastic processes

A combined model for revenues has thus far not been discussed. This section gives a brief introduction to haw such a model can be modeled. Let us assume that we have two stochastic processes following GBM, whose dynamics are given in the equations below:

$$dP = \mu_P P dt + \sigma_P P dz_P \tag{4.38}$$

$$dQ = \mu_Q Q dt + \sigma_Q Q dz_Q \tag{4.39}$$

Dixit and Pindyck (1993) showed that these two processes, when multiplied together, produced a new stochastic process following GBM, given as

$$dR = (\mu_P + \mu_Q + \rho_{PQ}\sigma_P\sigma_Q)Rdt + (\sigma_P dz_P + \sigma_Q dz_Q)R \tag{4.40}$$

In the above equation, *R=PQ* could be considered the revenues of a project. The corresponding log-process follows a generalized Wiener process:

$$d\ln R = (\mu_P + \mu_Q - \frac{1}{2}\sigma_P{}^2 - \frac{1}{2}\sigma_Q{}^2)dt + \sigma_P dz_P + \sigma_Q dz_Q \tag{4.41}$$

Hence, over any time interval t, the changes in *lnR* are normally distributed with mean and variance stated in the equations under:

$$\mu_R = (\mu_P + \mu_Q - \frac{1}{2}\sigma_P{}^2 - \frac{1}{2}\sigma_Q{}^2)T \tag{4.42}$$

$$\sigma_R{}^2 = \left(\sigma_P{}^2 + \sigma_Q{}^2 + 2\rho_{PQ}\sigma_P\sigma_Q\right)T \tag{4.43}$$

For commodity forward prices, the risk-neutral drift *r-c* replaces the former drift μ. Solving the differential equation with the appropriate boundary condition, the forward price of the product *R*, is given by the product of the two forward prices, correlated with a correlation factor ρ_{PQ}:

$$F_t\left(P_T Q_T\right) = F_t(P_T)F_t(Q_T)e^{\left(\rho_{xy}\sigma_x\sigma_y(T-t)\right)} \quad (4.44)$$

This is an expression of the future value of revenues in one time period. The present value of the forward contract for delivery at time T is then given as follows:

$$V_0\left(F_0(R_T)\right) = e^{-rT}F_0\left(P_T Q_T\right) = e^{-rt}\left(P_0 e^{(r-c)T}\right)\left(Q_0 e^{-\lambda_Q\sigma_Q T}\right)e^{(\rho_{PQ}\sigma_P\sigma_Q)T} \quad (4.45)$$

Equation (4.45) contains several parameters. The shadow price P_0, quantity Q_0, convenience yield *c*, correlation ρ_{PQ}, volatility σ_P and σ_Q of price and quantity respectively, and the risk-free rate *r* can be estimated from historical values. The market price of discharge risk λ_Q is slightly worse to estimate, as forward contracts are not traded in the market. The next section suggests how this still can be determined.

4.4.8.1 Market price of risk

Essentially, equation (4.45) has two unknown parameters, λ_Q and $V_0\left(F_0(R_T)\right)$. In other words, provided we can obtain an estimate for V_0 for another traded discharge power plant the market price of risk is given by the parameters. Since single period revenues are not traded, this means that the traded value of a power plant is required. The next section gives the value of such a power plant.

Once λ_Q is determined, equation (4.45) becomes risk free, and should therefore be discounted at the risk-free rate. Rearranging the equation, we can estimate the present value of the revenues in period T as

$$V_0(R_T) = P_0 Q_0 e^{(-c-\lambda_Q\sigma_Q+\rho_{PQ}\sigma_P\sigma_Q)T} \quad (4.46)$$

4.4.8.2 Value of power plant

The accumulated value of the plant revenues is then the integral of the forward contracts throughout the plant's life, discounted at the risk-free rate.

$$V_0(R) = \left[\frac{P_0 Q_0 e^{t\left(-c-\lambda_Q\sigma_Q+\rho_{PQ}\sigma_P\sigma_Q\right)}}{c-\lambda_Q\sigma_Q+\rho_{PQ}\sigma_P\sigma_Q}\right]_{T_1}^{T_2} \quad (4.47)$$

Finally, the corresponding plant value is the revenues minus running expenses and investment costs. These expenses are discussed in section 10.4. Note that the interest rate r does not enter equation (4.47) explicitly, because the forward prices are assumed to follow a risk free trend given as r-c, and the risk free rate thus cancels out. We will, however, need the risk-free rate, as the risk-neutral drift of the forward prices is given as r-c, in our case 3.15%.

4.4.9 Modeling the seasonal function

As explained above, the price process consists of a deterministic and a stochastic part. The deterministic part of the equation is often considered being a seasonal factor. Modeling the deterministic part of the equation is normally done using ordinary least square (OLS) regression techniques. In some cases non-linear least squares (NLS) methods are necessary. Other techniques include maximum likelihood methods and filters, such as the Kalman filter. We will mostly use OLS and NLS methods to obtain the desired estimates.

4.4.9.1 Selecting the deterministic model

Various models for modeling prices have been suggested. Lucia and Schwartz (2001) suggested and tested four models, two based on ln(P_t) and two based on P_t directly. The seasonal variations are modeled either through a sine function, or twelve single dummy variables. The four models are given as

Model 1 (4.48)

$$P_t = \alpha + \beta D_t + \sum_{i=2}^{12} \beta_i M_{it} + X_t$$

Model 2

$$P_t = \alpha + \beta D_t + \gamma \cos((t+\tau)\tfrac{2\pi}{365}) + X_t \tag{4.49}$$

Model 3

$$\ln P_t = \alpha + \beta D_t + \sum_{i=2}^{12} \beta_i M_{it} + X_t \tag{4.50}$$

Model 4

$$\ln P_t = \alpha + \beta D_t + \gamma \cos((t+\tau)\tfrac{2\pi}{365}) + X_t \tag{4.51}$$

In the above models, P_t is the price, α is a constant, β, β_i and γ are coefficients, D_t and M_{it} are binary variables indicating weekend or month, τ is a phase-constant and X_t a stochastic factor. As always, t denotes time.

Models 1 and 3 use twelve independent seasons, where the individual β_is are measured as a monthly factor. Models 2 and 4, on the other hand, use a cosine factor to model the seasonal behavior. Which approach is the best, depends on the conditions. The twelve dummies would be best if the months are clearly independent, or if there is more than one demand peak a year. The dummy model will, however, perform worse in the transition between two months, because the function would jump to a new level. The sine function is smooth between the months, but fails to capture a more sophisticated trend than patterns of one crest and one trough. Adding higher order sine functions does, however, solve this problem.

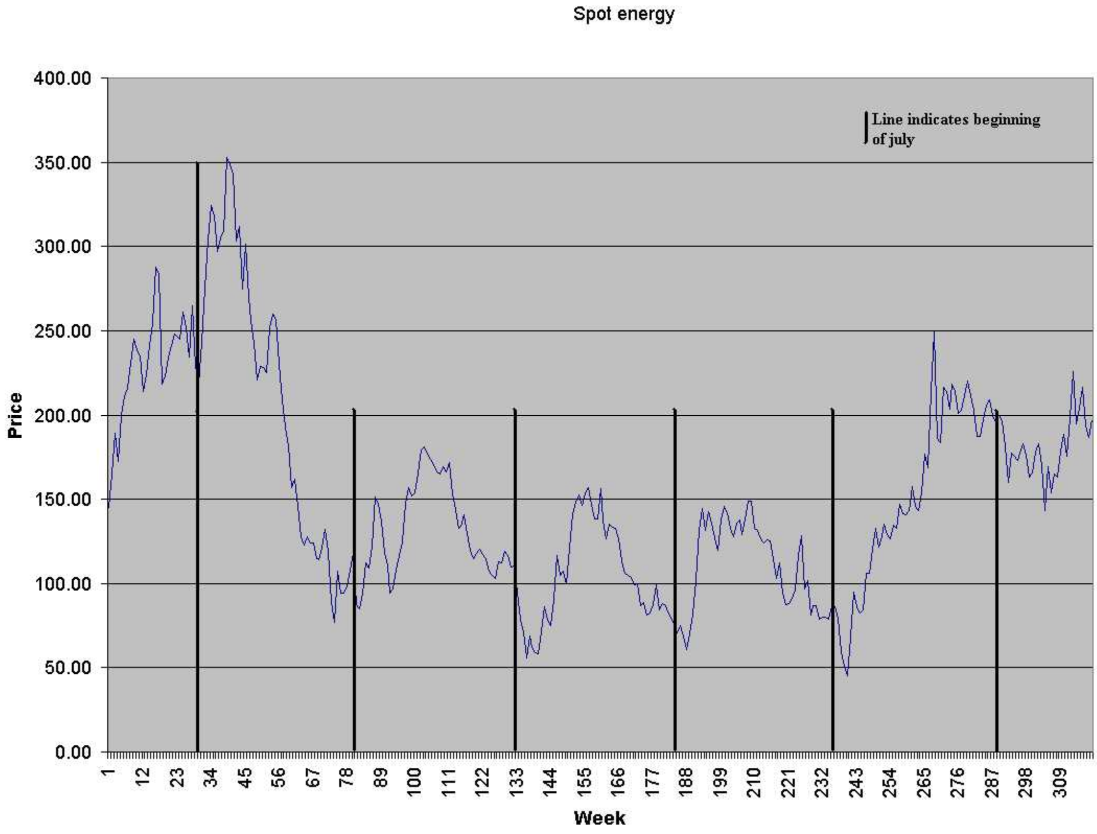

Figure 4.3: Spot prices with summer troughs.

In this thesis, we have based our analysis on one-week forward contracts, making the weekend factor obsolete in our analysis. Other data set anomalies are dealt with in the following manner:

- The last week of the year is combined with the first week of the next. Shifts in the seasonal function resulting from different week numbers are assumed minor and ignored.
- Easter week might have shorter demand. This is ignored.
- Summer holidays and public holidays are ignored.

As the Nordic energy market does not have an air condition peak in the summer, prices are assumed to be satisfactory modeled through a sine function. Looking at the graph in figure 4.3 we can see that this assumption is reasonable, as no large peaks appear in the summer. Looking at the forward curves, this seasonal pattern seems to be expected, at east in the near future. We hence model the seasonal factor as a cosine or sine function with a phase shift τ. The model presented under was tried both as a model for the spot price P_t, and its log price function $\ln(P_t)$.

$$f(t) = \alpha + \gamma \cos(\tfrac{2\pi}{52} t + +\tau) \tag{4.52}$$

4.5 Risk-free rate of return

The risk-free rate of return was estimated as a weighted average of Norwegian long-term bonds. The data was obtained from Norges Bank (www.norges-bank.no), and contained the bond rates over the last six years, divided into maturity of 3, 5 and 10 years. The risk-free rate was then estimated as the average of the three, using the bond length as weights.

4.5.1 Results

The procedure presented above resulted in a risk-free fate estimate of 5.88%. This is used throughout our analysis.

5 Parameter estimation of Spot Price Models

To find the value of the power plant by the use of risk-neutral valuation, we need to find a function for the risk-neutral expected forward prices. These prices will be used as strike price for the option valuation approach in section 11.1, and as part of a function for the implementation of the optimal exercise price. This section is therefore devoted to the determination of energy price and forward price parameters. Although we are aware there are more advanced models than the ones introduced in this section, we will, due to simplicity, only investigate one-factor models.

The section starts with a description of the data set used and a discussion on the possible candidates for a deterministic seasonal function. The expected spot prices are then estimated, and a discussion on risk-neutralization based on forward prices follows. The analysis is independent of model chosen, although the final regressions for the various spot models will be model specific.

5.1 Description of the price data set

To predict prices, an appropriate term structure model is necessary. In order to make the problem analytically tractable we desire a collection of energy forward contracts of equal length, for all desirable maturity dates. This is not readily available at Nord Pool, because of the block structure, resulting in forward contracts of unequal length and time of maturity. For planning problems with seasonal implications, this is unsatisfactory. The problem does not disappear when contracts approach maturity, although the number of short maturity contracts is greater. The second problem is that energy is a flow commodity. Since a contract might span several days, its price needs to reflect an average load during the contract delivery period. Lastly, long-term energy contracts such as seasons or years are highly illiquid. This means they might not reflect the actual prices at any given time.

To resolve the problems outlined in the above sections, and produce a sensible resolution for planning problems, Fleten and Lemming (2001) made a data set dividing energy contracts into one-week forward contracts. This is the data set we will be using in the pursuing analysis. The data set has been created based on the available contracts in the Nordic energy market, and smoothed over the whole time specter. We will briefly explain how Fleten and Lemming has created the data set.

5.1.1 The Fleten and Lemming data set

Fleten and Lemming (2001) used a data set of contracts in the Nordic energy market. By assuming that contracts were delivered at a constant rate throughout the year, the large contracts were split into weekly contracts. Then, these contracts were averaged over the delivery period. The method assumed a constant risk-free rate r to get the following expression for the averaged forward price:

$$F(T_0,T_1,T_2)=\int_{T_1}^{T_2}\frac{e^{-rs}}{\int_{T_1}^{T_2}e^{-rs}ds}f(T_0,s)ds \tag{5.1}$$

Here, the forward price $F(T_0,T_1,T_2)$ is an average of the series of weekly forward contracts $f(T_0,s)$ over the time period T_1 to T_2, appropriately discounted at the risk-free rate r.

Finally, Fleten and Lemming (2001) smoothed the prices over the whole interval to obtain a smoothed forward curve, utilizing the two contracts immediately preceding and following the actual contract.

We have decided to use only the parts of the data set that includes a full two-year forward curve of one-week contracts. To avoid overlapping, we use the closing price of every Friday for contracts delivering on a Monday-to-Sunday basis. The 1-week forward contract is thus three days from maturity, and is taken as our spot price. Using this criterion, the procedure resulted in our data set, consisting of 319 two-year (104 weeks) forward curves, or a total of 33176 points of data. The 319 data points hence represent 319 spot[11] contracts, i.e. approximately six years of data, ranging from December 1995 to January 2002.

5.1.2 Ten year contracts

The second data set used comprises ten-year contracts from 1992 to 2001, given at a weekly resolution. A ten-year contract is a contract for constant delivery for ten

[11] As mentioned, spot in this context, means a three-day forward contract for steady delivery in one week.

consecutive years, starting at the beginning of the next calendar year. The price shown is hence the expected average price during the next ten years.

5.1.3 One-year forwards

One-year forwards are contracts for constant delivery during a given year. Our third data set comprises the longest maturity contracts for Nord Pool from the years 2000 and 2001. The set from 2000 therefore includes one-year forwards for 2001, 2002 and 2003, whereas the data from 2001 include one-year forwards for 2002, 2003 and 2004. The data set includes prices with weekly resolution.

5.2 Estimating the long-term drift

The long-term risk-neutral price trend should be fitted to match the forward data available in the market, because forward prices represent a risk-neutral estimate of the expected price. It is common practice to adjust only the drift term to make the process risk-neutral. Hence, if the trend is risk neutral, then the prices based on the trend should also be risk-neutral.

Since our project has a life span of 40 years, we need to obtain contracts maturing as far as possible into the future. It will not be sufficient to estimate the trend over the next three years. The following points are considered supportive evidence for this conclusion:

- The unclear shape of the seasonal factor present in short term contracts makes the trend blur.
- Short-term deviations from the long term mean, due to for example reservoir levels will produce out-of-equilibrium prices.
- The general unobservable nature of the long-term mean in short-run prices

Instead of using the Nord Pool prices, the market's opinion on a future level of the energy prices can be found in prices so far into the future that temporary deviations are assumed to be zero. The problem with this approach is that the collection of forward data available at Nord Pool does not include long maturity contracts, as the longest contracts started trading three years before delivery. For longer maturity contracts, we need to turn to the OTC[12] market, to obtain price data for ten-year contracts.

5.2.1 Estimation approach[13]

To estimate the trend in the forward prices, we need as much information as possible. Assuming that a ten-year contract starts the delivery next year, we will use the next three one-year contracts to estimate the immediate prices. After year 3, Nord Pool does not provide longer contracts. We then assume that beyond the closest three years, forward prices grow at a constant rate, equal to *g%* annually, based on the forward contract in the market with the longest time to maturity. In other words, the forward contract for year 3 serves as a "shadow price" for the trend beyond this point. This forward growth rate is assumed to be risk free, as forward contracts are priced using risk-neutral measures. If our first year of forwards is 2002, the structure becomes as follows:

Year 1	Year 2	Year 3	Year 4	Year 5	Year 6
FWYR-02	FWYR-03	FWYR-04	(1+g)FWYR-04	$(1+g)^2$FWYR-05	…

Table 5.1: The growth structure of forward prices

Following this, we now have a series of actual and implied one-year contracts, up to ten years into the future. The (risk-free) discounted value of these contracts should be equal to the risk-free discounted value of the ten-year contract. Figure 5.1 summarizes the approach.

Y to mat	Year	One-year	Discounted		10-year	Discounted
0.19	2001	139.30	137.73		160.50	158.70
1.19	2002	144.32	134.55	Actual contracts	160.50	149.64
2.19	2003	149.41	131.35		160.50	141.10
3.19	2004	153.76	127.46		160.50	133.05
4.19	2005	158.25	123.70	Contracts based on trend	160.50	125.46
5.19	2006	162.86	120.04		160.50	118.30
6.19	2007	167.61	116.49		160.50	111.54
7.19	2008	172.50	113.04		160.50	105.18
8.19	2009	177.53	109.70		160.50	99.18
9.19	2010	182.71	106.46		160.50	93.52
				Difference		
		Sum	1323.8343	-1.91719E-08	Sum	1323.8343

Figure 5.1: The estimation procedure for the long-term trend.

[12] Over the counter.

[13] This approach is based on an approach in Dobbe and Sigmo (2002).

5.2.2 Results of the long-term trend estimation

The procedure described above was run on a data set containing prices from the years 2000 and 2001. It should be noted that the ten-year contracts are quite illiquid, and would hence probably produce an erroneous magnitude of the volatility. We therefore do not attempt to use these prices to find volatility estimates.

Using the Excel solver, an average drift of the forward prices was estimated by minimizing the squared difference between the discounted values of the ten-year contracts and the trended data for 47 different trading days in 2000 and 2001. The results are summarized in table 5.2. As can be seen, the geometric forward price trend was estimated to be 3.15% on average. This corresponds to a drift[14] in *ln*(P_t) was 0.61%. Looking at the last few years of actual spot prices, these are subject to a negative trend. This is, however, mainly due to the dry year of 1996. Looking at the two-year forward curves estimated from Nord Pool, an assumption of 3.15% growth seems plausible.

Parameter	Annual trend	Low	High	Standard deviation
Geometric trend in P_t	3.15%	1.42%	5.45%	1.24%
Arithmetic trend in P_t	3.38%	1.46%	6.05%	1.41%
Implied trend in $\ln P_t$	0.61%	0.25%	10.59%	0.27%

Table 5.2: Results of the trend estimation

5.3 Seasonal function

Section 4.4.9 outlined several candidates for the seasonal function for energy prices. The presence of a price crest in the winter and a corresponding trough in the summer suggests that a normal cosine function should have the ability to capture the long-term trend. This selection is backed by the findings of, amongst others, Lucia and Schwartz (2001). The seasonal function for prices is therefore given as

$$f(t) = c + \gamma \cos\left(2\pi\left(\frac{t}{52} + \tau\right)\right) \tag{5.2}$$

[14] The drift in $\ln(P_t)$ is $(r - c - ½\sigma^2)$. See for example Hull (2000). The trend was implied from the results of the geometric trend, and requires a starting point. For P_t=172, the drift is **Error!**.

5.4 Appropriate spot price models

There are several approaches and models possible to estimate the model parameters for the price. In section 4.4, we discussed several candidate models for the spot price. Since our planning project spans 40 years, a one-factor model without drift would be worthless. We therefore impose a deterministic drift in the model. For the sake of comparison, we will use the following two models:

Model 1: A one-factor mean-reverting price model, from Lucia and Schwartz (2001)

Model 2: A one-factor log-price model with short-term mean reversion, as given in section 4.4.2.

5.5 Regression results for the spot price models

The following tables include the results of the parameter estimation of the spot price models. The parameters are estimated using a non-linear regression approach. A total of 319 spot prices were used in the regression.

1-factor spot price model				
Expected spot price model:	$E_0(P_T) = f(t) + (P_0 - f(0))e^{-\kappa T}$			
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
CONSTANT	158.617	28.278	102.978	214.256
GAMMA	20.762	9.523	2.026	39.499
TAU	-0.020	0.072	-0.161	0.122
KAPPA	0.029	0.013	0.002	0.054
R^2	0.946			

Table 5.3: Parameters of the 1-factor model for price P_t

In table 5.3, the constant given reflects the starting value at the beginning of the data set. A price series starting at a later time T using this value must therefore be adjusted with a drift. This drift will be assumed constant, and was determined in section 5.2. For example, the constant one year later, is expected to be

$$158.62e^{0.0315} = 163.69 \tag{5.3}$$

The next model, is a one-factor log-price process. Here, the drift is dependent on the starting point P_0. The log-trend can then be determined as ln(1+0.315)/ln(P_0)=0.61. This drift is not explicitly modelled here, but used later to take into account the drift in energy prices.

1-factor log-price model with deterministic drift				
Expected spot price model:	$E_0(P_t) = \exp\left(f(t) + \left(\ln P_0 - f(0)\right)e^{-\kappa t}\right)$			
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
CONSTANT	4.975	0.161	4.657	5.293
GAMMA	0.196	0.065	0.068	0.324
TAU	-0.032	0.052	-0.135	0.070
KAPPA	0.035	0.015	0.006	0.064
R2	0.939			

Table 5.4: Parameters of the 1-factor model for $\ln(P_t)$

It should be noted that the phase angle τ was not significant for any of the models. This is mainly due to the peak in energy prices around January. A cosine function starts at a maximum value when (**Error!**+τ)=0, and τ would hence bee too close to zero to be significantly different from zero using only 319 data points. Another problem is that κ is quite small in these regressions. This might be explained by the fact that the model does not capture the dry/wet season differences in a satisfactory way. An alternative regression fitted to dry years, where the seasonla function fits better, is included in appendix B. Using this approach, the magnitude of κ increases to about 0.135, a more realistic level.

5.6 *Spot price volatility*

Spot price volatility can be estimated in a number of ways. For a simple oberview of the most common approaches, see Hull (2000). The methods include finding the differences between consecutive prices, or calculating the standard deviation of returns. The return can be calculated in different manners. One suggestion, used by Koekebakker and Ollmar (2001) is as follows:

$$return = \rho_t^{\,1} = \ln\left(\frac{P_{t+1}}{P_t}\right) \tag{5.4}$$

Assuming that returns are "small", this is an adequate approximation. A second possibility is to calculate the returns as

$$return = \rho_t^{\,2} = \frac{P_{t+1} - P_t}{P_t} \tag{5.5}$$

The above approaches only work if seasonal and trend effects are significantly smaller than the noise. If this condition is not appropriate, the data has to be de-trended before the procedure can be used. Investigating energy prices, an assumption of volatility much greater than the trend seems plausible. Having obtained the returns, an estimate for the volatility of spot prices, can be found as the standard deviation of the returns.

$$volatility = \sigma = \sqrt{\frac{1}{t-1}\sum_t (\rho_t - \bar{\rho})^2} \tag{5.6}$$

This gives the weekly volatility. An annualized measure can be obtained by multiplying by √52.

5.6.1 Results

The estimation of spot price volatility was done using the 319 one-week forward prices as proxies for the spot price. The results are presented below:

Model	**Weekly volatility**	**Annual volatility**
Model 1	10.22%	73.73%
Model 2	10.48%	75.58%

Table 5.5: Spot price volatility

The two estimates above are almost identical, as they should be, as model 1 is an approximation of model 2. We will nevertheless mostly use model 1, as this is most convenient.

Observe that this volatility reflects the average volatility throughout six years of data. If a more exact volatility pattern is needed, GARCH models or other autoregressive volatility estimators might prove to be useful. We will, however, not attempt to fit the volatility to the seasonal patterns.

5.7 Market price of risk

The model parameters presented in the previous section comprise an expected spot price model. It does not reflect a risk-neutral spot price, however, as the regression is done on spot prices. To find the market price of risk, we need to estimate the systematic difference between the model built and the actual forward contracts given in the data set. Sadly, the seasonal function, as presented in equation (5.2) does not appropriately capture the seasonal variations in the spot prices. The deviations caused by the seasonal function are hence greater than the risk. In effect, it has been impossible to develop a fair estimate for the market price of risk using these data.

We do, however, offer a solution to the problem. We have obtained a data set of expected spot prices. The difference between these data sets gives an estimate of the market price of risk.

5.7.1 The data set of expected spot prices

The data set includes expected spot prices at 48 selected dates over the same period as our original data set. A specific date includes an uneven number of spot price predictions, ranging from 105 to 254 weeks into the future. Subtracting the forward prices at these dates from this data set produces a new data set, explainable only through risk. By performing this subtraction, we obtain a new data set of 48 trading days, and a total of 4680 differences.

5.7.2 Estimation of the market price of mean-reversion risk

Using a one-factor price model of the type given in section 4.4.1 or 4.4.2, the difference between forward and expected spot price is given as

$$d = \frac{\lambda\,\sigma}{\kappa}(1 - e^{-\kappa T}) \qquad (5.7)$$

The λ of equation (5.7) represents the market price of mean-reversion risk, and is present in both the log model and the pure price model of Lucia and Schwartz (2001). The volatility is given as σ, and measured from the spot prices. Finally, κ is the mean reversion parameter, and T represents the weeks to maturity. We can now estimate the risk as the difference between the two observed values.

5.7.2.1 Results

An estimate for σ was obtained in section 5.6, using spot prices. The estimate of κ was estimated in section 5.5, and found to be 0.029 on a weekly basis for the one-factor mean reversion model. Using a non-linear solver, we regressed equation (5.7) against the systematic differences. The estimate for λ was hence found to be 9.82 NOK/KWh for the price model. Given this, a long-term measure of risk can be taken as the value of equation (5.7) with $T=\infty$. This yields an overall mean-reversion risk level for long-term contracts approaching 33.86 NOK, or 19.26% for the log model[15] described in section 5.5.

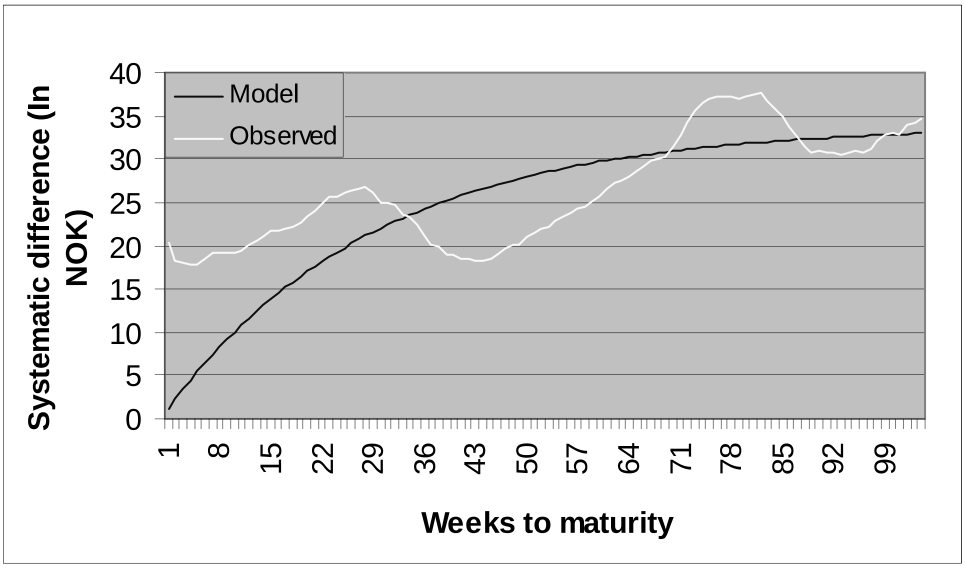

Figure 5.2: Systematic differences between expected spot prices and forward prices

As can be seen in figure 5.2, the function does not fit short maturity contracts. The shape of the observed data in the short end of the curve should be taken with a bit of caution, however. Since energy traders hedge their exposure in the short run, the market price of risk for short maturity contracts should in fact be expected to be negative. Hence, the expected forward prices seem to be over-estimated for short-run prices. Since short-term structures are not relevant for pricing long-term contracts, however, the long end of the curve is considered more important. This is also where the function from equation (5.7) fits the best.

5.7.2.2 Discussion

Lucia and Schwartz (2001) investigated the market price of risk for daily contracts. They found that the RMSE[16] of prices was approximately NOK 11.90 on a daily basis, and that λ was almost always positive. For longer contracts, a considerably higher risk premium could be expected if the relationship in equation (5.7) holds. Hence, a risk premium of 19.26% might prove to be a fair value for the long-run

---

[15] The log change is dependent on the initial value of P_t. The log-change is then found as the log of the actual price change due to the risk-neutralization, or as risk/price.

