Analysing Flexible Load Contracts in the Energy Market

Arne-Christian Lund and Fridthjof Ollmar *

November 20, 2002

Abstract

In this paper we analyse flexible load contracts (FLC), a type of "swing" option. This contract type has existed in energy markets for a long time and has proved to be challenging to value. The term swing refers to the flexibility in the quantity of energy that the holder of the contract can receive. We formulate the FLC as a stochastic optimisation problem. The price process, modelled as a time dependent Ornstein-Uhlenbeck process, is calibrated to the spot price on the Nordic electricity market. With this process the optimisation problem is solved numerically. The results of the algorithm are compared with the exercise policy for nine market participants. We find that our algorithm obtain the highest accumulated exercise revenue for a five year period.

*Institute of Finance and Management Science, Norwegian School of Economics and Business Administration, Helleveien 30, 5045 Bergen, Norway. Email arne-christian.lund@nhh.no and fridthjof.ollmar@nhh.no. We wish to thank Leif K. Sandal, Jostein Lillestøl and Gunnar Stensland for valuable comments. We would also like to acknowledge the helpful suggestions from the participants at the NHH's 2001 staff seminar.

Contents

1 Introduction

The first power plants built in the first part of the previous century were built to meet the nearby industry's demand for electricity. To reduce transmission losses or the cost of transporting the raw material used to make the electricity the industries were often located close to the power plants. Another characteristic of this early stage was that the same firms owned both the power plants and the industry. The value of electricity was consequently not exogenous calculated but endogenous valued as a part of the product costing. When it was possible to sell energy surplus, the need to formulate and value electricity contracts occurred. One of the first types of contracts to be traded was the contract that gave the owner the right to a certain amount of energy within a given period of time. To make it possible to deliver the electricity the seller restricted the maximum amount per hour (i.e. the effect) the buyer could withdraw. The buyer of the contract could then withdraw electricity, given the effect restriction, to cover his own electricity demand. This type of contract was the predecessor to the type of contracts we today call *flexible load contracts* (FLC).

Since the first flexible load contracts were traded, most electricity consumers and producers have interconnected themselves with a national or international power grid. In recent years many countries have also deregulated their electricity marked. These changes have influenced how we can utilise the flexible load contract. Before being connected to a power grid the owner of a flexible load contract had to withdraw the amount he consumed and any surplus energy was wasted. If there exists a liquid spot market the buyer of a flexible load contract can now withdraw energy from the seller of the contract and sell it in the spot market. To meet his own demand for electricity he can buy it directly from the spot market instead of exercising the contract. By incorporating the spot market the owner of a flexible load contract can fully utilise the flexibility of the contract. This effect has naturally increased the value of the contract. But it has also made the problem to value and decide when to exercise the contract more difficult.

Lets assume that we have a liquid spot market and that we have just bought an 8335 MWh flexible load contract[1]. The contract has a maximum effect of 5 MWh per hour and a delivery period from 1. May 1997 to 30. September 1997. The price we paid for this contract was 115 NOK/MWh or 958.525 NOK, and

[1] We will use this contract as an example throughout the paper

the delivery period of the contract consists of 3672 hours. Assuming we have a neutral attitude toward risk and we are profit maximising, then our target will be to exercise the contract during the 1667 [2] of the total 3672 hours with the highest spot price. Every day at 10 am we must inform the seller of the contract which hours the following day we want to exercise our right to buy for 115 NOK/MWh. The energy we buy will then be sold in the spot market, and our profit/loss will be the difference between 115 NOK/MWh and the price we manage to sell the energy in the spot market for. The flexibility of the contract is the ability to change our exercise policy during the delivery period. After buying the contract we may ask ourselves the following questions: How high should the spot price be before we start exercising the contract? What is its theoretical value? Which factors influence the value of the contracts and how do they influence the contract? All these questions and more will be answered in this paper.

The remainder of this paper is organised as follows. In the next section we formulate the FLC as a mathematical optimisation problem, and in section three we analyze the spot price and decide upon a spot price model. Then in section four we analyse how we can solve the optimisation problem numerically. In section five we describe our data-set and estimate the price process. The results and concluding remarks are given in section six and seven.

[2] 8335MWh/5MW=1667h

2 Mathematical formulation of the FLC

We will in this section show how we can find the optimal exercise policy and corresponding contract value by formulating the FLC as an continuous stochastic optimal control problem.

2.1 Background

Optimisation methods for hydroelectric power has been studied since the early 1960s. The first attempt to value a flexible load contract was by modelling the contract as a hydroelectric power plant with no inflow. In [4] Stage and Larsson developed one of the first optimisation methods for hydroelectric power plants. Their method was called incremented cost of waterpower and was based on finding the hydroelectric production that minimise the cost of the thermal power in a system where hydroelectric power is predominant. To implement this type of model one usually has to represent all hydroelectric power production as one representative hydroelectric power plant. If the individual hydroelectric power plants are significantly different from each other, representing them as one unit is both difficult and an inferior representation. Since there was no spot or forward market when the model was developed, they did not incorporate any information from these markets into the model. Instead they regarded the price as an endogenous function of the marginal production costs. This is a good approximation when there is no spot or forward market. If there exist a spot or forward market it is common to regard the price as exogenous. Despite several weaknesses, Larsson and Stages model is still used today.

Recent literature ([5], [2]) on valuing flexible load contracts is more based on contingent claims and derivative theory. If there exist a forward market with the same resolution[3] as the flexible load contract Øksendal shows in his PhD thesis [2] how to value it. He shows that it is possible to hedge the claim with a portfolio of forward contracts. The value of the FLC is equal to the value of the hedging portfolio. This method of valuing a flexible load contract only works when we have a forward market with equal or higher resolution than the FLC contract. If we try to use the method in practice we will discover that this assumption is not fulfilled. This can give an erroneous valuation of the flexible

[3] With "same resolution" we mean that if the FLC is based on an hourly resolution then the forward market must have one forward contract for each hour in the delivery period of the FLC.

load contract.

In this paper we will regard the FLC as a contingent claim on the spot price. We assume there is no forward market, or that the owner of the contract is unable to participate in the forward market. Since we in this setup cannot hedge the contract we need to specify the owner's risk attitude. We have decided to value the flexible load contract under the assumption of risk neutrality, and postpone further risk considerations to future work.

2.2 FLC as an optimisation problem

In this section we show how the flexible load contract can be formulated mathematically as a continuous stochastic optimal control problem. In the real world this optimisation problem is typically a combined discrete-continuous problem. It seems natural to think of the spot price as a continuous process. The control is however chosen on an hourly basis. Still one hour is a small time interval compared to the total contract length. A continuous model formulation is therefore natural. When we later implement a numerical scheme, one hour is used as the basic discrete time interval.

We study a control problem related to the optimal delivery of electrical power. We assume that a contract for a specified amount of energy over a period $[0, T]$ is given. The price of the electricity at a certain time $t \in \langle 0, T \rangle$ is given by a specified price process P_t. We assume that the 'producer' is a small participant in the market, so the price does not depend on the amount of delivered power. Further we assume that the contract puts restrictions on the delivery; At each instant the rate of delivered energy must be in a specified interval. Let Q_t denote the amount delivered up to time t. Our goal is to find the optimal control choice at each moment t, and for all levels of Q and P. This is a feedback form of the control. With this optimal policy in hand, the controller can choose the best delivery, given the current levels of the state variables. Further, the actual value of the contract is important when such contracts are bought or sold. We now show how this problem may be formulated as a stochastic optimal control problem with a terminal condition.

Suppose that we have agreed to deliver M units of a product (e.g. power) during the period $[0, T]$. The delivery rate is called u_t. Therefore

$$dQ_t = u_t dt$$

with $Q_0 = 0$. Obviously Q must satisfy $Q_T = \int_0^T u_t dt = M$. This is an *end* *constraint* on the variable Q_t. We assume that the control u_t must be in an interval $[u_0, u_1]$ for all t. Further, the contract specifies that the holder of the contract is paid a spot price P_t for the delivered amount of product. We assume that P follows a process

$$dP_t = \mu dt + \sigma dW_t$$

where W_t is a Brownian motion, μ and σ may be functions of t and P. At time t the price $P_t \stackrel{\Delta}{=} p$, and the amount $Q_t \stackrel{\Delta}{=} q$ are known by observation. The objective for the producer is now to maximise the net present value. Let the function Π represent the instantaneous profit of the delivery, and δ the discount factor. We want to find the *value function*

$$V(t, q, p) = \max_{u \in \mathcal{U}} E \int_t^T e^{-\delta s} \Pi(s, u_s, P_s) ds. \tag{2.1}$$

when $t < T$ and the corresponding control under the condition that $Q(T) = M$. This side condition calls for a control space $\mathcal{U}$ which is explicitly dependent of t and Q. In general such problems are hard to solve. In this case we may reformulate the problem to get a state independent control space. In subsection 2.4 we give a more precise formulation of the problem but first we need to study the structure of the problem more thoroughly.

2.3 Further observations

The function $V(t, q, p)$ is the value of the remaining period, given that the time is t, the delivered amount so far is q, and the current spot price is p. As seen above, there is an intimate relationship between the control and the level of the Q variable. When the problem is solved numerically, we take advantage of this. We would expect that the value V must be found for all $p > 0$, all $Q \in [0, M]$ and for all $t \in [0, T]$. Actually this is not necessary. Let us take a closer look on the condition $Q(T) = M$. The restrictions on u limits the Q-space that must be considered. See figure 1. For this problem to be well posed we must assume that $Tu_0 < M < Tu_1$. The problem is trivial if one of the two extremes is binding. The upper boundaries of the parallelogram are traced out by the policy

$$\begin{aligned} u &= u_1 \quad \text{for} \quad t \in [0, T_1] \\ u &= u_0 \quad \text{for} \quad t \in [T_1, T] \end{aligned}$$

where

$$T_1 = \frac{M - u_0 T}{u_1 - u_0}.$$

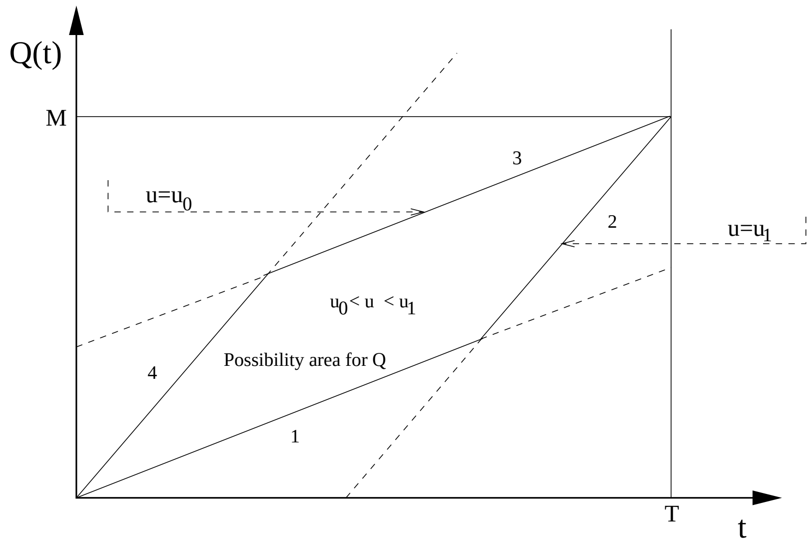

Figure 1: The possible values of $Q(t)$, given the restrictions on u_t.

The lower boundaries are on the other hand given by

$$\begin{aligned} u = u_0 &\quad \text{for} \quad t \in [0, T_2] \\ u = u_1 &\quad \text{for} \quad t \in [T_2, T] \end{aligned}$$

where

$$T_2 = \frac{u_1 T - M}{u_1 - u_0}.$$

Depending on the parameters of the problem we may have $T_1 < T_2$, $T_1 = T_2$ or $T_2 < T_1$.

To simplify the analysis and the numerical scheme we focus on a problem with control restrictions of the form $[0, u_1]$. This is no limitation since a contract with the limitation $[u_0, u_1]$ may be modelled as a flexible load contract with $[0, (u_1 - u_0)]$ combined with a contract with constant delivery u_0 in the same period.

2.4 Precise formulation

We can now formulate the optimisation problem precise without a state dependent $\mathcal{U}$. By defining the stopping times

$$\begin{aligned} \tau_1 &= \inf\{t; Q_t = M\} \\ \tau_2 &= \inf\{t; Q_t = u_1 \cdot (t - T_2)\} \\ \tau &= \min(\tau_1, \tau_2), \end{aligned}$$

the value function can be expressed as

$$
\begin{aligned}
V(t,q,p) &= \max_{u\in\mathcal{U}} E_t\Big\{ \int_t^{\tau} e^{-\delta s}\Pi(s,u_s,P_s)ds \\
&\quad + I_{(x=M)}(Q_\tau)F(\tau,P_\tau) + [1-I_{(x=M)}(Q_\tau)]G(\tau,P_\tau)\Big\}.
\end{aligned} \tag{2.2}
$$

Here the functions F and G is defined as

$$
\begin{aligned}
F(t,p) &= E\left[\int_t^T e^{-\delta s}\Pi(s,u_0,P_s)ds \,\middle|\, P_t = p\right] \\
G(t,p) &= E\left[\int_t^T e^{-\delta s}\Pi(s,u_1,P_s)ds \,\middle|\, P_t = p\right]
\end{aligned}
$$

Now $\mathcal{U}$ is the space of functions taking values in $[0,u_1]$. It is important to keep in mind that this is not an optimal stopping problem.

2.5 The Hamilton Jacobi Bellman equation

First of all we assume that the instantaneous profit is given by

$$\Pi(u,P) = \alpha u P,$$

and let $\alpha = 1$ for simplicity of notation. This turns the control problem into a problem which is completely linear in the control u. We therefore expect optimal controls of the so called 'bang-bang' type.

