Helsinki University of Technology

Systems Analysis Laboratory Research Reports E15, May 2005

HYDROPOWER PRODUCTION PLANNING AND HEDGING UNDER INFLOW AND FORWARD UNCERTAINTY

Erkka Näsäkkälä

TEKNILLINEN KORKEAKOULU
TEKNISKA HÖGSKOLAN
HELSINKI UNIVERSITY OF TECHNOLOGY
TECHNISCHE UNIVERSITÄT HELSINKI
UNIVERSITE DE TECHNOLOGIE D'HELSINKI

Distribution:
Systems Analysis Laboratory
Helsinki University of Technology
P.O.Box 1100
FIN-02015-HUT, FINLAND
Tel. +358-9-4513056
Fax +358-9.4513096
Systems.analysis@hut.fi

This report is downloadable at
www.e-reports.sal.hut.fi/pdf/E15.pdf

Series E- Electronic Reports
www.e-reports.sal.hut.fi

ISBN 951-22-7688-7
ISSN 1456-5218

Title: Hydropower production planning and hedging under inflow and forward uncertainty

Author: Erkka Näsäkkälä
Systems Analysis Laboratory
Helsinki University of Technology
P.O. Box 1100, 02015 HUT, FINLAND
erkka.nasakkala@hut.fi

Jussi Keppo
Department of Industrial and Operations Engineering
University of Michigan
1205 Beal Avenue, Ann Arbor, MI, 48109-2117, USA
keppo@umich.edu

Date: May 2005

Status: Systems Analysis Research Reports E15 May 2005

Abstract: We consider long- and medium-term production planning and hedging. The uncertainties are from electricity price process and from random inflow. The price uncertainty is modeled with a multidimensional forward curve model. We give a simple and intuitive parameterization for the optimal production strategy of a profit maximizing producer. The parameterization leads to a simple production hedging policy. The accuracy of the parameterization is analyzed by comparing its expected cash flows to an upper bound of expected cash flows. In our bench marking case the parameterization gives profits that are less than 2.6% from the optimal profits. Further, our bench marking illustrates that during winters 1997-2003 our method would have increased the profits of an actual hydropower producer by 4.2%. This implies that the actual hydropower producer has not fully utilized the information in electricity derivative markets.

Keywords: Hedging, electricity forward curve dynamics, hydropower production

Hydropower production planning and hedging under inflow and forward uncertainty

Erkka Näsäkkälä[1], *Jussi Keppo*[2]

Abstract

We consider long- and medium-term production planning and hedging. The uncertainties are from electricity price process and from random inflow. The price uncertainty is modeled with a multidimensional forward curve model. We give a simple and intuitive parameterization for the optimal production strategy of a profit maximizing producer. The parameterization leads to a simple production hedging policy. The accuracy of the parameterization is analyzed by comparing its expected cash flows to an upper bound of expected cash flows. In our bench marking case the parameterization gives profits that are less than 2.6% from the optimal profits. Further, our bench marking illustrates that during winters 1997-2003 our method would have increased the profits of an actual hydropower producer by 4.2%. This implies that the actual hydropower producer has not fully utilized the information in electricity derivative markets.

Key words: Hedging, electricity forward curve dynamics, hydropower production

[1] Corresponding author. Systems Analysis Laboratory, Helsinki University of Technology, P.O. Box 1100, FIN-02015 HUT, Finland. E-mail: erkka.nasakkala@hut.fi

[2] Department of Industrial and Operations Engineering, University of Michigan, 1205 Beal Avenue, Ann Arbor, MI, 48109-2117, USA. E-mail: keppo@umich.edu

1 Introduction

We study production planning and hedging of a hydropower plant. This is important since hydropower production has a significant effect in electricity markets. In 2001, in OECD countries, electricity generation from hydro reservoirs was 1229.1 *TWh,* which corresponds 13% of the total electricity production. For comparison, the total electricity production in 2001 from renewable sources was 1424 *TWh* given that about 86% of renewable energy was produced by hydropower systems. The largest hydropower generating countries are Canada (333.0 *TWh*), the United States (201.2 *TWh*) and Norway (120.4 *TWh*). Other big hydropower producers are Japan, Sweden and France (see, e.g., IEA, 2003). The hydropower plants' negligible response times and production costs enable significant variations in a hydropower plant's production strategy. Thus, in addition of being the largest renewable source in electricity production, hydropower is important as a regulating power reacting rapidly to changes in electricity supply and demand.

Generally, electricity is a non-storable commodity and thus seasonal variations in supply and demand cause seasonality into the electricity spot prices. However, the production capacity in a hydro reservoir is storable and thus a hydropower producer decides each moment of time, whether to use the reservoirs with the current spot price or wait for higher prices in the future. This means that the production decision made today is not based solely on the current spot price, but also on the expected future electricity prices and their uncertainties. Further, also the reservoir's water level and expected inflows affect the production decision. When the reservoir is full all new inflow will be spilled. Thus, if the water level and/or the expected future inflow are high, the owner of the plant is more eager to produce than when the water level is low and/or the estimated future inflow is small. Due to the uncertainty in the future rainfall and temperature, the inflow to the reservoirs is stochastic. When the water level is high an increase in the inflow uncertainty increases the risk of spillage making earlier production more favorable. Hydropower production is a dynamic problem, where today's decision made today affects the future production decisions. By producing today the estimated future water level decreases, which makes future production less attractive. Stochastic programming methods are suitable for problems with such a dynamic structure. Gjelsvik et al. (1992), Pereira (1989) and Pereira and Pinto (1991), among others, use stochastic programming in hydropower production planning in regulated electricity markets. In their models the price is deterministic and the demand is stochastic. Stochastic programming for deregulated electricity markets is studied, for example, in Gjelsvik and Wallace (1996), Fosso et al. (1999), Mo et al. (2001), and Fleten et al. (2002). A good survey of the stochastic programming models, both in regulated and deregulated electricity markets, is given in Wallace and Fleten (2002).

Since the beginning of the 1990's the electricity markets around the world have started to move from centrally planned monopolies towards deregulation. For example, the United Kingdom, New Zealand, the Nordic countries, and some states of USA have opened their electricity markets for competition. Also the European Union has adopted directives to deregulate electricity markets (directives 96/92/EC and 98/30/EC). In a deregulated market there is a publicly quoted reference price, set by the equilibrium between supply and demand. Even though the deregulated electricity markets remind us a lot of stock markets, there are several fundamental differences. From a view point of a market participant one major difference is that electricity usage/production is given by an exogenous consumption/production process, while a stock investor can decide asset holdings himself. In the case of hydropower the total production depends on the reservoir inflow, which, of course, depends on the amount of rainfall. The publicly quoted electricity spot price has created electricity derivative markets. For example, electricity forwards and options are traded in electricity exchanges [e.g., Nord Pool (Scandinavia), APX (USA and Netherlands), NYMEX (New York), and VicPool (Victoria, Australia)]. In Scandinavia, for example, around one quarter of the total physical demand is traded via Nord Pool and, therefore, the spot price is a credible reference index for the whole market. The Scandinavian derivative markets are also active the combined volume of forward and futures in the Nord Pool and in the OTC-markets is nearly ten fold over the size of the physical market.

We show how electricity forward markets can be used in hydropower production. First, we use the electricity forward curve dynamics in production planning. Second, we show how the production can be hedged by selling electricity forward contracts. Remember that the production decisions are also affected by the stochastic inflow process. If the inflow uncertainty is not perfectly correlating with the electricity derivative prices, then the load uncertainty cannot be perfectly hedged with those derivatives. Therefore, the market is incomplete (for incomplete markets see, e.g., Karatzas et al., 1991, Cuoco, 1997, Cvitanic et al., 1997, and Pennanen and King, 2004). The incompleteness now means that the production process can not be hedged perfectly, i.e. there is always some load risk that the producer is not able to hedge in the electricity derivative markets. If the correlation between the inflow uncertainty and derivative prices is perfect (1 or -1) then the market is complete and swing options can be used for hedging (see, e.g., Keppo, 2004, Jaillet et al., 2001 and Thompson, 1995).

If the forward markets are not used in the production planning, important information from financial markets is lost. The information in financial markets is important because the value of the production cash flows can be calculated in terms of the electricity derivatives. An expected spot price process omitting the forward prices can give reservoir values that are inconsistent with the law of one price, i.e. similar payoff with forward contracts may have different value. Even if the expected spot price process matches the forward curve the possibility to trade forward contracts should not be omitted. By selling forward contracts the uncertainty in the future cash flows can be decreased without changing the expected cash flows. For example, consider an electricity producer producing 16 MWh of electricity next week. By selling the 16 MWh in the forward markets today, the producer can guarantee that the cash flows from the next week's production are equal to the forward price. By not selling the forwards, the expected cash flows are dictated by next week's forward price, but it is possible that next week's forward price changes. Mo et al. (2001) and Fleten et al. (2002) study hydropower production planning when the possibility to trade forwards and options is included. In these studies the expected spot price process is first adjusted so that it is consistent with the forward curve. Second, the optimal production and trading strategy for a risk-averse producer is solved with stochastic dynamic programming. Numerical examples in Fleten et al. (2002) and in Vehviläinen and Keppo (2003) illustrate that by trading electricity forwards and options it is possible to reduce risks in the expected cash flows without affecting significantly the portfolio's value.

We take a slightly different approach. Instead of traditionally used stochastic programming methods we apply the martingale method of optimal consumption and portfolio selection (see, e.g., Merton, 1969, Merton 1971, Cox and Huang 1989, and Karatzas et al., 1987). We use electricity forward price dynamics as the price process. The electricity forwards are financial assets whose value converges to the spot price as the maturity decreases to zero. The electricity forward curve depends on several risk factors. The key features are the spot volatility curve, the volatility curve's maturity effect, and the forward curve's correlation structure (see, e.g., Audet et al., 2003 and Koekebakker and Ollmar, 2001). We derive an optimal production strategy for a firm that maximizes its value. The optimal production strategy is a bang-bang strategy, which can be characterized with a threshold function for electricity spot price. The probability that the future spot price is greater than the threshold function gives the optimal hedging strategy. Instead of trying to solve the threshold function analytically, we suggest a simple and intuitive parameterization for it. The parameterization is an approximation of the optimal production strategy, and thus it gives a lower bound for the optimal production strategy. We use Monte Carlo simulation to solve the parameters of the threshold function. The Monte Carlo simulation can also be used to obtain an upper bound for the expected discounted cash flows. The difference of the upper and lower bound gives an upper bound for the error of our parameterization.