[16] Root Mean Squared Error

mean-reversion risk. We note, however, that this is a high level of risk, and that uncertainties in the long-term mean might explain a considerable amount of risk.

5.8 Forward price models

This section is devoted to pure forward price modeling. The analysis above caused several problems when the market price of risk was to be determined. Instead of going through spot prices to obtain risk-neutral prices, we can model the forward prices directly, as forward prices are, per definition, risk free. This simplifies the estimation procedure considerably. The downside is that this approach leaves the market price of risk undetermined. The analysis in this section is very similar to the spot price analysis. The section will therefore be briefer in nature.

5.8.1 Data set

The data set contained 319 weeks of data, each week comprising 104 forward prices with maturity from 1 to 104 weeks. The set is the same as used in the spot price models above, only the whole data set is used to determine the shape of the forward curve instantaneously.

5.8.2 Seasonal function and models

The seasonal function is the same as of the expected spot price models discussed in section 5.3, and the forward price models are essentially equal to the expected spot price models from section 5.4.

5.8.3 Regression results

The following section lists the parameters estimated from the forward data. The forward curves used for the regressions are given in the tables. Observe that for the two-factor model, the market price of risk is not part of the regression, because the forward prices are already risk neutral.

1-factor forward price model				
Forward model: $F_0(P_0,T) = f(t) + (P_0 - f(0))e^{-\kappa T}$				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
CONSTANT	169.374	0.180	169.022	169.726
GAMMA	28.110	0.108	27.898	28.322
TAU	0.934	0.001	0.933	0.936
KAPPA	0.014	0.000	0.013	0.014
The long-term riskless drift μ^* is taken as 3.15% annually.				
R^2	0.869			

Table 5.6: Parameters of the 1-factor approximation of the forward price model

1-factor forward log-price model				
Forward model: $F_0(P_0,T) = \exp\left(f(T+t) + \left(P_{0,t} - f(t)\right)e^{-\kappa(T+t)}\right)$				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
ε_0	5.131	0.001	5.129	5.134
GAMMA	0.189	0.001	0.187	0.190
TAU	0.937	0.001	0.936	0.938
KAPPA	0.017	0.000	0.016	0.017
The long-term riskless drift μ^* is estimated to be 0.61% annually, based on the results of section 5.2.				
R^2	.833			

Table 5.7: Parameters of the 1-factor forward log-price model

The parameters presented above comprise a risk-neutral estimate of the forward curve, and do not, as spot prices, require risk-neutralization.

5.8.4 Forward volatility

How do we obtain an estimate for the forward volatility? Koekebakker and Ollmar (2001) suggested two different approaches for this purpose. Either the difference between two consecutive prices could be used, or the return between these. In section 5.6, we presented two ways of estimating spot price returns. In this section, we will use model 1 to find an estimate for the return, and then use two forward volatility models to fit the existing data.

In section 4.4.3, we introduced two different forward volatility models, and called them model A and B:

$$\sigma^A(t,T) = \sigma e^{-\kappa(T-t)} \tag{5.8}$$

$$\sigma^B(t,T) = \frac{a}{(T-t+b)} + c \tag{5.9}$$

Model A is consistent with a normal mean-reversion model, whereas model B is an empirical model fitted to observed forward prices. These models can be fitted to the observed volatility of the available price data. Koekebakker and Ollmar (2001) investigated the volatility of energy prices, and suggested using the forward contracts of length 1, 2, 3, 4, 5, 6, 7, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 70, 88 and 104 weeks. This captures the level of the long tem contracts, without over-fitting the long run smoothed prices. We will try to fit both the models to both the selection suggested above, and the whole data set, because the long-term perspective enforces our interest in long-term volatility behavior.

5.8.4.1 Method used

The data set used was the normal smoothed set of forward prices. For every (weekly) forward maturity T=1 to 104 the return was calculated using the formula given in equation (5.10). This is identical to model 1 of section 5.6.

$$return = \ln\left(\frac{F_{t+1}}{F_t}\right) \tag{5.10}$$

At each maturity *T*, the standard deviation of the return is a measure of the volatility. This standard deviation of return was then plotted against maturity. This was done both for the selection suggested by Koekebakker and Ollmar (2001), and for the whole data set. The resulting graph was fitted to model A and B using a non-linear regression optimizer in SPSS.

5.8.4.2 Estimation results

The regression parameters are reported in table 5.8 and table 5.9, and a plot of the functions is given in figure 5.3.

Parameter estimation for model A				
Model A using the selection				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
σ	.650	.000	.650	.650
κ	.064	.017	.027	.101
R^2	.95768			
Model A using all data				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
σ	.465	.000	.465	.465
κ	.012	.010	.009	.032
R^2	.82998			

Table 5.8 Model A parameter estimates

Parameter estimation for model B				
Model B using selection (annualized estimate)				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
A	2.511	.965	.476	4.550
B	5.030	.015	5.000	5.060
C	.211	.458	-.755	1.177
R^2	.96298			
Model B rerun with all the data (annualized estimate)				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
A	5.980	.571	4.843	7.109
B	12.324	1.275	9.795	14.852
C	.145	.008	.130	.1604
R^2	.93863			

Table 5.9: Model B parameter estimates

As can be seen from table 5.9, the results from the regression using the selection produces an insignificant long-term trend, whereas the regression using all the data produces a highly significant, but lower estimate.

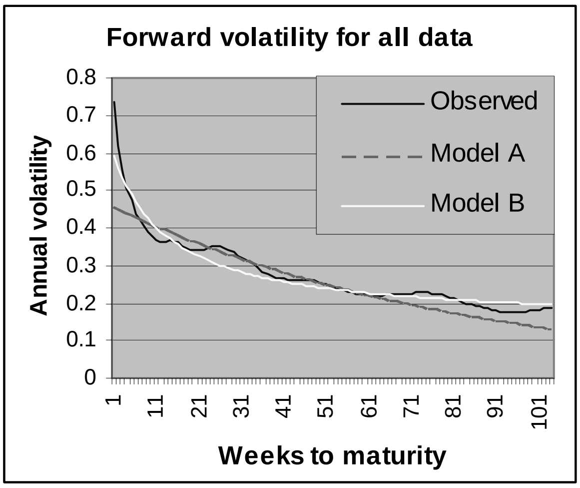

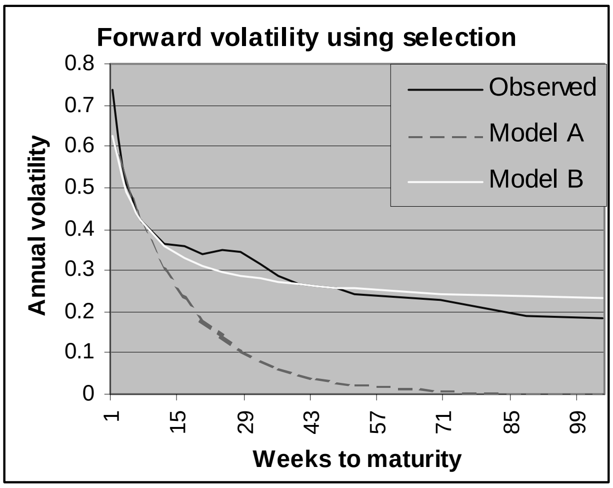

Figure 5.3: Volatility and volatility functions

5.8.4.3 Implied volatility

The implied volatility is calculated as the volatility of options contracts. It is found by observing the option price, forward price at maturity, strike and time to maturity. Using an option formula, such as a modification of Black and Scholes (1973), the volatility is fitted to the observed data. Using an excel spreadsheet obtained from APX, we found that an option maturing four years into the future has implied volatility of approximately 15%. This should be the approximate level of the long-term volatility. Investigating table 5.9 we find that model B fitted to all the data converges to a long-term volatility of 14.5%, very close to the implied volatility.

5.8.4.4 Discussion on forward volatility

Comparing our two models to the implied volatility, we find that model B produced the best results. As can be seen in figure 5.3, model A greatly underestimates the long-term volatility for the selection, and worsens for long maturity contracts. Although regressing over the whole data set improves the estimate, it still approaches zero for long contracts. Model B is considerably better, but the constant level c is not significant using only parts of the data set. This could be expected, however, as the analysis after 52 weeks contains a total of 3 data points. Consequently, the models fitted to all the data performed considerably better in the long run than the model fitted to the selection only. Although over-fitting might occur, it is important to have the volatility as good as possible in the correct interval. When pricing long-term contracts, the volatility structure should be made consistent with observed prices, as long as the model permits these changes. Besides, the results match the implied volatility observed in the market. In conclusion, model B is the better choice for a long-term price volatility model.

5.8.5 Discussion of the models estimated

Looking at the results of the spot parameter estimations, we see that the seasonal variations are large and significant, but not equal. Investigating forward contracts, we find the presence of the same kind of trend. We are therefore fairly confident about the future presence of a defined seasonal shape in prices.

For the long-term trend, our estimate of 3.15% annual price growth seems to be fair. The number is fairly uncertain, however, as the value ranges from 1.42% to 5.45%. Our estimate of 3.15% annual growth gives a price growth of approximately 5 NOK from a power price of 160 NOK. According to an anonymous[17] source, this assumption was fair, and close to the company’s own measures. The determination of price trends in the main data set obtained from Nord Pool, proved to be impossible. This is due to a variety of reasons. Firstly, the data set is too short, and too dependent

[17] The comment was considered privileged information, and he therefore listed anonymously.

on short-time deviations. Complex seasonal patterns and dry years complicate the picture considerably, and make the use of ten-year contracts necessary.

It should also be mentioned that the data set is relatively speaking far too short to estimate the parameters needed. Modeling the unobservable mean by the deterministic function might prove to give bad results. The mean itself becomes a function lying in the middle of dry and wet years, not necessarily reflecting the real mean.
Another approach would be to model energy prices using the wet years[18] only and then use a jump probability for dry years. In appendix B, the parameters for the process are measured on the basis of the wet years. The estimates of the mean-reversion and seasonal effects are far greater for this set. This suggests that it might be fruitful to estimate short-term variations on only parts of the data set.

In order to properly value energy prices, a higher order model, such as a two-factor model, might be considered. The addition of jump-diffusion and time-varying volatility will probably improve the forecasting abilities. For a discussion of these models, Knittel and Roberts (2000) investigated jump-diffusion processes, whereas Kellerhals (2001) discusses pricing electricity forwards using stochastic volatility. These additions might come at the expense of analytically tractability and simplicity, however.

The market price of risk has been very difficult to determine based on the available data. The problem has to a large degree been avoided by fitting models to the already risk-neutral forward curves instead. A positive long-run risk premium as reported by Pirrong (2000) for the PJM markets, should, however, be present. Due to the large amount of noise in the data set, however, it is extremely difficult to quantify.

[18] In the price data set from 1996 to 2001, 1997-2000 are considered wet.

6 Weather derivatives frameworks

Since the bulk of the weather derivatives market is made up of temperature-based contracts, such contracts have also been the main point of attention from the academic community. However, it is still quite a novel field of research, and only few articles have been published on the subject. Approaches for calculating a fair price of a derivative (typically an option) vary but can be broadly put in four categories: Stochastic models, Equilibrium approaches, pricing methodologies using only the statistical properties of the underlying and burn analysis. This section gives a brief summary of models developed.

6.1.1 Stochastic models

The most common types of models are stochastic models. A one-factor model is convenient for analytical tractability, and this approach has been taken by amongst others Dischel (1999), Geman (1999), Torro, Meneu and Valor (2001) and Alaton, Boualem and Stillberger (2001). Dischel (1999) also developed a two-factor model including stochastic volatility, while Brody, Syroka and Zervos (2001) extended the notion of Brownian motion to fractorial Brownian motion[19] to better capture long-term interdependencies in temperature time series.

6.1.2 Equilibrium Models

A second type of models is equilibrium approaches incorporating the agent's utility function. Frameworks for valuing general weather instruments in such a setting was developed by Cao and Wei (2000) and Davis (2001). Such models have little appeal outside the academic community, as they generally introduce hard-to-measure factors such as aggregated dividend on the market portfolio and risk preferences of agents.

6.1.3 Models using the distribution of the underlying

McIntyre (1999) introduced a model for valuing temperature-based options by calculating the expected loss of a contract by integrating over the probability distribution of the underlying. His methodology has several resembles to the Black-

[19] Fractorial Brownian Motion is a generalization of Brownian motion calculus and is considered as one of the simplest stochastic processes that exhibits long-rang dependence (Brody, Saroka and Zervos, 2001)

Scholes framework. Martin, Barnett and Coble (2001) used a similar methodology for precipitation derivatives, but recommended using the gamma distribution as it cold be fitted more accurately to distributions of precipitation observations than the normal or lognormal distribution.

6.1.4 Burn analysis

A second commonly employed pricing approach is burn analysis. The idea is to calculate the payoff from a given contract over a set of historical outcomes for the underlying and take the average payoff as the fair price of the contract. This method is popular due to its simplicity, and can be applied to virtually any contingent claim. Care has to be taken with data sets that exhibit some long-term trend (Dischel, 2001). Burn analysis applied on precipitation derivatives is discussed in the next chapter.

6.1.5 Applicability to precipitation derivatives

A survey of the available stochastic models applied to temperature-based weather derivatives did not find any frameworks that directly fit precipitation derivatives. The main reasons are that precipitation follows a different process than temperature – a negative value of precipitation is impossible, indicating that a model for the logarithm of the precipitation level might be appropriate. Furthermore, whereas temperature is continuos, precipitation is discrete in its nature, and this will affect the development of a stochastic model. The simpler approaches such as expected loss and burn analysis are easily adapted to alternative underlyings, and we return to burn analysis for pricing precipitation derivatives in chapter 8.

7 Precipitation models

This section discusses common properties of stochastic precipitation models and the data set we use as source for our analysis. Since our ultimate aim is to develop derivatives written on a precipitation index, we argue why standard hydrological models employed are inappropriate if the aim is to derive closed-form solutions. Finally, we develop a one-factor stochastic model for the Enron precipitation index and test how well this model performs.

7.1 Data used

As a basis for the analysis, we use a subset of the Scandinavian Precipitation Index developed by Enron Nordic. The index consists of 19 Norwegian and 8 Swedish measurement stations, and the index is constructed such that 1 index point equals 1 GWh of generation potential. The index we use was aggregated from precipitation time series for the Norwegian stations in the index. A detailed overview of the index can be found in Appendix C.

7.2 Modelling precipitation

7.2.1 Hydrological background

Following Shaw (1988), the basic idea behind stochastic hydrological models for precipitation is normally based on a two-step process where the first step consists of determining whether a given day at a given site is a *”wet day”* (a day with measurable precipitation) or a *”dry day”* (a day with no or not measurable amount of precipitation). In simple models, the distribution of wet and dry days is driven by a Bernoulli process where the probability of a wet day is higher if the previous day was a wet day than if the previous day was a dry day. If a day is a wet day, then the intensity of the rainfall is modelled by sampling from a common distribution with a minimum of zero and no upper bound such as the lognormal or the gamma distribution. (See i.e. Shaw, 1988 or Moreno, 2001) The gamma distribution is most commonly employed, since it can be fitted with greater accuracy to a precipitation model with varying time resolution (Marin, Barnett and Coble, 2001). The intensity of precipitation on a wet day can be dependent on the intensity on a number of previous days (Moreno 2000).

There are numerous refinements to this basic approach available (see i.e. Cameron et al, 2000) for a good discussion on stochastic precipitation models), and the most commonly employed in Norway is a variant of the Bartlett-Lewis model (Skaugen, 2002) developed by Onof and Weather (1994). It has seven stochastic parameters and is too complex for deriving analytical expressions for derivatives prices.

It is known that precipitation levels on different sites are inter-dependent, since weather systems tend to move over time, hence it would require a complex estimation of correlation parameters between the stations in the index if we choose to model stations individually (Killingtveit, personal communication). On this basis, we choose to model the index directly as opposed to modelling the individual stations in the index. This makes the standard hydrological approach outlined above less relevant, since it is unlikely that it will not rain on one out of the 19 stations in the index on a given day, hence the wet-dry distinction become meaningless. In fact, only 1% of the days are indeed dry days. This could be mended by choosing a cut-off value and define the days where the index is lower than this as dry days, but this would be arbitrary and unnecessarily complicate the model development.

From the preceding section, it is obvious that if the underlying stochastic model for precipitation aims to take research on stochastic hydrology into account, simulation is the only possible result. Our main goal, however, is to make our development analytically tractable and we therefore chose a simplified approach. The one-factor model suggested by Schwartz (1998) incorporates some of the key features of precipitation time series, namely seasonally dependent mean and mean-reversion. The mathematical properties of this model were discussed in section 4.4.2. This model is a gross oversimplification of reality, but as we shall see later, the model performs adequately compared to historically observed data.

7.3 A one-factor model for precipitation

7.3.1 Aggregating data

The purpose for hydropower generators is to hedge against continous precipitation shortfall over weeks if not months (Gustavsen, personal communication). Furthermore, AEP Energy, the company maintaining the index uses weekly-accumulated precipitation as their reference, so we want to keep the same resolution to compare our modelled prices to market quotes. On this basis, we choose to model weekly-accumulated precipitation. .

Aggregating to weekly totals further helps clarifying the seasonal trend and reducing the volatility of the data set. A plot of the weekly-accumulated index values and weekly standard deviation measured in percent is presented below

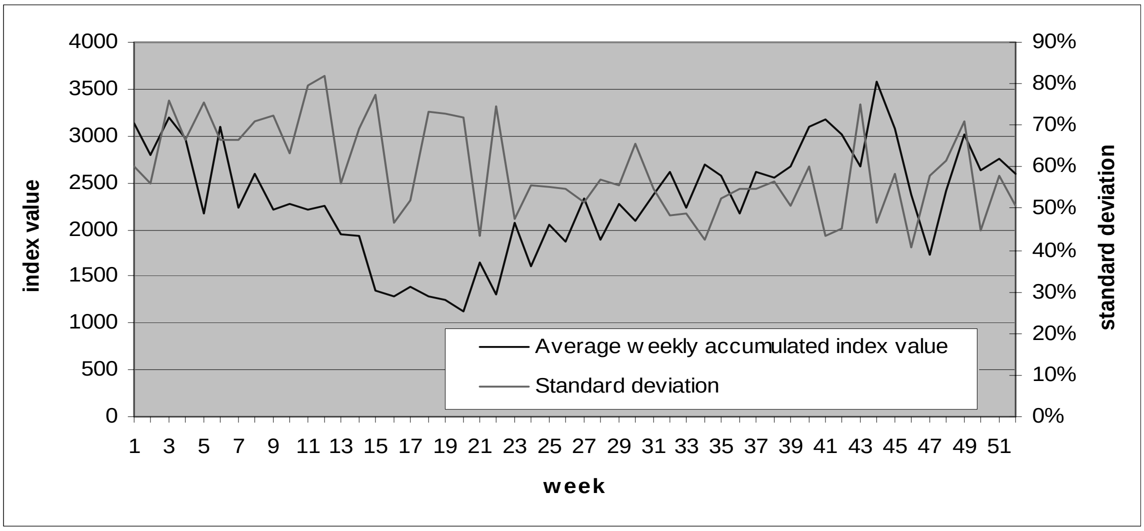

Figure 7.1 Average weekly accumulated index values and weekly standard deviation

Since precipitation has a logical lower bound of zero, a log-transformed model of the time series is convenient. For the rest of this analysis, we will take

$$I_t = \ln(Index_t) \tag{7.1}$$

7.3.2 The deterministic part

The deterministic part of the stochastic model serves two purposes, firstly it expresses the expected outcome at a given point in the future and it determines the long-term mean to which the stochastic values revert. By inspection, the seasonal pattern in the precipitation time series looks simple, although a higher-order transformation might be required to capture the shape. Consequently we seek a model of the form

$$I_t^m = \alpha + \beta t + \gamma \sin(\omega t + \tau) + \gamma_2 \sin(2(\varpi t + \tau_2)) \tag{7.2}$$

This function is determined by fitting the equation

$$Y(t) = a_1 + a_2 t + a_3 \sin(\omega t) + a_4 \cos(\omega t) + a_5 \sin(2\varpi t) + a_6 \cos(2\varpi t) \tag{7.3}$$

to the weekly accumulated precipitation index values yielding

Coefficient	Value	Std Error	P-value
a_1	7.47	0.0217	0.000
a_2	-0.0003	0.000	0.355
a_3	-0.274	0.02835	0.000
a_4	0.205	0.02873	0.000
a_5	0.127	0.02837	0.000
a_6	-0.001	0.02867	0.978
R^2	**9.5%**		

Table 7.1 Regression coefficients on the deterministic precipitation function

No significant linear trend is observalble in precipitation as opposed to temperature time series – that is, there is no significant evidence of increasing or decreasing overall precipitation level in the data set.[20]

The expression for mean precipitation can be rewritten to

$$I_t^m = 7.47 + 0.34\sin(\omega t + 0.648) + 0.13\sin(2\omega t) \qquad (7.4)$$

7.3.3 Fit to observed data

As shown in the plot of the deterministic function versus historical data below confirms the low fit obtained through the regression

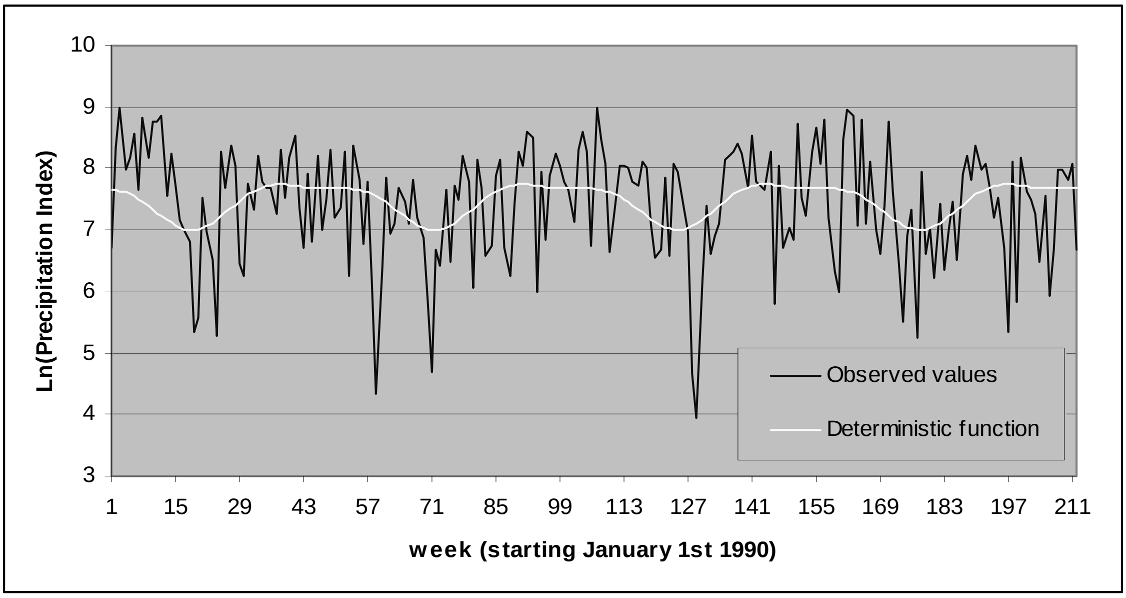

Figure 7.2 Deterministic component of precipitation model versus observed values

[20] However, Schieldrop (2002) noted that whilst the overall precipitation level is unchanged, there seems to be evidence that the discrete outcomes have changed. He found that extreme values occur more often the latest ten years than during the preceding decade.

The low R^2-value does not mean that the fit is a bad one, an alternative interpretation is that precipitation levels are highly volatile and the noise process is an important component of the model. Indeed, when run against weekly averages, the R^2 – value is 69,4% indicating a good fit to the seasonal fluctuations.

7.3.4 The noise process

We know from the model formulation in section 4.4 that the mean-reversion parameter κ determines how quickly the model reverts to its deterministic value. κ is found by analyzing the autocrorrelation function. The MINITAB printout is presented below

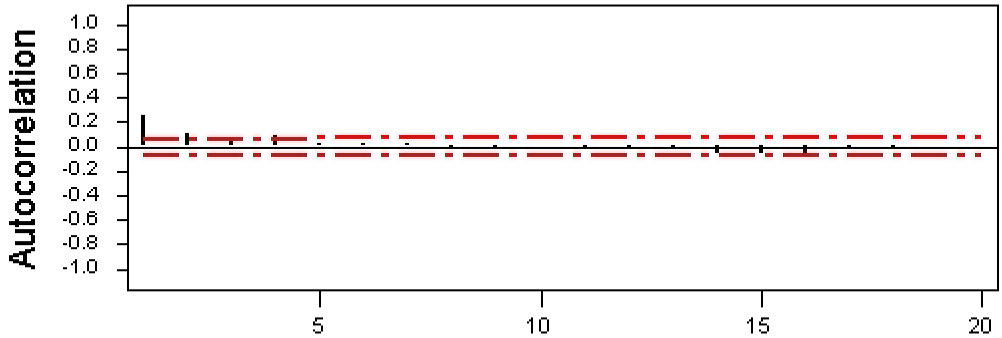

Figure 7.3 Autocorrelation pattern for precipitation time series residuals

The autocorrelation coefficients are

Lag	Autocorrelation (ϕ)	t-value[21]
1	0.25	7.85
2	0.11	3.30
3	0.07	2.07
4	0.09	2.70
5	0.03	0.90

Table 7.2 Autocorrelation coefficients for precipitation model residuals

We see that only the first two lags are significant, and the plot reveals that the coefficients die down in an exponential manner. This is the case in the one-factor model we use. An estimator for ϕ is the first-order coefficient yielding an estimate of $\kappa = 1 - \phi = 0.75$ for the mean-reversion parameter in the model.

This is a surprisingly high value of the mean-reversion parameter, indicating that the time series are almost memory-less. This is contrary to our assumptions, but a reason

[21] If the absolute value of the t-statistic is greater than 2, the lag is normally assumed significant (Fueller 1996)

might be that we model weekly precipitation levels, and precipitation on a given day is conditional on the past (or past few) days, and a weekly resolution in the data removes important information about the autocorrelation pattern. The autocorrelation factor when using daily resolution is found to be 0.42 with four highly significant lags. On this basis, our estimate of 0.25 for the weekly autocorrelation seems more reasonable.

7.3.5 Estimating σ

Hull (2000) suggested estimating volatility as the continuously compounded return volatility as given in equation (7.5)

$$u_i = \ln\left(\frac{I_i}{I_{t-1}}\right) \tag{7.5}$$

Taking the standard deviation of the returns yields an estimate *s* of σ as *s* = 0.90 or 90% on a weekly basis.

A plot of the continuously compounded weekly returns confirms the extremely large volatility present in the series.

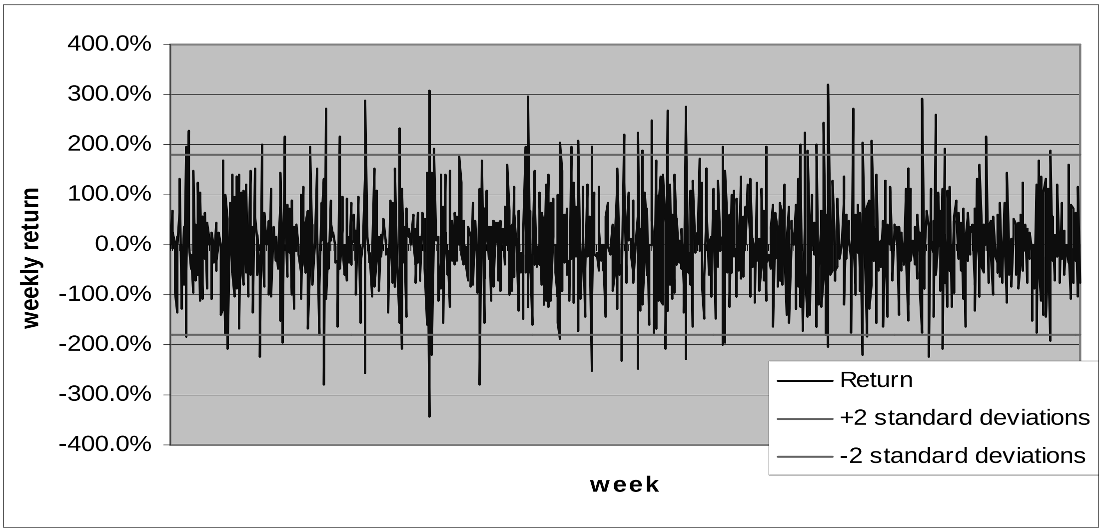

Figure 7.4 Weekly precipitation index returns

Note however that the high mean-reversion parameter estimated above will quickly draw the series back to the expected value, and precipitation will not reach extreme levels According to Shaw (1988), it is customary to use monthly or seasonal resolution for variance in precipitation time series. The Schwartz one-factor model assumes constant return volatility, and to check if the error introduced is large, we test the following hypothesis:

H_0 : Return volatility in week t is equal to week t' for all t ≠ t'

H_1 : Return volatility for some week t ≠ t' is not equal to week t'

Performing a Fisher-test for equal variances checks the null-hypothesis. The F-test matrix can be found in Appendix D[22]. As the results show, the evidence for constant return volatility is conclusive, and we keep the null hypothesis of constant volatility.