We want to find the value function $V(t,q,p)$. Define the space (see figure 2). $\Omega(t) \subset \mathbb{R}^2$ by

$$\Omega(t) = \{(q,t) \in \{M > q > 0\} \cap \{u_1 t \geq q > u_1 \cdot (t - T_2)\}\}.$$

The function $V : \Omega(t) \times \mathbb{R} \to \mathbb{R}$ can be found as the (viscosity) solution of the partial differential equation

$$V_t + \mu(t,p)V_p + \frac{1}{2}\sigma^2(t,p)V_{pp} + \max_{u\in\mathcal{U}}\{uV_q + e^{-\delta t}up\} = 0. \tag{2.3}$$

Here subscripts on V denotes the partial derivatives with respect to the subscript. This equation is called the Hamilton Jacobi Bellman (HJB) equation. The equation cannot be uniquely solved without proper boundary conditions. We know that the value is zero at time T, i.e.

$$V(T,q,p) \equiv 0 \quad \forall \quad q,p.$$

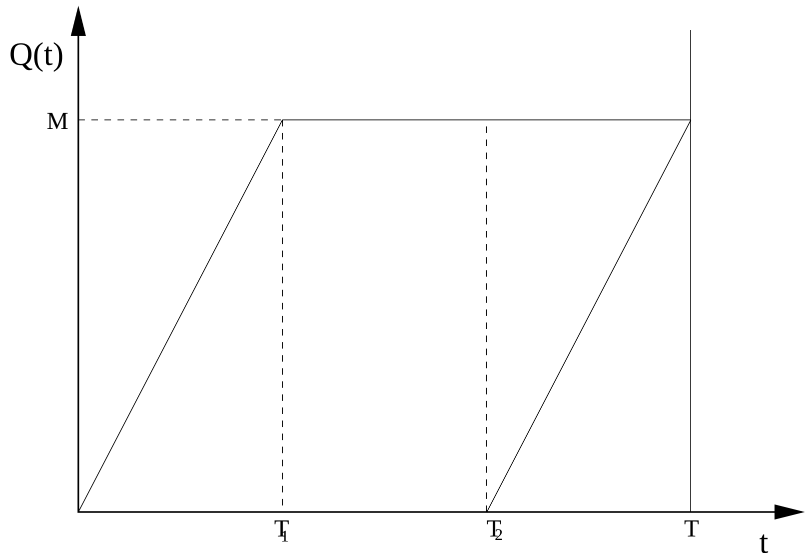

Figure 2: The (t, Q) projection of the parallelepiped, defining the space $\Omega(t)$.

Further, $V(t,q,p) = F(t,p)$ when $q = M$ and $V(t,q,p) = G(t,p)$ when $q = u_1 \cdot (t - T_2)$. From the definition of F and G we see that they can be found as solutions of the following partial differential equations[4]

$$W_t + \mu(t,p)W_p + \frac{1}{2}\sigma^2(t,p)W_{pp} = 0$$

$$W_t + \mu(t,p)W_p + \frac{1}{2}\sigma^2(t,p)W_{pp} + u_1 e^{-\delta t}p = 0 \tag{2.4}$$

both with end condition $W(T,p) = 0$. We see that this gives $F(t,p) \equiv 0$.

We now focus on the maximum operator in equation (2.3). Observe that

$$\begin{aligned} e^{-\delta t}p > -V_q \quad &\Rightarrow \quad u = \quad u_1 \\ e^{-\delta t}p = -V_q \quad &\Rightarrow \quad u = \quad ? \\ e^{-\delta t}p < -V_q \quad &\Rightarrow \quad u = \quad u_0. \end{aligned}$$

It can be shown that the optimal control only takes the extreme values, thus a bang-bang control. This is a consequence of the risk neutral formulation.

The flexible load contract is now formulated as a stochastic control problem. Observe that the equations in this section suggests that the value function may be found by backward induction, starting at time T. To solve the problem we need to specify a reasonable spot price process. We focus on this task in the next section.

[4]Alternatively, for a price process with simple structure, the functions may be calculated directly from the definitions. We chose however to keep the presentation general with respect to the process choice.

3 Modelling the spot price

We will in this chapter analyse the spot price to find a suitably stochastic differential equation to model it. After deciding upon a stochastic process we will show how we can calibrate the process parameters to the data.

3.1 Examining the spot price

The Nordic spot market for electricity is a market for physical delivery of electricity. Each day at noon, spot prices and volumes for each hour the following day are determined in an auction. The spot price is the clearing price that makes the demand for a given hour match the supply. Real aggregated supply and demand curves for hour 12 on 10. July 2000 are shown in figure 3. To understand the dynamics of the spot price it helps to understand the dynamics of the aggregated supply- and demand curve. Since a high degree of all energy used for heating in the Nordic countries is electricity, the demand for electricity is closely linked to temperature. The demand for electricity is also influenced by general work activity. Due to limited choice in alternative energy forms and lack of end users that actually observe real time price movements, the demand for electricity is highly inelastic (i.e. independent of market clearing price). The inelasticity of the demand curve can be seen from the steepness of the demand curve in figure 3. From figure 4 we see that the demand follows daily, weekly and yearly cycles. We also observe a small growth in electricity demand of approximately 1% to 1.5% per year. Induced by extreme weather conditions one can on several occasions observe temporary spikes in electricity demand. These spikes are not sustainable and the demand reverts back to normal levels within a short time.

In contrast to the nearly price independent electricity demand, the supply characteristics of the electricity producers are price responsive. The supply characteristic is mainly a function of *generation technology*, *fuel costs*, *availability of generation* and the possibility of *import/export*. The supply depends, in the long run, on the production cost for electricity. In the short run the supply is influenced by production outages and constraints in the power grid. Production costs for thermal based power depends mainly on the degree of utilisation and fuel costs. For hydroelectric power the production cost depends more on the reservoir filling, inflow and accumulated snow. The sum of deviation of reservoir filling and accumulated snow from the normal level is called the hy-

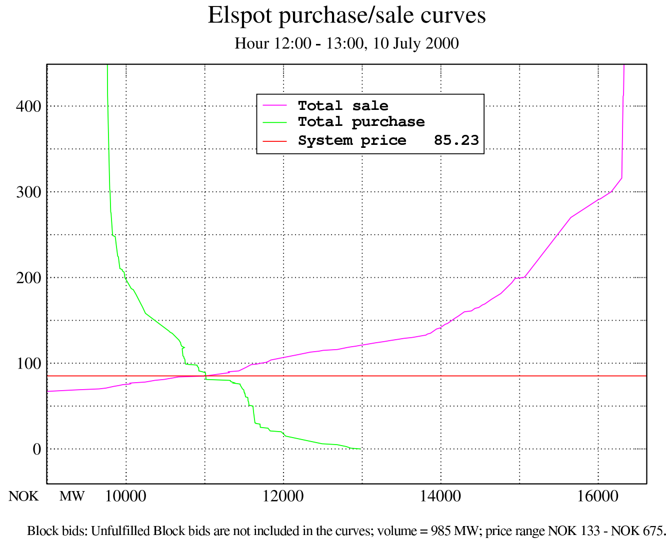

Figure 3: Supply and demand curves for hour 12 on 10. July 2002.

drological balance. Estimated hydrological balance together with spot price for the period 1996-2001 are shown in figure 5. From the figure we see a clear mean reversion in the hydrological balance. If we compare the hydrological balance with the spot price for the same period we see a strong negative correlation (Since we have inverted the scale in the figure it appears to be positive correlation). The empirical correlation coefficient is -0.72. The strong negative correlation is due to the fact that the hydroelectric power has a high alternative cost when the hydrological balance is below normal and a low alternative cost when the balance is above normal. The share of hydroelectric power in the Nordic electricity market is approximately 60%[5]. It is therefore no surprise that the supply curve is strongly influenced by the hydrological balance. Since the hydrological balance is so important it is crucial that the price process we choose is able to capture its effect.

[5]The total consumption of electricity for the Nordic countries were in 2000 384 TWh, and 234 TWh of this was hydroelectric power.

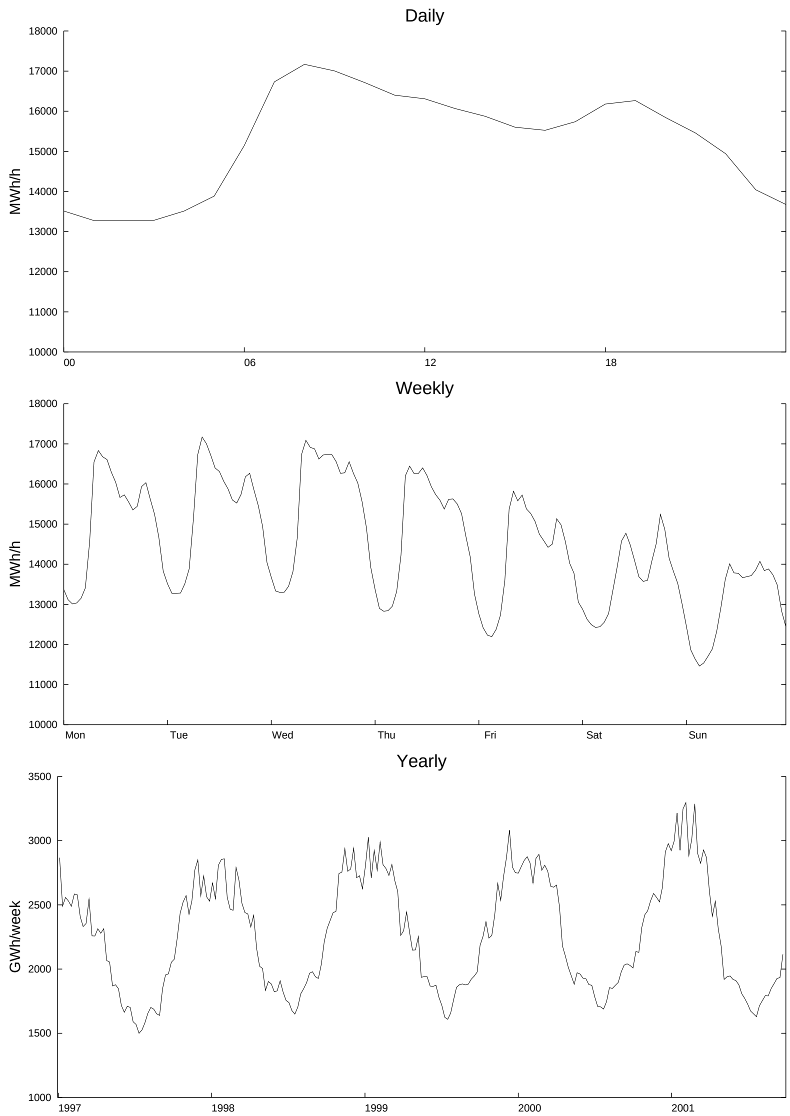

Figure 4: In the first figure we have total consumption of electricity in Norway during Tuesday 20. October 1998. The second figure is the total consumption during one week (19. October 1998 - 25. October 1998) and the third figure is the total consumption during a period of 4.5 years. As we can see the consumption of electricity follows daily, weekly and yearly cycles. Since the demand curve is highly inelastic we expect to find the same cycles in the electricity price.

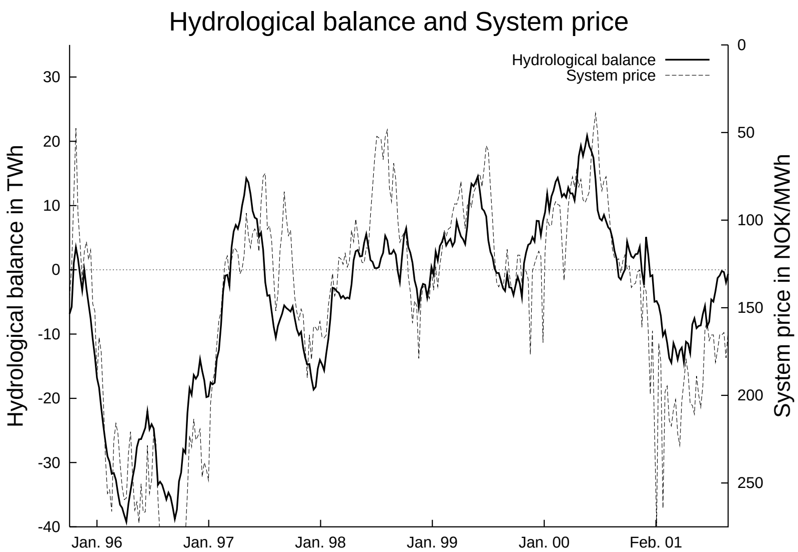

Figure 5: System price together with the hydrological balance. The hydrological balance is defined as the energy in snow and water minus their normal values. We have inverted the scale of the system price to illuminate the negative correlation.

3.2 Selecting a model

This section addresses the challenge of selecting a suitably stochastic process to model electricity prices in the Nordic electricity market. For reasons mentioned in the previous subsection the stochastic characteristics of electricity production and consumption are reflected directly in electricity prices. In addition to the lack of storability the cyclical patterns of electricity demand makes modelling the electricity price a challenge.

By analysing the Nordic electricity market we find following important factors influencing the spot price process:

- Cyclical patterns in demand over the course of the day, week and year.
- Price spikes or fast mean reversion due to unusual load conditions.
- A slow mean reversion in price caused by mean reverting hydrological balance.
- More long term factors such as fuel prices, currency exchange rates, emission costs and climate changes.

We are analysing flexible load contracts with a settlement period of approximately six to twelve months. We therefore focus on the spot dynamics within this time horizon, and long term factors such as inflation, fuel prices and climate changes can be ignored.

In the book “Energy modelling and the management of uncertainty” [1], B. Johnson and G. Barz analysed how the following four stochastic differential equations managed to model the spot price for different electricity markets:

Brownian motion:	$dP_t = \mu_t dt + \sigma dW_t$
Mean reversion, OU:	$dP_t = \kappa(\alpha_t - P_t)dt + \sigma dW_t$
Geometric Brownian motion:	$dP_t = \mu_t P_t dt + \sigma P_t dW_t$
Geometric mean reversion:	$dP_t = \kappa(\alpha_t + \frac{\sigma^2}{2} - \ln P_t)P_t dt + \sigma P_t dW_t$

where P_t is the spot price of electricity, μ_t is the drift term, σ is the diffusion term, W_t is a Brownian motion, κ is the speed of mean reversion and α_t is a sort of long run mean. They tried the above models with and without jump terms. The jumps where modelled with a Poisson arrival time, Bernoulli (positive or negative) jump direction, and exponential jump magnitude. The eight models where tested on four different electricity markets. They found the best model regarding sum of log-likelihood values for the Nordic electricity market to be the mean reversion with jumps followed by the pure mean reversion model. Since we chose not to incorporate jumps into our model, we use the mean reverting Ornstein-Uhlenbeck process to model the spot price.