We bench mark our method with the realized production strategy of a Norwegian hydropower producer during winters 1997-2003. During winters 1997-2003 the average upper bound of our model's error was 2.6%. On average our method would have increased the profits of the hydropower producer by 4.2%. Statistical analyses indicate that the improvement given by our model compared to the actually used production strategies is statistically significant. This implies that the actual hydropower producer has not fully utilized the information in electricity derivative markets.

This paper is organized as follows: The forward curve dynamics are introduced in Section 2, and Section 3 characterizes the optimal production strategy. In Section 4 parameterization for the optimal production strategy is given together with the upper and lower bounds of the value function. The optimal hedging strategy is solved in Section 5. In Section 6 the method is compared with the realized production strategy of a Nordic energy company and Section 7 gives a production planning and hedging example based on the estimated model parameters. Finally, Section 8 concludes the paper.

2 Forward Curve Dynamics

We consider an electricity market where spot and derivative instruments are traded continuously in a finite time horizon $[0,\tau]$. In describing the probabilistic structure of the market, we will refer to an underlying probability space (Ω, F, P), where Ω is a set, F is a σ-algebra of subsets of Ω generated from continuous and jump uncertainties, and P is a probability measure on F. The following assumption characterizes the derivative market.

Assumption 1. *There exist different maturity forward contracts on electricity spot price. The electricity derivative markets are complete and there is no arbitrage.*

The no arbitrage assumption states that all portfolios with the same future payoffs have the same current value. The no arbitrage condition and the completeness of the market ensure the existence of a unique linear pricing function, which can be described by an equivalent martingale measure Q. Under the martingale measure all the expected returns of traded non-dividend paying financial assets are equal to the risk-free interest rate r (see, e.g., Duffie, 2001 and Björk, 1998). Since electricity is not a financial asset its expected return under Q is not usually equal to the risk-free rate. However, the derivatives are financial assets and thus, at time t, the price of a T-maturity derivative on electricity spot price equals

$$f(t,T) = e^{-r(T-t)} E^Q\left[\phi(S(T)) \mid F_t\right] \quad \forall \ t \in [0,T], \tag{1}$$

where $\phi(\cdot)$ is the payoff function, $S(T)$ is the electricity spot price at time T, E^Q is the expectation operator under the martingale measure Q and the expectation in (1) is taken with respect to the information at time t, F_t. For simplicity, we assume that risk-free rate r is constant.

We denote the T-maturity forward price, at time t, by $S(t,T)$. The electricity forward $S(t,T)$ pays at the maturity the difference of forward price at time t and realized spot price at time T, denoted by $S(T)$. In other words, the payoff function of a T-maturity forward is $\phi(S(T)) = S(T) - S(t,T)$. The value of the forward contract when initiated is by definition zero, i.e. $f(t,T) = 0$. Thus, from (1) we get

$$S(t,T) = E^Q\left[S(T) \mid F_t\right] = E\left[M(T)S(T) \mid F_t\right] \ \forall \ t \in [0,T] \ , \tag{2}$$

where $M(t)$ is Radon-Nikodym derivative dQ/dP on F_t and E is the expectation operator under the objective measure P. Hence, the forward price is the risk-adjusted expected spot price. The risk adjustment (the Radon-Nikodym derivative $M(t)$) is common for all financial assets in the market, and thus it can be estimated from the derivative prices. Koekebakker and Ollmar (2001) report descriptive statistics of the one week, one year and two year forward prices in Nord Pool during the period 1995-2001. Their results indicate that forward prices have significant volatility, but the price processes do not have a clear trend (see Table 2 in Koekebakker and Ollmar, 2001). Therefore, in this paper we assume that $M(t) = 1$ for all t, i.e. $P = Q$. Thus, we assume that the forward prices are martingales under P. We will postpone the discussion about the situations when the forward prices are not martingales under P until the introduction of the hedging strategy in Section 5.

Electricity forwards with maturities from t to τ give an electricity forward curve $S(t,\cdot):(t,\tau]\rightarrow \mathbf{R}_+$. There are cycles and peaks in the forward curve due to the seasonality in electricity spot price. The electricity spot price is determined by supply and demand. Increase in the spot price attracts high cost producers to produce more which sets a downward price pressure. Conversely, when prices decrease some high cost producers will withdraw capacity temporarily, putting upward pressure on prices. The starting point of the forward curve equals electricity spot price $S(t)$ and we write $S(t) = S(t,t)$. Next, we describe the forward curve dynamics. Similar parameterization has been used, e.g., in Audet et al. (2003) and Koekebakker and Ollmar (2001).

ASSUMPTION 2. *T-maturity forward price follows an Itô process*

$$dS(t,T) = S(t,T)\,\mathrm{e}^{-\alpha(T-t)}\,\sigma(T)\,dB(t,T) \quad \forall\ t\in[0,T],\ T\in[t,\tau] \tag{3}$$

where α *is a strictly positive constant,* $\sigma(\cdot):[0,\tau]\rightarrow R_+$ *is a given bounded spot volatility curve, and* $B(\cdot,T)$ *is a standard Brownian motion on the probability space* (Ω,F,Q)*, along with the standard filtration* $\{F_t : t\in[0,T]\}$. *The correlation structure of the Brownian motions governing the different maturity forwards is given by*

$$dB(t,T_1)\,dB(t,T_2) = e^{-\rho|T_1-T_2|}\,dt \quad \forall\ T_1,T_2\in[0,\tau] \tag{4}$$

where ρ *is a strictly positive constant.*

Assumption 2 implies that the forward volatility is lower than the corresponding spot volatility. Parameter α models the exponential decrease in the forward volatility as a function of maturity. Electricity prices tend to revert toward some equilibrium price due to the high cost producers reactions to the price decreases. The decrease in the forward volatility can be seen as a consequence of this mean-reversion in electricity spot price (see, e.g., Schwartz, 1997 and Clewlow and Strickland, 1999). Electricity forwards whose maturity dates are close to each other are highly correlated. Parameter ρ captures this effect. The correlation decreases exponentially as the difference between maturities increases. Koekebakker and Ollmar (2001) report that the parameterization in Assumption 2 explains 75% of variations in the one week, one year and two year forward data in Nord Pool during the period 1995-2001.

Assumption 2 does not restrict the forward curve as a function of maturity, but it states that the forward prices are lognormally distributed and, therefore we have

$$\log\left(S(T)\right)-\log\left(S(t,T)\right)\sim N\left(-\frac{1}{2}\hat{\sigma}^{2}(T-t),\hat{\sigma}\sqrt{T-t}\right), \tag{5}$$

where $N(m,s)$ is a normal distribution with mean m and standard deviation s. The average volatility of $S(\cdot,T)$ between t and T is given by

$$\hat{\sigma}=\frac{\sigma(T)}{\sqrt{2\alpha(T-t)}}\sqrt{e^{-2\alpha t}-e^{-2\alpha T}}\ . \tag{6}$$

3 Production planning

Hydropower production planning is often broken down into long-term (3 - 5 years), medium-term (1 - 2 years) and short-term (1 week) planning. Basically the forward markets can be used in all these cases. However, often the forward curve has a minimum maturity and duration of one week/day. This means that the forward curve does not give information about the hourly variations in the spot price, which is an essential part of the short-term planning. In this section we make some assumptions that are often relaxed in the short-term planning and thus our method is suitable for medium- and long-term planning. In the bench marking Section 6 we consider production 21 weeks ahead, whereas in the example of Section 7 we use one year maturity. For short-term planning see, e.g., Gröwe-Kuska et al. (2002).

For simplicity, we concentrate on a single reservoir even though hydropower systems often consist of several linked reservoirs. The long- and medium-term scheduling of a multi-reservoir system can be approximated with one aggregated reservoir (see, e.g., Fleten et al., 2002 and Archibald et al., 2001). Thus, in a multi-reservoir case our single reservoir can be used as the aggregated reservoir. Multi-reservoir planning without an aggregated reservoir is studied, for example, in Salinger and Rockafellar (2003) and Keppo (2002).

Next, we will discuss the commonly used inflow estimation methods and give an overall description of the inflow process needed for the production planning. For more discussion about the different inflow estimation methods see e.g. Faber and Stedinger (2001), Tejada-Guibert et al. (1995), Gjelsvik et al. (1992), Røtting and Gjelsvik (1992), and Fosso et al. (1999). If there are derivative instruments that correlate with the inflow process they can be used in the inflow estimation. Unfortunately, there does not usually exist liquid markets for such instruments. The historical inflows provide valuable information for inflow predictions. However, when production less than one year ahead is considered the existing snow reservoirs, for example, provide better information. When the maturity shortens to less than one week, also the weather forecast is relevant for the inflow prediction. Thus, both the data and methods for the inflow prediction change as the planning horizon changes. This means that we will not restrict to any specific parameterization of the inflow process. The following assumption gives an overall characterization of the inflow process.

ASSUMPTION 3. *The expected inflow at time T based on the information* F_t *is* $\nu(t,T)$. *The expected inflow process* $\nu(\cdot,\cdot):[0,\tau]\times[0,\tau]\rightarrow R_+$ *follows right continuous process that is driven by jump uncertainties on the probability space* (Ω,F,P) *along with the standard filtration* $\{F_t : t\in[0,\tau]\}$.

In assumption 3 we assume that the inflow comes in random lumps at random times, however in reality the inflow comes with random flows over random time periods. Thus, we assume that the flow during a discrete time interval $[t,t+\Delta t]$ comes in a lump at time *t.* This simplification is standard in long- and medium-term hydropower planning, and the error is small (see, e.g., Tejada-Guibert et al., 1995 and Røtting and Gjelsvik, 1992). In the example of Section 7 we use weekly granularity for the forward price and inflow estimates, which means that we assume that all the inflow occurring during the upcoming week is available in the beginning of the week. For example, if the predicted inflow of week 52 is 5 Mm^3 we assume that the amount of hydro in the reservoirs increases by 5 Mm^3 on the Monday of week 52.