7.3.5.1 Are the residuals normally distributed?

A crucial assumption in the Schwartz one-factor model is that the noise parameter, dz, follows a Wiener process. To investigate this property, we check whether the residuals are normally distributed. A histogram of the residuals is presented below

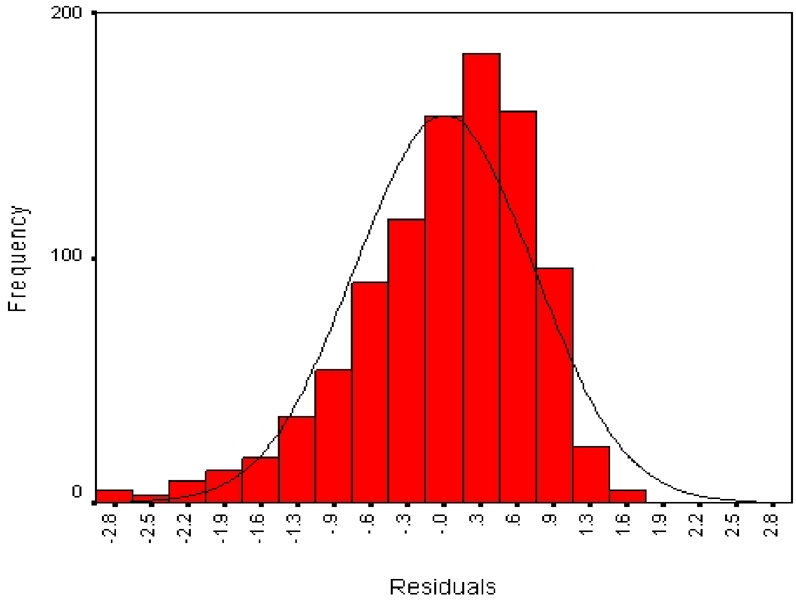

Figure 7.5 Histogram of precipitation model residuals and fitted normal curve

This plot reveals that the assumption of normally distributed residuals is not entirely supported by data. The descriptive statistics confirm this.

---

22 Since the approach involves testing for equal variance over 51 weekly data sets, it comprises over 1200 significance tests. The result table is too large to list here, but can be found in Appendix D

Parameter	Value	Standard error
Mean	0.000	
Weekly Standard deviation	0.747	
Skewness[23]	-0.896	0.149
Excess kurtosis[24]	1.102	0.157

Table 7.3 Descriptive statistics of precipitation model residuals

We observe that both skewness and kurtosis are significantly different from 0 and 3 respectively, and consequently, the residuals are formally not normally distributed. For analytical tractability, we nevertheless stick to the assumption that the noise term in the model follows a Wiener process.

7.3.6 How good is the model?

Lacking a predefined way of testing the precipitation model, we turn to two approches based on simulating precipitation index trajectories and compare to the data set.

7.3.6.1 Comparing to historical observations

As a first test of the goodness of the simple precipitation model developed above, we generate 1000 simulations using the stochastic differential equation in equation (4.9) and see how many of the index values generated fall within the boundaries of weather scenarios observed during the last 19 years. The results are presented below

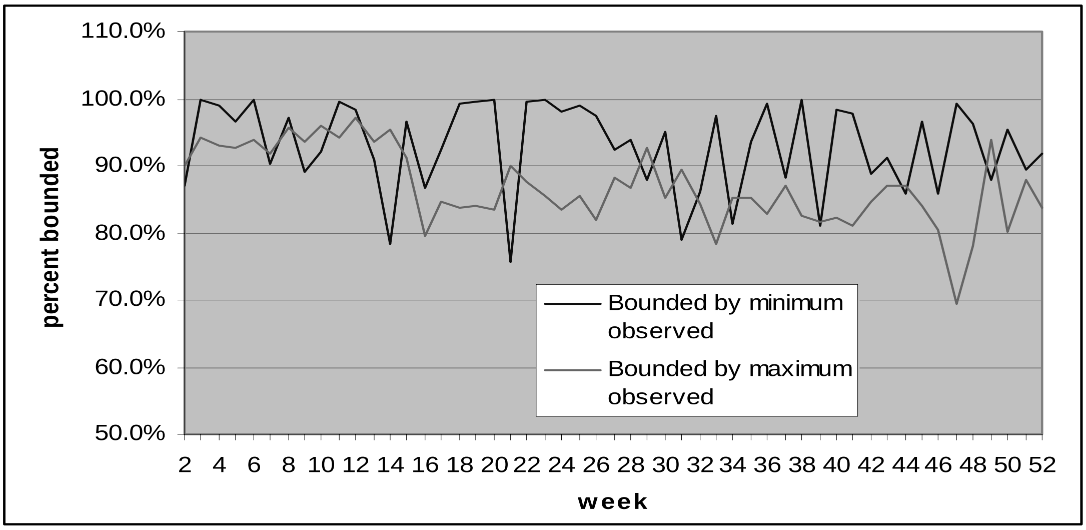

Figure 7.6 Simulations of precipitation index compared to historical observations

[23] Negative skewness indicates clustering to the right

[24] Positive excess kurtosis indicates that the distriution is “taller” than the standard normal distribution

Since the last 19 years contain unusually dry (i.e. year 1996) and wet (i.e. year 1990), we would expect the majority of scenarios generated to fall within historical bounds, while still being able to generate weather events not yet occurred. As the plot reveals, this is the case for the model for most weeks. There is no formal way of defining "adequate" in this context, but it is clearly a weakness that for some of the weeks, not a single of the 1000 scenarios fall outside the bounds. This is consistent with the histogram residuals in Figure 7.5, since the noise distribution used in the model

1. Does not generate enough low extreme values
2. Generates too many high extreme values

Consequently, too many simulation results will fall outside the observed upper bound and correspondingly too few ouside the observed lower bound.

7.3.7 Formal statistical test

An alternative way to test the performance of the model is to compute a confidence interval for the weekly outcomes based on historical observations. Assuming that the outcomes are independent and normally distributed, the simulation results can be standardized by computing (Walpole, Walpole and Myers, 1998)

$$-\alpha_{1/2} < \frac{x_{ij} - \bar{x}_j}{\sigma_j} < \alpha_{1/2} \tag{7.6}$$

where α denotes the confidence level, σ_j the standard deviation for week j. and x_{ij} the i-th simulation for week j. From the properties of the normal distribution, we would expect 95% of the standardized simulated values to fall in the interval [-1.96,1.96]. As summarized in the plot below, this seems to be the case for most weeks

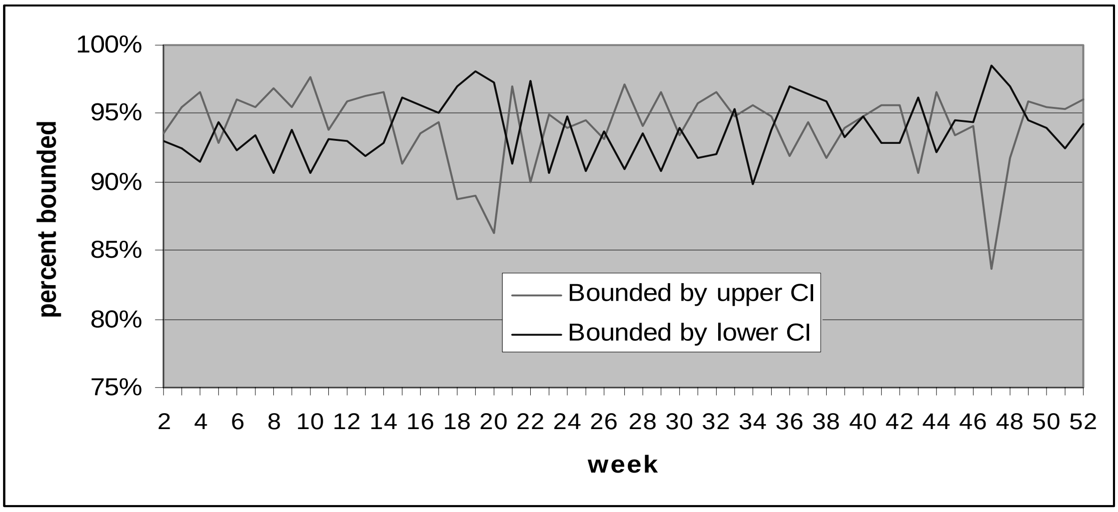

Figure 7.7 Statistical test of precipitation simulation results

The spikes in weeks 20 and 47 are attributable to unusually low values, which in turn makes the confidence interval lower than for the adjacent weeks.

7.3.8 Evaluation of the model

In sum, the model, albeit far from a correct replication of the precipitation index, seems to have the desired properties, and in the next chapter, we proceed to pricing derivatives written on the index using this model. Care should however be taken to use this model slavishly, but applied with care, it should give an indication of where the price of a derivative instrument should lie.

7.3.8.1 Volatility problems

Estimating the weekly standard deviation for the model is not straightforward. Experiments with varying estimates for σ suggest that the volatility structure in the time series is too complicated to be adequately captured by our simple model. An alternative is to use the Mean Sqared Error from the residuals (s=0.75), but firstly the one-factor model based on the logarithm of the underlying assumes geometric standard deviation, and secondly, this lower estimate is obviously less volatile than the actual observations – a crucial factor when valuing derivatives. Some of the simulation results are two or three times as large as the largest historical observations.

An explanation to this can be that the standard deviation of the index returns is as high as measured, but some physical constraint is involved that prevents the precipitation levels from reaching the high extreme values indicated by the noise distribution. A compromise can be to use the residual standard deviation for simulation purposes when the aim is to generate weather scenarios for a specific week, while the intra-week return volatility should be used for valuing derivatives. When pricing precipitation options in section 8.3.2, we use standard deviation of returns as model input.

7.4 Possible extensions to the model

Our model of weekly precipitation is simple to ensure analytical tractability. More realistic models can be created, some improvements are suggested below

- A stochastic process including jumps on the form (Clewlow and Strickland, 2000)

$$d\ln(I) = \kappa(\mu_t - \Phi K_m - \ln I_t)I_t dt + \sigma I dz + KIdq \qquad (7.7)$$

  might better capture the high volatility in the series.

  dq is a Poisson process

  K_m is the mean jump size (possibly also a function of time)

  K is jump size with lognormal distribution ln $(1+K)$~N[ln$(1+ K_m - v^2/2), v^2$]

  ϕ is the average number of jumps per year (possibly also a function of time)

  ν is the standard deviation of the proportional jump (jump volatility)

  The rest of the variables as defined in the chosen model. As spointed out by Clewlow and Strickland (2000), incorporating jumps into the stochastic model does not allow for closed-form formulas for derivatives pricing. For this reason, and time constraints, we choose not to pursue this approach further.

- Volatility is assumed constant across the year. This is not consistent with existing meteorological knowledge; variance is higher during the summer months. For simulation purposes, time dependent volatility would be a good extension. The problems are to Experience from previous modelling effort showed that it is difficult to obtain a daily model which.

  1) generates daily or weekly precipitation levels where the peak values are consistent with historical data while at the same time
  2) ensures the monthly and yearly accumulated precipitation levels adequately bounded by historically observed minimum and maximum.

- Standard stochastic precipitation models could be fitted to data from the individual stations in the index, and the index could be aggregated from these models. However, since the stations are distributed all over Norway, this would call for a complex estimation of inter-station correlation factors and would certainly require computationally intensive simulations.

8 Pricing precipitation derivatives

As shown in Chapter 6, several approaches have been attempted to value weather (read: temperature) instruments. In this section, we derive closed-form formulas for valuing precipitation derivatives. We will discuss two appraches: Burn analysis and arbitrage-free pricing using two derivatives.[25] We will also obtain an estimate for the market price of precipitation risk and finally compare how sample contracts obtain different values accoring to the pricing methodology employed.

8.1 A note on forecasting impact

In all cases, we assume the contract to be entered into long enough in advance so that no forecasts are available that will change the market's perception of expected precipitation during the contract period. Following Dischel (2000), it is not clear-cut whether i.e. seasonal forecasts are reliable, but it is obvious that the expected outcome for a wether phenomenon is less uncertain the closer into the future we look. Day-to-day forecasting beyond 8-14 days is considered impossible, but Dischel pointed out that it might be possible to predict a certain trend from the given state of the weather, comparaing with historical data to detect patterns. For our pricing, we assume the contracts to be long enough into the future for the outcome to be unaffected by any forecast.

8.2 Contract types

Since weather derivatives are meant to hedge against non-catastrophic weather events, contracts are normally structured to have a limited up- and downside. An example of this is a spead (or reverse spread) whose payoff prifile has an upper and lower bound.

[25] The arbitrage-free approach uses a forward contract and an option. However, there is no observable forward curve for precipitation, but given an estimate of the market price of precipitation rrisk, the teoretical curve can be derived.

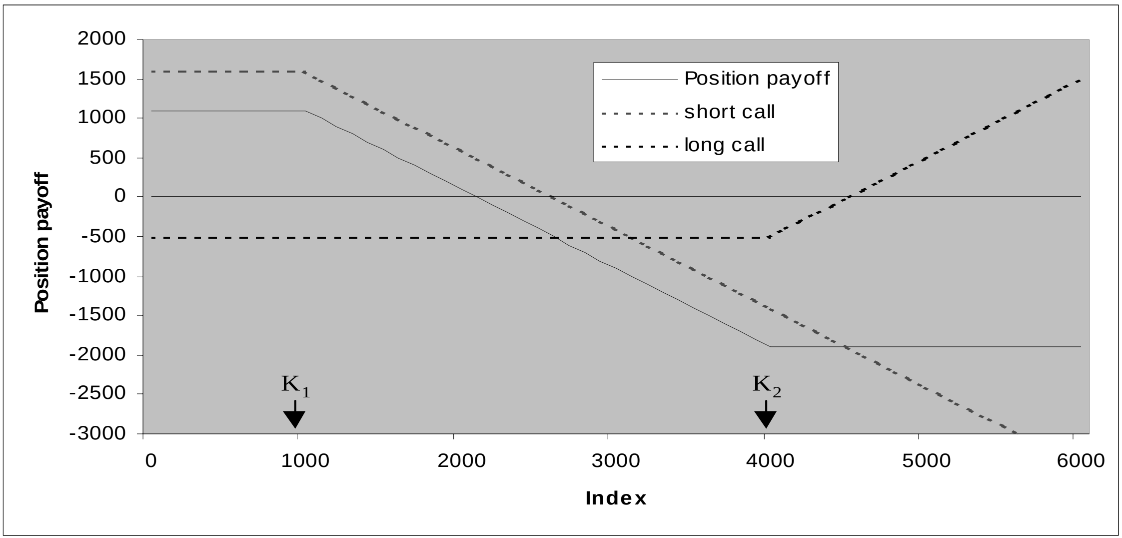

Figure 8.1 Enginering a spread or a reverse spread and their payoff profile

As shown in Figure 8.1 a reverse spread (known as a *bear spread* in the stockmarket) is easily constructed by selling a call with strike K_1 and buying a call with strike K_2. The corresponding position can also be engineered using put options. Since the premium of a call always decreases with increasing strike, this strategy will always incur a cost to set up when using call options. (Hull, 2000)

8.3 Pricing methodologies

This section discusses burn analysis in datail since this is a method employed by weather traders, then we move on to derive analytical pricing formulas based on the precipitation model from the previous chapter.

8.3.1 Burn analysis

The first providers of weather protection were insurance companies, and a common prcing method for insurance policies is to calculate the expected loss for a given policy. In a weather derivatives context, this is equal to answering the question: “Given a contract and a set of historical data, what would the average payout of the contract be when run over observed outcomes of the underlying weather phenomenon?” This is a nonparametric approach, since nothing is assumed about the distribution or the behaviour of the underlying.

8.3.1.1 Pricing example

To illustrate how burn analysis is done, we price a reverse spread with payoff profile as shown in Figure 8.1. In the example, we price the contract for week 42. The expected value of the accumulated index values for this week is 3000, and the contract is structured to pay one unit of currency for each point the index is below the expected value, limited to a payoff of 1000, and for each point the index is above the expected value, the contract buyer pays the contract seller one unit of curency, limited to a maximum of 500. The result of the burn analysis is shown in Figure 8.2

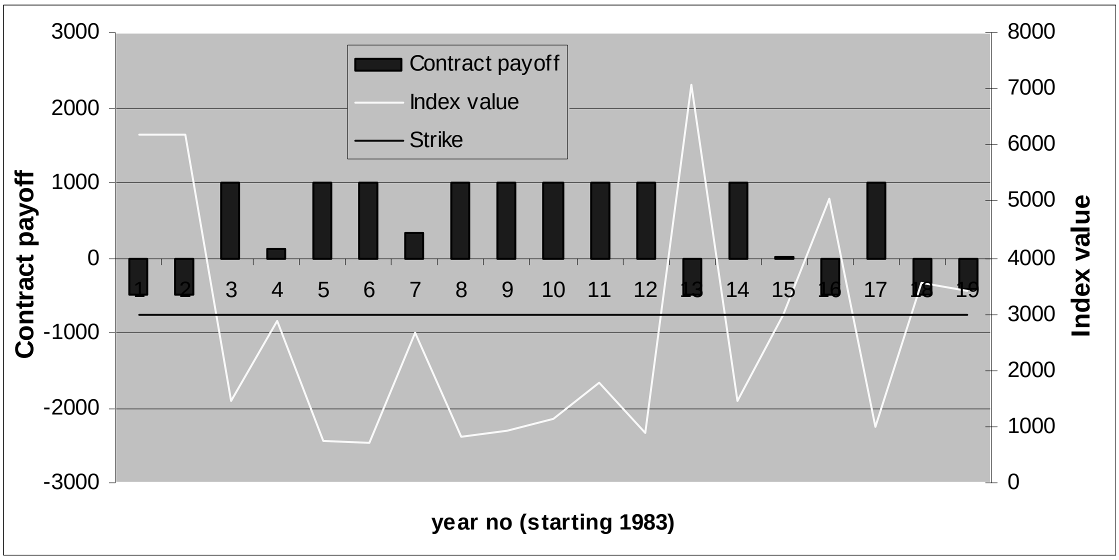

Figure 8.2 Burn analysis example

For this sample contract, the expected payout to the contract buyer is 372 units of currency, which has to be discounted at an appropriate discount rate. A problem is that this rate is unknown – it depends on the risk exposure and preferences of the contract buyer and seller and is not explicitly defined. The ambiguity of the discount rate is a major weakness of burn analysis.

8.3.1.2 Sensitivity to record length

A second problem is the sensitivity to the length of record available. To clearify this, we compute the fair value of a precipitatin derivative written on a single measurement station (Vardø) to have a longer record (52 years) to work on. The contract is is structured to pay the buyer one unit of currency for each point the accumulated precipitation is below the average of 13.5 and the buyer pays the seller the same ammount if precipitation is above the mean. Both payments are capped at five units of currency.

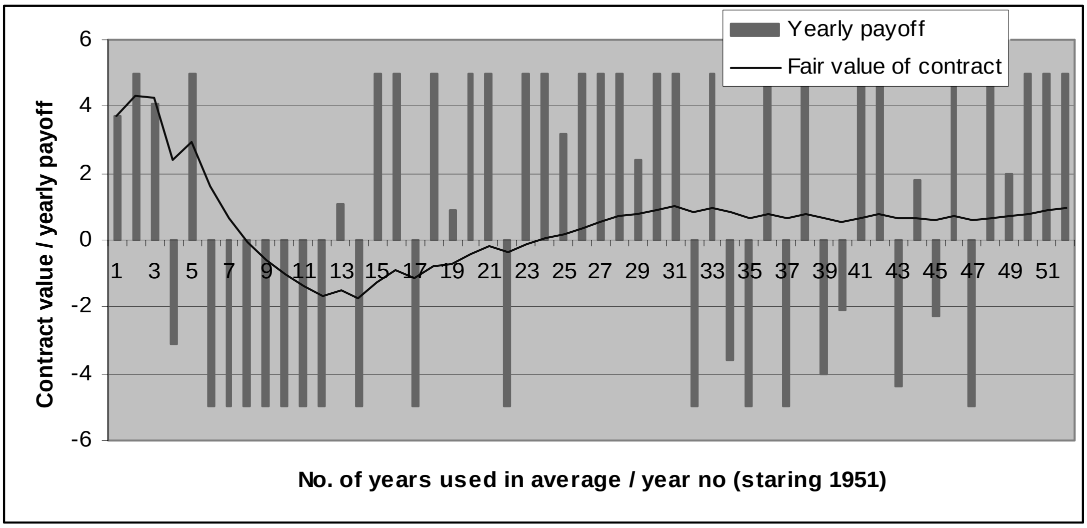

Figure 8.3 Burn analysis payoff and sensitivity to record length

As seen from the plot in Figure 8.3, at least 30 years of data is needed for the analysis to yield stable values.

8.3.1.3 Evaluation of method

The main virtue of burn analysis is its simplicity. It is emplyed by practioneers to get a first indication of the fair contrat price (Schieldrop, personal communication). Some of the serious shortcomings are

- It implicitly assumes the market price of precipitation risk to be zero. As we shall see later, this price differs from zero, at least in the short run.
- No hedge parameters can be calculated
- It fails to distinguish between weather patterns that might exhibit substantial differences when the observed values hovers around the strike level[26]
- With limided data available, burn analysis can price contracts that should be identical (i.e. precipitation levels in two adjacent weeks) significantly different. This shortcoming can be mended by using more than one week in the analysis.

8.3.2 Arbitrage-free pricing using the forward curve

From equation (4.20), we know that the forward price under the risk-natural probability measure for the one-factor model developed in section 4.4.2 is given by

[26] An extreme example: Constant precipitation level at a site can be priced equal to a site with extremely volatilite precipitation, provided the two sites have the same mean.

$$F_0(I_o,T)=E^*(I_t)=\exp\left(I_t^m+(\ln I_0-I_0^m)e^{-\kappa t}+\alpha^*(1-e^{-\kappa t})+\frac{\sigma^2}{4\kappa}(1-e^{-2\kappa t})\right) \quad (8.1)$$

where $\alpha^* = \lambda\sigma/\kappa$. We assume λ, the market price of precipitation risk, constant.

Given a forward curve, we know that the price of a European call option written on the same underlying is given by

$$c(t,T,s,K)=e^{-r(T-t)}[F(t,s)N(h)-KN(h-\sqrt{w})] \quad (8.2)$$

where w and h are defined in equations (3.32) and (3.33) respectivly. If the maturity of the option and the forward contract has the same maturities, the option becomes an option on the spot of the underlying. (Clewlow and Strickland, 2000) The price of a put option is easily found through the well-known put-call parity and is given by

$$p(t,T,s,K)=e^{-r(T-t)}[KN(-h+\sqrt{w})-F(t,s)N(-h)] \quad (8.3)$$

All parameters are observable from historical series, except λ.

Although we use constant volatility in our model, the option formulas are still valid if time-dependent volatility is used. As we showed in section 7.3.5, this is most likely not necceassay. The integral for accumulated forward return volatility, *w*, will also generally require numerical integration. (Clewlow and Strickland, 2000)

8.3.2.1 Estimating the market price of precipitation risk from traded contracts

As pointed out in section 2.2.4, there is no liquid market for precipitation derivatives from which to obtain prices. However, AEP Energy agreed to quoute prices on the Norwegian part of the Enron precipitation index (similiar to the data set used in chapter 7), and the quotes are presented in the table below. The contracts are for uncapped call options on accumulated precipitation during a given week with strike set 10% above the expected value and a payoff of 1000 NOK per index point

Contract week	Strike (GWh)	Actual call option premium	Implied λ from model[27]	Burn analysis premium
5	2900	NOK 650-700,000	5% - 7%	686,000 [28]
15	1400	NOK 190-210,000	15% - 18%	197,000
40	3300	NOK 500-540,000	5% - 7%	500,000

Table 8.1 Estimating market price of precipitation rinsk and model results

It has to be stressed that the prices obtained are not true market prices, but according to weather trader Lars Elmlund at AEP Energy, the price ranges quoted are prices "at which a deal would be likely to take place"

8.3.2.1.1 Discussion of results

We see from Table 8.1 that the market price of discharge risk is postive for the three contracts investigated, and λ is consistent for the winter and fall contacts desipte a difference in strike price. The high value of λ for the spring contract is most likely due to the fact that the determenistic part given in equation (7.4) of the forward curve does not adequately capture the drop in average precipitation level for weeks 15 to 20.

With only three data points available, there is a natural limit to the conclutions that can be drawn, but the findings do at least indicate that the market price of precipitation risk is positive and dependent on time of the year. A model involving time-dependent volatility might give more consistent results, and more contracts at true market prices would be desirable for a through analysis.

Lacking more information, and due to the simplicity of our model, we take the market price of discharge risk to be constant at 7%. This gives little weight to the spring contract for the reason pointed out above.

---

[27] Obtained by finding the value of λ that would equal the price given by the model and the price quote. Excel's Goal Seek function was used for the analysis, since the option expression cannot be solved analytically for λ

[28] Week 5 is a week that strongly deviates from the adjacent weeks due to two unusually low values. To smooth the results, we have taken the average of weeks 4,5, and 6 to calculate the price

8.3.2.2 Evaluation of method

Arbitrage-free pricing has the advantage of obtaining the same price for contracts that should theoretically be close to identical (i.e. precipitation spreads for two adjacent weeks). All hedging parameters can be calculated by straightforward differentiation of the pricing formulas. The major drawback is the simplicity of the underlying precipitation model. The contract prices are highly dependent on the forward curve provided by the model and the volatility estimate. As seen in Section 7.3.8.1, finding the correct estimator for volatility is less than straightforward. Given a proper estimate of the market price of risk, arbitrage-free pricing is indifferent to an agent's risk preferences and has a clearly defined discount rate – namely the riskless rate.

8.3.3 Which is the better pricing formula ?

As an example of which prices the three pricing methods yield for similar contracts, we have priced a European uncapped call for each week of the year. The strike is set at the weekly historical three-week moving average for each week.

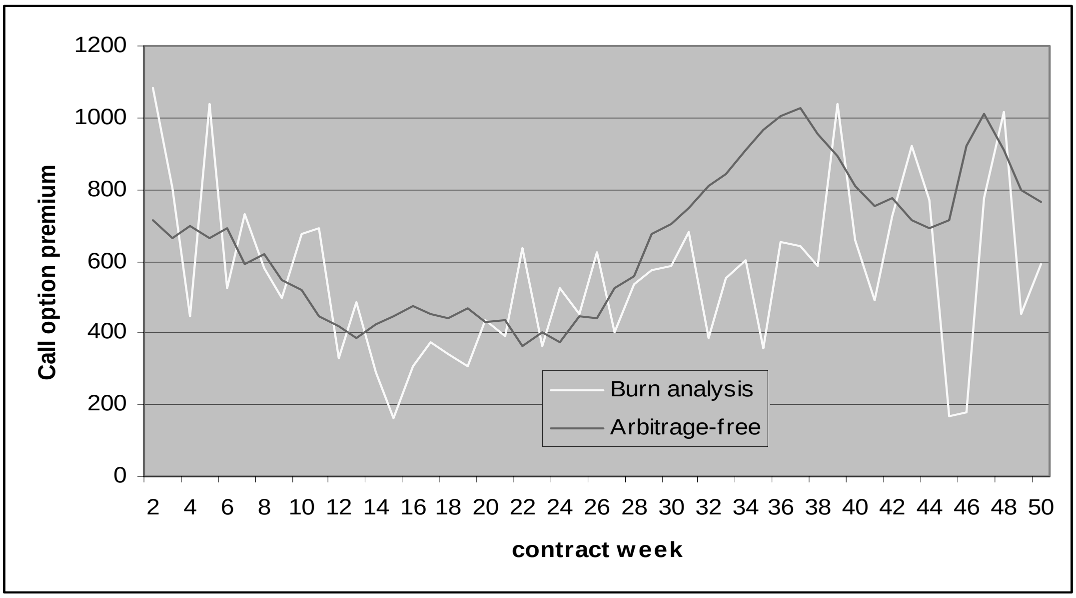

Figure 8.4 Precipitation call option premiums according to pricing method used

Burn analysis yields highly volatile prices – and as seen in section 8.3.1.2, 19 years of historical data is not sufficient to make prices converge due to the large volatility of weekly observations. If using the data from week t-1 and t+1 to price a contract for week t, the prices will be more stable.