To model the seasonal changes in the demand curve, we need a time dependent mean. In addition we want some kind of mean reversion to capture the effect of the hydrological balance. Since this reversion is slow compared to the mean reversion generated by price spikes, we need to separate them. If we do not separate them we get a mixture of fast and slow reversion. This will result in a volatility that is so high that the daily and weekly price patterns will vanish and a volatility that is too small to model the large deviation from the long run mean due to the hydrological balance. We specify the price process as $P_t = X_t + D_t$, where X_t represents the low frequent changes and D_t represents the high frequent changes. It is now possible to model the slow hydrological mean reversion together with annual seasons in X_t, and high frequent changes such as fluctuation in price over the course of the day or a week in D_t. We define the high frequent changes, D_t, as changes within one week and X_t as all other changes. Further we specify changes in X_t as

$$dX_t \quad = \quad a_t(b_t + \frac{b'_t}{a_t} - X_t)dt + \sigma_t dW_t, \qquad X_s = x_s,\ s < t \tag{3.1}$$

where a_t is the speed of mean reversion due to hydrological balance, b_t is the normal seasonal price, b'_t is the derivative and σ_t is the price volatility. We specify the normal seasonal price, b_t, and D_t as a sum of trigonometric functions.

$$
\begin{aligned}
b_t &= b_0 + \sum_{j=1}^{k} R_j^X \cos(\omega_j^X t + \phi_j^X) \\
&= b_0 + \sum_{j=1}^{k} \{A_j^X \cos(\omega_j^X t) + B_j^X \sin(\omega_j^X t)\} \qquad (3.2) \\
D_t &= d_0 + \sum_{j=1}^{l} R_j^D \cos(\omega_j^D t + \phi_j^D) \\
&= d_0 + \sum_{j=1}^{l} \{A_j^D \cos(\omega_j^D t) + B_j^D \sin(\omega_j^D t)\}
\end{aligned}
$$

where $A_j = R_j \cos(\phi_j)$, $B_j = -R_j \sin(\phi_j)$, ω is the frequency, ϕ is the phase, R is the amplitude and b_0 is a constant level. The parameter d_0 will later be used to ensure that the process D_t starts at zero every week. Choosing appropriate frequencies, phases and amplitudes we can model the daily, weekly and yearly price patterns. By specifying b_t and D_t as we did in (3.3) we could alternatively simplify P_t by incorporating D_t in b_t as an extension to the sum of trigonometric functions. Since we are going to use different sampling intervals for the estimation of the parameters in X_t and D_t, we will keep D_t separated from b_t.

The explicit solution to the price process, P_t, is given by

$$P_t = (P_s - D_s - b_s)e^{-\int_s^t a_u du} + D_t + b_t + \int_s^t \sigma_u e^{-\int_u^t a_r dr} dW_u \qquad (3.3)$$

If we let $a_t = a$, $\sigma_t = \sigma$ we can write P_t as

$$P_t = (P_s - D_s - b_s)e^{-a(t-s)} + D_t + b_t + \sigma\left(\frac{1 - e^{-2a(t-s)}}{2a}\right)^{1/2}\varepsilon \qquad (3.4)$$

where ε is a standard normal distributed random variable. From the above equation we see that the Gaussian process, P_t, has an conditional mean equal to $(P_s - D_s - b_s)e^{-a(t-s)} + D_t + b_t$ and a conditional standard deviation equal to $\sigma(\frac{1-e^{-2a(t-s)}}{2a})^{1/2}$. Since the expected value of P_t, when $t \to \infty$, is equal to $b_t + D_t$, we can interpret $b_t + D_t$ as the long run mean function for the price process.

3.3 Parameter estimation method

In the previous section we chose a stochastic differential equation with solution given by equation (3.4) to model the spot price. In this section we will show

how to estimate all the parameters in this process. Since the distribution of P_t is known, we can make use of the maximum likelihood estimation method. Let the parameter vectors $\alpha = \{a, \sigma, \omega_1^X, \ldots, \omega_k^X, \omega_1^D, \ldots, \omega_l^D\}$ and $\beta = \{b_0, A_1^X, \ldots, A_k^X, B_1^X, \ldots, B_k^X, A_1^D, \ldots, A_l^D, B_1^D, \ldots, B_l^D\}$. The reason for collecting the parameters into two vectors is to shorten our notation and, as we see later, we can use different methods to obtain the estimates of the different parameter vectors.

Let $\mathbf{P} = [p_{t_1}, p_{t_2}, \ldots, p_{t_n}]$ be a vector of observations of P_t at $t = t_1, t_2, \ldots, t_n$. The maximum likelihood estimates $\tilde{\alpha}$ and $\tilde{\beta}$ are the solution to the following maximisation problem

$$(\tilde{\alpha}, \tilde{\beta}) = \arg\max_{\alpha,\beta} \Psi\left(\mathbf{P}, \alpha, \beta\right) \tag{3.5}$$

where $P_t \sim \mathbb{N}\Big(m(p_{t_i}|p_{t_{i-1}}; \alpha, \beta)\ ,\ s(\alpha)\Big)$ and

$$\begin{aligned}
\Psi(\mathbf{P}, \alpha, \beta) &= \sum_{i=1}^{n} \log f(p_{t_i}|p_{t_{i-1}}; \alpha, \beta), \\
f(p_{t_i}|p_{t_{i-1}}; \alpha, \beta) &= \frac{1}{\sqrt{2\pi} s(\alpha)} \exp\left\{ -\frac{\left(p_{t_i} - m(p_{t_i}|p_{t_{i-1}}; \alpha, \beta)\right)^2}{2 s(\alpha)^2} \right\} \\
m(p_{t_i}|p_{t_{i-1}}; \alpha, \beta) &= (p_{t_{i-1}} - D_{t_{i-1}} - b_{t_{i-1}}) e^{-a(t_i - t_{i-1})} + D_{t_i} + b_{t_i} \\
s(\alpha) &= \sigma \Big(\frac{1 - e^{-2a(t_i - t_{i-1})}}{2a} \Big)^{1/2}.
\end{aligned}$$

This maximisation problem has a parameter space of $3(k + l + 1)$ dimensions.

As we discovered in the first section of this chapter the spot price has three distinct seasons. The seasons have periods of one day, one week and one year. To get a realistic representation of the spot price we need at least two trigonometric functions to represent each season. With $k = 2$ and $l = 4$ the maximisation problem given by (3.5) has a 21 dimensions parameter space. Numerically solving a maximisation problem with such a high degree of freedom can be difficult. To simplify the problem we fix the frequencies ω_j^X and ω_j^D to

$$\begin{array}{llllll}
\omega_1^X = 2\pi/8760 & \omega_2^X = 2\pi/4380 & \text{(year)} \\
\omega_1^D = 2\pi/168 & \omega_2^D = 2\pi/84 & \text{(week)} \\
\omega_3^D = 2\pi/24 & \omega_4^D = 2\pi/12 & \text{(day)}
\end{array}$$

We have here assumed an hourly sampling resolution of $\mathbf{P}$. The frequencies in the left column makes the long run mean follow cycles with a period of one year, one week and one day. The frequencies in the right column are set to one half of the frequencies in the left column. The reason for this is to make the long run mean able to model non-symmetric seasons. With fixed frequencies we only need to estimate 15 parameters with maximum likelihood. To further reduce the number of parameters to be estimated by maximum likelihood we reformulate (3.4) as

$$\begin{aligned} Y_{t_i} &= b_0 Z_0(t_i) + \sum_{j=1}^{k} A_j^X Z_j^{AX}(t_i) + \sum_{j=1}^{k} B_j^X Z_j^{BX}(t_i) \\ &\quad + \sum_{j=1}^{l} A_j^D Z_j^{AD}(t_i) + \sum_{j=1}^{l} B_j^D Z_j^{BD}(t_i) + \varepsilon_{t_i} \end{aligned} \tag{3.6}$$

where

$$\begin{aligned} Y_{t_i} &= (p_{t_i} - p_{t_{i-1}} e^{-a(t_i - t_{i-1})})/s(\alpha) \\ Z_0(t_i) &= (1 - e^{-a(t_i - t_{i-1})})/s(\alpha) \\ Z_j^{AX}(t_i) &= \{\cos(\omega_j^X t_i) - e^{-a(t_i - t_{i-1})} \cos(\omega_j^X t_{i-1})\}/s(\alpha) \\ Z_j^{BX}(t_i) &= \{\sin(\omega_j^X t_i) - e^{-a(t_i - t_{i-1})} \sin(\omega_j^X t_{i-1})\}/s(\alpha) \\ Z_j^{AD}(t_i) &= \{\cos(\omega_j^D t_i) - e^{-a(t_i - t_{i-1})} \cos(\omega_j^D t_{i-1})\}/s(\alpha) \\ Z_j^{BD}(t_i) &= \{\sin(\omega_j^D t_i) - e^{-a(t_i - t_{i-1})} \sin(\omega_j^D t_{i-1})\}/s(\alpha) \end{aligned}$$

The solution of the stochastic differential equation is now linear in the β parameters, and we can use ordinary least square to obtain an estimate of β. By specifying the frequencies and reformulating (3.4) we were able to reduce the number of parameter to be estimated by maximum likelihood from 21 to 2. Solving equation (3.5) is equal to solving

$$(\tilde{\alpha}, \hat{\beta}) = \arg\max_{\alpha} \Psi\left(\mathbf{P}, \alpha, \hat{\beta}(\alpha)\right) \tag{3.7}$$

where $\hat{\beta} = \{\hat{b}_0, \hat{A}_1^X, \ldots, \hat{A}_k^X, \hat{B}_1^X, \ldots, \hat{B}_k^X, \hat{A}_1^D, \ldots, \hat{A}_l^D, \hat{B}_1^D, \ldots, \hat{B}_l^D\}$ is the ordinary least squares estimate. The procedure used to solve this maximisation problem is as follows: First we start with an initial guess α, and find the corresponding $\hat{\beta}(\alpha)$ by OLS. The OLS-estimates together with α are then used to calculate the value of the log likelihood function. This procedure is then repeated until we find the α that maximises the log likelihood function and thereby solving the problem. We will later in section 6 use (3.7) on historical price data to estimate the parameters of the spot price process.

4 Numerical solution

In this section we show how the problem formulated in section 2 can be solved numerically on a discrete state space.

4.1 Discretisation

To solve the problem on a computer we need to discretise the time and state space. In this market the natural smallest time scale is one hour, and this is chosen as the basic time discretisation. In combination with the limitations on the control, this also gives the discretisation of the Q space, see figure 6. The parameters of the traded contacts are typically specified such that T_1, T_2 and T are all integers. The price space is truncated and divided into N uniform

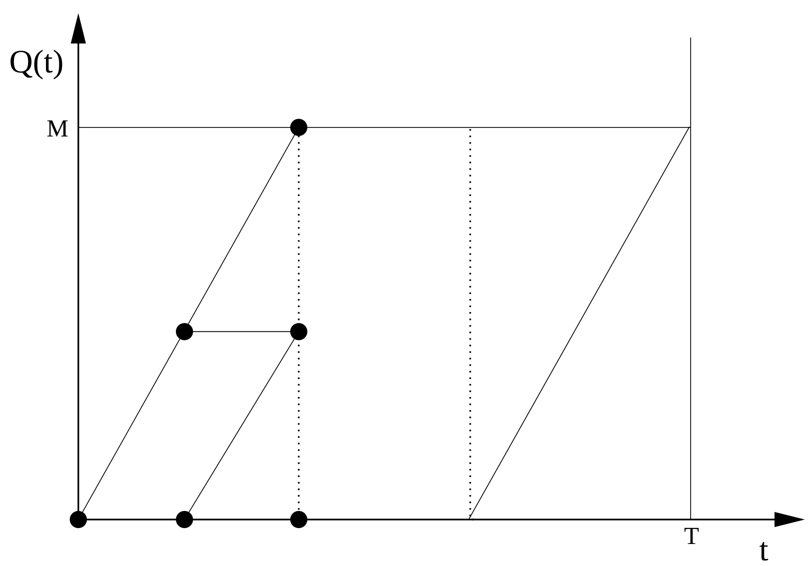

Figure 6: The natural nodes in the Q-space.

intervals. The value function is found in every node of the three dimensional parallelepiped in the (t, q, p)-space. We use backward induction, starting at time T.

The time horizon T is typically measured in whole hours. If $T_1 = \frac{M}{u_1}$ is an integer number of hours[6], it is natural to use one hour as the basic discrete time interval. In this case both T_1 and T_2 are reached after an integer number

[6]This is typically the case for the traded contracts. To increase numerical stability we may introduce e.g. four sub steps within each hour.

of periods. The initial time node is denoted 0, the last node is T, that is,

$$0 = t_0, \ldots, t_T = T$$

totally $T+1$ nodes. The control applied in the first hour is found in time-node 0.

The time discretisation combined with the control gives a natural discretisation of the Q space into $T_1 + 1$ nodes, see figure 7. Totally the (Q, t) space consists of $(T_1 + 1)(T - T_1 + 1)$ nodes[7].

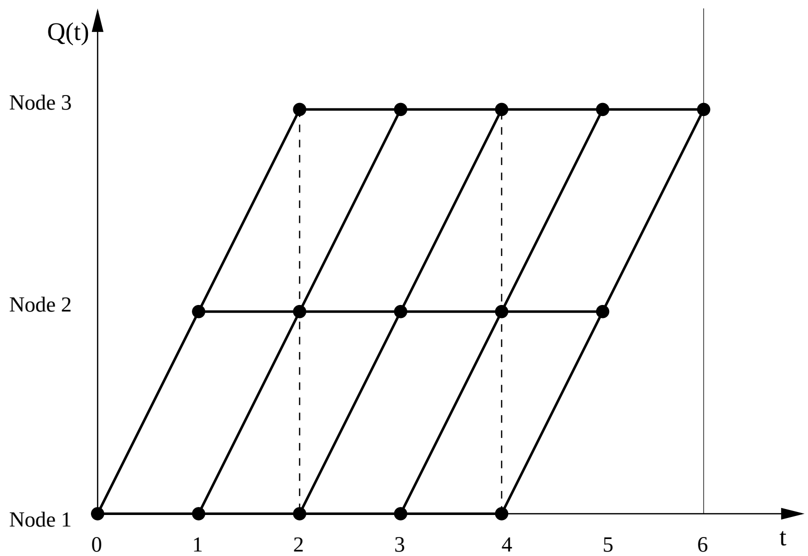

Figure 7: The discretisation of the (t, Q)-space.

The price process P_t studied in section 3 is unbounded. The infinite P-space must therefore be truncated before the optimisation problem can be solved numerically. Assume that the process can only take values in $[\underline{P}, \overline{P}]$. This interval is represented discretely as $\{\underline{P}, \underline{P} + \Delta P, \underline{P} + 2 \cdot \Delta P, \ldots, \underline{P} + N \cdot \Delta P, \overline{P}\}$ that is, $P_i = \underline{P} + i \cdot \Delta P$.