The inflow estimate $\nu(t,t+y)$ approaches the realized inflow $\nu(t)$ as y goes to zero and we write $\nu(t)=\nu(t,t)$. The realized inflows runs down to a hydro reservoir whose water level at time t is denoted by $x(t)$. The water level is constrained by the size of the reservoir $\bar{x}$ and some lower bound, which we assume to be zero, i.e. $x(t)\in[0,\bar{x}]\ \forall\ t\in[0,\tau]$. Inflow occurring when the reservoir is full, i.e. when $x(t)=\bar{x}$, creates spilling. Spilled water will not be used for electricity production and it will not be available in the future. The hydropower producer tries to maximize expected future cash flows by selecting for all $t\in[0,\tau]$ an optimal discharge strategy $u(t)$. This strategy satisfies $u(t)\in[0,\bar{u}]$ for all $t\in[0,\tau]$. In addition to these constraints the optimal production strategy is constrained by the fact that the water level can not be negative. In other words, the maximum production is the minimum of maximum discharge $\bar{u}$ and the amount of water available, i.e. $u(t)\in\left[0,\bar{u}\wedge\left(x(t)+\nu(t)\right)\right]\forall t\in[0,\tau]$. Thus, the upper boundary of the production strategy is stochastic.

The use of zero lower bound for discharge and water level is not restrictive. We can always rescale the variables so that the lower bound is zero. For example, consider a situation that the weekly production must be on the interval [5; 15] Mm^3 and the water level must be on the interval [100; 300] Mm^3. This means that the maximum weekly production $\bar{u}$ =10 Mm^3 and the size of the reservoir $\bar{x}$ = 200 Mm^3. The constant discharge can be subtracted from the reservoir by changing the inflow process. For example, in this example case the weekly inflow is decreased by 5 Mm^3. The following assumption formalizes the objective function of our profit maximizing producer.

ASSUMPTION 4. For each time $t\in[0,\tau]$ the producer maximizes expected future cash flows

$$V(t,x,S,\nu,u)=E\left[\eta\int_t^\tau e^{-r(s-t)}u(s)S(s)ds\mid F_t\right]\ \forall\ t\in[0,\tau], \tag{7}$$

by selecting production strategy $u(\cdot)$ that satisfies $u(s)\in\left[0,\bar{u}\wedge\left(x(s)+\nu(s)\right)\right]\forall s\in[0,\tau]$. The parameter η denotes the efficiency of the plant. The spot price $S(\cdot)$ and inflow $\nu(\cdot)$ dynamics are characterized by the forward and inflow dynamics in Assumptions 2 and 3. The water level has following dynamics

$$dx(t)=\nu(t)-u(t)-I\left\{x(t)+\nu(t)-u(t)>\bar{x}\right\}\left(x(t)+\nu(t)-u(t)-\bar{x}\right)\ \forall\ t\in[0,\tau] \tag{8}$$

$$x(0)=x_0\,. \tag{9}$$

In equation (7) the producer maximizes expected cash flows by selecting the optimal production strategy. Equation (8) is the balance equation of the reservoir. It states that all new water to the reservoir comes via the inflow process whereas the water can leave the reservoir either due to the discharge or spilling. The third term in the balance equation, i.e. $I\left\{x(t)+v(t)-u(t)>\bar{x}\right\}\left(x(t)+v(t)-u(t)-\bar{x}\right)$, states that water exceeding the size of the reservoir is spilled.

In assumption 4 there are no variable costs in production. We also omit time delays in the production, i.e. changes in the discharge will happen immediately once the decision to change strategy is made. The variable costs in hydropower production are negligible and the response times are often less than an hour. Thus, these assumptions are standard in long- and medium-term planning (see, e.g., Fleten et al., 2002). By using constant efficiency parameter η for the electricity production, we assume that the power generation of the plant is proportional to the discharge. Generally, the efficiency of the power generation varies nonlinearly as a function of water level. In long- and medium-term planning the constant efficiency is traditionally used (see, e.g., Fleten et al., 2002).

Often hydropower production models have also a "legacy" term in the objective function giving the value of the reservoir at the end of the planning period as a function of water level $x(\tau)$ (see, e.g., Fleten et al. 2002 and Mo et al. 2001). Due to technical reasons that will be clarified after Corollary 1 we omit this term in our model. Thus, we assume that the reservoir value after the planning horizon is zero independent of the water level. In reality hydropower producers operate with an indefinite planning horizon, and thus the water in the reservoirs has always some value. The zero value assumption can be made as realistic as possible by selecting the end date τ so that it is logical to have the water level close to some predefined level. For example, in Section 7 we use spring as the end date, because snow melts after that and, therefore the reservoir's water level should be as low as possible after the winter. Other possibility is to use some longer-term planning model to calculate the optimal usage with a larger granularity and then use our model to optimize the production in each period of the long-term planning model. For example, a longer term model can be used to plan the production for next three years with yearly granularity and then our model can be used to calculate the optimal weekly production for each year.

From the objective function and water level dynamics given in Assumption 4 it follows that the optimal production strategy is of the following form.

Corollary 1. The optimal production strategy is

$$\begin{aligned} u\left(K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right) &= I\left\{S(t)\geq K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right\}\left(\bar{u}\wedge\left(x(t)+\nu(t)\right)\right) \\ &+ I\left\{S(t)< K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right),x(t)+\nu(t)\geq\bar{x}\right\}\left(\bar{u}\wedge\left(x(t)+\nu(t)-\bar{x}\right)\right) \qquad \forall t\in\left(0,\tau\right), \end{aligned} \tag{10}$$

where $K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)$ *is a production threshold.* $S(t,\cdot):\left(t,T\right]\rightarrow R_+$ *and* $\nu(t,\cdot):\left(t,T\right]\rightarrow R_+$ *denote the forward and expected inflow dynamics.*

Proof: See Appendix A.

Corollary 1 states that the optimal production strategy is a bang-bang strategy which varies between minimum and maximum discharge. A bang-bang strategy can be characterized by production threshold $K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)$. The production threshold is a function of forward price dynamics $S(t,\cdot):\left(t,T\right]\rightarrow R_+$, expected inflow dynamics $\nu(t,\cdot):\left(t,T\right]\rightarrow R_+$, time *t*, and water level $x(t)$. There are two motivations for production. First, if the price obtained by producing today is higher than the expected price in the future then it is optimal to produce now. In other words, production is optimal if the current spot price is greater than the production threshold. The first term in (10) denotes this case. When water is spilled from the reservoir, it is lost. Thus, whenever spillage is about to happen it is optimal to produce regardless of the price. The second term in (10) models this case. If we had used a non-zero "legacy" function in the object function (7), then the production threshold would also be a function of this function.

Because of the forward and inflow dynamics in Assumptions 2 and 3, it is not possible to solve the production threshold $K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)$ analytically. Instead of solving the threshold analytically, we suggest a simple and intuitive parameterization for it. By maximizing the expected discounted cash flows given by the parameterized production threshold we get a lower bound for the optimal expected discounted cash flows. The parameterized production threshold can be used as an approximation of the true production threshold $K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)$. The quality of the approximation is obtained by comparing the lower and upper bounds of expected discounted cash flows.

4 Threshold estimation

Let us summarize some intuitive characteristics of the production threshold:

1. As the reservoir value at the end of the planning horizon is zero, the production threshold must converge to zero as the time approaches the end of the planning period.
2. The probability of spillage increases as the water level increases, and thus the threshold must decrease as a function of water level.
3. Similarly for the future inflow, if the future inflow estimates increase the threshold decreases.
4. If the forward curve increases, the value of waiting increases. Thus, the production threshold increases as the future electricity prices increase.

The future inflow estimate and the forward curve consist of several random variables and thus there is no exact definition for the increase. We use average future inflow estimate $\tilde{v}(t)$ and average forward curve $\tilde{s}(t)$ to characterize the changes. By summarizing the above mentioned four characteristics, we give the following parameterization for the production threshold.

ASSUMPTION 5: We assume that the production threshold is of the following form

$$\tilde{K}_S(t, x(t), \tilde{v}(t), \tilde{s}(t)) = \alpha_s \tilde{s}(t) \mathrm{e}^{-\alpha_x x(t) - \alpha_v \tilde{v}(t) - \alpha_t \frac{1}{\tau - t}} \quad \forall\ t \in [0, \tau], \tag{11}$$

where $\tilde{v}(t)$ *is the average future inflow estimate*

$$\tilde{v}(t) = \frac{\int_{t_+}^{\tau} v(t,s)\, ds}{\tau - t_+}, \tag{12}$$

and $\tilde{s}(t)$ *is the average future forward curve*

$$\tilde{s}(t) = \frac{\int_{t_+}^{\tau} s(t,s)\, ds}{\tau - t_+}. \tag{13}$$

We assume that the threshold decreases exponentially as the time t approaches the maturity τ. Parameter α_t gives the rate of decrease. Correspondingly, parameters α_x and α_v give the rate of decrease in threshold as a function of water level $x(t)$ and future inflow estimate $\tilde{v}(t)$. The linear decrease in the threshold as a function of average forward curve is given by slope α_s. By using the average future inflow estimate $\tilde{v}(t)$ and average future forward curve $\tilde{s}(t)$, we lose some effects from the forward curve and inflow dynamics on the production threshold. For example, consider two discrete time forward curves having equal means. In the other forward curve all except one forward price are lower than the current price, whereas in the other curve all prices are higher than the current price. The optimal production threshold is different for these forward curves even though the means are equal. Note that, the forward curve and inflow dynamics still affect the actual production decisions and, therefore, water level dynamics (8).

If we had not used the zero value assumption for the reservoir's legacy value then the threshold would decrease towards the marginal value dictated by the legacy function. The marginal value is a function of the water level at the end of the planning horizon and thus non-zero reservoir value after the planning horizon makes the parameterization more complicated.

There is no theoretical ground for the exponential form in (11) and for the use of average inflow and forward price. Thus, it is important to know how close the discounted expected cash flows given with the threshold in Assumption 5 are from the cash flows given with the optimal threshold. We can not calculate the optimal threshold, but we can calculate an upper bound for the expected discounted cash flows. Broadie and Glasserman (1997) suggest simulated trees based method to calculate upper and lower bounds for the value of American-style securities. In their method the upper and lower bound converge to the true value as the amount of nodes used to describe the stochastic variations is increased. Due to the complicated dynamic structure of the hydropower production decisions, we are not able to construct these kind of upper and lower bounds. Our lower bound gives a production strategy that models the optimal production threshold, and our upper bound gives an upper bound for the expected discounted cash flows that can be obtained by improving the lower bound model. Before going to the calculation of the upper bound, let us formalize the lower bound given with the parameterization in Assumption 5.