Two parts of the price curves deserve further attention. Weeks 28 to 39 and weeks 45-46. We see that for weeks 28 to 39, the differences between the analytical contract price and the burn analysis price are over 100% for some of the weeks. Again, this can be explained by the shape of the deterministic function from euqation X – the fit to the historical observations used in the burn analysis is poor, hence the pricing disparices. Burn analysis for weeks 45 and 46 yield extremely low option premiums – this is attributable to some extreme values in the lower end of the distriution drawing the average down.

Regardless of the method used, some contracts look overpriced and some look underpriced depending on the contract. This makes it difficult to clearly prefer one model over the other, and part of the explanation can be found in the ambiguity of how to estimate the crucial parameters in a precipitation model. Burn analysis has the virtue of being simple, but fails to take weather dynamics into consideration. To formally test the pricing methods, more reference prices are needed.

8.4 Summary

In this section, two different methods for pricing precipitation derivatives were evaluated. Our analysis showed that the theoretical price for contracts differ significantly depending on which method is used, but we do not have enough data to formally test the differnet methods. We also obtained an estimate of the market price of precipitation risk from quotes on contracts, and this value seems to be positive (5%-15% depending on the contract), but with only three contracts available, the data does not allow for through analysis or drawing clear conclutins.

9 Hedging volumic risk using precipitation derivatives

This section will explore two points. Firstly, we investigate how the precipitation index used in chapters 6 and 7 is correlated with forward prices to better understand how precipitation levels diverting from the expected value influence power prices. Secondly, we will explore whether it is possible to use derivatives written on the precipitation index as a volumic hedge for a run-of-river power plant on a single site.

9.1 Price correlation

We would expect long-term contracts to be negatively correlated with the precipitation index, since expected spot during the winter will be strongly dependent on the fill level of the reservoirs. To investigate this, we compute the correlation between the precipitation index and forward contracts maturing between 1 and 104 weeks from a trading date, yielding 104 data points – one for each maturity. The results are presented below

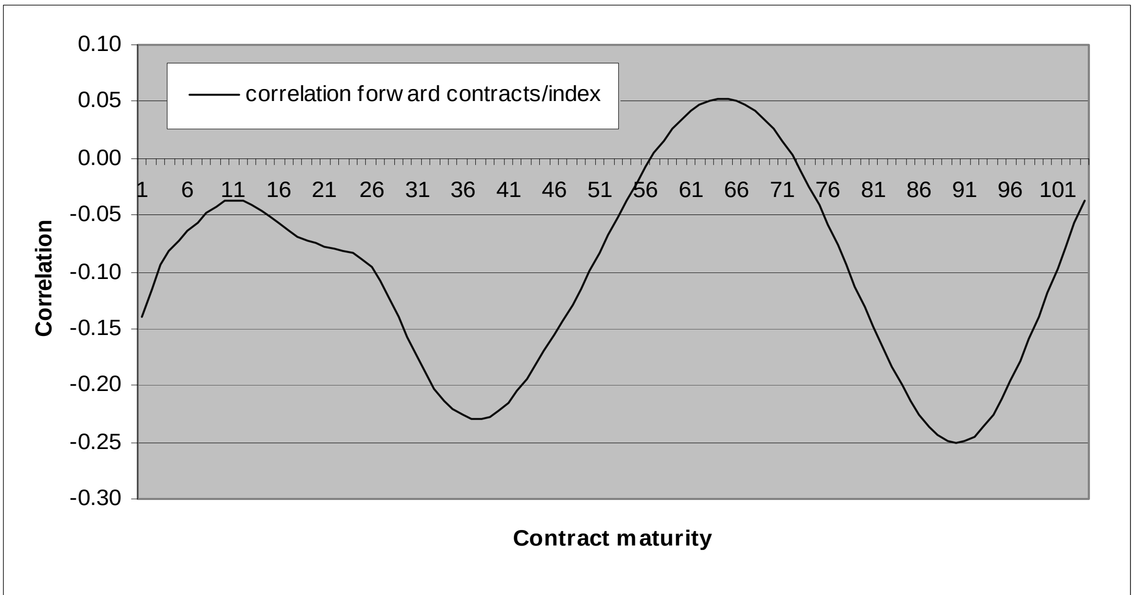

Figure 9.1 Correlation between forward contracts of different maturities and the precipitation index

This plot reveals a strongly time-dependent correlation. Given an efficient market, intuition would suggest a more constant correlation. An explanation can be found in the way the curve in Figure 9.1 is calculated – it is an average over six years, and the following relationship might be true

- For short-term contracts, buyers are active, delivery takes place and low precipitation forces the prices up

- For long-term contracts, the true correlation is dependent on the trading date, since there might be a period between the trading date and contract maturity during which the reservoirs are expected to be full. If such a date is present, correlation should be negliable.

Consequently, the correlation between precipitation and forward contracts with longer maturities is expected to be a function of trading date, and the plot in Figure 9.1 is influenced by seasonal variations. We therefore attempt a different methodology for investigating the correlation between precipitation and forward prices

9.1.1 Can precipitation explain short-time fluctuations?

An interesting question is to which extent changes in precipitation level explain short-time changes in forward prices. To investigate this, we detrend the precipitation data and forward curves for selected maturities through equations (9.1) and (9.2).

$$I_{detrend}(t) = I_{observed}(t) - I_{expected}(t) \tag{9.1}$$

$$F^{m}_{detrend}(t) = F^{m}_{observed}(t) - F^{m}_{expected}(t) \tag{9.2}$$

where $I(t)$ and $F^m(t)$ denotes the precipitation index value at time t and forward contract price for maturity m at time t

We use only data from 1997 to 2000, since as briefly commented on in chapter 5, the high prices in 1996 and 2001 make it impossible to get a good estimate for seasonal fluctuations. Results from the analysis are presented in Table 9.1

Contract maturity	Weekly correlation raw data	P-value	Weekly correlation detrended data	P-value
1st week	0.127	0.071	-0.225	0.001
5th week	0.148	0.035	-0.27	0.000
10th week	0.055	0.433	-0.167	0.017
12th week	-0.019	0.785	-0.202	0.004
20th week	-0.296	0.000	-0.273	0.000
26th week	-0.371	0.000	-0.164	0.020
28th week	-0.407	0.000	-0.169	0.016
36h week	-0.396	0.000	-0.216	0.002
48th week	0.178	0.011	-0.261	0.000
52nd week	0.265	0.000	-0.238	0.001
70th week	-0.198	0.005	-0.212	0.002
88th week	-0.343	0.000	-0.256	0.000
104th week	0.255	0.000	-0.200	0.004

Table 9.1 Correlation between detrended forward prices and detrended precipitation index

We see that when we remove the seasonal trend inherit in power prices and to a lesser extent in precipitation series, the correlation between precipitation during a given week and the forward contracts traded that week is significant at the 2% level and negative for all maturities investigated.

An interpretation of these findings is that an *unexpected* increase in overall precipitation level decreases forward prices. As expected, deviations from the expected precipitation level constitute a significant part of the information flow determining power prices. From the selected contracts, a clear trend cannot be extracted, but the correlation hovers around –0.2 and are significant for all maturities.

To further inspect this relationship, we plot the correlation between changes in forward contracts and the precipitation index as a function of trading data to see if seasonal patterns are still present. The plot is shown below

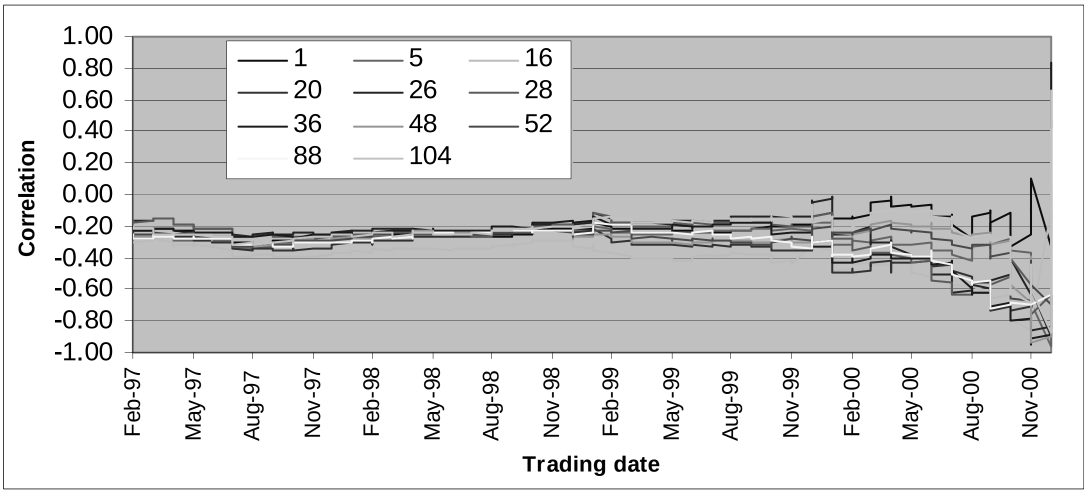

Figure 9.2 Correlation between precipitation index and forward contracts for different trading dates

We see that the correlation coefficient is constant throughout the year. The volatility towards the end of the curve is purely a result of the fact that fewer and fewer data points are used for calculating the correlation. Johnsen (2001) found that unexpected inflow had a greater effect on prices during the late winter weeks (week 10-16) than during the early winter (weeks 34-46). We have attempted a similar analysis on precipitation, but failed to find any significant pattern when calculating the correlation between unexpected precipitation and forward prices for each week of the year. Possible reasons for this are:

- Johnsen studied the impact of unexpected *inflow* while we study unexpected *precipitation*. These are not necessarily related to each other if precipitation comes as snow
- After detrending, there is still much noise left in the data, and only six data points for each week[29] of the year may be too little to extract a clear correlation.

9.2 Can precipitation derivatives hedge volumic risk ?

The correlation between power prices and precipitation was found to be significantly negative as expected. We now turn to the question of hedging volumic risk using precipitation derivatives – this will have implications for the last part of this thesis where we value a run-of-river power plant. For such a hedge to be effective, the underlying precipitation measure will have to be well correlated with the water discharge in the specific river in which the plant is to be built.

9.2.1 Assessing volumic risk

The volumic risk of a run-of-river plant is mainly dependent on two factors: The discharge capacity of the plant and the hydrograph of the river. Hydrographs for two typical Norwegian rivers (Orkla and Gaula) are presented in Figure 9.3. From this figure, several conclutions can be drawn

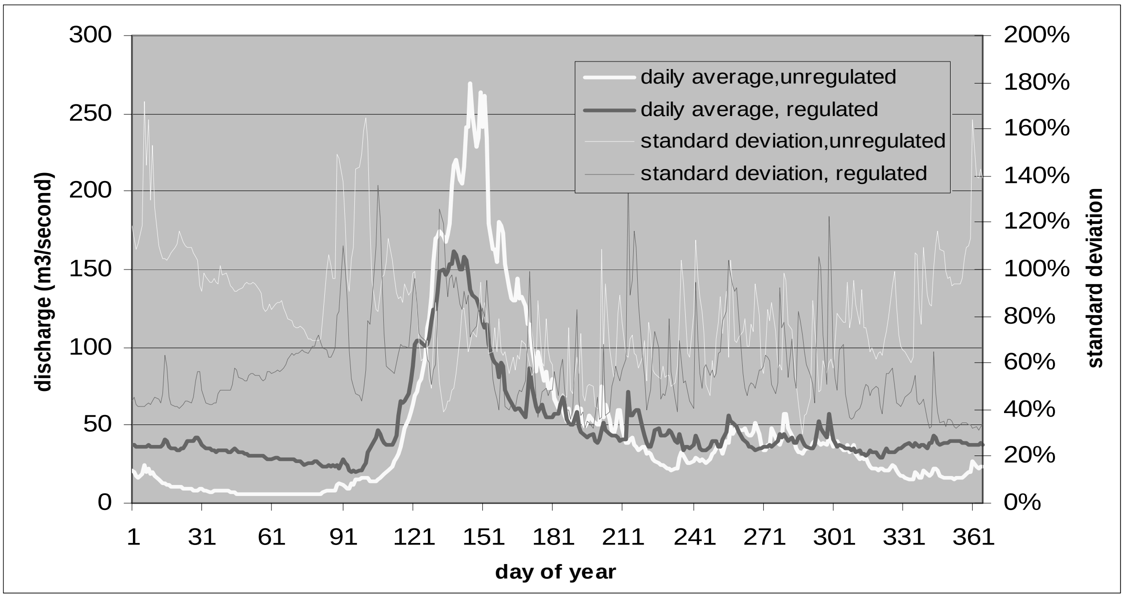

Figure 9.3 Average daily discharge and standard deviation for regulated and unregulated rivers

[29] We used weekly resolution, and while precipitation data are available from 1983 to 2001, we only have price data from 1996 to 2001

- Regulating a river (by exploiting it for power production) levels out the yearly flow and decreases flow volatility
- Water discharge in an unregulated river is very low during the winter
- Optimal discharge capacity is likely to be higher in a regulated river

The implications for run-of-river plants are that more even production can be expected in a regulated river, while the volumic risk depends on the discharge capacity of the plant. We return to the problem of determining the optimal size in chapter 10.

9.2.2 The problems of hedging

9.2.2.1 Hydrological reasons

The main problem in constructing a hedge using precipitation derivatives is that the relationship between precipitation and water discharge is less than straightforward. The reasons are

- During the winter, precipitation generally arrives as snow, which does not translate into water discharge until weeks or months later.
- Water discharge during the spring flood is mainly a function of accumulated snow in the mountains, not precipitation during the flood period (Killingtveit, 2000)

9.2.2.2 Basis risk

A further problem is the correlation between the underlying (the Enron precipitation index in our case) and water discharge in the selected river. The index consists of measurement stations scattered all over Norway, and Moreno (2001) pointed out that basis risk is present if the hedger’s buisiness is not located very close to the measurement station(s) used.

9.2.3 Estimating correlation for three rivers

To investigate the correlation between the precipitation index and water discharge, we use three sample rivers – Orkla, Gaula and Glomma. On the background of the abovementioned factors, we analyze the correlation for weeks 29-46, since this is the period when the spring flood is over and precipitation generally arrives as rain.[30] Consequently, this is the period during which the correlation can be assumed to the highest. Significant correlation is indeed found, as shown in Table 9.2

River	Station	Average weekly Correlation	P-value
Orkla	Syrstad	0.221	0.001
Gaula	Gulfoss	0.207	0.003
Glomma	Solbergfoss	0.133	0.057

Table 9.2 Correlation between precipitation index and water discharge for fall weeks

The correlation coeffecients in Table 9.2 are calculated over all 19 years in the data set simultaniously. For a more detailed breakdown of the weekly correlation coeffecients as well as the cross correlation between precipitation in week t and water discharge in weeks t-1 and t-2, we refer to a collection of figures in Appendix E. A detailed analysis reveal that as for the correlation between power prices and the precipitation index, large irregularities are hidden behind the average correlations found.

From a hydrological viewpoint, low correlation between an index such as the Enron precipitation index and water discharge in a specific river is as could be expected (Killingtveit, personal communication).

9.3 Concluding remarks

Based on the analysis conducted in this section, we conclude that derivatives written on the Enron Precipitation Index do not constitute an adequate hedge for the volumic risk inherent in a single run-of-river power plant. This appears to be the case in both regulated and unregulated rivers and irrespective of the geographical location of the river.

Another question is if a precipitation index can serve as an adequate volumic hedge for an operator with run-of-river plants located across Norway. We do not have the

[30] If the same analysis is run on all weeks of the year, the correlation is significantly negative(!). The explanation for this is probably that discharge is large during the spring despite (potentially) low precipitation, and that large precipitation during the winter does not mean large discharge in the river

data source required to pursue this question further, but preliminary tests do show that correlation increases when more than one river is used.

When correlating discharge for the Gaula weeks 29-45 with a weather station in the same area as the river, the correlation is found to be as high as 0.6. This indicates that for a precipitation hedge to be an efficient hedge for a run-of-river plant, the derivative has to be written on one or more geographically clustered weather stations in proximety to the river in which the plant is to be constructed. Any such derivative is unlikly to materialize because of no potential liquidity.

Due to the factors discussed in this chapter, we have to reject the idea of valuing a run-of-river plant as a portfolio of precipitation derivatives and power contracts. We therefore need to consider alternative real option approaches for valuing the plant, and this will be the topic of the coming chapters. We first turn to the problem of deciding how large the discharge capacity of the plant should be.

10 Power plant characteristics

For the analysis of a run-of-river power plant, we choose to use the time series for the Gaula river. This is due to the fact that the river is unregulated, making the revenues more volatile and the exposure to volumetric risk clearer, thus making the valuation problem more interesting. To get discharge levels in the range our cost data is valid for, we transform the series by taking

$$Q_{transformed}(t) = \frac{1}{8} Q_{original}(t) \qquad (10.1)$$

where Q denotes the water discharge measured in m^3/second. For the rest of this thesis, the transformed series are used. The transformation preserves the seasonal fluctations and the volatility.

10.1 Description of a run-of-river power plant

A run-of-river power plant channels a portion of a river through a canal or penstock. It may or may not require the use of a dam. Even if a dam is present, it does not have the multi-period storage facilities present in reservoir systems. Thus, the all-important factor for determining production is instantaneous water discharge. Water not immediately used for power generation is considered lost. A typical hydrograph for an unregulated river (Gaula) and the production is shown in Figure 10.1

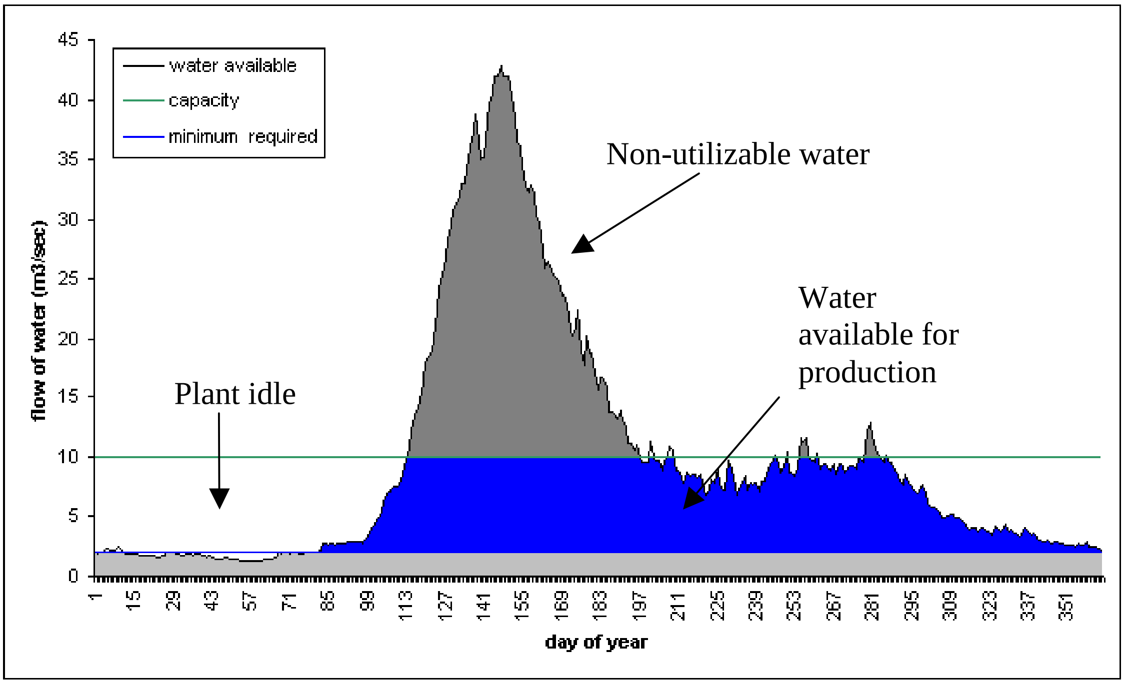

Figure 10.1 Production in a run-of-river power plant versus time-varying water discharge

As seen from Figure 10.1, the discharge capacity of a plant greatly influences

1. Fraction of the year it can be expected to run at full capacity
2. Fraction of the year the plant can be expected to be idle
3. How much water is expected “lost” during the flood period

A plant will have a minimum discharge needed – the reason for this can be an environmental constraint requiring a certain minimum discharge in the river or a technical constraint requiring a set amount of water for the turbines to operate at all. Regardless of the reason, no production will take place if the discharge is below this limit.

10.2 Determining the effect rate

The maximum effect in kW of a run-of-river power plant is calculated as:

$$N = \frac{g\rho\varepsilon QH}{1000} \tag{10.2}$$

where

N = Effect in kW,

g = gravaitional acceleration = 9.81 m/s^2,

ρ = density of water = 1000 kg/ m^3,

ε = energy conversion rate (normally assumed equal to 0.9),

Q = discharge capacity in m^3/s

H = net height (head) in meters.

Production in kWh is then equal to

$$E = Ndt \tag{10.3}$$

where N is as defined above and dt = time increment measured in hours

Assuming a head of 20 meters, an effect rate of 0.9 and weekly resolution, equation (10.3) yields E = 9.81*0.9*20/1000*24*7 = 29.66 MWh per week per unit discharge. This value will be used through the analysis.

10.3 Determining the value

If spot prices for and discharge at time t are taken as deterministic and marginal costs are assumed zero, the value of a plant can be expressed as

$$V_{plant} = \int_0^T e^{-\delta t} N(t)P(t)dt - I(\Theta) \tag{10.4}$$

where *N(t)* is the effect at time *t* given a water discharge, *P(t)* is the spot price at time *t*, *T* is the lifetime of the plant, δ is a risk-adjusted continuously compounded discount rate (7% is normally used) and *I(θ)* is the investment required expressed as a function of a vector θ of parameters determining the cost. We turn to valuation under uncertainty in chapter 11.

10.4 Determining the cost

To determine the optimal size for a plant in a river, it is essential to know the cost of construction. The cost for turbines, generators, transformers etc. will be a function of the generating capacity of the plant, while the construction work will be dependent, amongst other factors, on the rated discharge of the plant, the head etc.

10.4.1 Data used

As a basis for our cost analysis, we use a detailed cost analysis made by NVE (Norwegian Water Resources and Energy Directorate). The cost functions provided by NVE are valid for mini- and micro-sized plants with a discharge capacity up to 10 m^3/second and an effect of up to 5000kW.[31] The equations used are shown in Appendix F.[32]

10.4.2 Cost drivers

We choose to focus on the parts of the project whose costs are 1) independent of the actual site conditions and 2) constitute a significant part of the total cost. The main cost drivers fulfilling these requirements are: Turbine, power station generator, interface to power grid, dam and controlling equipment. The dam cost is a function of discharge capacity, while the remaining costs are directly dependent on the power capacity. We also add an overhead of 20% for labour and other costs.

---

[31] While we analyze the project for a discharge capacity greater than 10m^3/sec, this only affects the cost of the dam. We assume the dam cost function to be valid also for larger discharges. The dam only constitutes about 10% of the total cost, so the error is negliable

[32] We use a continous cost function for the coupling equipment. In the NVE data, this is a stepwise funxtion. See Appendix F for a detailed disussion of how and why this is done

10.5 Determining the optimal size.

The tradeoffs of plant dimensioning can be seen from Figure 10.1. The central question is if it will pay off to capture the parts of the year when there are large amounts of water in the river, but then risking the plant to run at low capacity for the bulk of the year. As the size of the plant is taken as an input parameter to the valuation approaches in the following chapter, we need some indication on the optimal size of the plant. Using a simple NPV analysis, we take weekly water discharge determinsitic as the observed average and we use forward prices. Since no stochastic factors are present, we discount at the risk-free rate. It must be stressed that this is *not* meant to yield a correct value of the plant, it is purely an approach for obtataining plant sizes for later use and study the robustness of the value to changes in price models and their parameters. Formally, we solve

$$\underset{Q_{\max}}{Max} \sum_{t=1}^{40*52} \left(W_t(Q_t, Q_{\max}) F_o(t)\right) - I(Q_{\max}) \tag{10.5}$$

$$s.t: W_t(Q_t, Q_{\max}) \in \begin{matrix} 0 & Q_t < 2k \\ 2k & 2k < Q_t < Q_{\max} \\ kQ_{\max} & Q_t < Q_{\max} \end{matrix}$$

Here, k is the conversion factor from discharge to MWh, Q_t represents the discharge at time t, and Q_{max} the maximum production possible. I is an investment cost as a function of the plant characteristics.

Finally, we discount these revenues back to time zero using a discount rate of 5.88%, and subtract the initial investment. Due to the form of the investment function and the upper and lower production boundaries, this is a non-linear problem solved using a numerical search technique. The project value as a function of the plant's maximum discharge capacity is shown in the plot below.

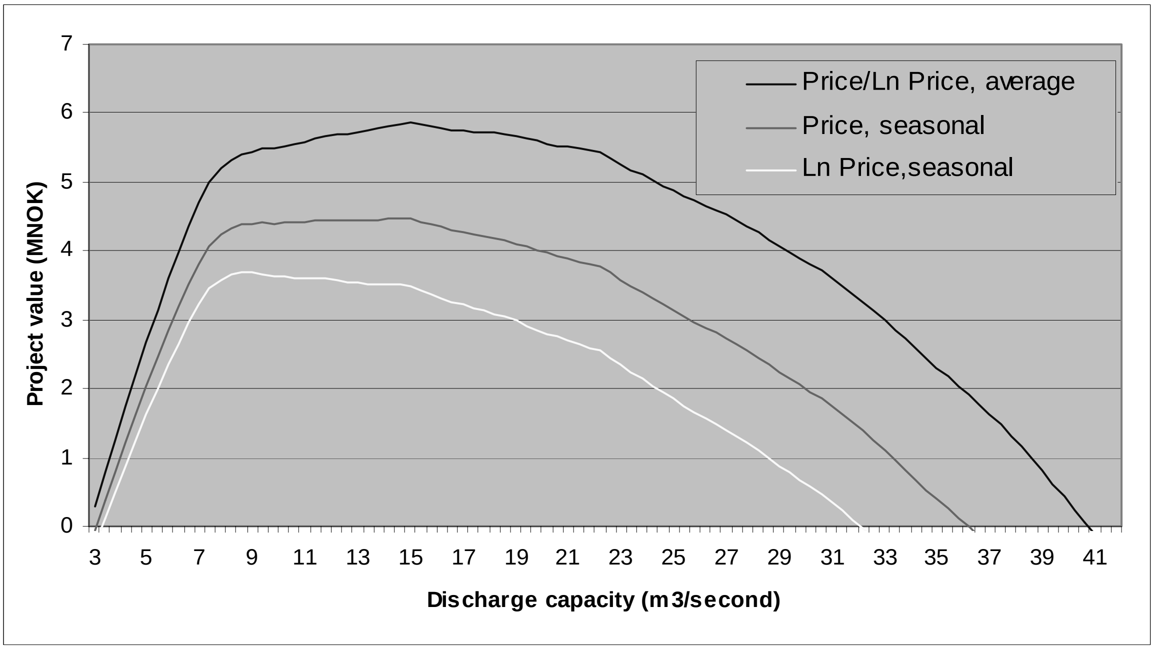

Figure 10.2: Optimal size of run-of-river power plant for different price models

10.5.1 Sensitivity to price model

From Figure 10.2, we see that the project value is highly sensitive to the price model used, while the actual sizing is fairly robust – the only exception being the Log model with seasonal variations yeilding a smaller plant.. To explain this, we have to take the characteristics of the different price models as well as the hydrograph into consideration. The model properties are summarized in table Table 10.1

	Seasonality	**No Seasonality**
Price	Trend: Exponential Amplitude: Constant	Trend: Exponential Amplitude: Zero
Ln Price	Trend: Exponential Amplitude: Increasing	Trend: Exponential Amplitude: Zero

Table 10.1 Properties of price models used in plant sizing optimization

As seen from Table 10.1, all models have exponentially increasing trend. This is the main factor for sizing – an increasing trend justifies a larger investment outlay today to have production to sell at higher prices in the future. Therefore, the plant size is fairly robust to the price model used.

The project value, however, is not. To understand this, we investigate the relationship between production in week t and the expected price during the same week. We know that initially, the model for Price and the model for Ln Price will have the same amplitude, while the seasonal amplitude in the Ln-model will increase as a function of time. The relationship is shown in Figure 10.3 for prices after 20 years of operation.