4.2 The numerical scheme

On the grid previously stated, we define

$$V_{i,j}^k = V(t_k, q_i, p_j).$$

This is a discrete approximation to the continuous value function[8] in equation (2.1), see page 7. After the choice of control the HJB equation (2.3)

[7] For the test case we have $T_1 = 1667$ and $T = 3672$. Therefore, with the P space divided into 100 intervals, we get 335 million nodes in the three dimensional parallelepiped.

[8] We denote the approximation and the true (continuous) function as V. When this is unclear, the continuous function is called $\tilde{V}$.

reduces to the partial differential equation

$$V_t + \mu(t,p)V_p + \frac{1}{2}\sigma^2(t,p)V_{pp} + \hat{u}V_q + e^{-\delta t}\hat{u}p = 0 \tag{4.1}$$

where $\hat{u}$ is either 0 or u_1.

Let us first focus on the interior of the P-space. We use finite difference to approximate the derivatives of V. At time t_k we use the known value function V^{k+1} to approximate V_q while the unknown V^k is used for V_p. Therefore the scheme is explicit in the q-variable and implicit in p. We use

$$\begin{aligned}
\frac{\partial}{\partial t}\left\{V_{i,j}^k\right\} &\approx \frac{V_{i,j}^{k+1} - V_{i,j}^k}{\Delta t} \\
\frac{\partial}{\partial q}\left\{V_{i,j}^k\right\} &\approx \frac{V_{i+1,j}^{k+1} - V_{i,j}^{k+1}}{\Delta q} \\
\frac{\partial}{\partial p}\left\{V_{i,j}^k\right\} &\approx \begin{cases} \frac{V_{i,j+1}^k - V_{i,j}^k}{\Delta p} & \text{when} \quad \mu \geq 0 \\ \frac{V_{i,j}^k - V_{i,j-1}^k}{\Delta p} & \text{when} \quad \mu < 0. \end{cases} \\
\frac{\partial^2}{\partial p^2}\left\{V_{i,j}^k\right\} &\approx \frac{V_{i,j+1}^k - 2V_{i,j}^k + V_{i,j-1}^k}{(\Delta p)^2}
\end{aligned}$$

This is a downwind-upwind discretisation of $\frac{\partial V}{\partial p}$. Observe that the approximation is done in the flow-direction of the underlying process. Define

$$\begin{aligned}
\mu^+ &= \max(\mu, 0) \\
\mu^- &= \max(-\mu, 0).
\end{aligned}$$

Observe that $\mu^+ + \mu^- = |\mu|$ and $\mu^+ - \mu^- = \mu$. Inserting the above approximations into (4.1) we get

$$\begin{aligned}
\frac{V_{i,j}^{k+1} - V_{i,j}^k}{\Delta t} &+ (\mu_j^k)^+ \frac{V_{i,j+1}^k - V_{i,j}^k}{\Delta p} - (\mu_j^k)^- \frac{V_{i,j}^k - V_{i,j-1}^k}{\Delta p} \\
&+ \frac{1}{2}(\sigma_j^k)^2 \frac{V_{i,j+1}^k - 2V_{i,j}^k + V_{i,j-1}^k}{(\Delta p)^2} + \hat{u}\frac{V_{i+1,j}^{k+1} - V_{i,j}^{k+1}}{\Delta q} + e^{-\delta t_k}\hat{u}p_j = 0
\end{aligned}$$

Using that $\Delta q = \Delta t u_1$, $\hat{u} \in \{0, u_1\}$ and collecting the terms, we get

$$\begin{aligned}
\frac{V_{i+1,j}^{k+1} - V_{i,j}^k}{\Delta t} &+ (\mu_j^k)^+ \frac{V_{i,j+1}^k - V_{i,j}^k}{\Delta p} - (\mu_j^k)^- \frac{V_{i,j}^k - V_{i,j-1}^k}{\Delta p} \\
&+ \frac{1}{2}(\sigma_j^k)^2 \frac{V_{i,j+1}^k - 2V_{i,j}^k + V_{i,j-1}^k}{(\Delta p)^2} + e^{-\delta t_k}u_1 p_j = 0
\end{aligned}$$

when $\hat{u} = u_1$. When $\hat{u} = 0$ the equation is

$$\begin{aligned}
\frac{V_{i,j}^{k+1} - V_{i,j}^k}{\Delta t} &+ (\mu_j^k)^+ \frac{V_{i,j+1}^k - V_{i,j}^k}{\Delta p} - (\mu_j^k)^- \frac{V_{i,j}^k - V_{i,j-1}^k}{\Delta p} \\
&+ \frac{1}{2}(\sigma_j^k)^2 \frac{V_{i,j+1}^k - 2V_{i,j}^k + V_{i,j-1}^k}{(\Delta p)^2} = 0.
\end{aligned}$$

Observe that this may be seen as discrete representations of (2.4) with the convention that we move up in the Q-grid over the time step when u_1 is chosen.

Since we use backward induction, V^{k+1} is completely known at time t_k. At the boundaries two and three (see figure 2 on page 10) the control must be u_1 and 0 respectively, that is, there is no choice here. At these boundaries the value function is equal to the functions G and F. At the boundaries one and four, and in the interior both control choices may be used.

The above scheme can be organised as

$$a_j^k V_{i,j-1}^k + b_j^k V_{i,j}^k + c_j^k V_{i,j+1}^k = e^{-\delta t_k} u_{i,j}^k p_j \Delta t + W_{i,j}^{k+1}(u_{i,j}^k) \tag{4.2}$$

where

$$\begin{aligned} a_j^k &= -\left((\mu_j^k)^- \frac{\Delta t}{\Delta p} + \frac{1}{2}(\sigma_j^k)^2 \frac{\Delta t}{\Delta p^2}\right) \\ b_j^k &= 1 + \left|\mu_j^k\right| \frac{\Delta t}{\Delta p} + (\sigma_j^k)^2 \frac{\Delta t}{\Delta p^2} \\ c_j^k &= -\left((\mu_j^k)^+ \frac{\Delta t}{\Delta p} + \frac{1}{2}(\sigma_j^k)^2 \frac{\Delta t}{\Delta p^2}\right) \\ W_{i,j}^{k+1}(\hat{u}) &= \begin{cases} V_{i,j}^{k+1} & \text{when} \quad \hat{u} = 0 \\ \\ V_{i+1,j}^{k+1} & \text{when} \quad \hat{u} = u_1. \end{cases} \end{aligned}$$

At the boundaries 2 and 3 of the (Q,t)-space (see figure 1 on page 8) $\hat{u}$ is known. At each time step t_k and for each Q-node q_i away from these boundaries we find the optimal control

$$u_{i,j}^k = \arg \max_{u \in \{0, u_1\}} (e^{-\delta t_k} u p_j \Delta t + W_{i,j}^{k+1}(u)),$$

and thereby also the righthand side of the linear system of equations defined by equation (4.2).

Before we focus on the discretisation on the boundaries of the P-space, we show how the above scheme may be linked to a Markov chain approximation of the underlying stochastic process.

With reference to the book by Kushner and Dupuis [3] we note that the scheme (4.2) may be written as

$$V_{i,j}^k = \sum_{l=\{j-1,j+1\}} p(i,l;k,k) V_{i,l}^k + p(i,j;k,k+1)\left[W_{i,j}^{k+1}(u_{i,j}^k) + e^{-\delta t_k} u_{i,j}^k p_j \Delta t\right]$$

with the following definition of the "probabilities";

$$
\begin{aligned}
p(i, j-1; k, k) &= -\frac{a_j^k}{b_j^k} \\
p(i, j+1; k, k) &= -\frac{c_j^k}{b_j^k} \\
p(i, j; k, k+1) &= \frac{1}{b_j^k}.
\end{aligned}
$$

Observe that $p(\cdot) \geq 0$ and $\sum p = 1$. With this representation we see that this scheme may be associated with a Markov chain approximation of the price process. The chain lives in the discrete (p, t) space, and time is treated as just another state variable. At each period there is only a certain probability that a time step is taken. See figure 8. This intuition proves useful when we study the boundaries of the P-space in the next subsection.

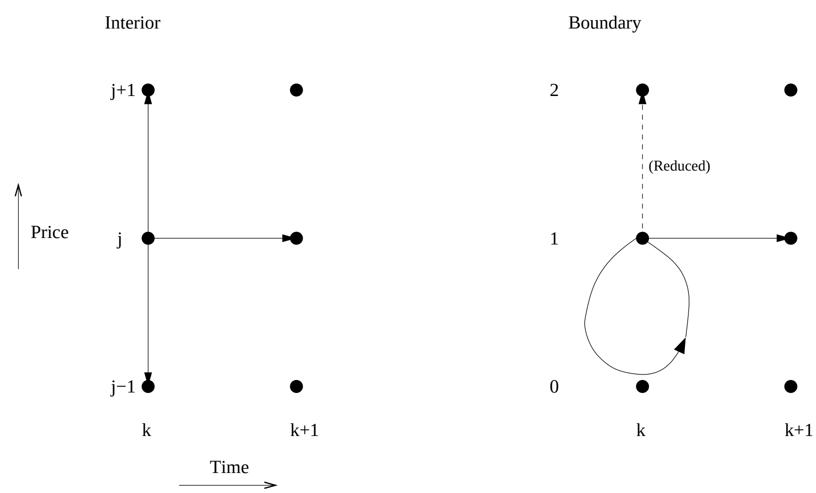

Figure 8: The Markov chain interpretation, with reflection on the boundary.

4.3 The boundaries of the price space

In this section we study the boundaries of the truncated P-space $[\underline{P}, \overline{P}]$. Two different types of boundary conditions are used. We first present the method called "Absorbtion". This type of boundary conditions typically arise when a process is absorbed in a boundary node, and a specified value is known in that node. In this case conditions are put directly on the value function. In the theory of partial differential equations such boundary conditions is called "Dirichlet" conditions.

The second method is called "Reflection", and must be used when we study a process that is reflected at a boundary. In this case it is important that the discrete Markov chain is reflected in the proper direction. This can be seen to correspond to conditions on the derivative of the value function, so called "Neumann" conditions.

We now show in detail how these boundary conditions affects the above scheme.

4.3.1 Absorbtion

We first focus on how absorbtion may be implemented. Suppose that the value $V_{i,j}^k$ is known (or approximated) for all k, i at the boundaries of the P-space, that is,

$$\begin{aligned} V_{i,0}^k &= \tilde{V}(t_k, q_i, \underline{P}) \\ V_{i,N+1}^k &= \tilde{V}(t_k, q_i, \overline{P}) \end{aligned}$$

for all i, k. Next to the lower boundary the equation (4.2) must be changed to

$$b_1^k V_{i,1}^k + c_1^k V_{i,2}^k = e^{-\delta t_k} u_{i,1}^k p_1 \Delta t + W_{i,1}^{k+1}(u_{i,1}^k) - a_1^k \tilde{V}(t_k, q_i, \underline{P})$$

and

$$b_N^k V_{i,N}^k + a_N^k V_{i,N-1}^k = e^{-\delta t_k} u_{i,N}^k p_N \Delta t + W_{i,N}^{k+1}(u_{i,N}^k) - c_N^k \tilde{V}(t_k, q_i, \overline{P})$$

at the upper. Remember that $p_0 = \underline{P}$ and $p_{N+1} = \overline{P}$.

In this subsection the value function were taken as given at the boundary. The problem is that it may be hard to say anything meaningful about this value in advance. The error done in this specification typically propagate towards the center of the grid. It is damped as it gets far away from the boundary, but still this may be a problem for the scheme, especially when the volatility (modelled by σ) is large. The solution is to truncate the price process at levels far away from the regions of interest. Further, we must keep an eye on the approximate solution near the boundaries, and adjust the specifications if it is clearly inconsistent with the real value function. Such methods are quite easy to implement, but is costly since the grid must be enlarged and the resulting value function inspected carefully.

4.3.2 Reflection

Reflection is an alternative to the method studied in the previous section. The idea is easier grasped when we think of our scheme as a (Markov chain) approximation of the movements of the price process. Instead of letting the process be absorbed at the boundaries as in the last section, we now assume that the process is reflected at the boundaries. This may be seen as a condition on the derivative of the value function, and as such, a weaker condition.

When the real process possess reflection, it is important that the reflection in the scheme is implemented in a consistent manner. The process we study has no natural reflection. We have therefore freedom to choose the approximation. What the most efficient reflection looks like is not obvious in advance, and we found a good approximation by experimentation.

At the boundaries the chain was reflected back into the grid, we here use the lower boundary as an illustration. When the chain goes from node 1 to node 0 it is immediately returned to node 1, i.e. the probability $p(1,1;t,t)$ is positive. Now the expected movement of the process is shifted upwards. To reduce this effect, we decrease $p(1,2;t,t)$ and increase $p(1,1;t,t)$ further. For our problem this procedure proved efficient.

We here present the chosen probabilities at node 1,

$$\begin{aligned} p(1,1;t,t) &= -\tfrac{2a_1}{b_1} \\ p(1,1;t,t+1) &= \tfrac{1}{b_1} \\ p(1,2;t,t) &= -\tfrac{c_1-a_1}{b_1}, \end{aligned}$$

where a, b and c (We have suppressed the time index) is defined on page 22. We see that they sum to unity. Further, if the drift is positive at node 1, they are all positive and less than one. Therefore we may interpret them as transition probabilities. The procedure is similar at node N.

4.4 Implementation of the scheme

Our problem is time dependent, with very explicit periodicity on a daily, weekly and yearly scale. Further, the mean reversion effect is small. This means that the drift of the process change sign during the day. We may therefore suspect that the reflection procedure at node 1 is a good approximation when the drift is positive, but poor when the drift is very negative. Opposite in node N. We have therefore implemented absorbtion when the drift is smaller than a chosen level (e.g. zero). The value associated with the absorbing node is approximated as the value at the previous time step. Then the unknowns $V_{i,1}^k, \dots V_{i,N}^k$ may be found as the solution of the system of linear equations defined in equation (4.3). Afterwards $V_{i,0}^k$ and $V_{i,N+1}^k$ can be approximated with interpolation of their neighboring values.

To illustrate the above discussion we present the scheme in a situation where reflection is used on the lower boundary, and absorbtion on the upper. We can find $V_{i,j}^k$ for all k, i, j by the following procedure

1. $V_{i,j}^T = 0$ is given from the end conditions.