PROPOSITION 1. *The lower bound of the discounted expected cash flows is given by*

$$H_L\left(t,x,S,\nu\right)=\sup_{(\alpha_x,\alpha_s,\alpha_\nu,\alpha_t)\geq 0} V\left(t,x,S,\nu,u\left(\tilde{K}_S(t,x(t),\tilde{\nu}(t),\tilde{s}(t))\right)\right) \quad \forall\ t\in\left[0,\tau\right], \tag{14}$$

where the electricity and inflow dynamics are given with Assumptions 2 and 3, and the water level dynamics are given with (8). *The parameterized threshold* $\tilde{K}_S(t,x(t),\tilde{\nu}(t),\tilde{s}(t))$ *is given in* (11).

PROOF: Corollary 1 defines that the all the production strategies $u\left(K_S\left(t, x(t), S(t,\cdot), \nu(t,\cdot)\right)\right)$ given with different threshold functions $K_S\left(t, x(t), S(t,\cdot), \nu(t,\cdot)\right)$ are admissible strategies. The parameterized threshold function $\tilde{K}_S(t, x(t), \tilde{\nu}(t), \tilde{s}(t))$ models the threshold function of the optimal production strategy, and thus lower bound is obtained. Q.E.D.

The inflow dynamics are driven by jump uncertainties and the forward curve dynamics by several correlated Wiener processes. When the optimization problem in (14) is solved the stochastic integral in (7) needs to be calculated each time object function value is needed. Due to the forward curve and inflow dynamics the stochastic integral is hard to solve analytically, and thus we will use a numerical integration method. We will use a Monte Carlo simulation based method. Simulation methods for asset pricing were introduced in Boyle (1977). A good survey of Monte Carlo methods in finance is given in Boyle et el. (1997). Lattice and finite-difference methods are not suitable for our model as the size of the grid describing the stochastic variations of price and inflow estimates increase exponentially. This means that also computing time increases exponentially. In Monte Carlo simulations an approximation can be recovered for higher dimensional integrals in a time which does not increase exponentially as a function of stochastic variables. For more discussion about the differences of finite-difference methods and simulation based methods see, e.g., Broadie and Glasserman (1997).

Let us summarize our solution method for (14). First, we select a group of possible parameter combinations. The selection of the possible parameter combinations is done by using intuition about the threshold level with different price, water level, time, and inflow values. Then we calculate the expected discounted cash flows for each parameter combination from Proposition 1. The expected value in (7) is calculated by simulating the forward curve and inflow processes and by selecting at each time step the production strategy given by the parameterized threshold (11). Once the discounted expected cash flows for each parameter combination are calculated, the parameters giving the highest lower bound are used as an approximation for the production threshold.

The highest lower bound does not describe much about the quality of the parameterization. Therefore, we analyze how far this lower bound is from the expected discounted cash flows given by the optimal production threshold. The following proposition gives an upper bound for the expected discounted cash flows. Thus, with Propositions 1 and 2 we get an upper bound for the error, in Assumption 5 in terms of the expected discounted cash flows.

PROPOSITION 2. *The upper bound of the expected discounted cash flows is given by*

$$H_U\left(t,x,S,\nu\right)=\sup_{u(s)\in\left[0,\bar{u}\wedge(x(s)+\nu(s))\right]\forall s\in[t,\tau]} E\left[\eta\int_t^\tau e^{-r(s-t)}u(s)S(s)ds\mid F_\tau\right], \tag{15}$$

where the electricity and inflow dynamics are given with Assumptions 2 and 3, and the water level dynamics are given with (8).

PROOF: This strategy is not any more admissible because the production strategy $u(t)$ is not measurable with respect to $F_\tau\ \forall\, t\in\left[0,\tau\right]$. This means that the optimization is done over a larger set than in the actual case and thus an upper bound for the expected discounted cash flows is obtained. Q.E.D.

The upper bound in Proposition 2 can be numerically estimated by Monte Carlo simulating the stochastic processes and by calculating the optimal solution for each simulated path. Clearly, it is better to know the value of the random process before making a decision than to make a decision with uncertain future. The difference between the upper bound $H_U(\cdot)$ and the actual reservoir value $H(\cdot)$ shows how much one could expect to win if one were told what will happen after the decision. Thus, the difference is called the value of perfect information. The introduced upper bound is called wait-and-see solution in Kall and Wallace (1997). In Broadie and Glasserman (1997) a similar upper bound for the value of an American option is called perfect foresight solution.

The difference in expected cash flows given by Propositions 1 and 2, gives an upper bound for the gains that could be obtained by selecting better parameterization and/or better parameter values. Naturally, if this maximum error is large there is a need to try some new parameter combinations. If this does not help then the parameterization for the threshold should be improved. In our bench marking case, in Section 6, the optimal threshold gives expected cash flows that are on average 2.6% from the upper bound given by Proposition 2. Note that the upper bound for the error gives only the difference in estimated cash flows. We do not have any information how the optimal production strategy differs from the approximated production strategy.

5 Hedging

The general idea of real asset hedging is to find a portfolio of standard derivative instruments that replicate the asset. Shorting this portfolio provides the hedge. To introduce a simple and practical hedging strategy we approximate the optimal production strategy, given in Corollary 1, as follows.

ASSUMPTION 6. When the hedging is considered we assume that the optimal production strategy is

$$\tilde{u}(t) = I\left\{S(t) \geq \tilde{K}_S\left(t, x(t), \tilde{s}(t), \tilde{v}(t)\right), x(t) + d\nu(t) > 0\right\}\bar{u} \quad \forall t \in [0, \tau] \tag{16}$$

where $\tilde{K}_S(\cdot)$ *is the parameterized production threshold given in* (11).

Assumption 6 states that when the hedging is considered the production to prevent spillage is omitted, i.e. the second term in (10) is not included in the hedging strategy. It is never optimal to spill and thus in reality spilling rarely occurs. Equation (16) also states that hedging is done as if the maximal discharge is always used. In reality, when there is less water in the reservoir than the maximal discharge, only the water available can be used. Thus, we are over hedging the production when the water level is less than maximum discharge. Note that even without these simplifications the production can not be hedged perfectly. Due to the uncertainty in the inflow process, there is always some load risk that the producer can not hedge.

The following lemma gives the expected discounted cash flows as a function of forward prices when the production strategy in Assumption 6 is used. Thus, following Lemma gives the amount of forwards that need to be sold to obtain the hedge.

LEMMA 1. When the production strategy in Assumption 6 is used, the expected cash flows are

$$V(t, S, x, \tilde{u}) = \eta\bar{u}\int_t^\tau e^{-r(s-t)} S(t,s) Q_{S(t,s)}\left(S(s) \geq \tilde{K}_S\left(s, x(s), \tilde{s}(s), \tilde{v}(s)\right), x(s) + d\nu(s) > 0\right) ds \quad \forall\ t \in [0, \tau], \tag{17}$$

where $Q_{S(t,s)}\left(S(s) \geq \tilde{K}_S\left(s, x(s), \tilde{s}(s), \tilde{v}(s)\right), x(s) + d\nu(s) > 0\right)$ *is the probability under the forward martingale measure* $Q_{S(t,s)}$ *that the spot price* $S(s)$ *will be greater than parameterized threshold* $\tilde{K}_S\left(s, x(s), \tilde{s}(s), \tilde{v}(s)\right)$, *and that the reservoir is not empty, i.e.* $x(s) + d\nu(s) > 0$.

PROOF: Following applies for the expected cash flows, for all $t \in [0, \tau]$, when the forward price is used as a numeraire (see, e.g., Duffie, 2001 and Björk, 1998) and the production strategy is given by Assumption 6.

$$
\begin{aligned}
V(t,S,x,\tilde{u}) &= \eta\int_t^{\tau} \mathrm{e}^{-r(s-t)}\, E\left[S(s)\tilde{u}(s)\,|\,F_t\right]ds \\
&= \eta\int_t^{\tau} \mathrm{e}^{-r(s-t)}\, S(t,s)E^{Q_{S(t,s)}}\left[\tilde{u}(s)\,|\,F_t\right]ds \\
&= \eta\bar{u}\int_t^{\tau} \mathrm{e}^{-r(s-t)}\, S(t,s)E^{Q_{S(t,s)}}\left[I\left\{S(s)\geq \tilde{K}_S\left(s,x(s),\tilde{s}(s),\tilde{v}(s)\right),x(s)+d\nu(s)>0\right\}|\,F_t\right]ds
\end{aligned}
\quad . \qquad (18)
$$

Taking the expectation of the indicator function gives (17). Q.E.D.

Lemma 1 implies that the expected cash flows can be replicated with forward contracts paying the forward price at maturity. Standard forward contracts pay the difference of current forward price and realized spot price. The forward contracts paying the forward price at maturity can be replicated by the standard forward contracts and risk-free instruments. For instance, in Lemma 1 $\mathrm{e}^{-r(s-t)}S(t,s)$ is equal to the value of one *s*-maturity forward contract plus $\mathrm{e}^{-r(s-t)}S(t,s)$ in the risk free asset at time *t*. Note that the probabilities dictating the forward holdings change constantly and thus the hedging strategy is dynamic, i.e. the amount of forwards that need to be sold change constantly. We will give a numerical example of the hedging strategy in Section 7.

The replication argument in Lemma 1 holds only under the pricing measure *Q*. Thus, if we had not assumed in Section 2 that the pricing measure *Q* is equal to the objective measure *P*, the production decisions in Proposition 1 could not be hedged with Lemma 1 unless the production decisions are done under *Q*. In other words, the cost of having a production strategy which can be hedged is the difference of the objective measure *P* and pricing measure *Q*. Because of this, the Radon-Nikodym derivative *dQ/dP* is often called market price of risk.

6 Bench marking

Driva Kraftverk, a hydropower producer located in the middle of Norway, has given us their production statistics with weekly granularity. Their biggest hydro reservoir, Gjevilvatnet, has a capacity of about 280 Mm^3. In this section we compare the production strategies suggested by our method to the actual production strategies used in Gjevilvatnet during winters 1997-2003.

Let us first discuss the characteristics of hydropower production in Scandinavia and introduce the used data. In Scandinavia the precipitation during winter accumulates as snow and thus the inflow to the reservoirs comes mainly during spring and summer when the snow melts. The start of the smelting varies year to year depending on the temperature. The inflow statistics to the Gjevilvatnet, given in Figure 1, illustrate how both the timing and total amount of inflow varies year to year. In Figure 1 the inflow during winter (weeks 46-16) is always negligible compared to the inflow during other weeks of the year.