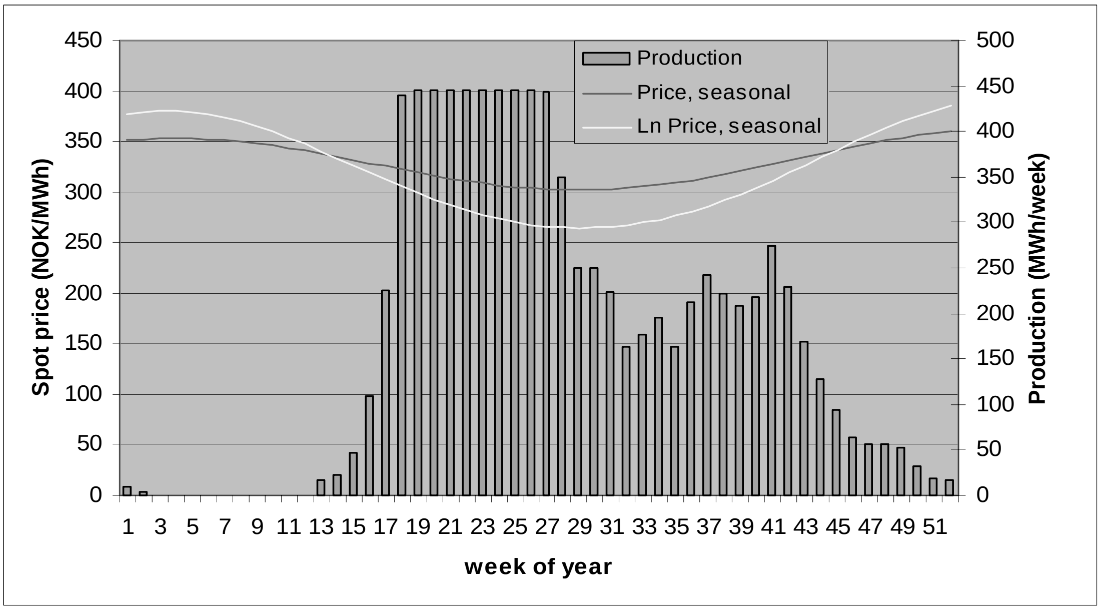

Figure 10.3 Weekly production and weekly price according to model used

We see that almost all production is sold at below-average prices. Hence, introducing seasonal variations significantly impacts the project value. As the difference between high and low prices increase in the Ln-model, the Ln-model yields a lower project value. The optimal capacity according to the price model employed is summarized below

Price model	Optimal discharge capacity[33]
Price, seasonal variations	15.0 m3/second
Price, no seasonal variations	15.0 m3/second
Ln Price, seasonal variations	9.0 m3/second
Ln Price, no seasonal variations	15.0 m3/second

Table 10.2 Optimum discharge capacity according to price model used

We also observe from Figure 10.2 that the marginal benefit of a larger initial investment to get a discharge capacity above $9m^3$/second is small.

[33] We here assume a constraint requiring a discharge of $2m^3$/sec in the river. Hence, a discharge capacity of 15.0 m^3/sec requires an actual discharge of 17.0 m^3/sec in the river for the plant to run at full capacity.

10.5.2 Sensitivity to trend

An interesting question is how robust the sizing decision is to small changes in the price parameters, notably the upward sloping price trend of 3.15%. To investigate this relationship, we calculate the modeled project value and optimal rated discharge using price trends ranging from 1.42% to 5.45% spanning the trend estimates found in section 5.2.2. We repeat the optimization in equation (10.5), but adding trend as a second variable. The optimal discharge capacities for various price trends using the Ln Price model with seasonal variations are shown below

Price trend	Project value	Optimal discharge capacity
1.42%	-1,700,000	7.4
2.00%	-600,000	7.8
2.50%	500,000	7.8
3.00%	1,800,000	8.4
3.50%	3,200,000	9.0
4.00%	4,900,000	14.6
4.50%	7,300,000	15.0
5.00%	10,000,000	22.2
5.45%	13,100,000	22.2

Table 10.3 Optimal maximum capacity with varying forward price trends

From this analysis, we see that an increasing price trend is highly important for both the sizing of the plant and the value of it. The explanation for this is that if prices increase in the future, it will pay off to make a larger initial investment today in order to have higher peak capacity in the future. We also see that a price trend of at least 2,5% per year is required for the project to become profitable (the exact break-even trend is 2.28%). A further question is if it is profitable to delay the investment, we return to the optimal investment timing in chapter 12.

10.5.3 Optimal size and the hydrograph

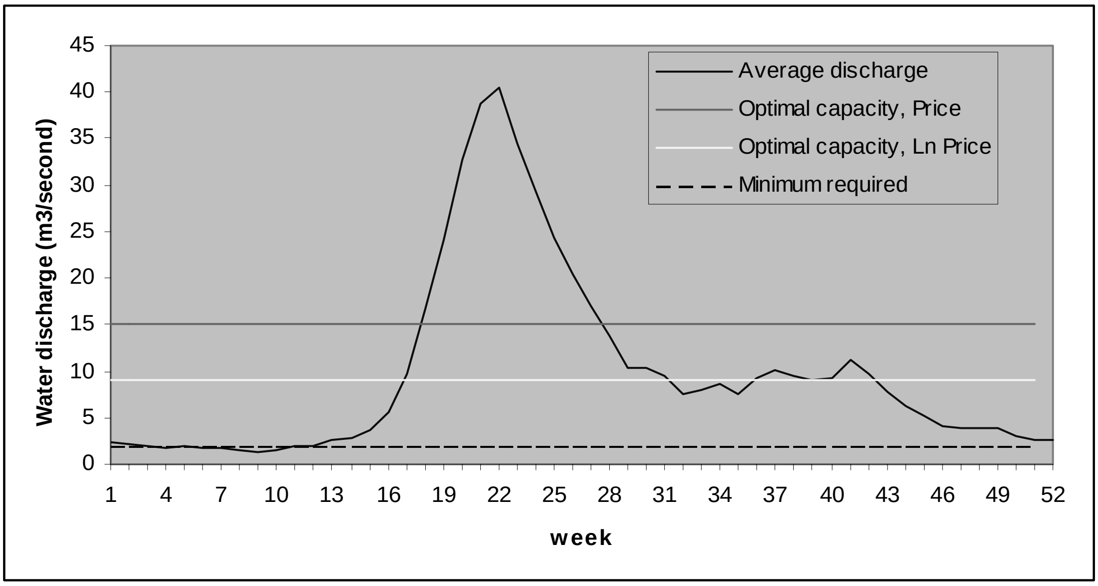

Figure 10.4 Optimal maximum capacity and hydrograph

From Figure 10.4, we see that if the seasonal variations do not increase, it pays off to increase the capacity to capture parts of the spring flood. However, if we include increasing seasonal variations in the price model, the optimal capacity becomes equal to the discharge level *after* the spring flood, allowing the plant to run at full capacity for more than half of the year. Due to the minimum discharge of $2m^3$/sec required, the plant will be idle during the bulk of the winter.

Sensitivity analysis finds the optimal size insensitive to the minimum discharge required, allthough the minimum water needed to run the plant has huge impacts on profits – which is reasonable, since this can be seen as production that is (almost) guaranteed to take place.

10.6 Summary

Using forward models from section 4.4 and a deterministic weekly water discharge defined as the average over historical observations, we find that the optimal size of the plant is a discharge capacity of 9 m^3/second when using a price model with a trend of 3.15% and increasing seasonal variations. If the seasonal variations do not increase, the optional discharge capacity becomes 15 m^3/second. This difference can be explained by the fact that production mostly takes place when prices are below average. We further showed that this dimension is highly sensitive to the trend in power prices, and that a trend of at least 2.5% is necessary for the project to become profitable.

11 Power plant pricing

In this section, we consider two alternative risk-neutral approaches for power plant valuation. The first approach determines the parameters for a stochastic process for revenues. The resulting revenue process is then used to evaluate the value of a power plant with possibility of production halts. The second approach considers the discharge level to be constant with Brownian motion, whereas prices follow GBM. The chapter is concluded by a discussion on the model tractability and market price of risk.

11.1 Modeling a power plant with production constraints using revenues

Revenue is defined as the product of price and volume. This valuation approach uses real data for discharge and prices to generate a time series for revenues, using a weekly resolution on the data. We start by investigating underlying assumptions and properties of the data set, and find a suitable stochastic process that describes it. Based on this, we estimate the parameters for the process, and suggest how to make the model risk-neutral. Then, revenue derivatives are introduced, and finally, the power plant is priced using these derivatives.

11.1.1 Assumptions

Since our river is not regulated, there is a real possibility of production halts, especially during the winter. A few assumptions about the running of the power plant are therefore necessary before we develop the model. These assumptions are listed below:

- The lower limit for production is assumed to be 2 m^3/sec. Discharge below this level means that the production is suspended until the level increases.
- Whatever the reason, the discharge below 2 m^3/sec can never be utilized in production. Hence, discharge of 3 m^3/sec will only produce 1 m^3/sec equivalent of power. (=29.66 MWh/week).
- The plant can be instantaneously started and stopped at no extraordinary cost.
- The plant has a predetermined maximum capacity equal to 9 or 15 m^3/sec. This part of the thesis will only consider these sizes.
- All production is sold spot, because the process multiplies spot prices together with production, and no storage is possible.

11.1.2 Data set

The data set comprises the level of water for the Gaula River and the spot[34] power prices in the period from 1996 to 2001. These sets are then multiplied together to form a new data set, of revenues. This can be illustrated as follows:

$$R_t = Q_t P_t k \tag{11.1}$$

Here, R_t, Q_t and P_t represent the revenues, discharge flow and prices at time t, whereas k is a conversion factor, converting river flow into MWh. This conversion rate k was found to be 29.66 MWh/m^3 on a weekly basis, and was discussed in section 10.2.

11.1.2.1 Resulting data sets

The resulting data sets comprise a data series representing the revenues for the power plant over the discussed period. The set comprises approximately six years of weekly observations, and a total of 319 values. Figure 11.1 shows the revenue function for the river, together with its log value.

[34] The ideal situation would be to multiply the production to a risk-neutral forward contract. This is not possible using this approach, as the river production is not storable. The product of spot prices and current volume is thus the only correct representation of the current production.

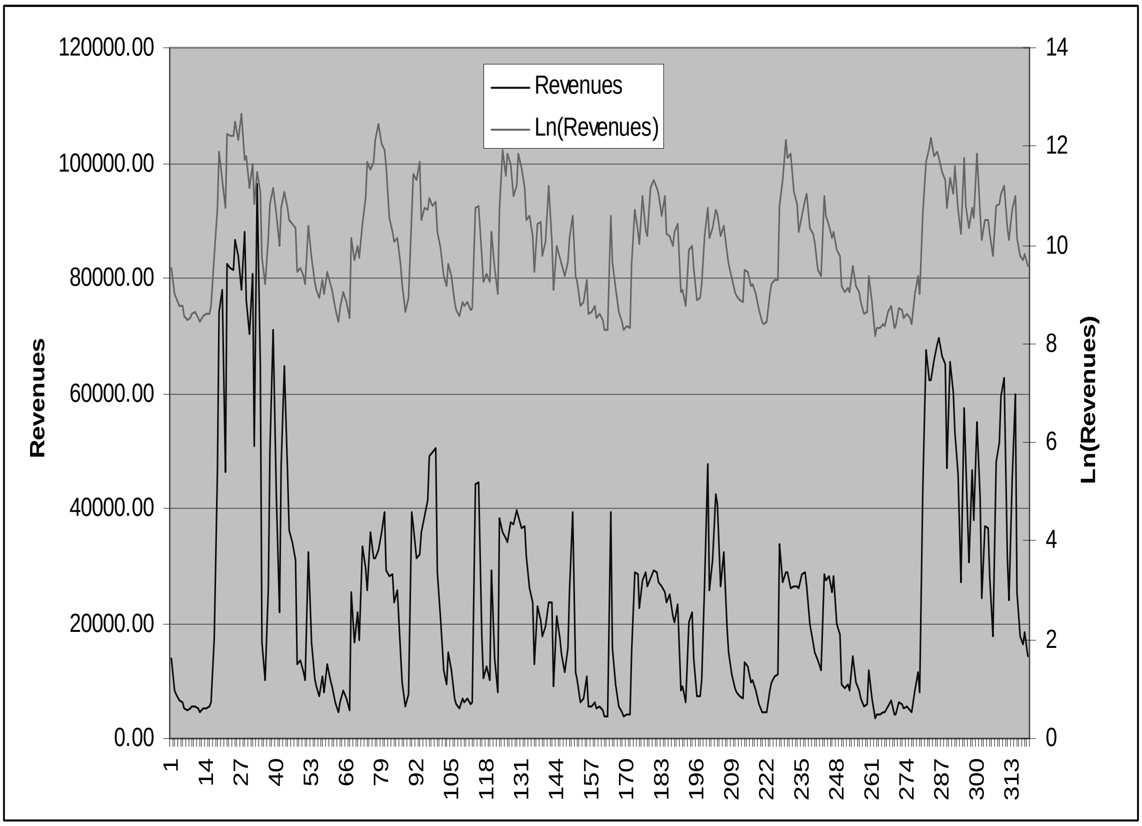

Figure 11.1: The revenue process and its log value

11.1.3 Seasonal function and revenue models

It is not easy suggesting an appropriate seasonal pattern from figure 11.1. Although the presence of some kind of seasonal factor might be a fair assumption, a clear seasonal effect is not present. Nevertheless, we choose to use a cosine function as suggested in section 4.4.9, represented as

$$f(t) = c + \gamma \cos(2\pi(t+\tau)) \tag{11.2}$$

Looking at figure 11.1, we detect a floor-reversion tendency. Our revenue function is therefore assumed to be following a mean-reversion process. Since the revenues are based on price, the data is expected to have a long-term risk-neutral drift equal to the forward price drift. We therefore add the forward price drift to the revenues, equal to the price drift determined in section 5.2. This drift will be considered deterministic, and included in *f(t)*. The discharge, on the other hand, is assumed to have zero drift, as there would be no reason to expect the water level in the river to increase in magnitude[35].

[35] Performing a simple regression analysis confirmed that there was no significant trend in the data set.

In section 4.4.2, we investigated a mean-reverting one-factor log-model for prices, described in Lucia and Schwartz' (2001). The process, originally from Schwartz (1997), is simple, and has the following solution for the risk-neutral expected spot price:

$$F(P_0,T)=\exp(f(T)+(\ln P_0-f(0))e^{-\kappa t}+\alpha^*(1-e^{-\kappa T})+\frac{\sigma^2}{4\kappa}(1-e^{-2\kappa t})) \tag{11.3}$$

By replacing prices *P* by revenues *R*, this model was then fitted to the revenue data set, and the results will be presented in the next section.

11.1.3.1 Regression results

The two data sets were first multiplied together in Excel. Then, the parameter fitting was performed in SPSS, using the non-linear regression solver The parameters of the process were fitted to the model in equation (11.3). Finally, the resulting process represents an analytic expression for expected spot prices, not adjusted for risk. Since the model is fitted to spot revenues only, the last term in equation (11.3) is removed, and included in the two first factors, whereas α^* is removed because this is an expression for forward price risk, not determinable from spot prices. We then estimate the parameters of the following equation:

$$E(R_T)=\exp\left[f(T)+(\ln P_0-f(0))e^{-\kappa t}\right] \tag{11.4}$$

This parameter estimation yielded the following results:

Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
CONSTANT	10.002	0.127	9.751	10.253
GAMMA	-1.015	0.164	-1.337	-0.693
τ	-0.040	0.026	-0.081	.012
κ	0.238	0.037	0.166	0.310
R^2	0.762			

Table 11-1: Parameters of the 1-factor log-revenue model

As can be seen from the above tables, all parameters except τ, the phase angle of the cosine function, were significant at the 5% level. The fact that that τ is not significant does not disrupt the problem, as this is just the phase angle, and must be seen together with gamma, the amplitude. This relationship basically tells us that the revenues are at the bottom in January, when τ=0.

Observe that the revenues have a very high mean-reversion tendency. The degree of mean-reversion reveals a lower degree of "memory" in the process, compared to energy prices. Therefore, temporary deviations should have less impact on long-term contracts.

11.1.4 Revenue volatility

The revenue volatility can be found by using the same approach as used for prices. The revenue return is first found by the use of equation (11.5):

$$return = \ln\left(\frac{R_{t+1}}{R_t}\right) \tag{11.5}$$

Here, R_t is the revenue at time t, and will have to be calculated for the whole data set. Taking the standard deviation of the returns then yields the weekly volatility. Since no revenue forwards are available, we only have the spot volatility to determine. The result of the volatility estimation for data set is presented below. The discharge and price volatility are included for comparison purposes

Variable	Weekly volatility	Annualized
σ_R	57.91%	417.59%
σ_W	59.62%	429.95%
σ_P	10.23%	73.73%

Table 11-2: Spot revenue volatility

The reason for the extremely high revenue volatility is due to the great variations in discharge flow. This volatility is far from constant throughout the year, though. The discharge has its highest volatility around the spring flood, and its lowest during the winter. We will, however, assume that the volatility is constant throughout the year.

11.1.5 Market price of risk and forward prices

After building the expected spot revenue models, we can develop revenue forward contracts by risk-neutralization of the revenue process. This process requires the market price of risk, however, and since no forward or option contracts on revenue are traded, this estimate is difficult to obtain.

Reinvestigating the forward expression of equation (11.3), this can be risk neutralized by subtracting the amount given in equation (11.6) from the expected log-values of the revenue function. Here, α^* is replaced by $\lambda\sigma/\kappa$.

$$d = \frac{\lambda\sigma}{\kappa}(1 - e^{-\kappa T}) \tag{11.6}$$

As always, the short-term volatility σ and mean-reversion κ are measured in the process whereas the maturity T is known. Hence, the market price of mean-reversion risk λ is the only unknown variable. This gives us several options for risk-neutralization of the revenues:

- The market price of risk can be determined by investigating the value of other traded power plants. After obtaining λ from this separate power plant, we can then risk-neutralize the revenue forwards.
- We can assume that, due to no uncertainty in drift, and the long run nature of the investment, the market price of risk for water is zero. Alternatively, the market does not price discharge risk. This would make price risk the sole driver of risk in the power plant. Any short-term uncertainties in production are hence reflected in prices.

11.1.5.1 A first approach to risk-neutralization of the revenue process

Our initial approach for pricing the mean-reversion risk in revenues used a certainty equivalent approach. The integral under the risk-adjusted discounted[36] revenues of another traded power plant should be equal to the trade price of the power plant. Hence, for a power plant whose terminal value is zero, its revenues could be expressed as

$$\int_{T_1}^{T_2} e^{-rt} F(R_t) dt \tag{11.7}$$

This is a contingent value, assuming zero variable costs, of all future revenues from time T_1 to T_2, the termination of the project. Mark that $F(R_t)$ represent risk-neutral revenue forwards, given in equation (11.3) .Consequently, it can be discounted at the risk free rate r.

The power plant's traded value is a certain amount, called a certainty equivalent (CE). This should reflect the risk-less income of the future income. Since forward contracts are considered to be risk-neutral expected spot prices, the expression in equation (11.3) can be substituted into equation (11.7). Hence, once the parameters for the spot

[36] Discounting at the risk free rate.

revenues of the traded power plant have been determined, λ is the only undetermined parameter of equation (11.8), and should explain the difference between the certainty equivalent and the plant value.

$$\int_{T_1}^{T_2} e^{-rt} F(R_t)dt = CE \qquad (11.8)$$

Due to the nature of our seasonal function, this integral would have to be approximated by a sum of all future revenues. Most likely, numerical methods would be needed to find an appropriate value of λ.

This approach had to be discarded, however, as no discharge power plants have been traded in the Nordic energy market the last couple of years. We therefore note that although the approach should be working, we will have to try a different approach.

11.1.5.2 A second approach to risk neutralization

Let us now assume that discharge risk is zero. This would reduce the risk left in the equation, to price risk only. In economic terms, this would mean that he market does not put a price on the uncertainty in future river flow. This can be explained by the fact that as long as the plant is run for a sufficiently long time, temporary variations in discharge should even out and hence be less important in the long run.

Lars-Ove Skorpen at Pareto Securities agreed that this assumption seemed fair but added that the regulation of a river would increase the price of the plant. Thus, at some point, steady production seems to have added value. We will nevertheless try this approach.

Investigating equation (11.6), we assume that the mean reversion risk approaches a set value as T increases. For long-term decision problems like investments in power plants, it is therefore fair to assume a constant risk level. In section 5.7, we found the systematic long-run difference between expected spot and forward contracts to be approximately NOK 33.86 of the price, or 19.3% of the value. Since we have multiplied real data together, the correlation structure is preserved in the resulting data set. We can therefore risk-neutralize the revenues by lowering our initial revenue by the same percentage. We do ths by subtracting .193 from the constant in the exponential function. We now have a risk-neutral revenue process.

Note that risk neutralization is model specific. We use the one-factor model in Schwartz (1997) as our model, hence the risk neutralization should follow equation (11.6).

11.1.6 Revenue forwards

Once a sensible market price of risk can be estimated, it is relatively easy to obtain the measure for the revenue forward contracts. We already presented the analytical solution for forward prices in the equation (11.3).

The adjustment approach assumes that the following conditions hold:

- Revenues are priced sufficiently far into the future to adjust mean-reversion risk by subtracting a lump sum. This can be done because equation (11.6) approaches a limit for long maturities.
- The expected revenues have a long-run forward risk profile similar to prices. They can therefore be risk-neutralized by removing 19.26% of the value of the revenues.

This way of risk neutralizing the revenues will most probably mean that the short-run risk premium would be too large. Assuming that the power plant will take two years to construct, we will have a fair risk-neutralization by the time the plant starts producing.

11.1.7 Revenue options and applications to revenue prices

The option pricing in this section is based on Clewlow and Strickland (2000). They showed that the mean-reverting log-process in Schwartz (1997) could be priced using through the use of a variation of Black and Scholes' (1973) option pricing formula, given in section 3.3.3.

The purpose of the revenue option approach is to build a spread[37] of two options, to price the revenues above the minimum production level constraint of the discharge power plant, but below the maximum level of discharge. The capacity of the plant was

---

[37] The payoff structure of a spread is shown in Figure 8.1.

in section 10.5 set to 9 and 15 m^3/sec, requiring a total of 11 and 17 m3/sec of discharge water. Since the minimum level is 2 m3/sec, and profits grow for revenues above this level, we have a call option payoff structure. The revenues earned can be seen in as the dark shaded area in figure 11.2.

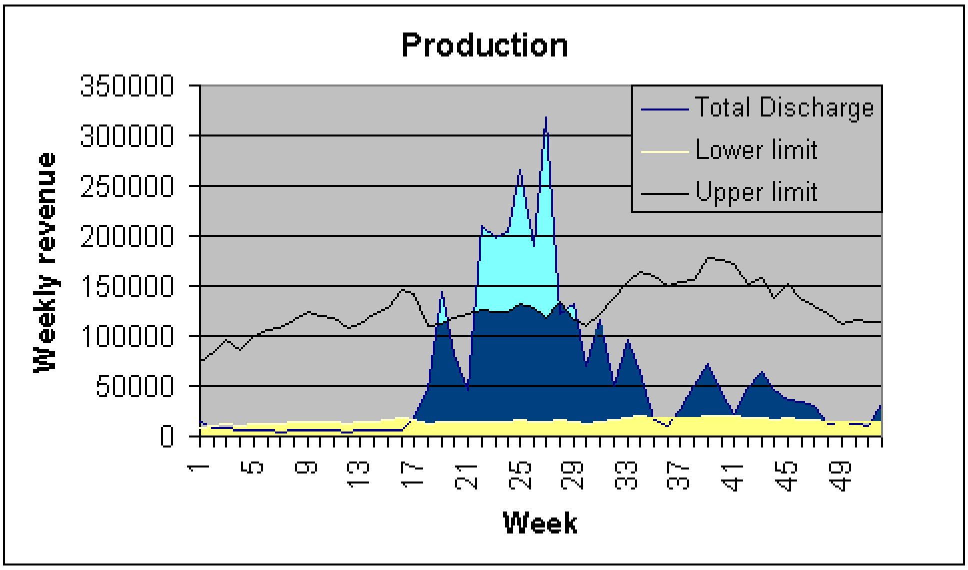

Figure 11.2: Power plant revenues for a random year

11.1.7.1 Strike

This model has two important levels. Firstly, since no production is assumed possible under a minimum level of 2 m^3/sec, a natural strike would be the value of the revenues lost due to this constraint. This strike X_t can be expressed as

$$X_t = 2kF_{Pt} \qquad (11.9)$$

Here, k is the conversion factor of the power plant, found to be 29.665 MWh/m^3 per week in section 10.2, and F_{Pt} is the expected forward price at time t. The forward price model is used to preserve risk-neutrality in the final option price. The parameters for the forward price were estimated in section 5.8, and will not be repeated here. Finally, the model has a strike on the maximum capacity of the turbines. This strike is equal in form to the lower strike, with the lower limit of 2 m^3/sec replaced by the upper limit of 11 or 17 m^3/sec.

11.1.7.2 Volatility

When pricing options, we need the cumulative variance. Variance is defined as the square of the volatility, and an expression for the accumulated variance at the maturity of the revenue option, is given as

$$w = \frac{\sigma^2}{2\kappa}\left(1 - e^{-2\kappa T}\right) \tag{11.10}$$

This expression is a natural consequence of the mean-reversion model, and is discussed in section 3.3.3. In section 11.1.4, the volatility of the revenues was found to be 57.91% on a weekly basis. Finally, κ was determined to be 0.238 in section 11.1.3.1.

11.1.7.3 Option valuation

We now have everything necessary to price revenue options. These options are used to repliate payoff structure of the power plant. As discussed above, Clewlow and Strickland (2000) showed that a variation of the model introduced by Black (1976) could be used for option pricing. This model is an extension of Black and Scholes (1973) using forward contracts. The only difference between the mean-reversion formula and Black (1976), is the expression for the *w(t)*, the cumulative variance. The expression for the option model is given in section 3.3.3.3.

Revenue options can now be priced at any maturity beyond the immediate future. A call option on production revenues would hence give the value of the production above the exercise level in the given time period.

Exercise date	Revenue Forward	Strike	w(t)	Call	Discounted
11/04/05	23487.67	10655.62	.33	16361.70	13094.36
01/11/19	37523.26	19221.47	.34	26621.48	9786.34
05/24/24	82166.35	17222.77	.35	69751.95	18682.72
09/28/40	101896.21	34588.64	.36	82926.15	8466.00

Table 11-3: Call options on revenues for the lower bound. All values except w(t) in NOK

Notice that because we are using an increasing trend of 3.15% for both strike and revenue forward, the future call value seems to grow steadily. On the other hand, the price trend is lower than the risk-free rate of 5.88%, and the resulting present value is therefore low.

11.1.8 Power plant value

Now, we have a set of options on the future weekly revenues of the power plant. We now set up a portfolio of the two types of options to price the power plant. For any given maturity, this portfolio contains one long call option on revenues above the minimum level, and one short call option for the value above the maximum capacity. The expected present value of production at time t can then be expressed as

$$V_0(R_t) = e^{-rt}\left[\max\left(R_t - X_t^{low}, 0\right) - \max\left(R_t - X_t^{high}, 0\right)\right] \qquad (11.11)$$

Now, we have an expression for the value of the revenues at a given time t. This can be used for the valuation of the power plant. Assuming that there are no variable cost of running the power plant, the value can be expressed as

$$V(R) = \sum_{t=T_1}^{T_2}\left(e^{-rt}\left[C_t^{low} - C_t^{high}\right]\right) - I \qquad (11.12)$$

Here, C_t^{low} and C_t^{high} represent the revenue option values at time *t*, and *I* is the initial investment, assumed to be constant. Assuming that the power plant can be built in two years and run for the next 40, the present value of the investment is found in the table below, using the two capacity choices.

Capacity	Revenues	Investment	Present Value
9 m^3/sec	16,591,382	15,065,000	1,526,382
15 m^3/sec	18,011,279	20,067,000	-2,055,721

Table 11-4: Value of the power plant, Revenues are risk adjusted by 19.26%. All values in NOK

As can be seen in table 11-4, the present value of the small plant is positive, whereas the present value of the large plant is not. This would mean that the peaks in the summer would not be large enough or persistent enough to generate the needed revenues to pay for the additional investment in equipment.

11.1.9 Discussion on the model and sensitivity analysis

The model presented in this section tried to utilize the option structure of the expected payoff of production. We built a theoretical spread of revenues, by developing options on the revenues in the future. In order to use the option approach, we needed the risk-neutralized forward prices. This was done by assuming that discharge risk has no market risk premium. This approach resulted in a reduction of 19.26% of the initial revenues. This is substantial, compared to normal risk adjustments, but nevertheless a consequence of the data analyzed. Consequently, we get very low estimates for the plant value.

The mean-reversion process used for pricing also has a mean-reverting volatility function. Because the mean-reversion parameter κ is high, the forward volatility soon approaches zero, making the cumulative variance estimate too low, compared to observed price volatility. Since revenue is a product of price and quantity, price volatility would be expected in long-run revenue prices. This could be included by changing the volatility model or implementing a two-factor model for revenues.