2. When $V_{i,j}^{k+1}$ is given, find $\bar{V} = [V_{i,1}^k, \dots, V_{i,N}^k]$ as the solution of the following linear system of equations (For simplicity of notation, we suppress the sub- and superscripts of V, A and G.)

$$AV = G \tag{4.3}$$

where

$$A = \begin{bmatrix} \tilde{b}_1 & \tilde{c}_1 & 0 & 0 & \dots \\ a_2 & b_2 & c_2 & 0 & \dots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & a_{N-1} & b_{N-1} & c_{N-1} \\ 0 & \dots & 0 & a_N & b_N \end{bmatrix}$$

and

$$G = \begin{bmatrix} e^{-\delta t_k} u_{i,1}^k p_1 \Delta t & + & W_{i,1}^{k+1}(u_{i,1}^k) \\ e^{-\delta t_k} u_{i,2}^k p_2 \Delta t & + & W_{i,2}^{k+1}(u_{i,2}^k) \\ & \vdots & \\ e^{-\delta t_k} u_{i,N-1}^k p_{N-1} \Delta t & + & W_{i,N-1}^{k+1}(u_{i,N-1}^k) \\ e^{-\delta t_k} u_{i,N}^k p_N \Delta t & + & W_{i,N}^{k+1}(u_{i,N}^k) - c_N V_{i,N+1}^k \end{bmatrix}$$

Here

$$\begin{aligned} \tilde{b}_1 &= b_1 + 2a_1 \\ \tilde{c}_1 &= c_1 - a_1. \end{aligned}$$

This system of equations is tridiagonal and can be solved efficiently by Gauss elimination.

3. Iterate from step 2.

Observe that the coefficients a, b, c of the A matrix are independent of the control and the Q-level. Therefore, at a given instant t_k, the A matrix is the same for all Q-levels. The three-diagonal system of equations is solved once, with a loop calculating the solutions corresponding to the different righthand sides. This improves the efficiency of the algorithm considerably. Also observe that $b_i \geq 1.0$, and that the matrix A is strictly diagonal dominant. This secures the stability of the scheme.

4.5 The control matrix

The algorithm in the previous section calculates the value and the optimal control in each node of the grid. As previously pointed out, the grid may typically have more than 300 million nodes. Consequently it is inefficient to store all the information. We chose to store the value only at the first time step. This gives an estimate for the initial value of the contract. The value of the contract may be interesting at later time steps if the contract is re-traded, but we put this aside at the present.

The optimal control is however needed at each node of the grid. Still the structure of the problem gives a limited demand for storage. The point may be explained by the following argument.

Suppose the time is t_k and the current price is p_l. Focus on the amount delivered up to this point, i.e Q_t. If we choose to deliver u_1 for $Q_t = q_j$, then we chose u_1 for all q_m where $m < j$. Therefore we need only keep the critical $q_{\hat{n}}$ such that

$$u = \begin{cases} u_1 & \text{for} \quad n < \hat{n} \\ 0 & \text{for} \quad n \geq \hat{n}. \end{cases}$$

This critical level must be found for all t_k and all p_l, thus giving a $N \times T$ matrix. See figure 9. If the process parameters are fixed for the whole period, this matrix is generated only once. This can be quite time consuming, especially when the contract horizon is long.

The real observed price is now used to find the optimal control for each hour, and to calculate the realised value of the contract. In section 6 we use this algorithm to analyze several different contracts. The resulting control policy is compared to the strategy of competitors in the market.

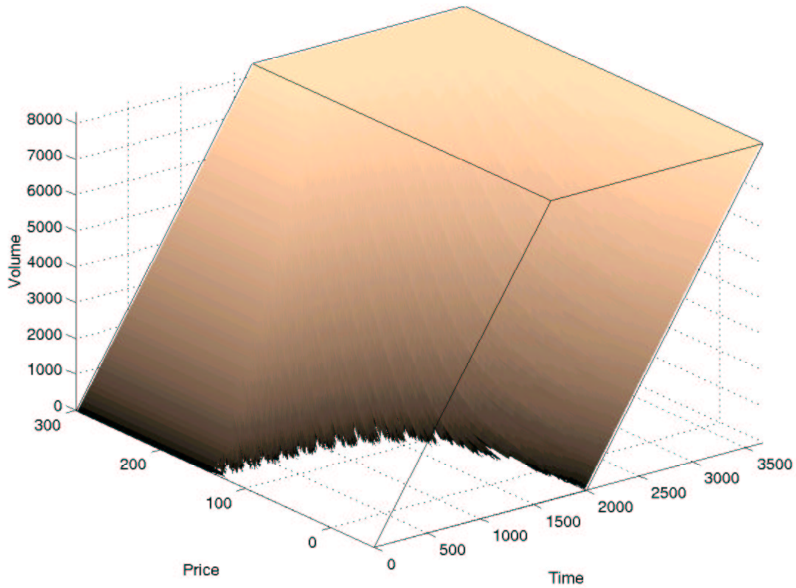

Figure 9: The control matrix.

4.6 Deterministic test

It is important to try to check the results generated by the numerical algorithm. For this problem we have no explicit solution to compare with. Still, if we let $\sigma \equiv 0$, we may test the algorithm by the following method.

Suppose we study a contract over the period[9] [0,168] and that we have to find the 100 hours with the highest price. If the price process is purely deterministic, the price is known for the whole period at time 0. It is therefore a simple task to find the hours to exercise the contract. The time 0 value we achieve (called explicit solution below) is compared with the value calculated by our algorithm. The price grid is $[-50, 300]$ with $\Delta p = 3.846$. In the time space we use $\Delta t = 1$. The results are presented in table 1. Relative error is the absolute error divided by the explicit solution. Observe that the numerical scheme is good in the middle of the grid but worse close to the boundaries, especially at the lower boundary. The error close to the boundaries is expected. We can however not explain the asymmetry in the error.

In this version of the paper we do not include the proof for convergence of the algorithm. This proof is rather technical and does not give any new intuition for a reader interested in applications.

[9]For simplicity, we study a contract with a short settlement period.

Table 1: Deterministic test of the algorithm

Price time 0	Explicit solution	Algorithm	Absolute error	Relative error
-23	-14666	-6988	7678	-0.52
0	7597	8821	1224	0.16
50	55840	54311	-1529	-0.027
100	104104	102461	-1643	-0.016
150	200789	198949	-1840	-0.0092
250	249205	245088	-4117	-0.017

4.7 Remarks

The above scheme has transition only to neighboring[10] nodes. This limits the possible movements of the process from hour to hour. The weakness of this implementation may be dealt with in different ways. One possibility is to use non-local finite difference approximations. Another is to introduce intermediate time steps, where the control is inherited from the large time step of one hour.

An easier way to more flexible movements of the process is to introduce intermediate time steps. At each small step the optimal control is found. The control for the present hour is the accumulated controls for the sub-steps. We have promising results using this method, but the full study of this extension is left for future work.

[10]From one time step to the next the Markov chain may move to other nodes. This is because the probability that a time step is actually taken is less than 1.0. If a fully explicit scheme was used, the chain had been limited to the neighboring nodes. This motivates, from a Markov chain perspective, why implicit schemes are more stable than explicit schemes.

5 Data and estimation

We have chosen a stochastic process for the spot price and developed a numerical algorithm to find the value and optimal policy for a flexible load contract. The next step is to implement our algorithm. To do this we need to estimate the price process from historical prices.

5.1 Price data

To estimate the parameters in the price process we use historical spot prices obtained from Nord Pool. The spot price is called *system price*, and is the price in NOK for one MWh of electricity for a given hour. Our data sample consists of 76 608 hourly prices from 4. January 1993 to 1. October 2001. See figure 10 for a graphical illustration of the data sample. There where no missing data but the prices were in a standard time format. Since cyclical patterns of electricity demand over the course of a day mostly depends on the time shown by the clock and not the time implied by the sun, we need to adjust for daylight saving time. To adjust for daylight saving time we inserted one fictitious price observation in the spring and removed one in the autumn. The observation we inserted in the spring was the average of the price value before and after. If we do not adjust for daylight saving time we will get a phase shift between the daily patterns on a winter day and the daily patterns on a summer day.

Another characteristic of our data sample is that it includes several price spikes due to unusual load conditions. Since we chose a price process without a jump term we are unable to model price spikes or fast mean reversion directly. We must therefore be careful not to let the spikes influence the parameter estimation too much. By closer inspection it seems that the price spikes mainly occurs in the morning or in the afternoon, with a duration of one to six hours. Fortunately the data sample used to estimate the parameters in the weekly process, X_t, does not include many spikes. The reason for this is that the data sample consist of the first hour on every Monday and at this time of the night the demand is low and price spikes rarely occurs. Since the intra-weekly process, D_t, is deterministic an occasional spike does not influence the estimation much. If we look at the descriptive statistics in table 2 the price spikes shows up as increased skewness and kurtosis. We can also see from figure 10 that the occurrences of price spikes has increased dramatically the last three years. The descriptive statistics also indicate that the spot price is

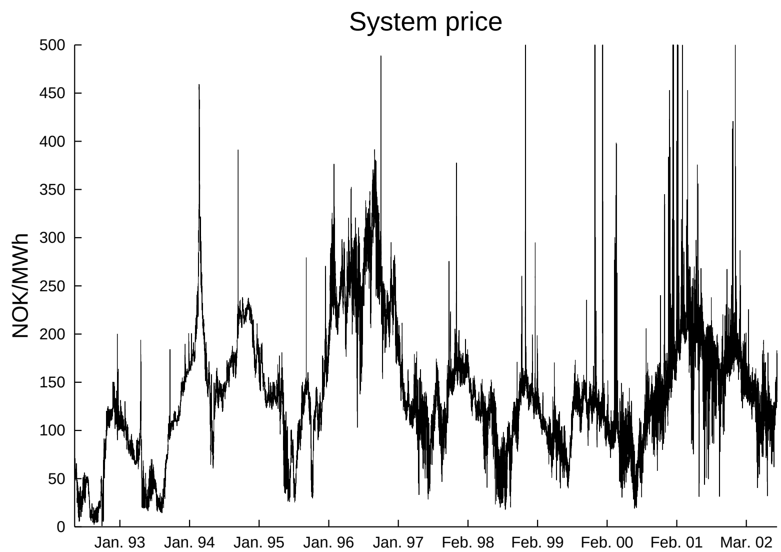

Figure 10: System price for the period 4. January 1993 to 1. October 2001, a total of 76 608 hours. Since we have plotted hourly prices we can more clearly see the occasional price spike. We can also see periodicity in the price.

lower and more volatile in the summer than the rest of the year. The low price is due to the seasonal pattern of electricity consumption, and the high volatility is because of deviations in hydrological balance.

5.2 FLC data

To be able to compare our algorithm to real market participants we have managed to get hold of a very unique data-set. The data-set consist of historical FLC policies for nine real market participants. The policies are for two kinds of flexible load contracts:

- Summer FLC: With a settlement period from 1. May to 30. September. The flexibility is to exercise in 1667 of 3672 hours (45.4%).
- Winter FLC: With a settlement period from 1. October to 30. April. The flexibility is to exercise in 3333 of 5088 hours (65.5%).

Together these two contracts make up a flexible load contract called "5000 hours FLC with 2/3 of the volume in the winter and 1/3 in the summer." For three of the participants we have policies from 1. May 1997 to 1. May 2002, and for three other participants we have policies from 1. May 1999. Due to

Table 2: Descriptive statistics

	Nominal prices					
	Avg	Min	Max	SD	Skewness	Kurtosis
1993*	80.04	14.27	193.75	41.10	0.1896	2.1258
1994	182.65	60.81	459.35	68.49	1.0679	7.0252
1995	117.67	25.38	210.89	60.59	-0.5807	2.7852
1996	253.63	102.97	391.62	79.16	0.2880	3.0466
1997	134.99	28.40	377.80	73.61	0.7945	4.4389
1998	116.35	17.97	735.28	70.47	0.5873	17.7840
1999	112.11	39.99	654.98	67.29	1.9075	28.3528
2000	103.33	19.01	1808.66	66.06	12.7962	410.4873
2001*	188.46	31.21	1951.76	67.86	11.6510	218.7751
Full sample	142.42	14.27	1951.76	67.86	2.2028	33.5961
W1	154.56	20.36	1951.76	65.90	5.4542	100.3643
SO	124.54	14.27	391.62	73.05	1.0170	3.7532
W2	157.53	29.45	735.28	51.37	1.1702	6.0739

	Deseasonalised prices					
	Avg	Min	Max	SD	Skewness	Kurtosis
1993*	-34.41	-94.98	90.41	28.15	0.31	2.64
1994	57.05	-51.22	318.14	40.31	1.30	8.77
1995	-7.93	-107.76	75.60	28.87	-0.49	3.23
1996	128.55	-7.91	271.32	53.04	0.03	2.38
1997	9.06	-74.10	234.89	28.50	1.02	5.85
1998	-9.30	-97.56	583.46	27.80	1.99	46.66
1999	-13.48	-61.04	509.17	23.27	3.81	60.41
2000	-22.25	-74.80	1658.95	36.79	19.46	724.19
2001*	66.58	-93.85	1801.81	68.14	10.55	195.98
Full sample	18.35	-107.76	1801.81	63.78	2.74	39.26
W1	14.32	-88.16	1801.81	63.26	5.99	113.10
S	20.48	-97.56	271.32	70.63	1.09	3.68
W2	20.09	-107.76	583.46	50.44	1.31	6.29

Descriptive statistics conducted on yearly and seasonal subsamples. W1 denotes the period 1. January to 30. April, S denotes the period 1. May to 30. September and W2 denotes the period 1. October to 31. December. The deseasonalising is performed by subtracting $E_t[P_s]$, $s = \{1, \ldots, 8760\}$ from the prices at the beginning of each year. The main results from the statistics is that the average prise has decreased and the skewness and kurtosis has increased. We also see that the skewness and kurtosis is highest in the W1-period, and the S-period has the highest volatility.
*not all prices for this year is included in the calculation of the statistics

incompatibilities we could only use the summer FLC policies for the remaining three participants.

The FLC data was obtained from Skagerak Energi AS - one of Norway's leading power companies. To get hold of the data set we had to anonymise the data by scaling the contracts and by naming the participants as $C1, C2, \ldots, C9$.

5.3 Parameter estimation

With the price data we can now begin the estimation of the parameters in the spot price process. As we recall from section 3 it is possible to estimate the spot price parameters by solving the maximisation problem given by (3.7) on page 18. To separate the fast mean reversion generated by large and sudden changes in the demand or supply from the more slowly mean reversion generated by the hydrological balance, we used a two stage estimation procedure. First we estimated the parameters in X_t from hourly prices with a weekly sampling interval. By construction D_t will start out at zero every week, meaning that D_t will be zero in the weekly data sample. Assuming an hourly sampling resolution of $\mathbf{P}$, we pick every 168'th value and use this data sample to estimate the parameters a, σ, A_1^X, b_0, A_2^X, B_1^X and B_2^X by solving (3.7) on page 18. The second stage is to estimate the parameters in D_t from the full data sample $\mathbf{P}$. To get an estimate for the parameters in D_t we insert the parameters estimated from the first stage into (3.6) and solve the maximisation problem given by (3.7). To ensure that D_t start at zero at the beginning of every week, we set d_0 equal to the value of -ε at time t_1.