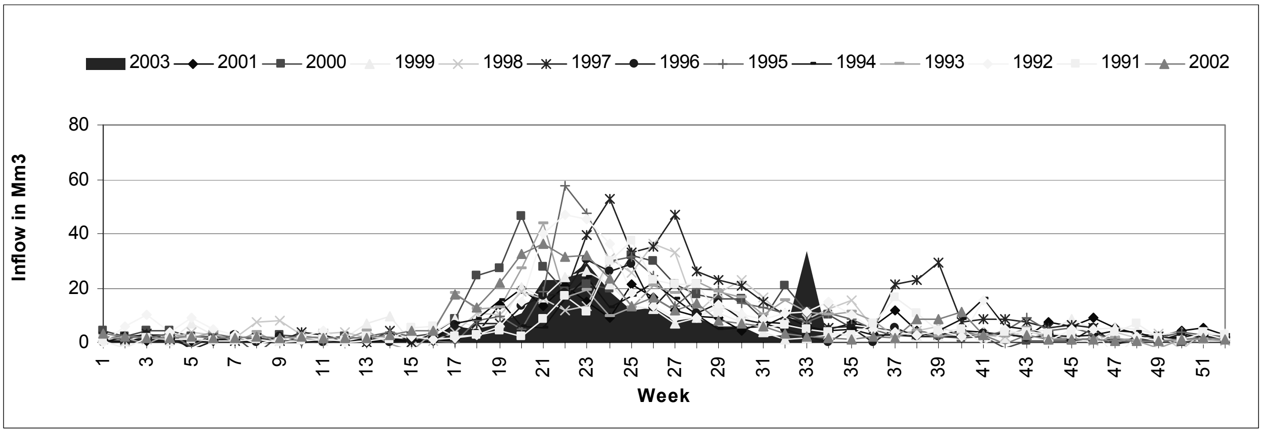

Figure 1: Inflow to the Gjevilvatnet during 1991-2003

Electricity spot and derivative markets in Scandinavia are operated by Nord Pool. The Nord Pool spot market has developed gradually from a national Norwegian power market, which started in 1991, to a joint Scandinavian market. Sweden joined the market 1996, Finland in June 1998, Western Denmark in July 1999, and Eastern Denmark in October 2000. The trading of forward contracts with physical delivery started in 1995. Since 1997 the forward contracts have been settled in cash. The change to cash settlement increased the liquidity of the forward markets. In 1995 the total volume of financial contracts traded on Nord Pool was 40.9 TWh whereas in 2000 it was 1611.6 TWh. Due to the low liquidity of the forwards with physical delivery we use only cash settled forwards, i.e. we use price quotes since 1997. The power generation in Nord Pool area is mixed: Denmark uses 85-90% fossil fuel-based generation and 10-15% wind power. Norway has nearly 100% hydropower production. Sweden and Finland rely on a mix of hydropower, nuclear power, and conventional thermal generation.

Due to the heating and etc., the variations in the weather cause seasonality into the electricity demand. The excess demand and lack of price competitive peak load production facilities makes the electricity prices increase during winter. The Nord Pool's average weekly spot prices during 1997-2003 are illustrated in the upper picture of Figure 2. The prices in Nord Pool are quoted in Norwegian krones, whose exchange rate is about 7 *NOK/USD.* In the lower picture of Figure 2 the total amount of water in the Norwegian hydro reservoirs, as a percentage to the total reservoir capacity, is given. The dashed vertical lines indicate the last weeks of June, whereas the first weeks of January are indicated by solid vertical lines. The summer prices are lower than the winter prices but the timing when the prices start to decrease or increase varies between years. For example, in 1998 and 2001 the cheapest electricity was available in August, whereas in 1999 and 2000 the lowest price was quoted already in July. Similar characteristics can be seen in winter prices. During winter 1998-1999 the prices spiked already before New Year, whereas during winter 2000-2001 the highest electricity price was quoted in February. The water level, in the lower picture, has an opposite seasonality. The hydro reservoirs are emptied during autumn and winter, whereas the filling takes place in spring and summer.

Figure 2 illustrates that changes in yearly price levels can be partly explained by the variations in yearly water levels, i.e. by the yearly rainfalls. For example, in 2000 the prices were considerably lower than in 2001, because the year 2001 was drier than the year 2000. Also the temperature affects the electricity price. For example, the high prices in December 2002 and January 2003 were caused by the combination of extremely cold winter and dry autumn. The highest quote of average weekly price (2nd week of 2003) was as high as 750 *NOK/MWh.* Hence, the period between December 2002 and January 2003 is sometimes called the Scandinavian power crisis. In the lower picture of Figure 2 the dry autumn and cold winter can be seen as a sharply decreasing water level.

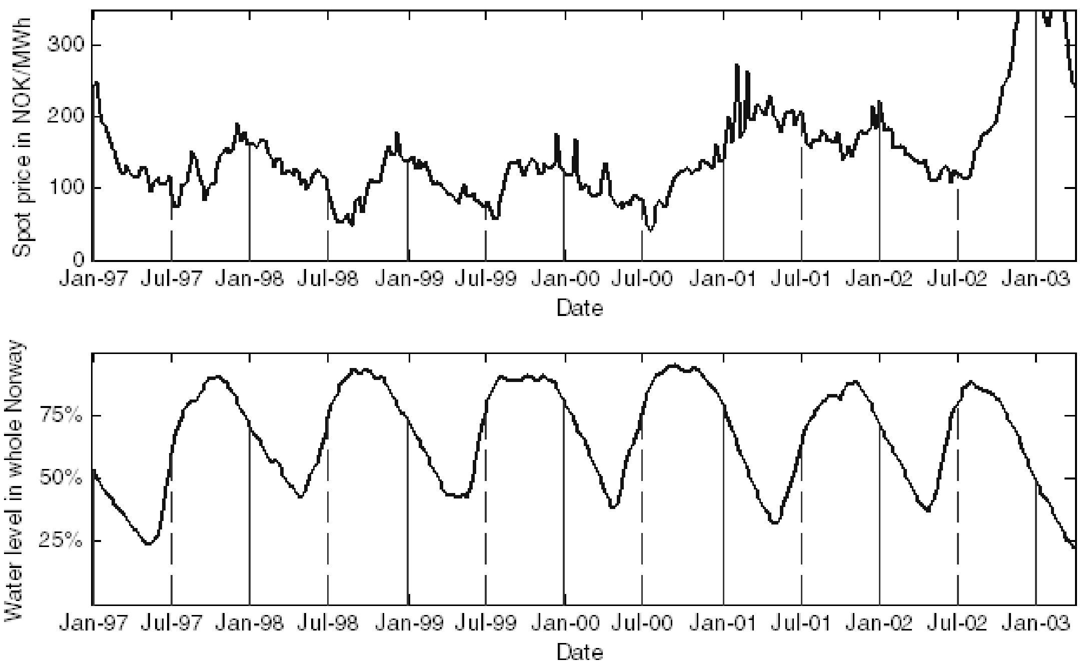

Figure 2: Weekly spot prices and water levels 1997-2003

The high prices in winter and the inflow to the reservoirs coming in spring and summer simplify the hydropower production planning into following: In spring and summer the reservoirs are filled. During the high prices in the winter, the reservoirs should be emptied so that when the smelting starts again, in the spring, the new water can be stored.

The inflow estimates that Driva Kraftverk was using when the production during years 1997-2003 was planned are not available. In this bench marking section we concentrate on the weeks when the inflow to the reservoirs is negligible, i.e. winter weeks 46-14, and assume that the inflow estimate is zero through out the planning period. This is done to minimize the bias caused by the different inflow estimates. An example of our model with a non-zero inflow estimate is given in Section 7. The results in Section 7 will not be compared to any actual production strategies. The main purpose of this section is to bench mark our method to a currently used production method, whereas the main purpose of Section 7 is to illustrate the capabilities of our method.

The area around Gjevilvatnet is a nature conservation area and thus there are some additional limits, varying between seasons, on the water level and discharge. To be sure that our method does not give strategies that could not have been implemented in reality, we assume that the lower and upper bound for the discharge are the minimum and maximum that Driva Kraftverk has had during the production period. In addition, we assume that the size of the reservoir is equal to the total discharge of Driva Kraftverk between 46th and 14th week. For example, during winter 2001-2002 the total discharge was 279.2 Mm^3 and the weekly discharge varied between [5.6; 16.5] $Mm^3/week$. The actual restrictions might have been looser, but these assumptions guarantee that the results of the bench marking are not biased towards our model due to the non-realistic production strategies. The discharge and total production for winters 1997-2003 are summarized in Table 1. The maximum weekly discharge 16.5 $Mm^3/week$ in 2001-2002 gives, when the efficiency of the power station η=1360 MWh/Mm^3, a maximum weekly electricity production of 22 440 MWh.

Table 1: Weekly discharge and total production during winters 1997-2003

Winter	*Discharge* (Mm^3/week)	*Total production* (Mm^3)
97-98	[4.5; 17.5]	282.9
98-99	[0; 16.6]	270.3
99-00	[0.1; 17.2]	292.7
00-01	[5.7; 16.2]	228.9
01-02	[5.6; 16.5]	279.2
02-03	[6.3; 13.4]	206.9

For the forward dynamics we use the price quotes of one week forwards from Nord Pool. In the upper picture of Figure 3, the forward curve for the winter 2001-2002 at the start of the planning period, i.e. at the beginning of week 46, is given with the dotted dashed line. The expected value of the Christmas week 52 is considerably lower than the value of the surrounding weeks, because during Christmas the industry consumption is lower. Similar forward curves are used for each week of each planning period. The grey lines in the upper picture of Figure 3 are the 68% confidence levels calculated with (5). For the forward curves uncertainty structure, spot volatility curve, and parameters α and ρ are needed. The spot volatilities for weeks 46-14 are illustrated in the lower picture of Figure 3. The spot volatilities are calculated from the weekly spot prices during 1995-2003. We use the same spot volatilities for all winters.

Audet et al. (2003) estimate the decreasing forward volatility parameter $\alpha = 4.02$, and the forward price correlation parameter $\rho = 4.51$ from the 1998-2000 forward data. We use their decreasing volatility and correlation parameters for all years. The correlation coefficient ρ gives that forwards whose maturity dates are two weeks from each other have a correlation of 0.85, but forwards whose maturity dates are 10 weeks from each other correlate only with 0.42. The decreasing forward volatility implies that, if the spot volatility is 40% the 4-week maturity forward has volatility of 28%. In Figure 3 the realized forward prices for winter 2001-2002 are given with the dashed line. Due to the cold days at weeks 51 and 52 the prices did not decrease during Christmas as was expected in the beginning of November. On the other hand, as the late winter turned out to be warmer than expected, the prices at the end of the winter (6-14 weeks) were lower than what was expected at the start of the planning period.