The major source of uncertainty in a power plant value would lie in the long-term price equilibrium. If the price trend of the power plant is changed to 1.42%, our lowest value in the trend analysis, the effect on revenue is enormous. This can be seen in table 11-5. The table shows the power plant value for different price trends within the interval obtained in section 5.2:

Price trend	Small plant (Value in NOK)	Big plant (Value in NOK)
1.42%	-3000000	-7000000
2%	-1700000	-5600000
2.5%	-400000	-4200000
3%	1100000	-2600000
3.5%	2700000	-800000
4%	4600000	1300000
4.5%	6700000	3600000
5%	9100000	6300000
5.45%	11800000	9400000

Table 11-5: Sensitivity of power plant value to the trend in prices

Knowing this, it seems like the risk adjustments done to adjust for mean-reversion risk are quite minor compared to the long-term effects of price changes. Hence, as a mean-reversion approach with seasonal patterns might be fruitful in the short run, the long-run focus must be on the price trend effects.

11.2 A joint stochastic process for revenues

Section 4.4.8 discussed a model for generating revenues using two dependent stochastic processes. The model considers the production Q and prices P as stochastic variables following GBM. This relationship can be utilized in pricing power plants by assuming that the power price follows GBM with positive drift, whereas the river discharge follows GBM with zero drift. To find the value of the plant, we first need to determine its parameters, adjust for risk, and integrate over the plant's lifetime.

11.2.1.1 Data set

The data sets used for this valuation, comprises 88 consecutive years of discharge data from the Gaula river from 1908 to 1995, together with the ten-year contracts and the

1-year forward contracts from 2000 to 2001. Most of the parameters have been determined in other sections of this thesis, and will only be restated in this section.

11.2.1.2 Model

Starting with the river flow, assume that the discharge flow in a river for a given year is assumed to be constant. We multiply this water by the conversion factor, to obtain a value Q, rpresenting the annual production. This production level will be assumed to have zero drift, but following Brownian motion, This can be described through the dynamics in equation (11.13):

$$\frac{dQ}{Q} = 0dt + \sigma_W dz \tag{11.13}$$

We will assume an annual resolution in the model. The average production level Q_{AVG} can be estimated as the average production throughout the year, capped at the plant's capacity, subtracting the minimum level of 2 m^3/sec required to run the plant. An estimate of the water flow volatility can then be obtained by calculating the annual percentage change based on the 88 years of data.

For prices, Schwartz' (1998) one-factor GBM approximation for long-term commodity prices seems to be a natural choice for a process. This is a risk-neutral process, using the risk-neutral forward trend as the drift, and the forward volatility as the model volatility. The process was introduced in section 4.4.7, and can be stated as

$$dZ = \frac{1}{F}\frac{dF}{dT} Zdt + \sigma_F Zdz \tag{11.14}$$

Assuming long-run prices are more important than short-run prices, long maturity forward contracts found the basis of the long-term risk-neutral trend. A natural choice of trend will then be to continue using the sample of ten-year forward contracts as an estimate for the risk-adjusted drift.

The annual revenues will now be the product of the two processes, given as

$$V_0\left(F_0(R_T)\right) = e^{-rT} F_0\left(Z_T Q_T\right) = e^{-rt}\left(Z_0 e^{(r-c)T}\right)\left(Q_0 e^{-\lambda_Q \sigma_Q T}\right) e^{(\rho_{ZQ}\sigma_Z\sigma_Q)T} \tag{11.15}$$

Here, Z is the shadow energy price, Q the average production, and r-c the risk-free forward price drift. The parameters ρ_{ZQ}, σ_Z and σ_Q are the correlation between prices and production, and the volatility of price and production, respectively.

11.2.1.3 Parameter estimation

The price trend was determined in section 5.2, and is taken to be 3.15%. We then need to estimate a starting price, the shadow price Z_0. This is the risk-neutral starting point of the forward price trend. Assuming we start our investment on the last day of our long-term price data, we use the longest contract on that given day as a proxy for the shadow price. Looking at the price forward data, the forward contract for 2004 is the longest contract at Nord Pool, and its value is 170.63 NOK/MWh at the last day of trading, the 27th of December 2001. This value is discounted back to the present, using the trend of 3.15%.

For the average annual production, we use discharge data for the last 88 years to estimate a sensible production rate. The production average is calculated from water flow series, capped at the top and bottom to account for over-flow and production suspension. The parameters of the model are then given as follows:

Parameter:	Explanation	Estimate:	Unit
$Q_{AVG\ 9}$	Average water for production[38], small plant.	8944	MWh/year
$Q_{AVG\ 15}$	Average water for production, large plant.	7083	MWh/year
Z_T	Original shadow price	170.63	NOK
Z_0	Discounted shadow price[39]	160.36	NOK
r-c	Price drift	3.15	%
λ_Q	Market price of discharge risk	?	%

Table 11-6: Model parameters for the revenue model

11.2.1.4 Volatility

The combined process introduced above can take any time dependent volatility as input. For prices, the long-term forward price volatility of 14.5% annually is used. This value is both estimated and implied from options, and should give a fair reflection about the plant's value. For the discharge level, we used the same 88 years as we used to calculate the average discharge, to find an estimate of the annual

[38] This is the average production capacity for the last 88 years.

[39] This is the discounted value of the longest one-year contract from Nord Pool. In our case, we start the investment problem 1/1/2002. For this purpose, Z_0 is estimated as the forward contract for 2004 discounted at 3.15% annually. The approach is taken from Schwartz (1998).

volatility. As usual, volatility was found as the standard deviation of returns, this time using annual averages. The volatility estimates for the two plants are included in table 11-7:

The instantaneous correlation, ρ_{PW} was estimated by comparing returns of prices and discharge over the same time interval. This was estimated from prices and discharge data from 1996 to 2001. This can be estimated by defining u_P and u_W as the price and discharge returns, respectively:

$$u_P = \ln\left(\frac{P_{t+1}}{P_t}\right), u_W = \ln\left(\frac{W_{t+1}}{W_t}\right) \tag{11.16}$$

Then, the correlation is found as

$$\rho_{PW} = \frac{Cov(P,W)}{\sigma_P \sigma_W} = \frac{1}{N\sigma_P \sigma_W} \sum_{i=1}^{N} \left(u_P{}^i - \bar{u}_P\right)\left(u_W{}^i - \bar{u}_W\right) \tag{11.17}$$

Performing the analysis, SPSS returned the correlation coefficients listed in table 11-7. They were both significant on the 5% level.

Parameter	Annual estimate
σ_{P-9}	14.5%
σ_{W-9}	19.11%
σ_{W-15}	20.45%
ρ_{PW-9}	-.203
ρ_{PW-15}	-.226

Table 11-7: Volatility parameters

11.2.1.5 Market price of discharge risk and long-term behavior

The forward price will have a long-term drift of (r-c), and is therefore already risk-neutral. The river flow, however, will follow the risk-neutralized drift $-\lambda_Q\sigma_Q$ because the unadjusted drift is given as zero. A long-term drift in water level would simply not be realistic[40]. Following this, the long-term volatility should also approach zero, as the uncertainty in this lack of trend is considered low, and the long-term market price of risk should be expected to be very small. Looking at the volatility of the annual production, this is fairly high at 20%, indicating that, the variance between years is considerable. Since we have not been able to find any power plants to estimate the

[40] This is actually an argument for modeling the discharge as a mean-reversion process with zero drift.

market price of risk from, we will base our analysis of the power plant value on various levels of the market price of discharge risk.

11.2.1.6 Resulting process and value of power plant

The resulting process used for revenues was described in section 4.4.8, and found to be

$$V_0\left(F_0(R_T)\right)=e^{-rT}F_0\left(P_TQ_T\right)=e^{-rt}\left(P_0e^{(r-c)T}\right)\left(Q_{AVG}e^{-\lambda_Q\sigma_QT}\right)e^{(\rho_{PQ}\sigma_P\sigma_Q)T} \quad (11.18)$$

for any given maturity T. Here, ρ_{PQ} is assumed to be constant over time. Finally, by integrating the revenues over the plant's lifetime, the value of the power plant can be shown to be:

$$V_0(R)=\left[\frac{P_0Q_0e^{t\left(-c-\lambda_Q\sigma_Q+\rho_{PQ}\sigma_P\sigma_Q\right)}}{c-\lambda_Q\sigma_Q+\rho_{PQ}\sigma_P\sigma_Q}\right]_{T_1}^{T_2}-I \quad (11.19)$$

As discussed above, we assume zero drift in the production. The plant building time is assumed to be two years and production is assumed to terminate after 40 consecutive years Variable cost are assumed negligible. The final value of the power plant, assuming the long-term discharge risk to be zero, is found in table 11-8, and a discussion of the plant value for various values of λ is given in table 11-9 in the next section.

Max. production	**Value of revenues**	**Investment cost**	**Value of plant**
9 m^3/sec	23,132,000	15,065,000	8,067,000
15 m^3/sec	28,695,000	20,067,000	8,628,000

Table 11-8: Value of power plant using a joint stochastic process

11.2.2 A discussion on the market price of risk

In section 11.1.5, we outlined several possibilities for how to risk-neutralize the forward revenue contracts. This discussion is to a large extent valid for the joint stochastic process. The joint process uses prices fitted to the long-run forward price term structure, and is therefore assumed to be risk-neutral. The same can not be said about the market price of discharge risk, λ_Q. We therefore assumed the discharge risk to be zero.

Our original idea was to value a traded power plant using one of the models above, since the analysis is easy to transfer to a different river. As mentioned in section

11.1.5.1, the traded value of a power plant can be taken as the current certainty equivalent of the power plant. Solving the value of the plant model with respect to λ_Q we could hence obtain the value of λ_Q by comparison to the certainty equivalent, because the two measures should be equal in the absence of risk. The value of λ_Q would then be the solution of equation (11.20).

$$\min_{\lambda_Q} \left(CE - V_0(R)\right)^2 \tag{11.20}$$

This method is particularly easy to apply to the analysis above, since an analytical solution for λ_Q can be determined. Once λ_Q has been determined, the value of our power plant is easy to determine. Since we did not manage to obtain the value of a traded power plant, the value of the plant for various measures of λ_Q is listed in the table below:

λ_Q	Model A (Value in NOK)	Model B (Value in NOK)
0%	8600000	8100000
2%	6700000	6600000
4%	5000000	5300000
6%	3400000	4100000
8%	2000000	3000000
10%	600000	1900000
12%	-600000	1000000
14%	-1700000	100000
16%	-2700000	-700000

Table 11-9: Plant value for different choices of λ

11.2.2.1 Traded Norwegian Power Plants

After asking several actors in the market, we did not succeed locating any traded discharge power plants during the past couple of years. Orkla Borregaard is at the moment considering selling eight discharge power plants, but none of the deals are finalized yet. Other power plant transfers have been parts of larger complex deals, involving more than one power plant, and a parent company. An example of this is the transfer of Pasvik Kraft to Varanger Kraft in 2000. The method is still considered relevant methodically.

11.2.3 Sensitivity to changes in the price trend

The model has assumed that the term structure of forward prices increases by an annual 3.15%. The revenue model was very sensitive to changes in this trend. Of course, the actual trend is unobservable, and 3.15% will never be more than our best scientific guess based on the analysis in section 5.2. The table below assumes that the market price of discharge risk is set to zero, while varying the price trend between 1.42% and 5.45%, as we did for the revenue options approach.

Price trend	Model A	Model B
1.42%	2400000	1600000
2.00%	4100000	3700000
2.50%	5700000	5700000
3.00%	7500000	7900000
3.50%	9500000	10400000
4.00%	11800000	13200000
4.50%	14400000	16400000
5.00%	17300000	20000000
5.45%	20300000	23700000

Table 11-10: Sensitivity to price trend. All values in NOK

As for the revenue options, the value of the plant is very sensitive to changes in the long-term trend. The value is positive for all values, however, which might not be realistic.

11.3 Comparison of the two models

The previous sections have discussed two models for power plant valuation. The first approach looked at a data set of "real" revenues, and the other looked at discharge water and price processes separately, generating a correlated GBM process. Looking at table 11-5 and table 11-10, we see that the value of the power plant is always positive for the GBM model, whereas the present value for the revenue model is considerably less. The main reason for this should be the major risk adjustment done for the first model. Removing the 19.26% risk adjustment in the first model creates values of similar magnitude to the GBM model.

Secondly, we have used slightly different data sets to measure the parameters. The average water in the joint stochastic process model was based on an average over 88 years, whereas the prices were based on long-maturity forward contract data from 2000 and 2001. The average production level during 1996 to 2001, which comprises the volume and price data set of the first model, was about 5% lower than the average over the 88 years. This is not enough to explain all the differences, however.

The seasonal factor of the mean-reversion model does not adequately capture the real seasonal variation in the revenues. It is obvious from the graph presented in figure

11.1 that a simple cosine function could not capture the whole truth. In the long run, the mean-reversion process will converge to this inexact seasonal function. The peaks of the revenue function might then be too low to capture the adequate structure of the revenues. The seasonal function would also partly explain why the two models produce different advice with respect to plant size. Hence, the averaged annual production might be a better estimate than the method of modeling the whole revenue process on a weekly basis. The average production model is also considerably easier to communicate.

None of the models adjust adequately for the market price of discharge risk. This is because we could not obtain a power plant traded in the market. It is not sure, however, that this is priced in the market at all. On the other hand, as seen in section 8.3.2.1, we found an implied market price of risk equivalent to 7% for precipitation derivatives. Although not equal, they represent a similar entity, and might to a certain extent be comparable.

Finally, both models were very sensitive to changes in the long-term risk-free trend. The net change in plant value compared to a trend change has the same magnitude for both models, indicating consistency between the models.

11.3.1 Summary

This chapter has introduced two different valuation approaches for a power plant, one using weekly revenues and option prices, and one using an annual average. The results of the two approaches were different, for a variety of reasons. In the long run, the second approach might turn out to be the best, due to both simplicity and ease of communication, and due to the fact that temporal differences even out over time. We have not seen, however, whether an actual investment can be delayed, to increase the present value of the project. This will be discussed in the following chapter.

12 Timing the investment

This section explains the implementation of the trinomial tree for finding the optimal investment timing for the (American) option to build a run-of-river plant. We investigate this using a mean-reverting and a GBM price process. Firstly, the valuation approach is described, secondly the tree building procedures are explained, then the methodology for valuing the option to invest and finding the early exercise boundary is commented on. We finally find the early exercise boundary and comment on the findings.

12.1 Necessary preconditions

Since our aim is to find the limit for spot price above which investment becomes more profitable than holding the option to invest, we have to use a price model without seasonal variations, otherwise the boundary will be dependent on the time of the year the option holder decides whether to invest or not. Consequently, we use a constant deterministic drift in the prices, from section 5.2, we know that the interval for this drift is [1.42%,5.45%] with 3.15% as the mean value. To keep the model internally consistent, we value the power plant using the same price model. The deterministic trend present in the spot prices is assumed to sustain trough the lifetime of the plant and there are no seasonal variations in prices after the investment is made. With this in mind, the value of the plant is given as

$$\sum_{t=1}^{40} e^{-rt}(P(t)Q) - I \tag{12.1}$$

where *P(t)* denotes the average spot price in year *t*, *Q* is a constant, deterministic production rate taken as the histroical average yearly production of a run-of-river plant with a discharge capacity of 15m^3/second. in section 11.2.1.3, this quantity was found to be 8944 MWh anually. *t* is measured in years.

12.2 Construction the trinomial tree for the mean-reverting process

For the tree construction, we follow Clewlow and Strickland (1998) for implementing an efficient procedure for the Hull-White model outlined in section 3.3.4.3. Firstly, we construct a simplified tree for the noise process as given in equation (3.43).

12.2.1 Determining the up and down moves.

At each time increment, the price can move either up or down a space step given by

$$dP^* = \sigma(i)\sqrt{3\Delta t} \tag{12.2}$$

Observe that the standard deviation is a function of the time increment. The volatility function is choosen to fit the observed volatility structure in spot prices, and is given from the Price model in equation (3.32). In the tree setting, it takes the form

$$\sigma^2(i) = \frac{\sigma^2}{2\kappa}(1 - e^{-2\kappa i \Delta t}) \tag{12.3}$$

12.2.2 Constructing the noise tree

With the volatility function and the up and down increments at each time step defined, the trinomial tree for the noise process is constructed according to the branching scheme described and with the risk-neutral transition probabilities defined in section 3.3.4.3.4

12.3 Fitting the tree to the observed term structure

12.3.1 Determining the displacement function

Following Clewlow and Strickland (1998), we first calculate the state prices at each node in the tree. The state price $Q_{i,j}$ is defined as the value today of a security that pays one unit of cash if node (i,j) is reached and zero otherwise. The state prices are obtained through forward induction by computing

$$Q_{i+1,j} = \sum_{j'} Q_{i,j'} p_{j',j} e^{-r\Delta t} \tag{12.4}$$

where $p_{j',j}$ is the probability of moving from node (i,j') to node (i+1,j). That is, the summation is taken over all nodes j, at time step i which branch to node (i+1,j). Each internal node at time i can be reached from three nodes at time t-1, while the nodes at j_{max} -2 and $-j_{max}$ +2 can be reached from four nodes provided the tree has stopped branching. After all state prices have been calculated, we find the displacement function *a(i)* by taking

$$a(i) = \mathrm{Ln}\left(\frac{e^{-ri\Delta t} F(0, i\Delta t)}{\sum_j Q_{i,j} e^{j\Delta P^*}} \right) \tag{12.5}$$

where $F(0,i\Delta t)$ is the observed price today for a forward contract maturing at time $i\Delta t$.

Given the dipsplacement function *a(i)*, the final price tree is constructed by calculating

$$P_{i,j} = a(i) + P^*_{i,j} \tag{12.6}$$

yielding a trinomial price tree consistent with the observed forward price structure. For the term structure, we use the trend estimated in section 5.2. The model handles any term structure, but we avoid seasonal variations due to the risk of obtaining a seasonally dependent exercise boundary.

12.4 Constructing the tree for the GBM process

The trinomial tree used in this approach is simpler than the corresponding tree for a mean-reverting model. In this case, the transition probabilities are the same for each node regardless of its location, and each node at time step i always branches to three nodes at time step i+1, and each node at time step i can be reached from three nodes at time step i-1. The heigth of the tree is always equal to the number of time steps. The implementation of the tree does not support a volatility function, and using the estimated yearly spot volatility will most likely result in a very high exercise boundary. To evaluate this, we will conduct two runs, one where we use the estimated spot volatility, and one where we use the estimated long-term forward volatility.

12.5 Timing the investment

The value of the American option at each node is found by backwards induction. We know that at maturity, its value is given by its European value (Hull,2000)

$$C_{i,j} = \text{Max}(V(P_{i,j}) - I, 0) \tag{12.7}$$

where V(P) denotes the value of the power plant given a power price P at the time of the investment. I denotes the investment.

12.5.1 The expected value of waiting one period

For a node $i < i_{max}$ we calculate the expected value of the investment if it is delayed one period. The expected value in the next period is given as

$$E(C_{t+1}) = e^{-r\Delta t} \sum_{j'} p_{i,j'} V_{i+1,j'} \tag{12.8}$$

where the summation takes place over all nodes j' that can be reached from node (i,j)

12.5.2 Checking for early exercise

At each node (i,j), the option value given as the expected value of postponing the investment one period is known. Furthermore, the price in the same node (i,j) allow us to obtain the value of exercising the option and build the plant. If the value of making the investment is greater than the value of the option, the option will be exercised.

12.5.3 Determining the early exercise boundary

After the steps outlined above, we will for each time step i have one or more nodes (i,j) where the it is optimal to exercise the option. Since the volatility function we use is a decreasing function of time, the early exercise boundary will be an increasing function of time.

12.6 The early excersise boundary using a mean-reverting model

The input parameters to the model are

Parameter	Value
Start Price	160 NOK/MWh
Volatility	Given in equation (12.3)
Mean-reversion factor (κ)	0.13 (annually)
Time steps	120 (monthly resolution)
Lifetime of option to invest	10 years

Table 12.1 Input parameters to mean-reverting trinominal tree

The choice of mean-reversion factor (κ) deserves a comment. In section 5.5, we estimated κ to 0.035. If we insert this value into equation (3.45) for finding the optimal branching size of the tree, we get a tree that never stabilizes ($j_{max} > i_{max}$) Consequently, we an estimate κ, from the price data from years 1997-2000[41], yielding a maximum branching of 18 with 120 time steps. The results are shown in the plot below.

[41] These years are choosen since years 1996 and 2001 display extremely high prices. A curve fitted to the whole data set will yield a mean to which the prices never revert, thus the value for the mean-reversion parameter becomes too low. See Appendix B

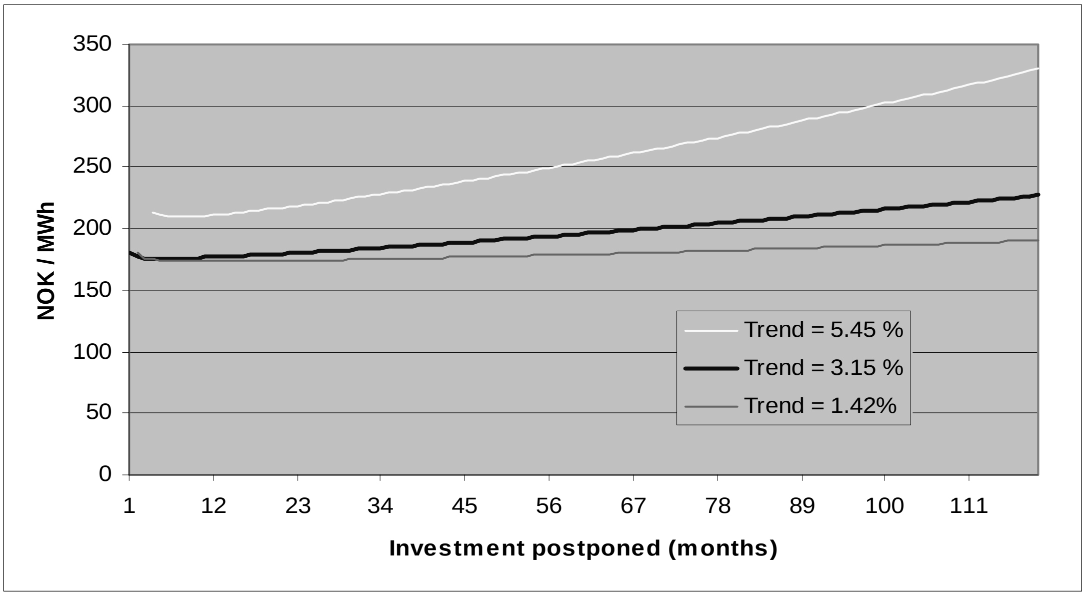

Figure 12.1Early excersise boundary for the option to invest using a mean-reverting price model

12.6.1 The early excersise boundary using Geometric Brownian Motion

Parameter	Value
Start Price	160 NOK/MWh
Volatility	73,5% (annually)
Convenience yield	Variable
Time steps	120 (monthly resolution)
Lifetime of option to invest	10 years

Table 12.2 Input parameters to GBM trinomial tree

The early excersise boundaries for differnet price trends are shown in the plot below. Note that for a trend of 5.45%, it is never optimal to excersise the option early.

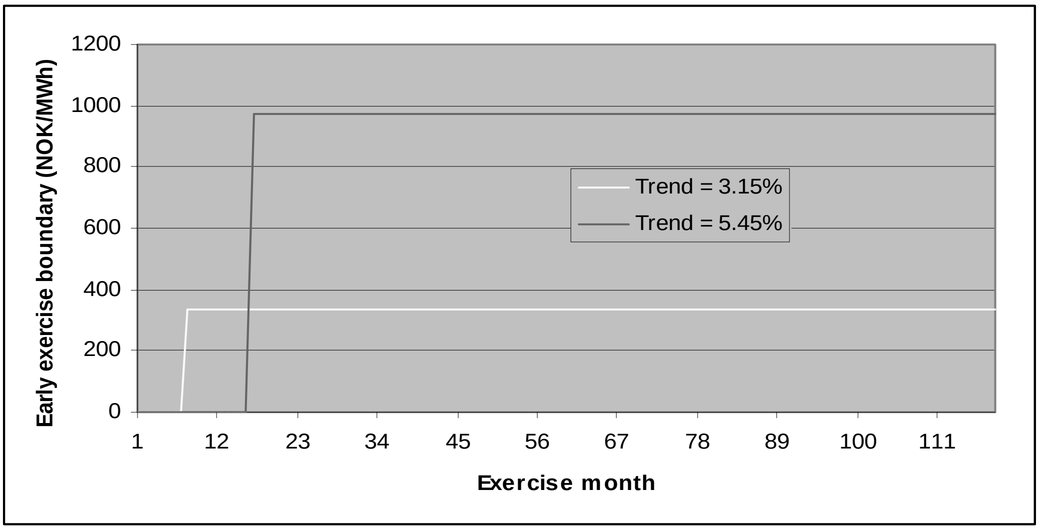

Figure 12.2 Early exercise boundary using a GBM process with annual spot volatility

Using these input parameters, we see that a high price is necessary for early exercise to become optimal. With the drift values used, the spot price will not reach the excersise boundaries of 336 and 970 NOK/MWh for a trend of 1.42% and 3.15% respectivly. This is due to the high volatility estimate. The implementation for the GBM tree only supports constant volatility. We conduct a second run where the volatility is changed to the long-run estimate of 14.5% annually. The early exercise boundaries obtained are presented below. Note that at a trend of 5.45%, it is still not optimal to exercise early.

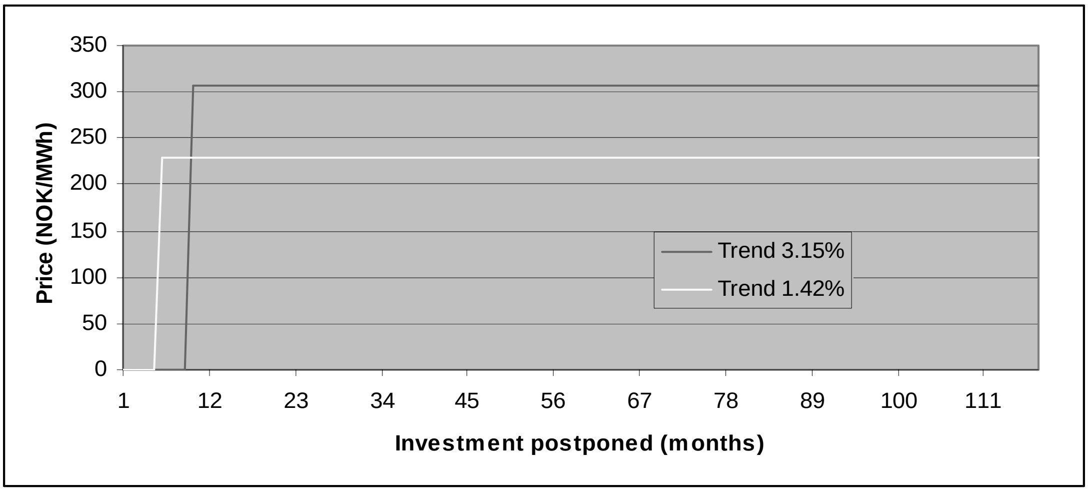

Figure 12.3 Early exercise boundary using a GBM process and forward volatility

12.7 Discussion of results

We see that when we use a mean-reverting model, the exercise boundary is an upward sloping function of time. This is due to the volatility function used. As we move closer to maturity, accumulated volatility approaches a fixed value, hence it becomes less and less profitable to postpone the investment. Furthermore, we observe the higher the trend, the higher the price must be for the investment to be worth more than the option. This is consistent with intuition – an increase in trend translates into an higher expected price in the future, and the probability of a favourable price development increases. It is not optimal to invest immediately, using a trend of 3.15%, the price must move close to 180 NOK during the first two years. This is about 6% above than the drift indicated by the observed term structure.

For the GBM model, the results look somewhat different. Firstly, the exercise boundary is a straight line once excercising becomes optimal. This is due to the design of the tree – in the GBM model, the drift is reflected in the (constant) transition probabilities, and we use a constant long-term volatility. If the trend is 5.45%, the expected profit of waiting one period is so high that early exercise never becomes optimal. The prices at which exercise takes place are also considerably higher than for the mean-reverting model. This is due to the fact that there is no parameter pulling the prices back to the long-term mean in the GBM tree, and hence the probability of an upwards move is equally high regardless of the position in the price tree.

Despite the fact that delaying the investment is optimal both for the mean-reverting model and the GBM model, the upwards price moves necessary are not large. This means that investment might well take place before the option to invest expires.