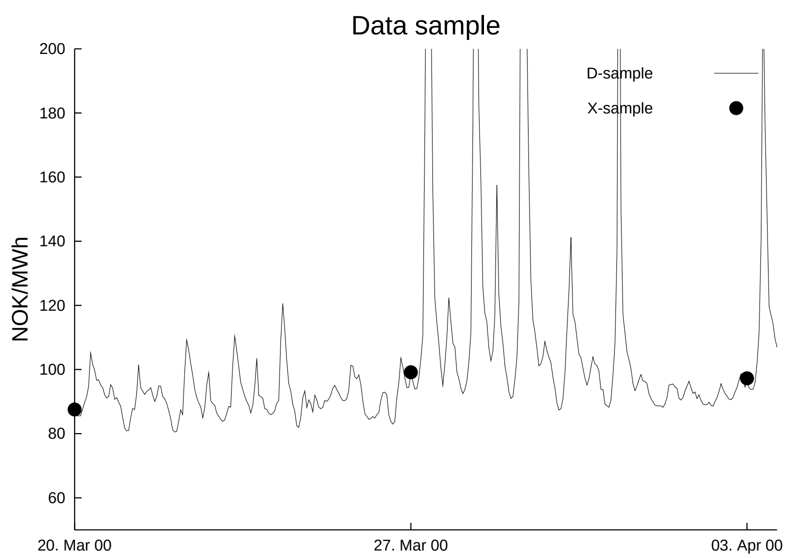

Figure 11: We can here see the relationship between the D_t sample and the X_t sample. We have in this figure on purpose picked a period with price spikes to show that X_t is usually not effected by spikes.

To incorporate new spot prices into our sample we re-estimated the parameters 1. May and 1. October each year. The re-estimation of the parameters made it possible to adapt to changes in dynamics of the spot price and use the largest available sample to get more accurate estimates. The results are given in table 3. As we can see from the estimated parameters the speed of mean reversion, a, is equal in all sub samples. This indicates that the mean reversion property of the spot price dynamics is unchanged over the last eight years. This is however not the case for the volatility parameter, σ, which has decreased. The deseasonalised long run mean, $b_0 + d_0$, has also decreased during the data period. Since we are operating with nominal prices we expected an increase, but the effect of the deregulation of the electricity market and several years with more than normal precipitation must have counteracted this.

The remaining parameters in the table determines the shape of the seasonal patterns. The day and weekly price patterns are quite stable throughout the sample period. The parameters A_1^X, A_2^X, B_1^X and B_2^X which control the yearly price cycle on the other hand seems to be more varying. This may indicate that the yearly price pattern is influenced by other factors than just the deviation from long run mean and the time of the year.

Table 3: Estimated parameters

	Sub samples									
	93-97S	**93-97W**	**93-98S**	**93-98W**	**93-99S**	**93-99W**	**93-00S**	**93-00W**	**93-01S**	**93-01W**
a	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004	0.0004
σ	1.9872	1.9836	1.8991	1.8922	1.8238	1.7819	1.7435	1.7126	1.6674	1.6655
b_0	159.97	152.52	150.93	146.04	142.03	141.03	136.05	134.98	138.43	138.39
d_0	-10.06	-10.12	-10.11	-10.48	-10.65	-10.59	-10.77	-11.06	-12.49	-12.85
A_1^X	32.19	31.16	33.80	33.70	34.44	33.70	32.55	32.94	28.87	27.14
A_2^X	-8.40	-6.41	-5.36	-3.36	-2.68	-2.84	-3.57	-3.40	-5.07	-3.68
A_1^D	-3.07	-3.18	-3.16	-3.30	-3.29	-3.27	-3.13	-3.31	-3.42	-3.59
A_2^D	-0.24	-0.24	-0.13	-0.13	-0.16	-0.18	0.00	0.02	0.07	0.09
A_3^D	-5.31	-5.56	-5.54	-5.87	-5.87	-5.93	-6.24	-6.51	-7.39	-7.82
A_4^D	1.58	1.96	1.76	2.00	1.80	1.90	1.78	2.01	1.88	2.28
B_1^X	-3.65	1.44	-2.42	3.66	-0.09	0.94	-2.00	-0.54	2.59	2.99
B_2^X	12.32	12.99	10.82	7.59	6.63	6.33	5.79	4.50	4.14	4.33
B_1^D	4.09	4.50	4.44	4.67	4.69	4.73	4.98	5.11	5.49	5.76
B_2^D	2.78	2.99	2.94	3.00	3.00	2.97	3.13	3.17	3.82	3.93
B_3^D	-4.75	-5.01	-4.92	-5.17	-5.07	-5.03	-4.88	-4.99	-5.19	-5.56
B_4^D	-4.24	-4.25	-4.23	-4.34	-4.34	-4.30	-4.61	-4.72	-5.54	-5.70

This table shows the estimated parameters obtained by solving (3.7). Each column represents a sub sample. "93-97S" represents the period 4. January 1993 to 30. April 1997, "93-97W" represents the period 4. January 1993 to 30. September 1997 and so on. We see a decrease in the long term average price ($b_0 + d_0$) and the volatility (σ). The parameters determining the daily and weekly price patterns is quite stable, but the parameters determining the yearly patterns is more unstable. In addition we use a risk free interest rate of 5% pro anno in the optimisation algorithm.

6 Results

The purpose of this paper is to study how a flexible load contract should be exercised optimally when only historical spot price information is to be used. So far we have expressed the FLC as an optimisation problem, found a good process to model the spot and estimated the parameters of this process. We are therefore ready to implement our model and compare it with historical FLC-data.

6.1 Results from case

In this section we focus on the contract defined in the first section. This was a FLC contract for the summer 1997 from 1. May to 1. October, totally 3672 hours. In our case we paid 958.525 NOK for the right to withdraw 8335 MWh, with a maximum of 5 MWh per hour. Therefore our target is to exercise the contract during the 1667 hours with the highest spot price.

In figure 12 we show how our algorithm exercised the contract during the summer period of 1997. The plot shows the accumulated control (i.e. the Q-variable) at each instant. We compare this with the aposteriori best path which picks exactly the best 1667 hours. We also show how the contract was utilised by a market participant. Even though the historical contract is closer to the aposteriori best, it is not necessarily better than the model. This is because there is no monotonicity in the value of the policies as we get closer to the ex post optimal curve. This can be illustrated by the policy picking the same hours as the optimal, but with a 12 hours lag. This policy will normally perform poorly (because of the low price levels in the night) but it will be very close to the aposteriori optimal curve. In figure 12 we have also plotted the frequency plot of all the prices together with the distribution of the prices for the exercised volume. As we can see our model managed to exercise most of its volume on the "right side" of the price distribution. This indicates an ability to distinguish a high-price state from a low-price state. The actual performance of the model is difficult to measure from the frequency plot or the accumulated control. Therefore we need to calculate the realised revenue from the different policies. The results are given in table 4. We see that the model based on an one hour update gets a fourth place, but the difference from the other competitors is relatively small. The model has an advantage since the control can use hourly

Table 4: Value of exercised FLC (Summer-1997)

	Total revenue from FLC	Revenue excess base load	Excess revenue per MWh
Model:			
1h	1 000 559	98 353	11.80
24h	982 141	79 935	9.59
Competitors:			
C1	1 010 281	108 075	12.97
C2	1 000 506	98 300	11.79
C3	1 007 316	105 110	12 61
C4	986 806	84 600	10.15
C5	**1 013 166**	**110 960**	**13.31**

Value of the exercised FLC obtained by our model and 5 competitors. Since the total revenue mainly consist of the value of the base load, and this base load is often hedged when a FLC is bought, it is common to look at the total or per MWh excess revenue. For this particular FLC we see that competitor C5 managed to obtain the highest revenue, with our model obtaining a 1.2% lower total revenue.

price information. We can adjust for this and use the (possibly[11]) more realistic model where a 24 hours deterministic development of the observed price is used to find the control. We see that the result is worsened, as expected.

Can we from this conclude that our model is inferior? The answer is no! It is important to keep in mind that this problem is of a stochastic nature. Therefore even though we knew the real stochastic process (which of course is impossible) the optimal control could give bad results when only one season (i.e. one replicate) is studied. But since the expected value is maximised, the long run accumulated value should be good. In the next subsection we introduce a new FLC for the winter period and again show how our model performs compared to real life competitors during the 1997 - 2001 period.

6.2 General results

To supplement the FLC for the summer period we introduce a new type of FLC for the winter period 1. October to 30. April, totally 5088 hours. The new contract has a total volume of 16665 MWh and a maximum effect of 5 MW. Our goal is therefore to pick the 3333 hours with the highest price. With this

[11] The market participants does have good estimates for price development the following days. If we use this information the model with an hourly update may be realistic after all.

contract we are able to show how the model performs over the whole 1997 - 2001 period. Finding a good method to compare different contracts is not straight forward. Contracts with the same degree of flexibility and with equal delivery period must be used. In addition competitors may have different risk attitudes. We decided to first focus on the excess revenue obtained for the period 1997-2001, and as a second stage see if there were any differences in the volatility of the revenue among the participants.

The results are presented in table 5 and figures 13 to 21. We may draw some conclusions from the results. First of all our model manage to obtain the highest accumulated revenue during the period. The model also demonstrates that it has the courage to pick many hours early if the prices are sufficiently good. Opposite, the model waits for a long time if the prices are poor. This can be seen as a risky behavior, and may be a consequence of the risk neutral model[12] formulation. The results also shows that the results vary substantially from extremely good (as in W2001) to extremely bad (as in W2000), but with a good average performance. This may also be seen as a materialisation of risk neutrality. We will in the next subsection take a closer look at the FLC for the winter period 2000, and try to analyse the result.

Another observation is that our model seems to perform better for the winter contracts than for the summer contracts. One reason can be that the winter contract has a lower degree of flexibility than the summer contract. For the winter contract we have to exercise $3333/5088 \approx 65.5\%$ of the hours against only $1667/3672 \approx 45.4\%$ for the summer contract. Another reason may be that the process is best calibrated to the winter data. The reason for this is that since we have only used two trigonometric functions to model the changes through the year, the process can not model the summer vacation and all the holidays in May properly. The low prices in the summer is typically expected to appear 6 months after the highest winter prices. This is not necessarily the case in the real world. Normally the lowest prices appear in the vacation weeks of July. Our spot process does not expect collapse in the July prices, and therefore the routine has a tendency to pick too few hours in May and June. Then, when the really poor prices appear in July these hours cannot be exercised either. Now the routine is basically forced to take all the hours in August and September. This scenario is broken if the early summer prices are sufficiently high as in

[12] On the other hand, we believe that this routine does the correct trade off between the different effects of the model such as interest rate, volatility, reversion and periodicity.

the summer of 2001. We believe that the performance in the summer contracts could be improved with a more representative process. Several different ideas can be followed.

- We could include more trigonometric functions into the spot process.
- We could estimate separate summer and winter processes.
- We could include a drift term into the process such that the holidays are placed properly.

We feel that we have demonstrated that the model works quite good with this level of precision, and leave the process of refinements for future work.

RESULTS SUMMER-1997

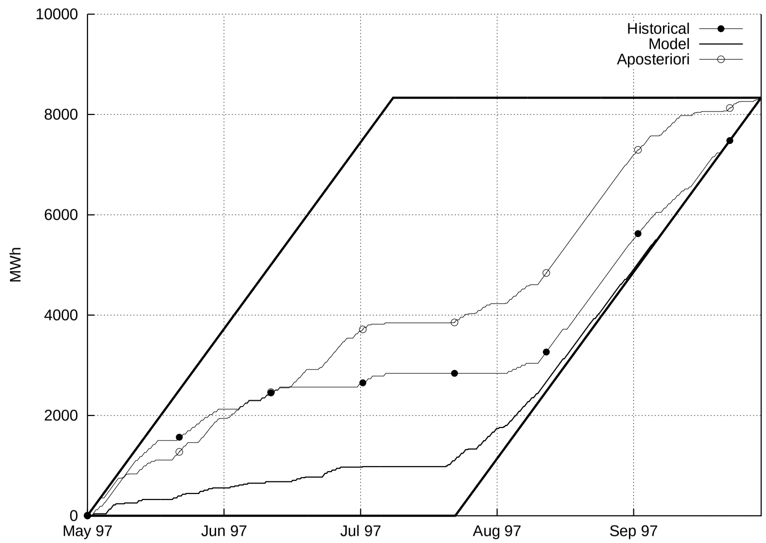

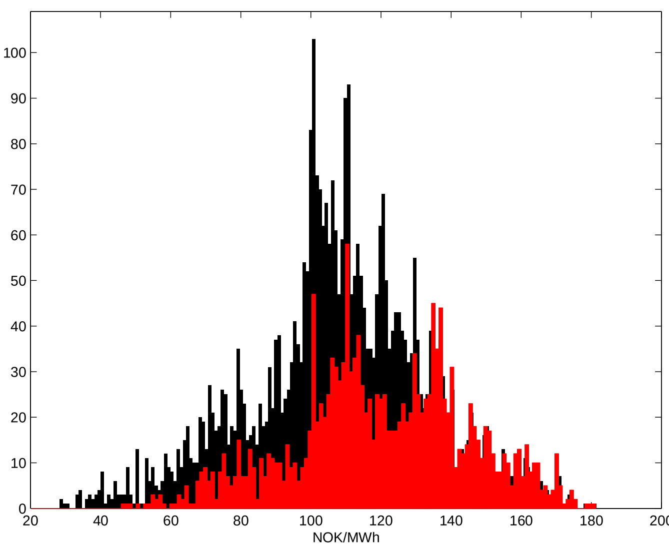

Figure 12: On the left hand side we have plotted the algorithm's exercise policy together with a competitor's (denoted as C1 in subsequent tables) exercise program. The a posteriori best exercise is also plotted as a reference. On the right hand side we have a figure of the empirical density function for all spot prices and spot prices for the exercised volume.