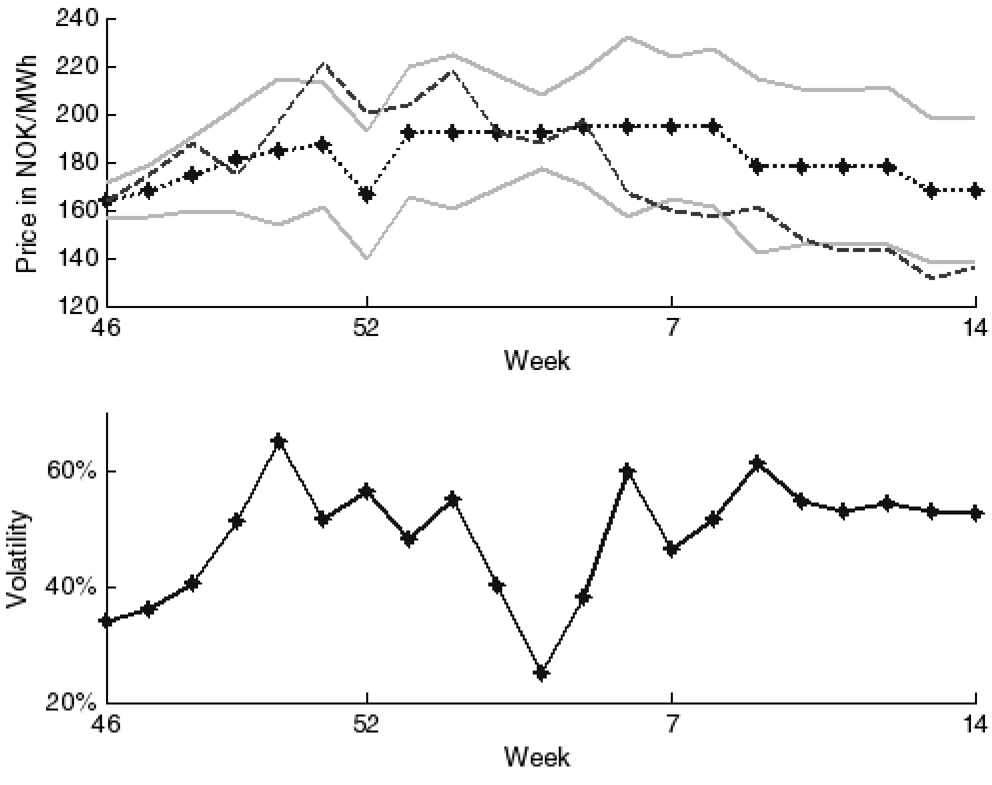

Figure 3: Forward curve and spot volatility curve for winter 2001-2002

As the inflow process is deterministically zero, the production threshold $\tilde{K}_S(\cdot)$ is a function of water level $x(t)$, average future forward curve $\tilde{s}(t)$, and time t. We estimate the threshold parameters α_t, α_x, and α_s for each winter by calculating the expected cash flows for about 150 possible parameter combinations. The expected values were calculated by using 50 000 simulations. The used 150 parameter combinations were chosen from a larger parameter set by using basic intuition about the realistic threshold level. Often the intuition was supported with some, around 1000, simulations. The parameters maximizing the expected cash flows are given in Table 3. The max error, i.e. the last column in Table 3, states how much larger the upper bound of expected cash flows, calculated with Proposition 2, is compared to the expected profits calculated with Proposition 1. For example, for winter 2001-2002 the Proposition 1 gives an expected profit of 47.47 *MNOK*, whereas the upper bound calculated with Proposition 2 is 48.87 *MNOK*. Thus, the optimal production strategy will not improve the average profits by more than 1.4 *MNOK*, i.e. 2.9%. The average of the max errors in Table 3 is 2.6%. Thus, on average our parameterized threshold gives a production strategy whose profits are less than 2.6% from the profits of the optimal production strategy. The parameters and the max error vary year to year due to the different forms of the forward curve. We have tested the statistical significance of α_t and α_x parameters by calculating the average profits when the corresponding parameter is set to zero. In all cases the changes in the average profits were considerable. The least significant was the α_t parameter for winter 1999-2000. When α_t for winter 1999-2000 was set to zero the expected profit decreased with 6.2%.

Table 2: The threshold parameters

Winter	α_S	α_t	α_x	*max error*
97-98	1.1	0.011	0.0010	3.4%
98-99	1.2	0.013	0.0015	2.4%
99-00	1.2	0.010	0.0020	2.9%
00-01	1.2	0.010	0.0015	2.2%
01-02	1.1	0.012	0.0010	2.9%
02-03	1.2	0.013	0.002	1.9%

In Figure 4 the production threshold as a function of time and average future forward curve is illustrated for three different water levels for winter 2001-2002. The lowest plane corresponds to the situation when the reservoir is full, i.e. x= 161.6 Mm^3, in the middle plane half of the available reservoir is used, i.e. x= 80.3 Mm^3, and in the highest plane there is one fourth of the available capacity left, i.e. x= 40.15 Mm^3. In Figure 4 the threshold increases linearly as a function of future forward curve, and decreases exponentially as the end of the planning horizon draws closer.

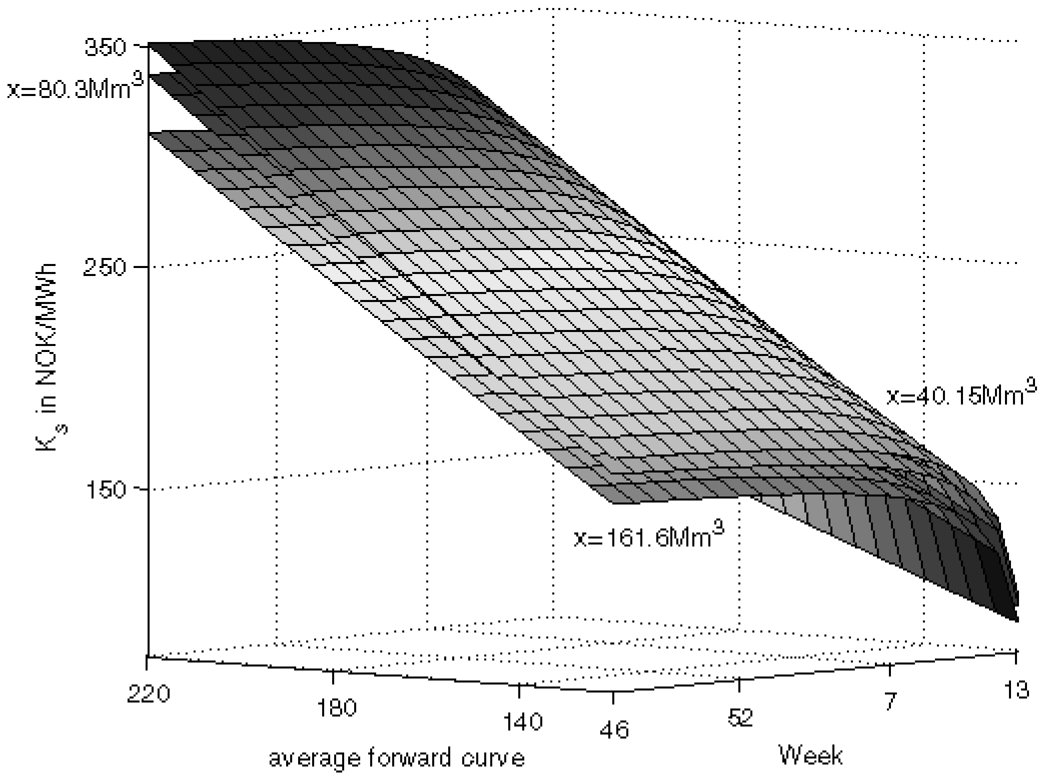

Figure 4: Threshold functions for winter 2001-2002

In the upper picture of Figure 5 the total profits of our method and Driva Kratverk's method are illustrated for winters 1997-2003. The available production capacity, given in Table 1, varies between winters, and thus the total profits are affected by changes in the total production. To remove the variations in the profits caused by the variations in the yearly production we compare the profits per used production capacity. In the lower picture of

Figure 5 the profits divided by total production are illustrated for winters 1997-2003. The variations in the available production capacity affect considerably the total profits. For example, during the winter 2000-2001 the profits per production capacity were on the same level as during winter 2001-2002, however the total profits were about 20% lower. The profits are summarized in Table 3. The last column in Table 3 gives the improvement that Driva Kraftverk would have gained if they had used our method. The average improvement is 10473 *NOK/Mm³*, and the 95% confidence level for the improvement is [5199; 15748] *NOK/Mm³*. The amount of observations is small, i.e. 6, and thus we use the student's t-test to test the statistical significance of the improvement. The t-test gives a p-value of 0.005, which means that there is a 0.5% probability that the improvement is from a distribution whose mean is zero or less. This means that based on the sample in Table 3 improvement gained with our model is statistically significant. In percentage terms our method improved the profits per used production capacity by 4.2%.