13 Suggestions for further work

We have approached the valuation problem from some angles, while others still remain unexplored. Some suggestions for further work on the topics discussed in this thesis are presented below

- Power prices display large peaks. It would be interesting to see how well a model incorporating a jump-diffusion process performs compared to the models we studied
- Although we have devoted space to discuss price volatility, there are still models available that can shed new light on the volatility structure. Notably two-factor models treating volatility as a stochastic process would be a welcome extension
- For precipitation derivatives, we developed a simple model. A better understanding of the underlying weather variable could help us develop a more sophisticated model
- While we found that the precipitation index is a poor volumetric hedge for a run-of-river power plant, it is reasonable to assume that it would work better for a production system with storage capacity. We know that inflow is largely dependent on accumulated snow, and preliminary analysis do show that accumulated precipitation and the fill rate of Norwegian power reservoirs are well correlated
- Some interesting correlation patterns between precipitation and prices and precipitation and water discharge has been commented on. To further study these relationships, longer time series would be needed. It is also likely that a more systematic statistical analysis or other techniques for detrending the data could provide more insight. Specifically, we did not discover how the trading date affects correlation for long-term forward contracts
- We have discovered that the price trend is important for the optimal size of a run-of-river power plant. An extension to our model could be to value the option of adding a second turbine if prices should indeed increase more than expected.

14 Conclusion

The aim of this thesis was to study models for power prices and power derivatives as well as models for precipitation with the aim of developing precipitation derivatives. Armed with these models, we planned to value a run-of-river power plant as a portfolio of power- and weather derivatives through a real options framework

A review of the spot and forward models for energy prices revealed a solid base of models of various complexities. Energy price behavior is driven by a number of factors, and the inclusion of all these in an analytical tractable model, seems very difficult. Seasonal variations, jumps, mean reversion and time-varying volatility characterize prices. Furthermore, the data set of power prices is too short to detect basic characteristics such as a price drift from the data set. We chose to explore the mean-reversion and seasonal dynamics, and soon found it very difficult to determine parameters such as market price of risk. In that respect, modeling forward prices proved to be easier, as these are already risk neutral.

A survey of the literature on weather derivatives did not yield any frameworks suitable for developing preference-free price models for precipitation derivatives. Furthermore, existing stochastic models for precipitation are not developed with financial applications in mind, and are consequently too complex to develop analytical solutions for derivatives pricing. On this background, we suggested using a one-factor mean-reverting process to model a precipitation index covering Norway. Our simple model seemed to perform adequately with respect to historical observations and allowed closed-form solutions to derivatives pricing. Due to the illiquidity of the weather market in Norway, it is difficult to test the models against market prices, but quotes obtained allowed us to estimate the market price of short-time precipitation risk to approximately 7%, thus making the model theoretical complete. Our model was compared to a nonparametric approach for derivatives valuation, but due to the lack of reference data, we cannot draw any clear conclusions on how well the different models perform. The extreme weekly volatility, estimated to 90% in precipitation and the relatively short data series further complicated the analysis.

Armed with precipitation derivatives, we wanted to use these as a hedge to remove the volumic risk inherent in a run-of-river power plant that is larger than the minimum water discharge of the river in which the proposed plant is built. Correlation analysis of discharge time series and precipitation data found the correlation low at around –0.2 depending on the river investigated during the period it would be expected to be at its highest. This was explainable, firstly by the fact that the index consists of measurement stations scattered all over Norway, and secondly by the relationship between precipitation and water discharge being complicated due to snowfall and the spring flood.

This forced us to consider alternative approaches for valuing the run-of-river plant, and we proceeded by finding the economically optimal size of a plant. Our findings indicated that the solution to the sizing problem and the plant value is highly sensitive to the estimates for trend in power prices as well as the underlying price model used.

Since the findings indicated the valuation to be dependent on the methodology employed, we explored two very different approaches to the plant valuation problem. First, we considered the seasonal structure of the power plant revenues, and included operational characteristics to model revenues as a portfolio of spread options. The second approach considered an annual resolution on the input, and no seasonal variation. Revenues were assumed to be composed of two correlated Brownian motions of price and quantity. Due to the lack of power plants traded in the market, this volume risk was assumed to be zero to keep the model consistent. Using the forward term structure and forward shadow price instead of spot price, the revenues were made risk neutral. The comparison of the two yielded different, but explainable differences.

Finally, we investigated the optimal investment-timing problem, and found that regardless of the price trend assumed and the underlying stochastic process for power prices, it was never optimal to invest immediately. The early exercise boundary at which investment becomes optimal varies considerably depending on assumptions about the volatility function and the method used in the analysis.

15 References

Alaton, P, D Boualem and D Stillberger (2001) ”On Modeling and pricing weather derivatives”. *Research paper*, Fat Tails Financial Analysis AB

Bjerksund, P, H Rasmussen and G Stensland (2000) ”Valuation and Risk Management in the Norwegian Electricity Market” *Working paper*, The Norwegian School of Economic and Business Administration (NHH).

Bjørkvoll, T, S-E Fleten, M Nowak, A Thomasgard, S W Wallace (2001) ” Power generation planning and risk management in a liberalised market” *IEEE Porto Power Tech Proceedings*, pp. 426-431.

Black, F (1976) ”The pricing of commodity contracts”, *Journal of Financial Economics* 3 (1976) pp 167-179.

Black, F and M Scholes (1973) ” The Pricing of Options and Corporate Liabilities” *The Journal of Political Economy*, Vol. 81, No. 3. (May - Jun., 1973), pp. 637-654.

Brody, D.C, J Syroka and M Zervos (2001) ”Dynamical Pricing of Weather Derivatives”, *Working Paper submitted to Quantitative Finance*

Cameron, D, K Beven and J Tawn (1999) “An evaluation of three stochastic rainfall models”, *Journal of Hydrology*, 228 pp. 130-149

Cao, M and J Wei (2000) ”Equilibrium valuation of weather derivatives”, *Working Paper,* Department of management, University of Toronto.

Clewlow, L and C Strickland (1998) ”Implementing Derivatives Models” Wiley & Sons, Chichester.

Clewlow, L and C Strickland (2000) ”Energy Derivatives: Pricing and Risk Management” Lacima Publications.

Clewlow, L, C Strickland and V Kaminski (2000) ”Making the most of mean reversion” *Energy Power Risk Management*, Risk Waters Group, November 2000, www.lacimagroup.com, visited 04/05-2002.

Cooper, V (2001) ”Singin' in the rain”, *Risk Management*, February 2001 pp 27-29

Cox, J, S. Ross and M. Rubinstein (1979) ”Option pricing: A simplified approach”, *Journal of Financial Economics*, 7 (October 1979), pp 229-64.

Dalsbø, G (2002) Email to Bjørdal, N received 06/03-2002

Davis, M (2001) ”Pricing weather derivatives by marginal value”, *Review Article, Quantitative Finance* vol 1-4

Dobbe, T and S Sigmo (2002) ”Realopsjonsanalyse av et gasskraftverk”, *Final thesis*, Institute of Industrial Economics and Management of Technology, Norwegian University of Science and Technology.

Dischel, R (1998) ”Options Pricing – Black-Scholes Won't Do”, *Energy and Power Risk Management – Weather Risk Special Report*, October 1998

Dischel, R (1998) ”Weather Data for Pricing Weather Derivatives” *Applied Derivatives Trading* 33, December 1998.

Dischel, R (1999) ”The fledging weather market takes off, Parts 1-4”, www.erivativesreview.com, accessed February 12th, 2002

Dischel, R (1999) ”A Weather Risk Management Choice”, in ”*Insurance and weather derivatives : from exotic options to exotic underlyings*”, ed. Geman, H, Risk Books, London

Dixit, A and R S Pindyck (1994), ”Investment under uncertainty”, Princeton Univ Press.

Element Re (2001), ”Insurance versus derivatives”, *Energy & Power Risk Management – Weather Risk Special Report* page S2, September 2001.

Elmlund, L (2002) Weather trader at AEP Energy.Phone conversation with Bjørdal, N on 26/06-2002

Eydeland A and H Geman (1999) "Pricing Power Derivatives," *Risk*, October, 1998.

Fleten, S-E (2000) ”Portfolio Management emphasizing electricity market applications. A stochastic programming approach” *Dr. Ing. Thesis*, Norwegian University of Science and Technology (September 2001)

Fleten, S-E and J Lemming (2001) ”Constructing Forward Price Curves in Energy Markets” *Working Paper*, Norwegian University of Science and Technology

Fuller, W.A (1996) ”Introduction to Statistical Time Series 2nd edition”, Wiley Series in Probability and Statistics, John Wiley & Sons, Inc, New York

Geman, H (1999), ”The Bermuda Triangle”, in *” Insurance and weather derivatives : from exotic options to exotic underlyings*”, ed. Geman, H, Risk Books, London

Gibson, R and E Schwartz (1990) ”Stochastic Convenience Yield and the Pricing of Oil Contingent Claims” *The Journal of Finance*, Vol. 45, No. 3. (July), pp. 959-976.

Gjølberg, O and T Johnsen (2001) ”Electricity Futures: Inventories and Price Relationships at Nord Pool” *Working paper*, Department of Economics and Social Sciences, Agricultural University of Norway.

Hull, J.C. (2000) Options, Futures & other Derivatives”, Prentice-Hall International Inc, 4th edition

Johnsen, T A (2001) ”Demand, generation and price in the Norwegian market for electric power” *Energy Economics* 23 pp227-251.

Kashyad, R.L and A Ramachandra Rao (1976) Dynamic Stochastic Models from Empirical Data, Academic Press, New York

Kellerhals, B P (2001) ”Pricing Electricity Forwards under stochastic volatility” *Working paper*, Department of Finance, College of Economics and Business Administration, Eberhard-Karls-University Tübingen.

Killingtveit, Å (2000) ”Tilsigsprognoser - lang sikt”, *Short paper* written for course on inflow modeling for NITO course on hydrology

Killingtveit, Å (2002), Meetings attended by Bjørdal, N and Skogen, A on 03/04-2002 and 14/05-2002

Knittel, C R and M Roberts (2000) ”Financial Models of Deregulated Energy Prices I: Discrete Time Models” *Working paper*, Department of Finance and Economics, Boston University and University of California Energy Institute.

Knittel, C R and M Roberts (2001) ”Financial Models of Deregulated Energy Prices II: An application to the California Market” *Working paper*, Department of Finance and Economics, Boston University and University of California Energy Institute.

Lucia J.J and E Schwartz (2001) ”Electricity prices and power derivatives: Evidence from the Nordic Power Exchange”

Lo, A W and A C MacKinlay (1999) “A non-random walk down Wall Street” Princeton University Press, Princeton, New Jersey.

Markovitz, H (1953) ”Portfolio selection” *The Journal of Finance*, Vol. 7, No. 1. (Mar, 1952), pp. 77-91.

Martin, S.W, B J Barnett and Coble (2001) ”Developing and Pricing Precipitation Insurance”, *Journal of Agricultural and Resource Economics* 26(1) pps 261-274.

McDonald, R and D Siegel (1996) ”The Value of Waiting to Invest” *Quarterly Journal of Economics* (November): 101, 707-728.

McIntyre, R (1999) ”Black-Scholes will do” *Energy and Power Risk Management*, November 1999

Modigliani, F and M H Miller ”The Cost of Capital, Corporation Finance and the Theory of Investment” *The American Economic Review*, Volume 48 (3) (June 1958)

Moreno, M (2001) Rain Risk, *Speedweel Weather Derivatives*, www.weatherderivs.com, accesed on March 13th, 2002

Onof, C. and H.S. Weather (1994) “Modeling of British rainfall using a random parameter Bartlett-Lewis Rectangular Pulse Model” *Journal of Hydrology*. 149, 67-95

Pindyck, R S (1991) ”Irreversibility, Uncertainty, and Investment”, *Journal of Economic Literature* 29(3) (September 1991) pp 1110-1148.

Pilipovic, D (1997) ”Energy Risk. Valuing and Managing Energy Derivatives” McGraw-Hill, New York

Pirrong, C (2000) ”The Price of Power: The Valuation of Power and Weather Derivatives”, Olin School of Business, Washington University

Ramamurite, S (1999) ”Weather Derivatives and Hedging Weather Risks”, in ” *Insurance and weather derivatives : from exotic options to exotic underlyings*”, ed. Geman, H, Risk Books, London

Schieldrop, B (2002) Phone conversation with Bjørdal, N on 25/02-2002

Schwartz, E (1997) ”The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging” *The Journal of Finance* Vol 52 (3) (July 1997) pp 923-973.

Schwartz, E (1998) ”Valuing Long-Term Commodity Assets” *Journal of Energy Finance & Development*, Volume 3 (2) pp 85-99.

Schwartz , E and J E Smith (2000) ”Short-Term Variations and Long-Term Dynamics in Commodity Prices” Management Science Vol 46 (7) (July 2000) pp 893-911.

Skaugen, T (2002) Email to Bjørdal, N received 04/03-2002

Sharpe, William F., “Capital Asset Prices: A theory of Market Equilibrium under conditions of risk” *Journal of Finance* 19 (September 1964) 425-442.

Shaw, E. M (1988) Hydrology in Practice, second edition, Van Nostrand Reinhold (International)

Smidt, P.C.T (2001) ”Weather Derivatives – A Look at the Practical Issues of Using Derivatives to Hedge Temperature-Sensitive Business”, *MSc Thesis*, Heriot-Watt University, Edinburgh

Steele, J M (2000) ”Stochastic Calculus and Financial Applications” Springer-Verlag, New York

Strickland, C (2002) ”Derivative modeling Approaches used in Real Option Valuation” Lacima Group, www.lacimagroup.com, visited 21/7-2002.

Telser, L G (1958) ”Futures Trading and the Storage of Cotton and Wheat” *The Journal of Political Economy*, Vol. 66, No. 3. (Jun., 1958), pp. 233-255.

Tigler, J and T Butte (2001), Weather Derivatives – a quantitative analysis. *Research Thesis* at Darmstadt University of Technology

Trigeorgis, L (2000) ”Real Options”, MIT Press, Cambridge, Massachutes

Torró, H, V Meneu, and E Valor (2001) ”Single Factor Stochastic Models with Seasonality Applied to Underlying Weather Derivatives Variables”, *working paper*, Facultat d’Economia, Universitat de València, Spain

Walpole, R.E, R.H Myers and S.L Myers (1998) ”Probability and Statistics for Engineers and Scientists”, Prentice Hall International Inc. New Jersey

Appendix A Derivations of mean-reverting processes

Normal mean-reverting process

Let Pt follow the equation

$$P_t = f(t) + X_t \tag{A.1}$$

$$dX_t = -\kappa X_t dt + \sigma dz \tag{A.2}$$

In equation (A.1), P_t is the price at time t, given as the deterministic function f(t) and a stochastic factor X_t following the dynamics given in equation (A.2). Equation (A.2) tells us that a random shock X_t will be smoothed away over time at the rate of κ. New random shocks may occur, but they will all be smoothed over time.

Now, rearranging equation (A.1), we obtain

$$X_t = P_t - f(t) \tag{A.3}$$

This can be put into equation (A.2) to obtain

$$d\left(P_t - f(t)\right) = -\kappa\left(P_t - f(t)\right)dt + \sigma dz \tag{A.4}$$

Using Ito's lemma and rearranging, we obtain

$$dP_t = \kappa\left(\frac{df}{dt}(t) + f(t) - P_t\right)dt + \sigma dz \tag{A.5}$$

This is a process reverting to a long-term mean of *f(t).*

Log mean-reverting process

Let $\ln(P_t)$ follow the equation

$$\ln P_t = f(t) + X_t \tag{A.6}$$

$$dY_t = -\kappa Y_t dt + \sigma dz \tag{A.7}$$

As for the process in P_t, equation (A.6) can be rearranged and inserted into equation (A.7) to obtain the expression in equation (A.8).

$$d\left(\ln(P_t) - f(t)\right) = -\kappa\left(\ln(P_t) - f(t)\right)dt + \sigma dz \tag{A.8}$$

Now, by using Ito's lemma on P_t, we get the following expressions

$$\frac{d\ln(P_t - f(t))}{dt} = \frac{df}{dt}, \frac{d\ln(P_t - f(t))}{dP_t} = \frac{1}{P_t}, \frac{d^2\ln(P_t - f(t))}{dP_t^2} = -1/P_t^2 \tag{A.9}$$

Following this, we put the results into equation (A.10)

$$dF = \frac{dF}{dt}dt + \frac{dF}{dP_t}dP_t + \frac{1}{2}\frac{d^2F}{dP_t^2}(dP_t^2), F = \ln P_t + f(t) \tag{A.10}$$

This result can then be written as equation (A.11):

$$d(\ln P_t + f(t)) = \frac{df}{dt}dt + \frac{1}{P_t}dP_t - \frac{1}{2}\sigma^2 dt \tag{A.11}$$

This expression is replaced into the left-hand side of equation, and rearranged to the result presented in equation (A.12) and (A.13).

$$dP_t = \kappa(b(t) - P_t)dt + \sigma P_t dz \tag{A.12}$$

$$\frac{1}{\kappa}\left(\frac{\sigma^2}{2} + \frac{df}{dt}(t)\right) + f(t) \tag{A.13}$$

This can easily be extended into a proof for a risk neutral process in P_t or $ln(P_t)$.

Appendix B Regression over wet years

The below table includes the parameters determined using the mean-reversion process in prices. The process was described in chapter 5, and was introduced by Lucia and Schwartz (2001).

Parameters for the mean-reversion price process estimated over the wet years, 1997-2000				
Parameter	Estimate	Std. Error	Lower 95% CI	Upper 95% CI
Constant	118.580	6.089	106.577	130.582
Gamma	36.087	6.657	22.965	49.210
Tau	0.978	0.029	0.921	1.036
Kappa	0.135	0.033	0.071	0.200
R^2	0.896			

Table B.1: Parameters of the mean-reversion price process estimated over the wet years 1997 to 2001

Appendix C The Enron Scandinavian Precipitation Index

The Norwegian Stations

ID	LOCATION	IN OPERATION SINCE	SHUT DOWN	COUNTY	BACKUP	WEIGTH
7010	RENA - HAUGEDALEN	Jan-58		HEDMARK	VENABU	6.9
13420	VENABU	Aug-80		OPPLAND	BERKÅK-LYNGHOLT	8.0
23420	FAGERNES	Jul-82		OPPLAND	VENABU	15.0
25590	GEILO - GEILOSTØLEN	Sep-66		BUSKERUD	FAGERNES	3.2
25610	GEILO - STRAND	Jan-51	Jul-66	BUSKERUD		
31610	MØSSTRAND	Dec-63	Apr-76	TELEMARK	GEILO-GEILOSTØLEN	11.5
31620	MØSSTRAND II	Nov-80		TELEMARK		
36560	NELAUG	Jul-66		AUST-AGDER	MØSSSTRAND II	2.0
36580	NELAUG - ØYNES	Aug-60	Jun-66	AUST-AGDER	MØSSTRAND II	
42920	SIRDAL - TJØRHOM	Sep-74		VEST-AGDER	SAUDA	8.4
46610	SAUDA	Mar-28		ROGALAND	EIDFJORD – BU	5.6
49580	EIDFJORD - BU	Jul-78		HORDALAND		3.9
49630	EIDFJORD	Jan-20		HORDALAND		
52290	MODALEN II	Jun-80		HORDALAND	EIDFJORD – BU	3.3
52300	MODALEN	Jul-85	May-80	HORDALAND		
54120	LÆRDAL - MOLDO	May-96		SOGN OG FJORDANE	FAGERNES	19.3
54130	LÆRDAL - TØNJUM	Jun-48	Apr-96	SOGN OG FJORDANE		
82290	BODØ VI	Jan-53		NORDLAND	LEKA	10.8
89350	BARDUFOSS	Jan-41		TROMS	BODØ VI	3.0
93140	ALTA LUFTHAVN	Dec-63		FINNMARK	BODØ VI	3.7
93150	ALTA ELVEBAKKEN	Sep-40	Nov-63	FINNMARK	BARDUFOSS	
98550	VARDØ	Jan-1867		FINNMARK	ALTA LUFTHAVN	1.7
60500	TAFJORD	Jan-30		MØRE OG ROMSDAL	LÆRDAL – MOLDO	2.5
66700	BERKÅK	Dec-29	Aug-67	SØR-TRØNDELAG	SELBU – STUBBE	
66710	BERKÅK II	Sep-67	Nov-80	SØR-TRØNDELAG	SELBU – STUBBE	
66730	BERKÅK - LYNGHOLT	Jul-82		SØR-TRØNDELAG	SELBU – STUBBE	7.0
68300	SELBU	Jan-21	Jun-76	SØR-TRØNDELAG	STORLIEN	4.8
68310	SELBU - BOGSTAD	Nov-76	May-79	SØR-TRØNDELAG	STORLIEN	
68340	SELBU - STUBBE	Sep-79		SØR-TRØNDELAG	STORLIEN	
75600	LEKA	May-40		NORD-TRØNDELAG	SELBU	3.7

Table C.1 Norwegian measurement stations in the Enron Scandinavian Precipitation Index

Source: Eliassen (2002), Schieldorp (2002) and DNMI

The weights are listed for the exact stations in the index. Stations without weights are older stations located close to a measurement station included in the index, and as seen from the “in operation since” column, these older stations will have to be used to construct longer series.

The Swedish stations

ID	LOCATION	IN OPERATION SINCE	SHUT DOWN	COUNTY	BACKUP	WEIGTH
N/A	GUNNARAN	N/A	N/A	N/A	HEMAVAN	4.3
N/A	GÄDDEDE	N/A	N/A	N/A	GUNNARN	16.4
N/A	HEMAVAN	N/A	N/A	N/A	RITSEM A	11.9
N/A	KATTERJÅKK	N/A	N/A	N/A	BARDUFOSS	6.0
N/A	MALUNG	N/A	N/A	N/A	RENA – HAUGEDALEN	15.4
N/A	RITSEM A	N/A	N/A	N/A	KATTERJÅKK	23.9
N/A	STORLIEN	N/A	N/A	N/A	GÄDDEDE	13.6
N/A	SUNDSVALL FLYGPLATS	N/A	N/A	N/A	MALUNG	1.8

Table C.2 Swedish measurement stations in the Enron Scandinavian Precipitation Index

Source: Eliassen (2002) and Schieldorp (2002)

Underlying data source

The underlying data source is the time series for the listed stations operated by the state-run meteorological agencies in Norway and Sweden. However, only the Norwegian data series are available for our analysis, since the cost of obtaining the Swedish data is too high as opposed to the Norwegian which are provided free of charge for educational purposes.

Quality of the data

According to G. A. Dalsbø at the Market Department of the Norwegian Meteorological Agency, quality control is conducted on all data series, so we have no reason to doubt the accuracy of the data or have any ambition to improve the series.

Missing data points

For various reasons (maintenance, vandalism, displacement etc) a measurement station can be unavailable for short periods of time (Dalsbø, 2002). The Enron Index defines back-up stations to use in such cases. The only problem is the 33 consecutive missing observations from the Selbu station (backed up by Storlinen, a Swedish station) We have used Berkåk instead, which is the closest Norwegian station for which we have data. As this is only 33 out of 6935 data points, the error introduced is negligible.

Length of record available

There are considerable variations in the length of the time series. The youngest station in the index has been in operation only since July 1982, while daily values are recorded back to 1867 for the longest-lived station used in the index. The data sets we obtained started January 1951 if the station had been in operation since before 1951, otherwise we received the complete time series for each statio, thus the index is complete from July 1982 until December 2001 if we back-up the few missing values.

Appendix D Precipitation return volatility

The plot below shows the standard deviation of the precipitation index returns. By insepection, the assumption of constant return volatility seems supported.

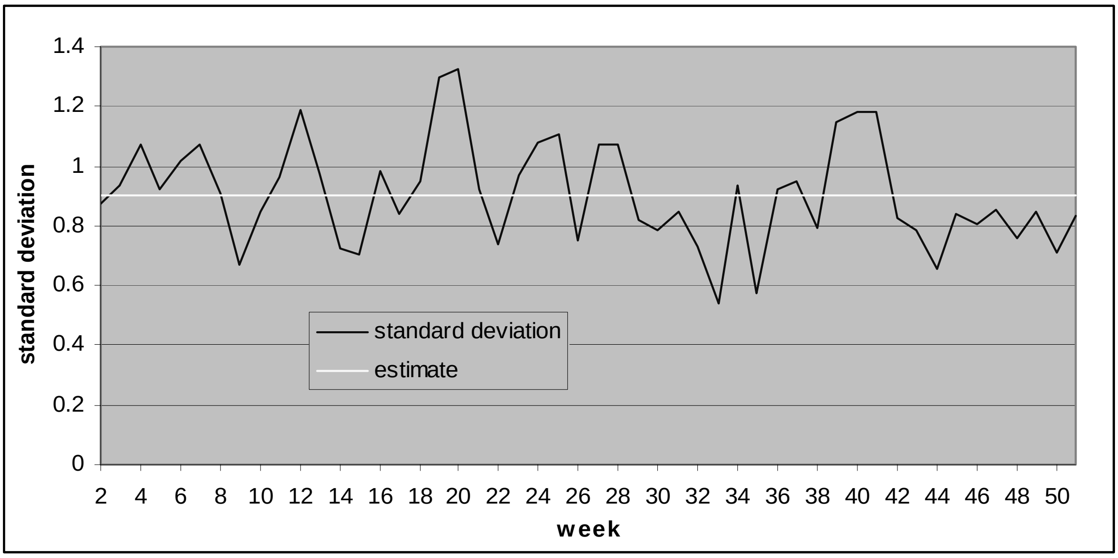

The results from the test of the hypothesis in chapter 7 are presented in the table below. If cell (i,j) is shaded, it means that we can be 95% confident that the return variance of week i is different from week j. The lefmost and upper columns denote week of year, while table values show the Fisher-test p-value in per cent.