Table 5: Results for the period 1997 - 2001

	Model		**Historical**								
			C1	C2	C3	C4	C5	C6	C7	C8	C9
	1h	24h	24h	24h	24h	24h	24h	24h	24h	24h	24h
S-1997	98 353	79 935	108 075	98 300	105 110	84 600	**110 960**				
W-1997	179 315	**185 980**	59 595	54 740	38 765						
S-1998	-8 670	-35 590	49 245	**133 155**	90 620	-25 090	90 605				
W-1998	**186 315**	185 315	122 885	183 210	51 650						
S-1999	174 035	155 200	177 770	175 450	**181 755**	165 450	181 735	168 205	171 080	164 520	
W-1999	**167 315**	157 985	67 340	40 175	7 565			60 740	73 125	69 590	95 935
S-2000	144 195	99 020	157 430	155 425	**180 775**	141 015	181 655	160 290	160 975	150 685	
W-2000	-75 990	-64 995	**125 820**	113 525	49 450			113 515	109 030	17 200	107 590
S-2001	**173 620**	152 530	111 875	101 030	24 870	121 305	44 145	112 755	116 605	71 790	
W-2001	**220 645**	207 815	162 452	143 243	*			82 733	83 271	82 688	*
Sum all summer	581 450	451 095	604 395	**663 360**	583 130	487 280	609 100				
Sum all winter	**677 600**	672 100	538 092	534 893	*						
Sum all	**1 259 050**	1 123 195	1 142 487	1 198 253	*						
Sum from S-1999	**803 820**	707 555	802 687	728 848	*			698 238	714 086	556 473	

This table shows the revenue excess base load for all the FLC for the period 1997 - 2001. The FLC with the highest exercised value is marked in bold. *The table is not yet updated with all the data from the W-2001 period