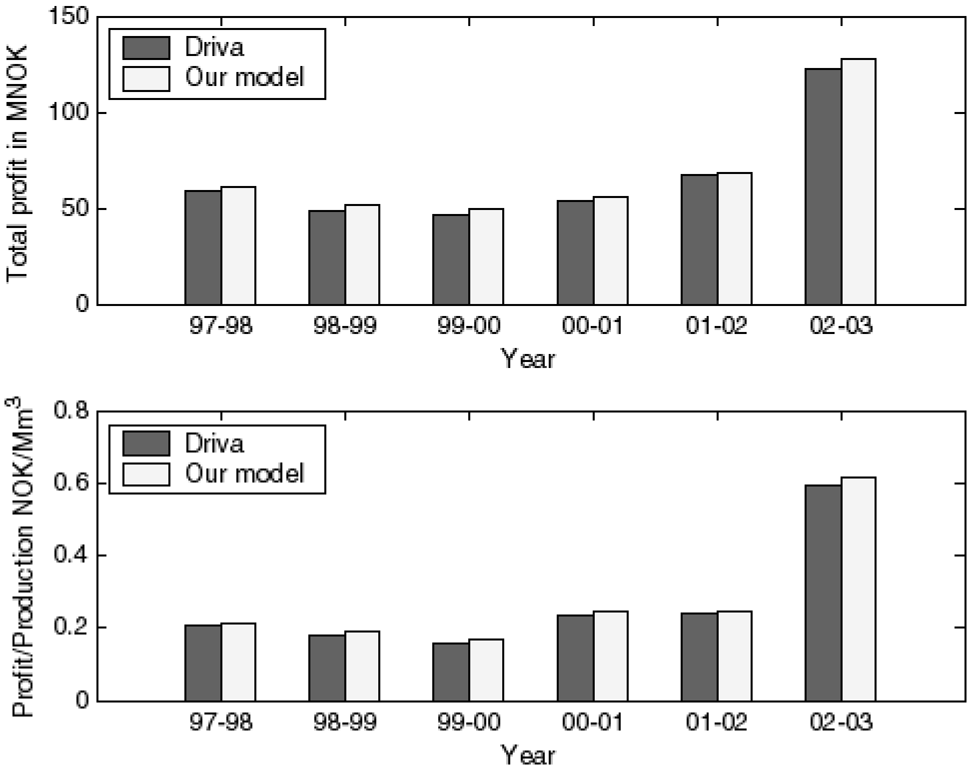

Figure 5: Profits for winters 1997-2003

Table 3: Profits

	Profit (MNOK)		**Profit/Production (MNOK/ Mm^3)**		**Improvement (NOK/ Mm^3)**
Winter	*Driva*	*Our Model*	*Driva*	*Our Model*	
97-98	58.5	60.8	0.207	0.215	8101
98-99	48.5	51.1	0.179	0.189	9600
99-00	46.4	49.1	0.158	0.168	9264
00-01	53.7	56.2	0.235	0.246	10837
01-02	66.8	67.9	0.239	0.243	3865
02-03	122.8	127.2	0.594	0.615	21174

Let us have a more thorough look on the production during winter 2001-2002. In the upper picture of Figure 6 we illustrate the forward curves with dashed lines. The dotted grey lines illustrate the realized prices. In the lower picture of Figure 7 the Driva Kraftverk's production strategies are illustrated with black bars, whereas the grey bars are the production strategies given with our method. The forward curves in the upper picture show how the market's expectations about the 52nd week's price change as the week 52nd gets closer. At the week 51 the estimate increases rapidly as weather forecasts start to predict cold weather for next week. Due to the mild January and early February the market started to expect considerably lower prices for weeks 9-14 already at week 6. In the lower picture we can see how our method reacted to these changes in the forward curve. During week 52 the production took place with maximum capacity even though the forward curve at the start of the planning period, i.e. at week 46, suggested lower production. Our method also used maximum production during the early winter so that most of the production capacity was used before prices started to decrease.

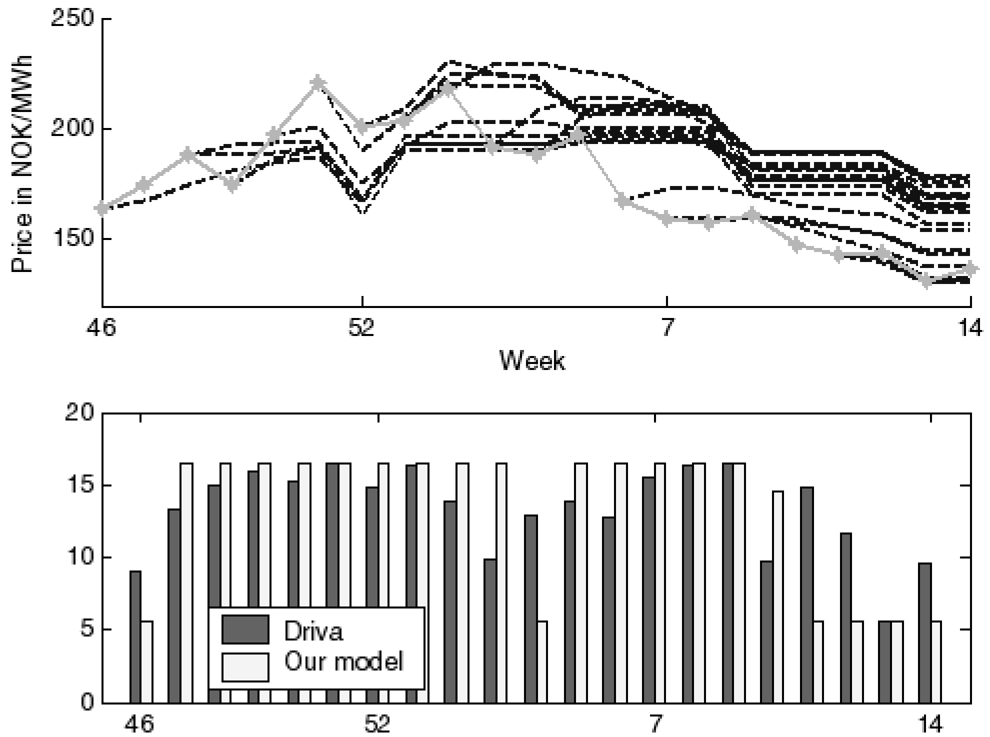

Figure 6: Forward curves and production strategies for winter 2001-2002

The production strategy calculated with our model is a bang-bang strategy, where the production varies between upper and lower bound. In weekly granularity the Driva Kraftverk's production strategy is not a bang-bang strategy. There are two explanations for this. In practice the production is optimized with smaller granularity and thus the Driva Kraftverk's strategy can, for example with hourly granularity, be a bang-bang strategy. For instance, in 2001 at the 47th week Driva Kraftverk has a total production of 13.3 Mm^3, which can be obtained by discharging with full capacity 135 hours out of the possible 168. Due to this also the profits in Figure 5 can be higher or lower than the true earnings. If Driva Kraftverk has managed to pick the hours, when the price is higher than the weekly average, the true earnings are greater and vice versa. Another explanation is that the objective function used in the production planning is such that the optimal production strategy is not a bang-bang strategy. The current industry practice seems to be towards using risk aversive objectives which have different target levels for profits, and where different target levels are progressively penalized (see, e.g., Fleten et al. 2002 and Mo et al. 2001). A nonlinear power production as a function of the water level, i.e. η depending on the water level, also gives production strategies which are not bang-bang strategies.

Our production plan is based on the information on the forward markets. Our bench marking indicates that the profits obtained with our method are better than the profits obtained with Driva Kraftverk's current production planning method. We do not know how much Driva Kraftverk uses information in the derivative markets. The current industry practice seems to be towards modeling future spot price dynamics with large complicated equilibrium models where each producer is considered. However the equilibrium models are used so that the forward markets are not totally omitted in the production planning. It is common to compare the results of the equilibrium model with the forward curve. If the difference is large the initial values of the equilibrium model are changed so that the estimates better match the forward curve. Whatever the Driva Kraftverk's price prediction method is, the results suggest that they could obtain valuable information from the forward markets.

7 An example

In this section we illustrate our model when a non-zero inflow estimate is used. We also show how our method can be used to hedge the future production. We concentrate on the time interval between the 16th week of 2001 and 15th week of 2002. The production planning and hedging are done with weekly granularity.

For each weeks inflow estimate we use a normal distribution fitted into the Gjevilvatnet's weekly inflows, given in Figure 1. The serial correlations in the inflow process are omitted. Note that also negative inflows are possible. For example, in 1996 at week 5 the inflow to Gjevilvatnet was -2.9 Mm^3. The negative inflows can be explained by variations in the measured water levels. For example, wind and/or ice can hold the water so that it can not spread equally around the reservoir. Thus, at some parts of the reservoir the water level can decrease even without discharge. Also evaporation from the reservoirs causes decreasing water levels. The expected future inflow at the start of week 16 is illustrated with a dark dotted line in the upper picture of Figure 7. The gray dashed lines indicate the 68% confidence intervals. As discussed in the previous section the majority of the inflow is assumed to come during spring and summer, i.e. between weeks 17-45.

Nord Pool does not quote forwards for each week of the following year. Thus, the forward curve for the upcoming year needs to be estimated from the existing longer term forwards, i.e. from month and season forwards. For example, Fleten and Lemming (2003) present a method to estimate a weekly forward curve based on longer term forwards and forecasts generated by bottom-up models. In the middle picture of Figure 7 our forward curve at the start of the planning period is presented with a dotted black line. The 68% confidence intervals, illustrated with dashed grey lines, are calculated by using the spot volatilities in the lower picture of Figure 7, volatility parameter $\alpha = 4.02$, and correlation parameter $\rho = 4.51$. The spot volatilities are calculated from the Nord Pool's average weekly spot prices during 1995-2003. The spot volatility is higher in spring than in autumn, because the temperature varies highly in spring. In winter the prices are often "high" whereas in spring the prices can be either "high" or "low".