	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28
2		92	70	90	11	75	93	76	41	57	96	23	35	67	37	33	82	13	85	49	97	91	**2**	49	83	27
3		57	94	73	57	89	16	67	91	32	86	28	24	83	65	95	18	15	95	31	88	56	50	36	57	58
4			52	83	100	48	5	33	65	67	70	10	8	73	31	61	42	38	54	12	68	98	90	14	99	99
5				67	52	95	19	73	85	29	80	31	27	77	71	89	15	13	98	35	82	51	45	41	52	53
6					83	62	8	44	81	52	86	15	13	89	42	77	31	27	68	18	84	81	74	21	83	84
7						48	5	32	65	67	69	10	8	72	31	61	43	38	53	12	68	99	91	14	99	98
8							21	78	80	26	75	34	30	72	75	84	14	12	93	38	77	47	41	44	48	49
9								33	13	**2**	12	75	83	11	34	15	**1**	**1**	18	70	12	5	**4**	62	5	5
10									59	16	55	50	44	53	97	63	8	7	72	55	57	32	27	62	32	33
11										38	95	23	19	92	57	95	21	19	86	26	97	64	57	31	65	67
12											41	**4**	**3**	44	15	35	71	65	30	**5**	40	69	76	6	67	66
13												21	18	97	53	91	24	20	82	24	98	68	61	28	69	71
14													92	20	52	25	**2**	**1**	30	94	22	10	8	86	10	11
15														16	46	21	**1**	**1**	26	86	18	8	6	78	8	9
16															51	88	25	22	79	22	95	71	64	26	72	74
17																61	7	6	69	57	55	30	26	65	31	32
18																	19	17	91	29	93	60	53	33	61	62
19																		93	16	**2**	23	44	50	**3**	43	41
20																			14	**2**	20	39	45	**2**	38	37
21																				34	84	52	46	39	53	55
22																					25	11	9	92	12	12
23																						66	59	29	67	69
24																							92	14	99	97
25																								11	91	89
26																									14	15
27																										98

	29	30	31	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47	48	49	50	51
2	61	64	85	4	86	43	78	6	81	69	10	71	58	40	45	30	4	68	15	85	46	60	31
3	58	46	68	30	3	100	4	95	97	47	40	34	34	60	45	14	65	52	69	38	66	24	62
4	27	20	33	12	1	57	1	54	60	20	77	69	69	28	19	4	31	23	34	15	32	9	30
5	64	51	74	34	3	94	5	98	90	52	36	30	30	66	50	16	71	57	75	43	72	27	68
6	37	28	45	17	1	72	2	68	76	29	62	54	54	38	27	7	42	32	46	22	43	13	40
7	27	19	33	11	1	57	1	53	60	20	78	70	69	28	19	4	31	23	34	15	32	8	29
8	68	55	79	38	4	89	6	93	85	57	32	27	27	70	54	18	76	61	80	46	76	30	73
9	39	51	32	71	38	16	52	18	15	49	3	2	2	38	52	92	34	45	32	60	34	82	36
10	90	75	99	54	7	68	11	72	64	77	21	17	17	92	74	28	98	82	98	65	98	45	95
11	51	39	60	25	2	91	3	86	94	41	46	40	40	52	39	11	57	45	61	32	58	20	55
12	13	9	16	5	0	32	0	30	35	9	89	98	98	13	9	2	15	11	17	7	16	3	14
13	47	36	56	23	2	86	3	82	90	38	50	43	43	49	36	10	53	41	57	30	54	18	51
14	59	72	50	95	23	28	34	30	26	71	6	4	4	57	74	68	52	66	49	83	52	93	55
15	52	65	44	87	27	24	39	26	22	63	5	4	3	51	66	75	46	59	43	75	45	99	48
16	45	34	53	22	1	83	3	79	87	36	52	46	45	46	34	9	51	39	54	28	51	17	48
17	92	78	97	56	7	65	11	69	62	79	20	16	16	94	76	30	100	85	96	67	99	47	97
18	55	43	64	28	2	95	4	91	99	44	43	37	37	56	42	12	61	48	65	35	62	22	58
19	6	4	8	2	0	17	0	16	19	4	61	68	69	6	4	1	7	5	8	3	8	1	7
20	5	3	7	2	0	15	0	14	16	3	55	62	63	5	3	0	6	4	7	2	6	1	6
21	62	50	72	33	3	96	5	82	92	51	37	31	31	64	49	15	69	56	73	41	70	26	67
22	64	78	55	99	21	31	30	34	29	76	7	5	5	62	79	62	57	71	54	89	57	87	60
23	49	38	58	24	2	88	3	84	92	39	48	42	42	50	37	10	55	43	59	31	56	19	52
24	26	19	32	11	1	56	1	52	59	20	79	71	71	27	18	4	30	22	33	15	31	8	29
25	22	16	28	9	0	49	1	46	52	17	87	78	78	23	15	3	26	19	28	12	26	7	24
26	72	86	62	91	17	37	26	39	34	84	8	7	7	70	88	55	65	79	61	97	64	79	67
27	27	19	33	11	1	57	1	53	60	20	78	70	69	28	19	4	31	23	34	15	31	8	29
28	28	20	34	12	1	58	1	55	61	21	76	68	68	29	20	4	32	23	35	16	33	9	30
29		85	89	63	9	59	14	62	55	87	17	14	13	98	84	34	92	92	88	74	91	53	95
30			74	77	13	46	19	50	43	98	12	9	9	83	99	44	77	93	73	89	77	66	80
31				54	6	68	11	72	65	76	21	17	17	91	73	28	97	81	99	64	98	44	94
32					21	31	31	33	28	75	6	5	5	61	78	64	56	70	53	88	56	88	59
33						3	81	3	2	12	0	0	0	8	13	43	7	10	6	16	7	27	8
34							4	96	96	48	39	34	34	60	45	14	65	52	69	38	66	24	63
35								5	4	19	1	0	0	13	20	59	11	16	10	24	11	38	12
36									92	51	37	31	31	64	49	15	69	56	73	41	70	26	67
37										45	42	36	36	57	43	13	62	49	66	36	63	22	59
38											12	10	10	85	97	43	79	95	75	87	78	64	82
39												91	91	17	11	2	20	14	22	9	20	5	18
40													94	14	9	2	16	11	18	7	17	4	15
41														14	9	2	16	11	18	7	16	4	15
42															82	33	94	90	90	73	93	51	97
43																45	76	91	72	90	75	67	79
44																	29	39	27	53	29	74	31
45																		84	96	67	99	47	97
46																			80	82	84	60	87
47																				63	97	44	93
48																					66	76	70
49																						46	96
50																							49

Table D.1 Results for Fisher test for equal variance in precipitation return volatility

Appendix E Correlation between the precipitation index and water discharge

The two following figures show the correlation coeffecients between the precipitation index and Gaula (Figure E.1) and Glomma (Figure E.2). While the plots indicate strong correlation compared to the values found in chapter 9, however, the correlation coeffecients shown below are not statistically significant. This is due to few data points (12 for each week) and clear noise in the data. The large values for the winter weeks are probably explainable by the fact that discharge is initially very low, and if a large increase in precipitation translates into a slight increase in water discharge, the correlation will come out as high.

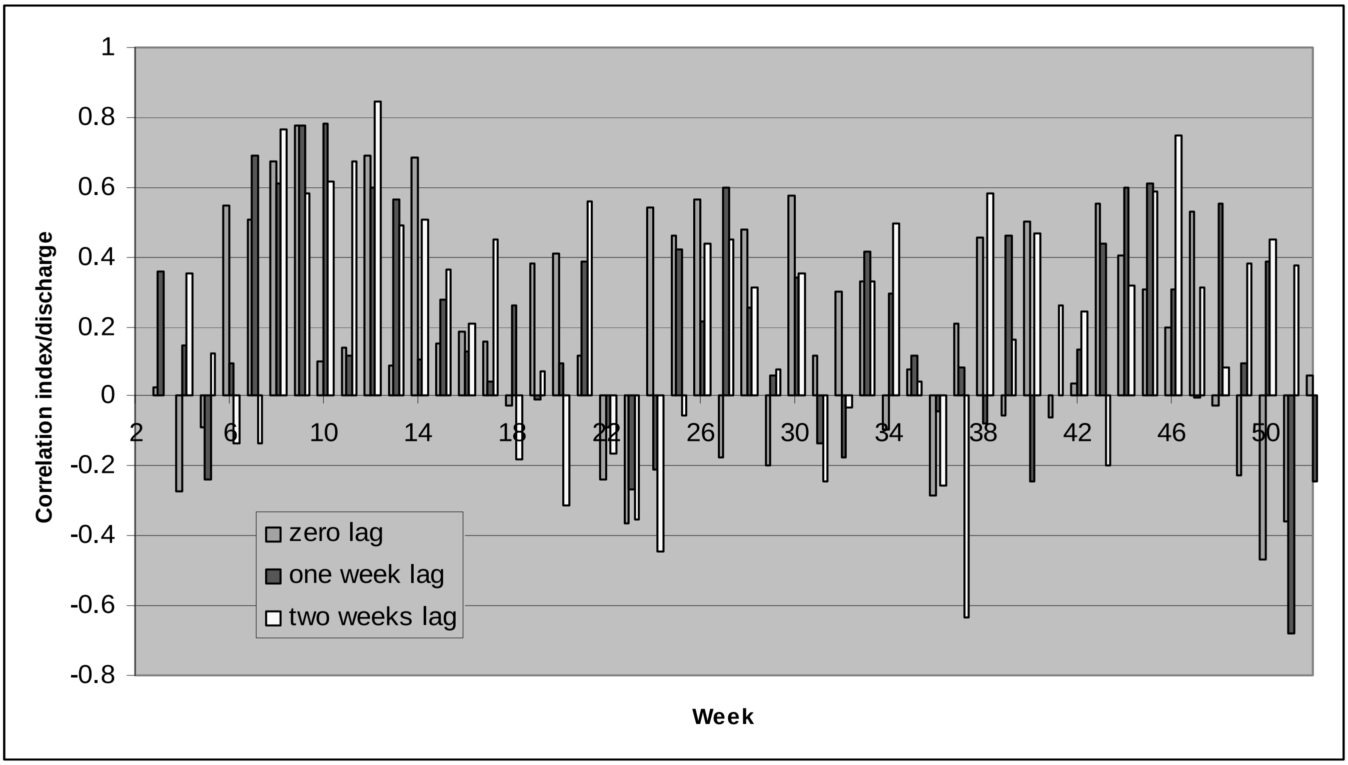

Figure E.1 Correlation between the precipitation index and Gaula water discharge on a weekly basis

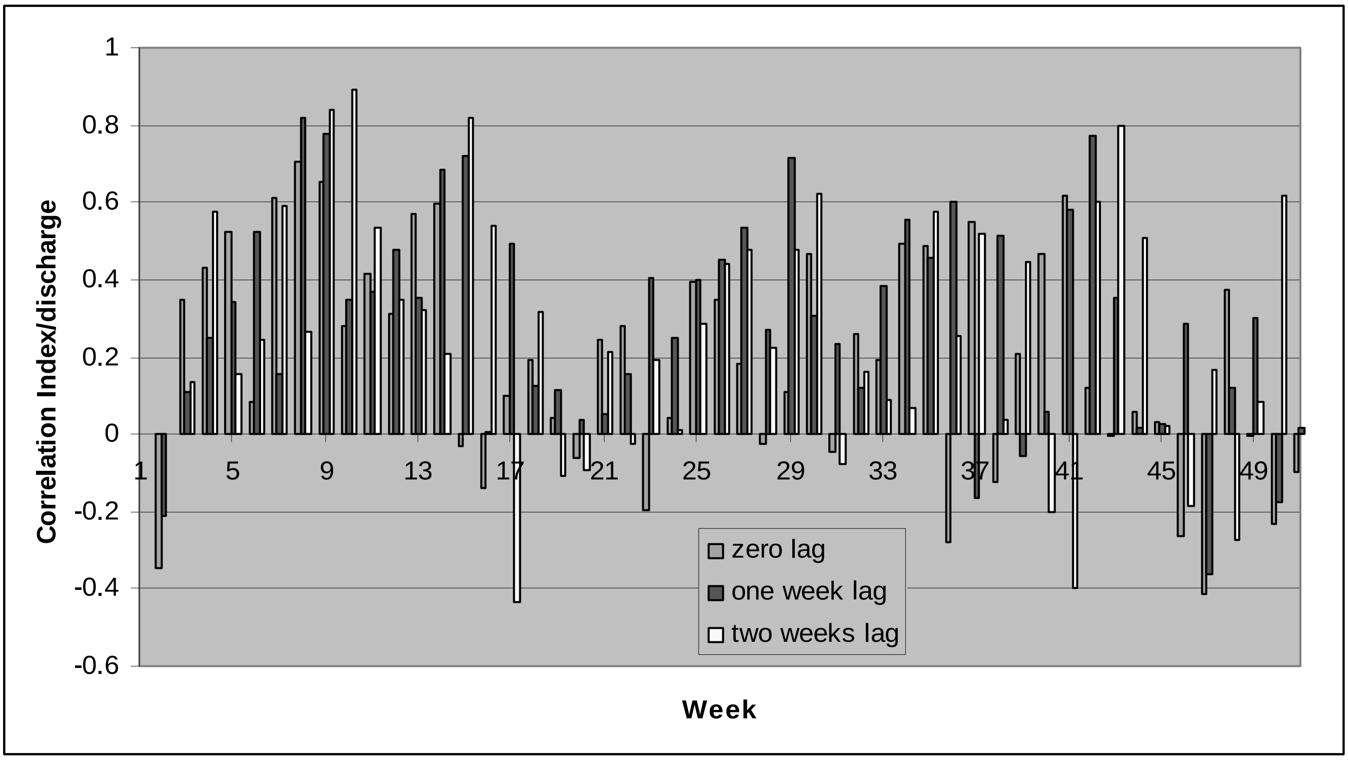

Figure E.2 Correlation between the precipitation index and Glomma water discharge on a weekly basis

In chapter 7, we also discuss correlation between discharge and the precipitation index during the late summer and fall weeks. While the correlation is found to be significant and negative, the correlation pattern shows strong deviations from the average values. The values are significant at the 5% level.

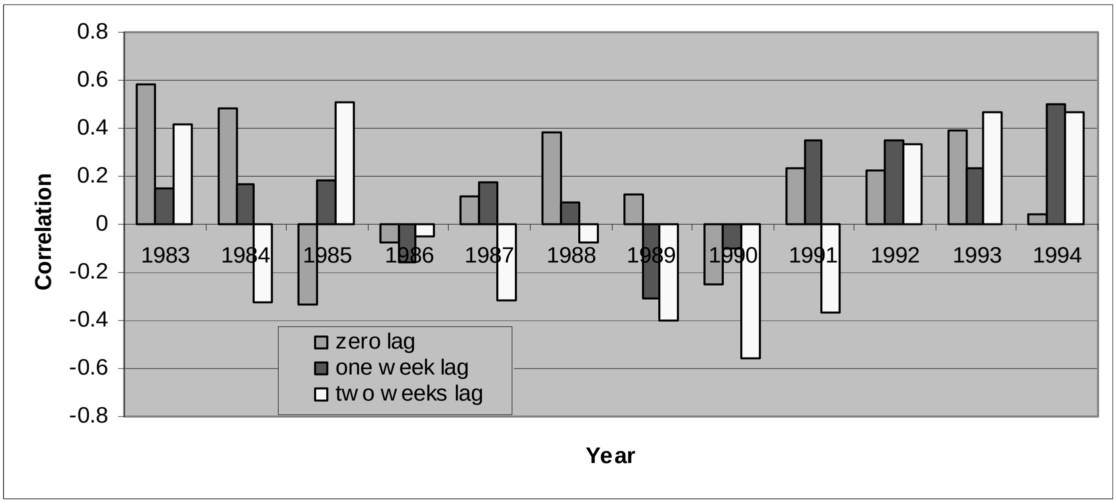

Figure E.3 Average weekly correlation between the precipitation index and Gaula discharge for weeks 29-45 broken down by year

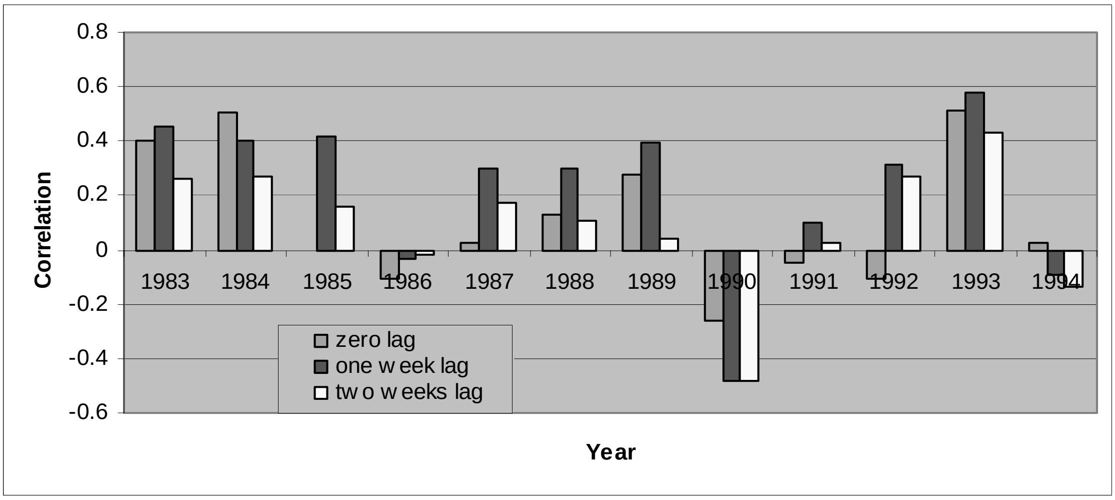

Figure E.4 Average weekly correlation between the precipitation index and Glomma discharge for weeks 29-45 broken down by year

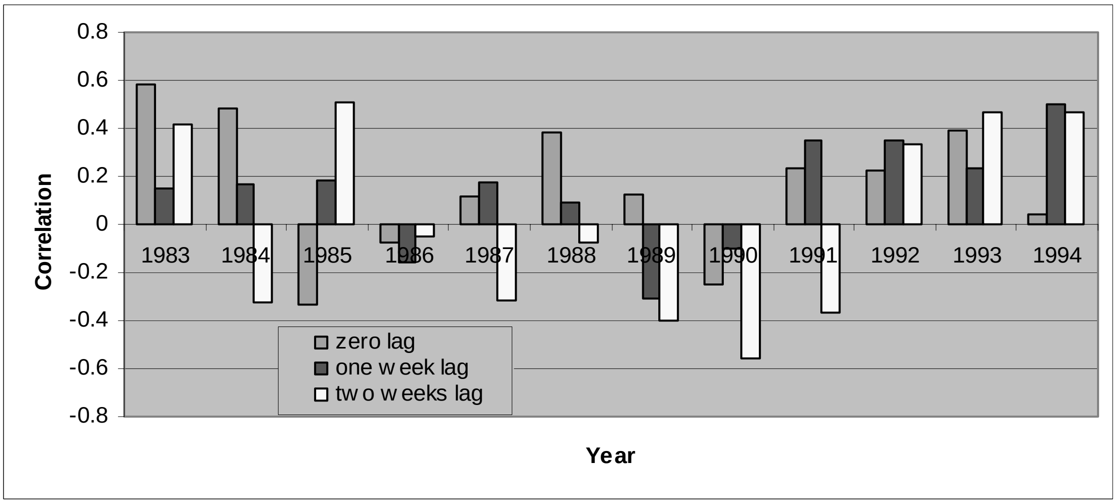

Figure E.5 Average weekly correlation between the precipitation index and Orkla discharge for weeks 29-45 broken down by year

Appendix F Cost of a run-of-river plant

The following cost equations are based on "Kostandsgrunnlag for mindre vannkraftanlegg" (Costs for small hydropower plants) publised by The Norwegian Water and Energy Resources Directorate. We got access to the data through the study "Mijøtilpasset Energiproduksjon ved små vannkraftverk i distrikts-Norge" (Environmental energy production through small hydropower plants in rural Norway) published by Gauldal municipaly. This publication gives detailed costs for most aspects of plant construction, but since our aim is mainly to determine the optimal size of a plant in a given river, we choose to include only the costs that 1) Constitute a significant part of the total and 2) are directly dependent on either discharge capacity or plant effect. All costs are in NOK unless otherwise stated. In the following, Q denotes discharge capacity and P effect.

Cost of dam

The cost of a dam for water is given by the following equations

$C_D = -0.1111Q^2 + 0.32222\,Q - 0.0111$ (MNOK) $Q \in [0.1\text{–}1\ m^3/sec]$

$C_D = -0.00001\,Q^3 - 0.0012\,Q^2 + 0.0772\,Q + 0.2409$ (MNOK) $Q \in [1\text{–}10\ m^3/sec]$

Cost of power station

The cost of a power station, will vary a lot from project to project, but the mean estimated cost is given by the equation

$C_{PS} = 0.44\,Q^{0.74}$

Cost of pipeline

This cost is dependent on what types of pipes are used and how the pipeline is constructed (dug down or supported above ground), the terrain in which it is build and so forth, so we ignore this cost, since it does not constitute a significant part of the total cost.

Cost of turbine

For the head we have chosen and the relevant discharge capacities, a turbine of the Kaplan type is appropriate. and the cost equations for a Kaplan turbine in the effect range 500-5000 kW are given by

$C_{turbine} = 27426\ Q^{-0.630\ NOK/kW}$ (Head 5 meters)

$C_{turbine} = 16106\ Q^{-0.648\ NOK/kW}$ (Head 10 meters)

$C_{turbine} = 9744\ Q^{-0.634\ NOK/kW}$ (Head 20 meters and above)

Cost of dam hatch / valve

A hatch is needed to allow for drenation of the dam. This cost constitutes a small ammount of the total cost, and is therefore neglected.

Cost of generator

The cost equation is for air-cooled generators with an effect range of 500-5000 kW. The cost is given by

$C_{generator} = 2134\ P^{0.8434}$

Cost of transformer

$C_{transformer} = 67.1\ P + 29258 \quad P \in [0.05 – 1.6\ kW]$

$C_{transformer} = 75.4\ P + 52206 \quad P \in [1.6 – 5\ kW]$

Cost of controlling equipment

$C_{control} = 1{,}500{,}000 \quad P \in [3.0 – 5.0\ kW]$

$C_{control} = 950{,}000 \quad P \in [2.0 – 3.0\ kW]$

$C_{control} = 550{,}000 \quad P \in [1.0 – 2.0\ kW]$

$C_{control} = 200{,}000 \quad P \in [0.5 – 1.0\ kW]$

Since this yields a stepwise continous cost function, we have used a continous function fitted to the above data. It is given as

$C_{control} = 1.573+0.0562*Q^{0.51}$ MNOK

Cost of coupling equipment

It is assumed that delivery ends at the power station, that is, the plant owner will not be responsible for high voltage power distribution. This cost is negliable

Appendix G Estimating the water discharge mean-reversion rate

An estimation for the mean-reversion parameter for the discharge data we use when valuing a run-of-river power plant is conducted below.

The deterministic part

Due to the shape of the average weekly discharge, it is evident that a higher-order transformation is necessary to obtain a good seasonal fit. We try 1st, 2nd, 3rd,4th and 5th order transformations to see how the goodness-of-fit increases.

We fit this model by fitting the function

$$f(t) = k + \beta t + a_1 \sin(\omega t) + a_2 \cos(\omega t) + ... + a_{2n-1} \sin(n\omega t) + a_{2n} \cos(n\omega t) \qquad \text{(G.1)}$$

Where n is the order of the transformation to the watershed data.

The results are

Order (n)	R-squared	P(sin nwt)	P(cos nwt)
1	45.9	0.000	0.000
2	58.2	0.000	0.000
3	63.5	0.000	0.000
4	64.1	0.000	0.000
5	64.3	0.000	0.000

Table G.1 Regression of order n fit to water discharge time series

We see that the increase in explainatory power of the models increases most steeply from the first to the second, and from the second to the third, while including terms of order higher than three gives little improvement to the model, despite all coefficients being significant at any significance level. We therefore choose the third-order transformation to capture the shape of the model. The most relevant residual plots from the MINITAB analysis are presented below:

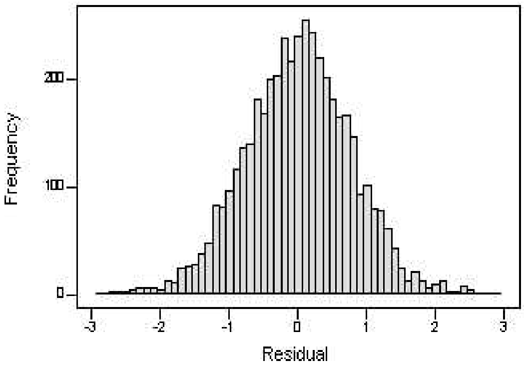

Figure G.1 Histogram of residuals

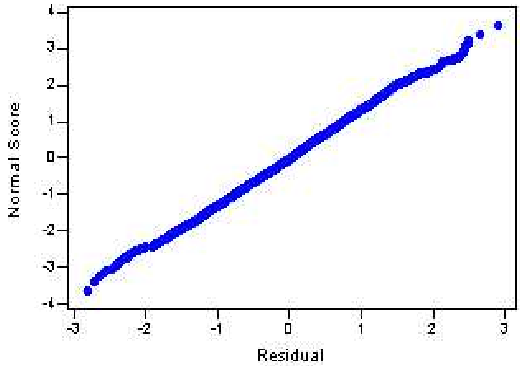

Figure G.2 Normal probability plot of residuals

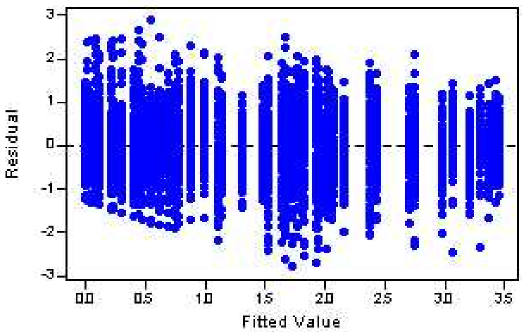

Figure G.3 Residuals versus fitted values

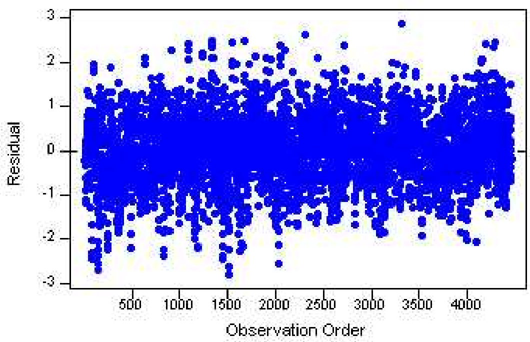

Figure G.4 Residuals versus order of data

From inspection, our assumption about normally distributed error terms seems to be justified. Regression yields the following parameters for the deterministic function

Predictor	Coeffecient	Standard Error	T-value	P-value
Constant	1.551	0.0115	134.66	0.000
trend (β)	0.000	0.0000	4.93	0.000
sin(ωt)	-0.457	0.01630	-28.05	0.000
cos(ωt)	-1.089	0.01629	-66.85	0.000
sin(2ωt)	-0.621	0.01629	-38.17	0.000
cos(2ωt)	0.344	0.01629	21.09	0.000
sin(3ωt)	0.364	0.01629	22.34	0.000
cos(3ωt)	-0.117	0.01630	-7.21	0.000

Table G.2 Regression results from deterministic water discharge function

The deterministic part of the watershed function can therefore be written as

$$w(t) = 1.55 + 1.18\sin(\omega t - 1.91) - 0.68\sin(2(\omega t - 3.33) + 0.41\sin(3(\omega t - 0.05) \quad (G.2)$$

As shown in the plot below, this is a very good fit to the average. The R2-value is 98.9%

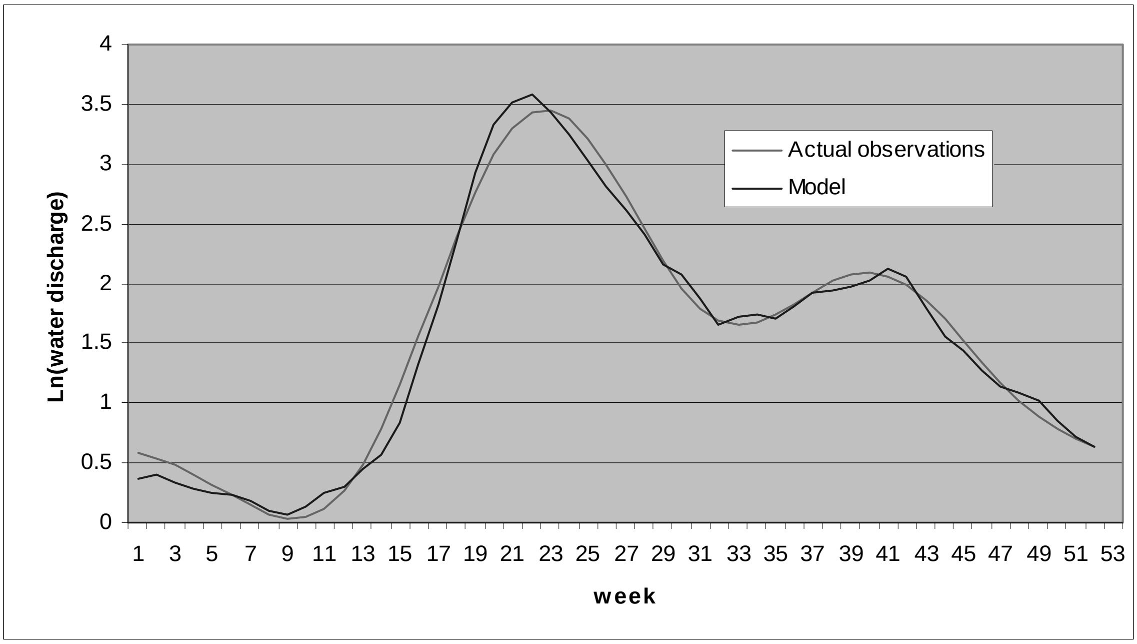

Figure G.5 Deterministic discharge curve versus historical average

Analysis of residuals

With a deterministic function and the residuals from the regression analysis, we are now ready to find an estimate of the mean-reversion parameter κ. This is taken as 1-ϕ where ϕ denotes the first-order autocorrelation coeffecient. The autocorrelation function is presented below:

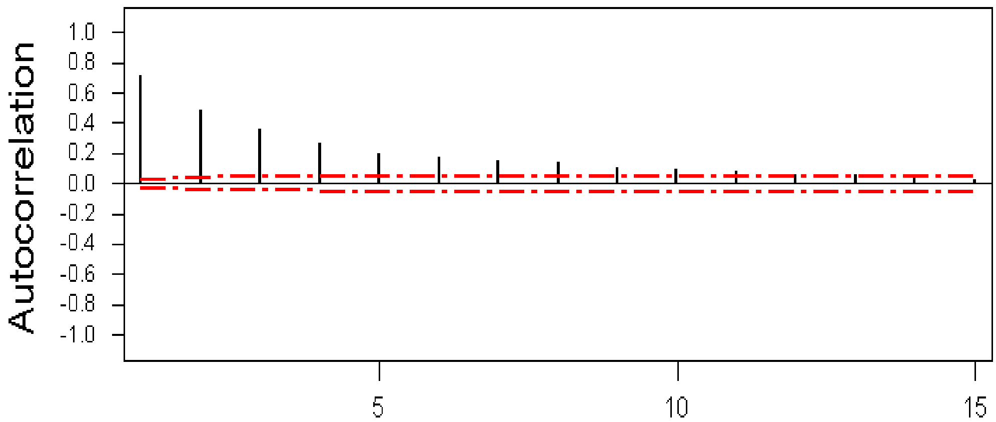

Figure G.6 Autocorrelation function for water discharge residuals

The first-order autocorrelation coeffecient is 0.71, hence our estimate for the mean-reversion parameter κ becomes κ_{est} = 1-0.71 = 0.29