RESULTS WINTER-1997

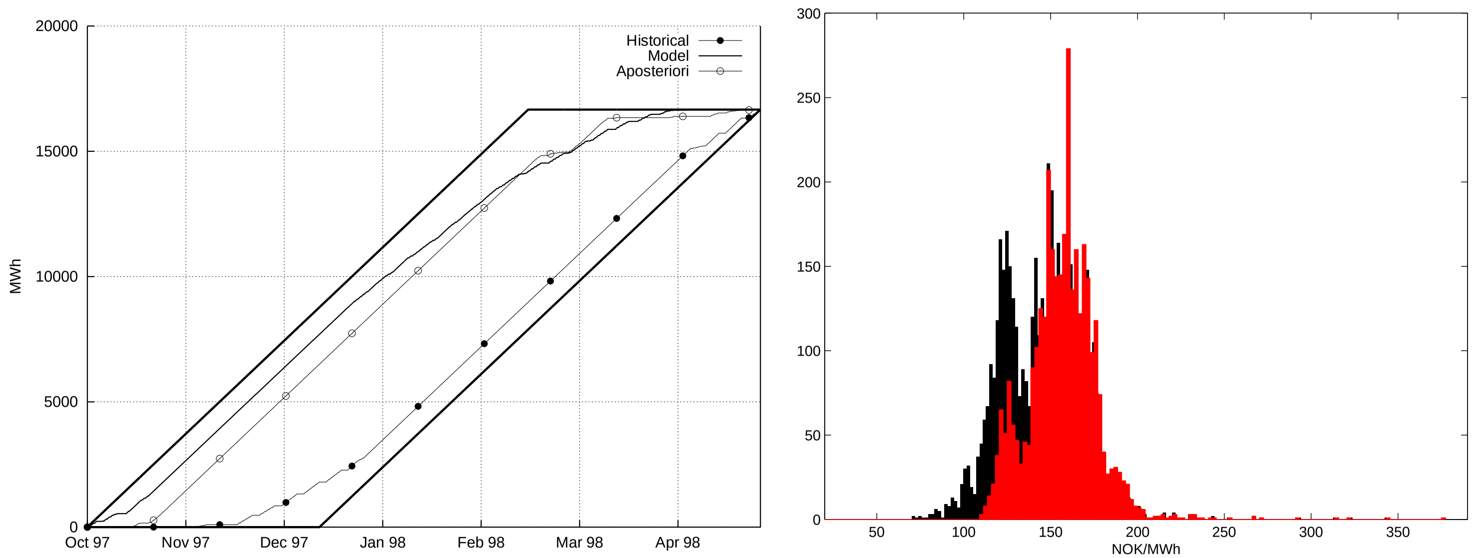

Figure 13: Flexible load contract: Winter 1997

RESULTS SUMMER-1998

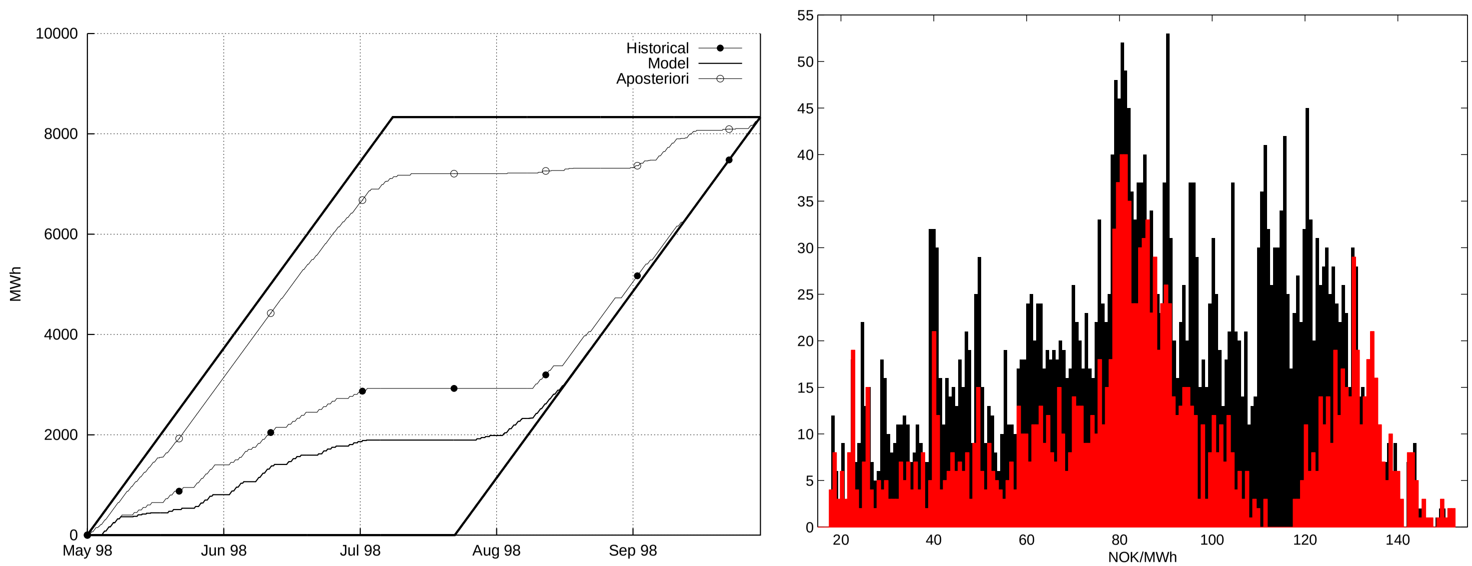

Figure 14: Flexible load contract: Summer 1998

RESULTS WINTER-1998

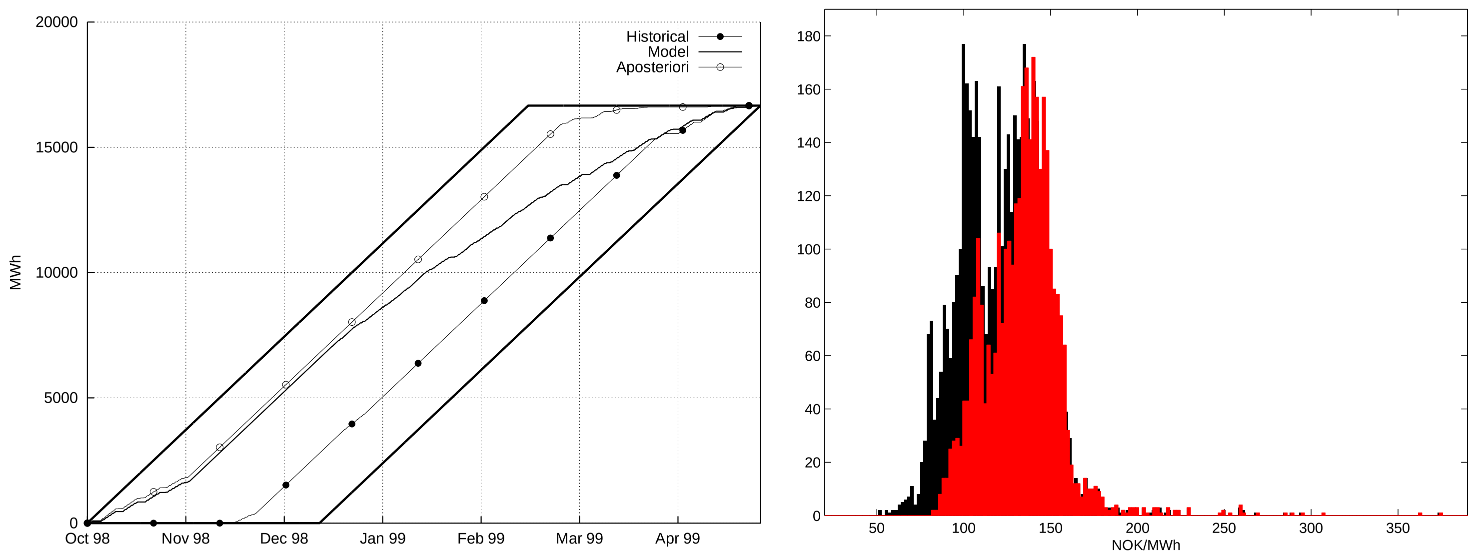

Figure 15: Flexible load contract: Winter 1998

RESULTS SUMMER-1999

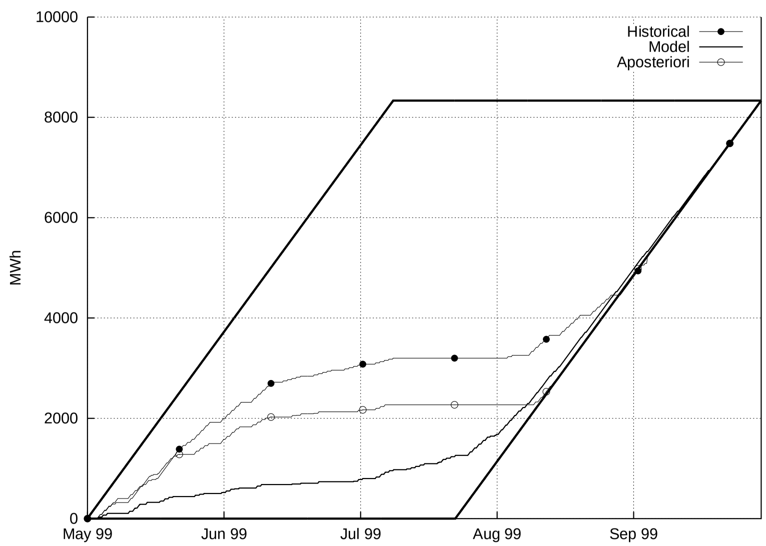

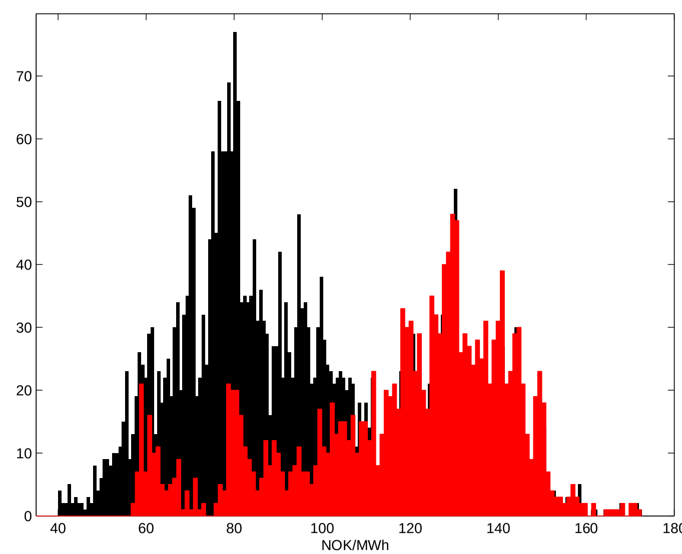

Figure 16: Flexible load contract: Summer 1999

RESULTS WINTER-1999

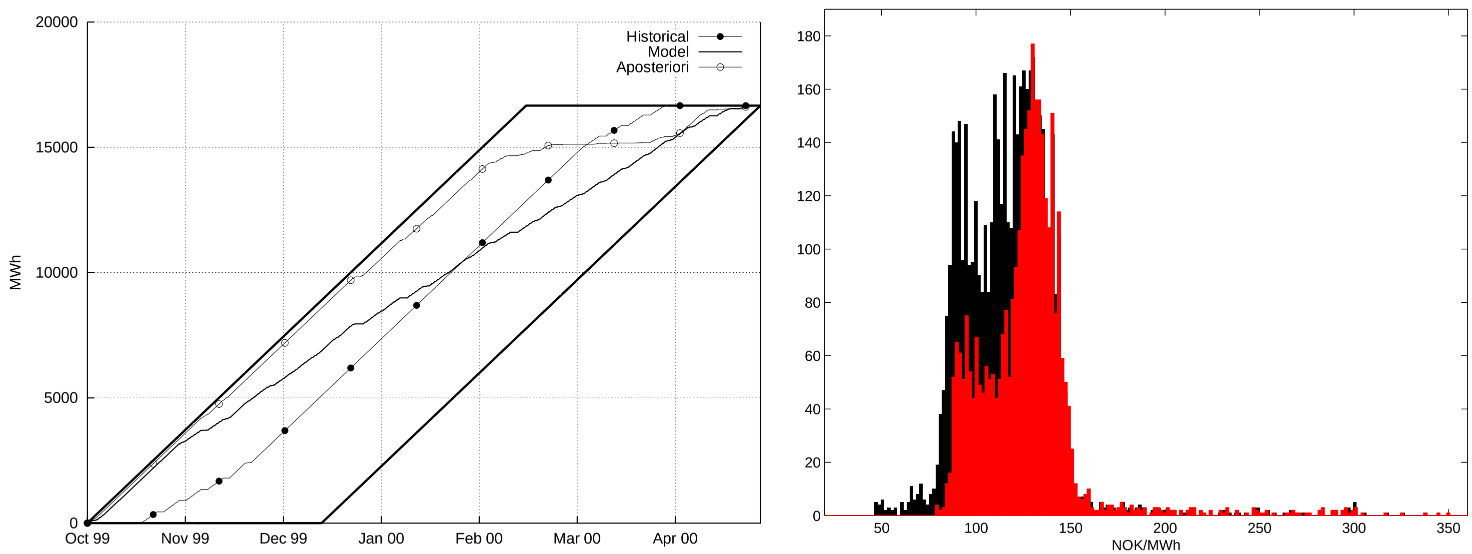

Figure 17: Flexible load contract: Winter 1999

RESULTS SUMMER-2000

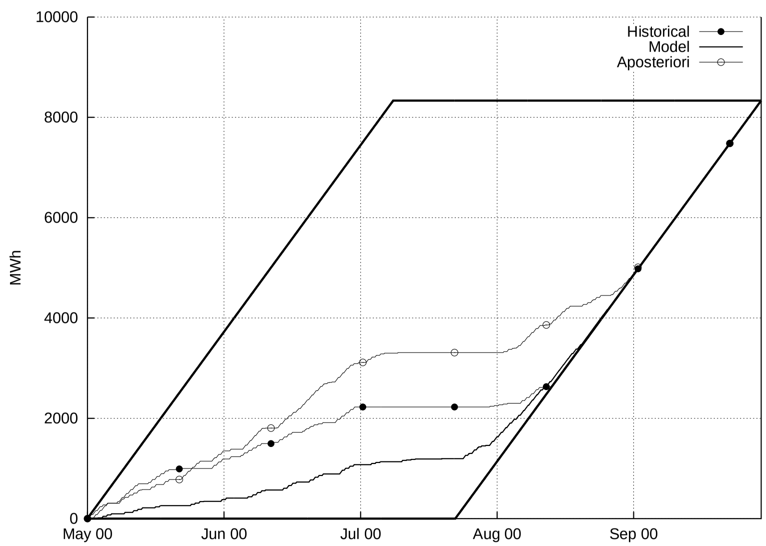

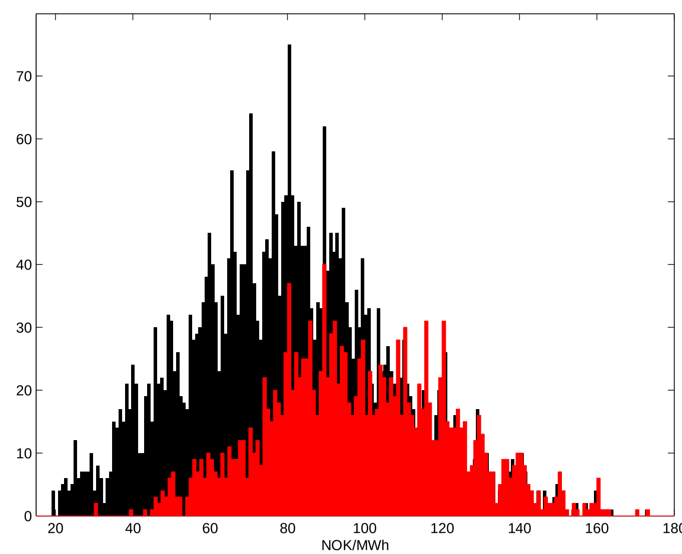

Figure 18: Flexible load contract: Summer 2000

RESULTS WINTER-2000

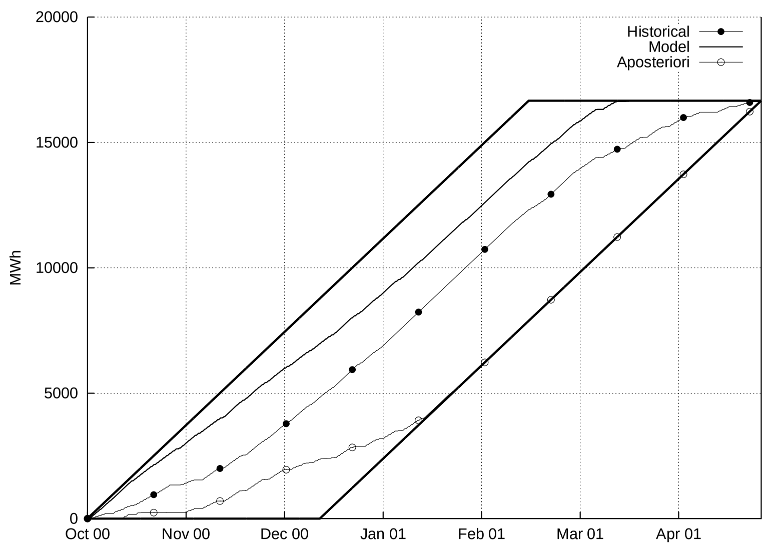

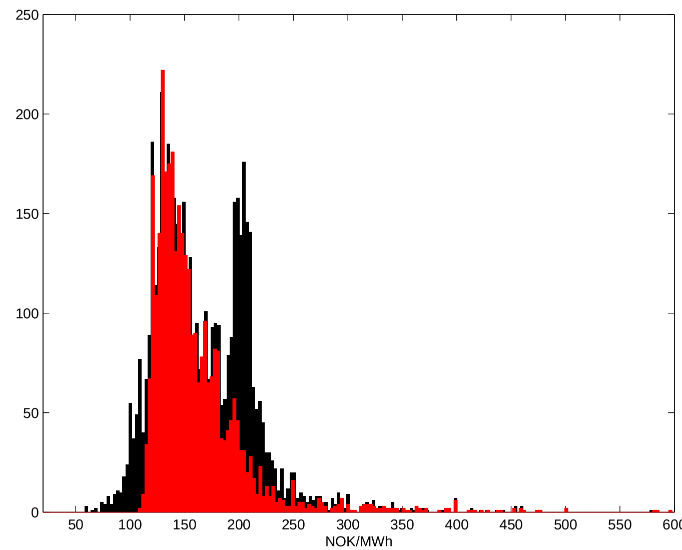

Figure 19: Flexible load contract: Winter 2000

RESULTS SUMMER-2001

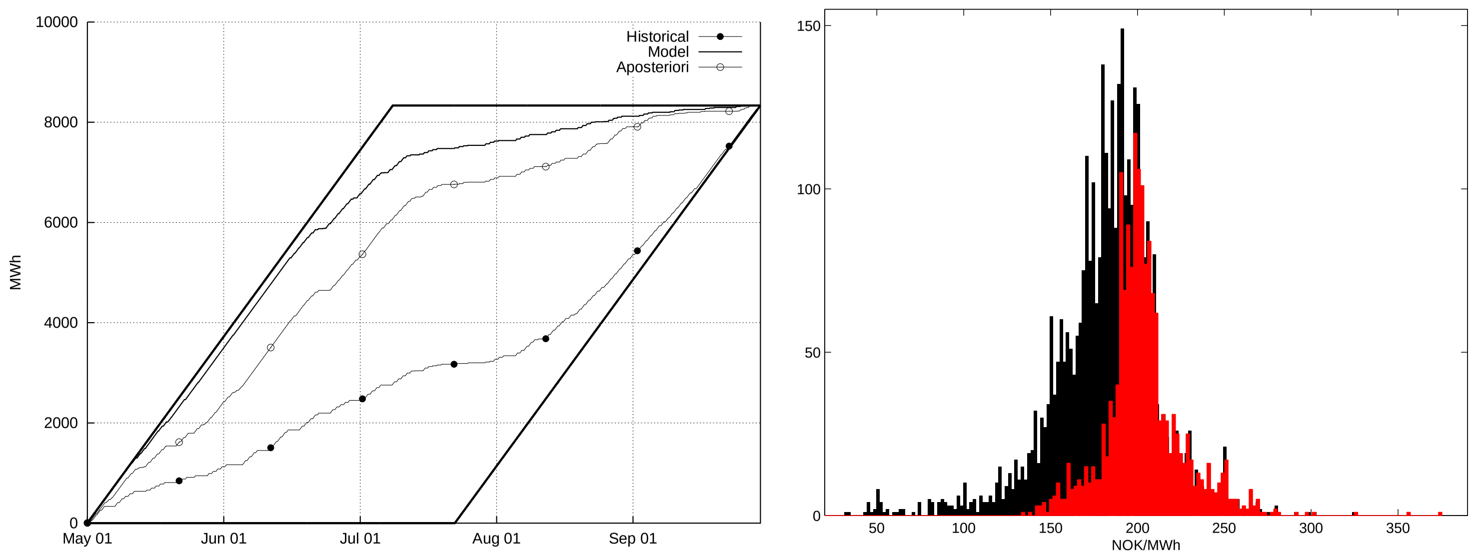

Figure 20: Flexible load contract: Summer 2001

RESULTS WINTER-2001

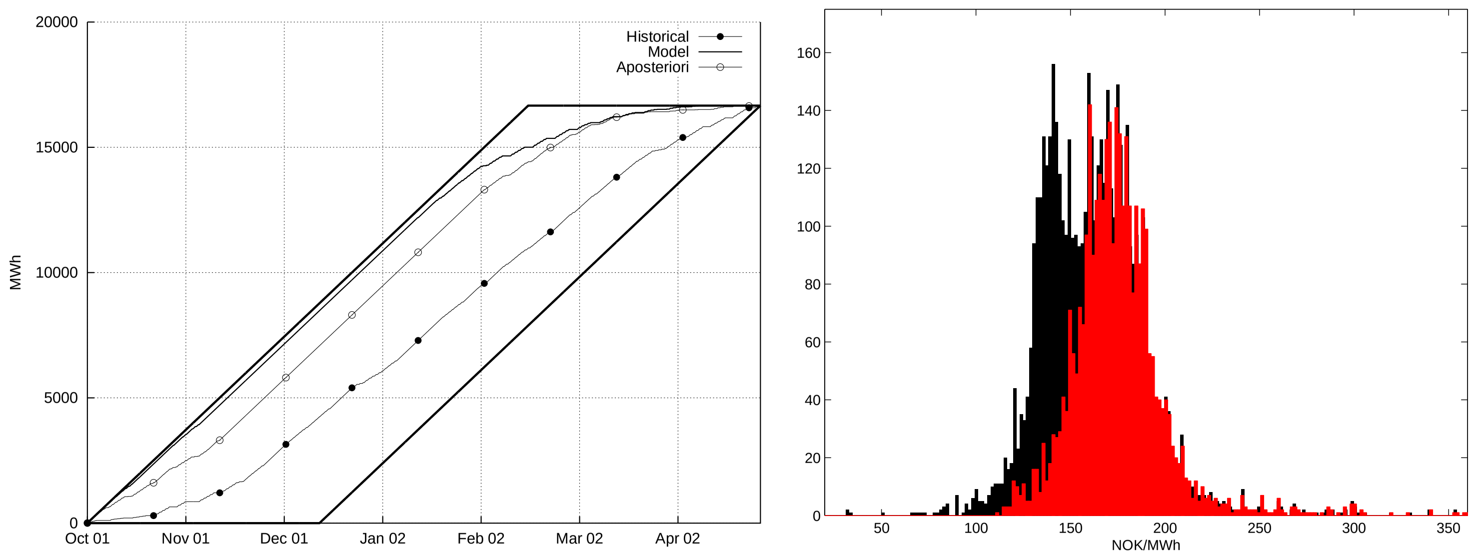

Figure 21: Flexible load contract: Winter 2001

6.3 A closer look at the winter 2000 FLC

The revenue from the FLC during the winter 2000 period was very low. In fact the revenue was lower than the value of a base load contract. We have two different types of explanations for this. Firstly, there may be weaknesses in the routine, especially due to the limitations of the movements of the discretised price process from hour to hour. Secondly and more importantly there are clearly an informational asymmetry since the market has access to information that the model does not have.

From figure 19 on page 48 we see that the algorithm starts out with an exercise policy close to maximum. This was a result of the higher-than-normal prices for this time of year. The degree of exercised volume was later reduced some, and the difference between the competitors policy and ours was reduced. During February our model saw sufficiently high prices to exercise the remaining volume, thereby missing several price spikes in March and April. This is reflected in the density plot as the low volume exercised in the 180 - 240 price range. If we had used a price process that was able to model spikes, the algorithm would not have exercised the remaining volume so soon. The forward prices did capture large parts of the price spikes and if incorporated would have helped. On the other hand the FLC for the winter 2001 period did very well since the routine overlooked the predictions given by the forward market. So the effect of information from the forward market is not clear. We therefore believe the main reason for the poor winter 2000 results was the algorithm inability to capture the possibility of future price spikes.

7 Concluding remarks

In this paper we have analysed flexible load contracts by formulating the contract as a stochastic optimisation problem. The value function is expressed as the solution of the Hamilton Jacobi Bellman equation in which the optimal control takes only the extreme values. By carefully examining the dynamics of the spot price in the Nordic electricity market we decided to use a time dependent mean reverting Ornstein-Uhlenbeck process. The process modelled daily, weekly and yearly price cycles. In addition it captures mean reversion due to deviations in the hydrological balance. The process has 21 parameters which was estimated from historical price data by a mixture of OLS and maximum likelihood. Estimation was conducted partly on a weekly data sample and partly on an hourly data sample. This to distinguish the short range factors from medium range factors.

To be able to solve the optimisation problem we discretised the time and state space and derived an algorithm to find the value function and optimal control in each node. To dampen the effects of a truncated price space we combined absorbing and reflecting boundary conditions.

We implemented the algorithm and calculated the optimal control for the five year period 1. May 1997 to 30. April 2002. The accumulated revenue from this control was compared to the revenue for nine market participants. We find that our algorithm obtains the highest accumulated exercise revenue for this period. The model also demonstrates that it has the courage to pick many hours early if the prices are sufficiently good. This can be seen as a more risky behavior, and may be a consequence of the risk neutral assumption. Another observation is that our model seems to perform better for winter contracts than for the summer contracts. We believe the performance for the summer contracts can be improved with a more representative process.

In our opinion this model demonstrates a great potential for utilisation of contracts of this type. The methods can be developed further to improve the results even more. The introduction of a jump process is important in this respect. We stress that the methods are fully operational, and can be implemented by practitioners, for instants as a tool for benchmarking (or just to improve their profits).

References

[1] B. Johnson and G. Barz. *Energy Modelling and the Management of Uncertainty.* Risk Books, second edition, 2001.

[2] A. Øksendal. *Mathematical models for investment under uncertainty.* Dissertation, University of Oslo, Faculty of Mathematics and Sciences, 2001.

[3] H. Kushner and P. Dupuis. *Numerical Methods for Control Problems in Continuous Time.* Springer, second edition, 2001.

[4] S. Stage and Y. Larsson. Incremental cost of water power. *AIEE Transactions*, pages 361–365, August 1961.

[5] A. C. Thompson. Valuation of path-dependent contingent claims with multiple exercise decisions over time; the case of take-or-pay. *Journal of Financial and Quantitative Analysis*, 30(2):271–293, June 1995.