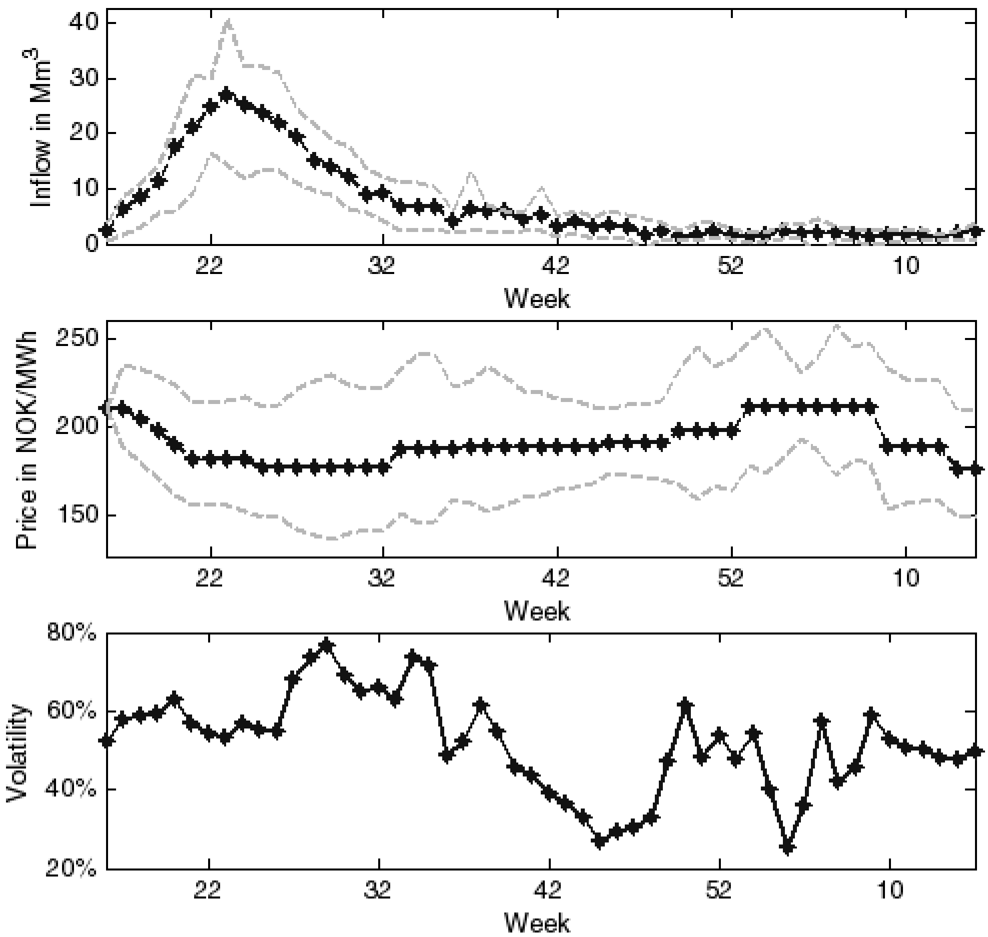

Figure 7: Estimated inflow, forward curve, and spot volatility curve

Let us assume that the production during the planning period must vary between [0; 16.3] $Mm^3/week$, the current water level is 86.6 Mm^3, and the lower and upper bound for the water level are [0; 220] Mm^3. By using following estimates for the threshold parameters: $\alpha_t = 0.011$, $\alpha_x = 0.00075$, $\alpha_s = 1.075$, and $\alpha_v = 0.00075$, we get the production probabilities under forward measure given in Figure 8. The probability of production is lower in summer, because the forward curve decreases during summer. In the early weeks of the planning period the probability of production is high as the reservoir is emptied before the expected inflows. Lemma 1 states that the production probabilities, illustrated in Figure 8, give the amounts of forwards that need to be sold to hedge the production. For example, the probability of production at week 22 is 48.1%, which gives, when the efficiency of the power station $\eta =$ 1360 MWh/Mm^3 that 10663 MWh of forwards need to be sold.

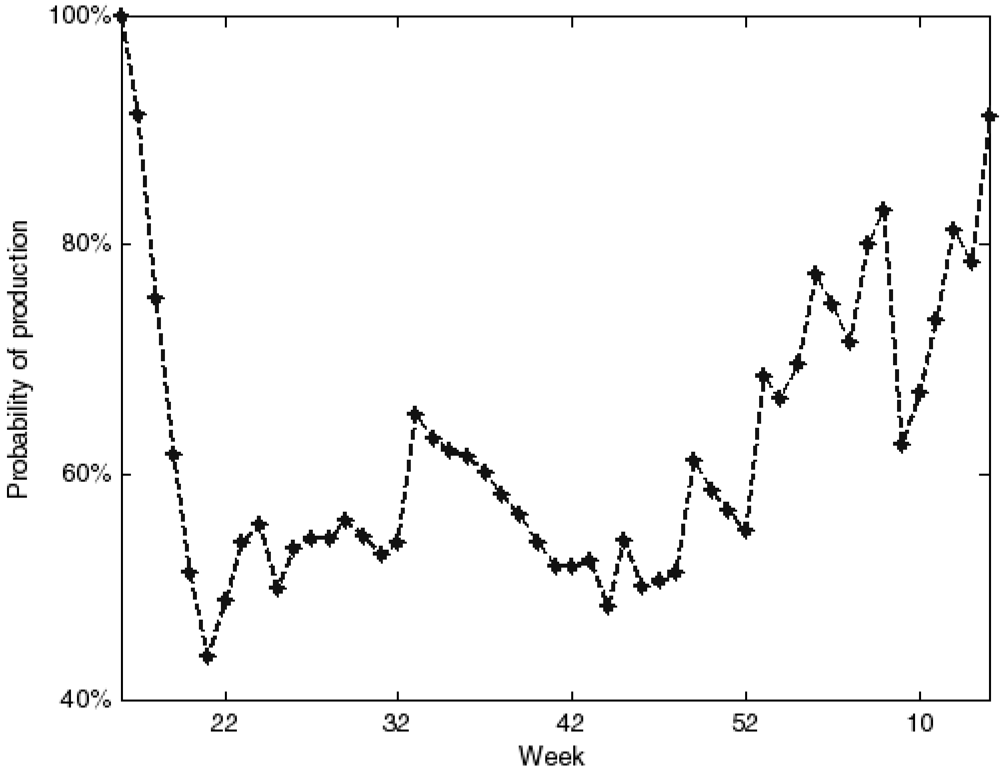

Figure 8: Production probabilities

In Figure 9 we illustrate how our method works during the planning horizon. In the upper picture the weekly production is given with dark bars. The inflow realization is given with white bars and the price realization is given with dashed line in the lower picture of Figure 9. The solid line in the lower picture is the realized threshold. The water level in the reservoir is given in the middle picture of Figure 9. Note how the reservoir is first emptied to make space for the expected future inflow. Once the reservoir is filled the minimum discharge is mainly used until the prices start to increase. When the reservoir is almost full the risk of spillage increases, which can be seen as a decrease in the production threshold. Few times the decrease in the threshold makes it optimal to produce with full capacity during low prices. Once the prices start to increase in the autumn, the production with full capacity will start. When the prices are high and the water level starts to decrease, the threshold first increases, as the forward curve states that there are plenty of weeks with high prices left. Towards the end of the winter the threshold starts to decrease, as the expected future price decreases and the end of the planning horizon draws closer.

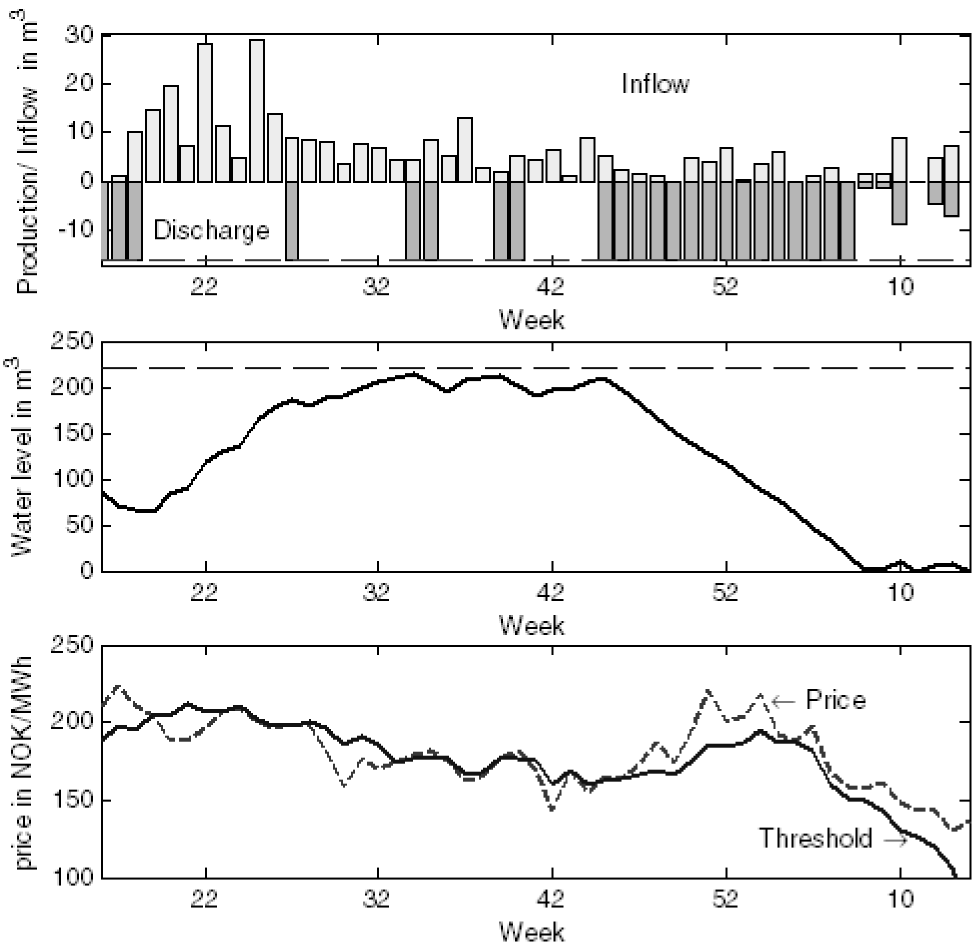

Figure 9: Inflow, water level and price realization

8 Conclusions

We have studied hydropower production planning based on forward curve dynamics. Our model can be used in medium- and long-term planning. We assumed that the hydropower producer maximizes expected cash flows. We also showed how the production strategy can be approximated with standard forward contracts. By selling the forward contracts the production strategy can be partly hedged. Our bench marking illustrates that during winters 1997-2003 our production strategy would have increased the profits of an actual hydropower producer by 4.2%.

Acknowledgements

We would like to thank Stein-Erik Fleten, Audun Botterud, and seminar participants at the Norwegian University of Science and Technology for helpful comments on earlier versions of the paper. Driva Kraftverk is gratefully acknowledged for providing bench marking data.

References

Archibald, T.W., Buchanan C.S., Thomas L.C., McKinnon K.I.M. (2001) Controlling multi-reservoir systems, European Journal of Operational Research 129, 619-626

Audet, N., Heiskanen P., Keppo J., and Vehviläinen I. (2003) Modeling Electricity Forward Curve Dynamics in the Nordic Market, forthcoming in the Modeling Prices in Competitive Electricity Markets, D. W. Bunn (ed.), Wiley Series in Financial Economics.

Björk, T. (1998) Arbitrage theory in continuous time, Oxford University Press, New York

Boyle P.P. (1977) Options: A Monte Carlo approach, Journal of Financial Economics 4, 323-338

Boyle P.P., Broadie M., Glasserman P. (1997) Monte Carlo methods for security pricing, Journal of Economic Dynamics and Control 21, 1267-1321

Broadie M., Glasserman P. (1997) Pricing American-style securities using simulation, Journal of Economic Dynamics and Control 21, 1323-1352

Clewlow, L. and Strickland, C. (1999) Valuing Energy Options in a One Factor Model Fitted to Forward Prices, Quantitative Finance Research Group, Working paper, University of Technology, Sydney, Australia.

Cox, J., and Huang C. -F. (1989) Optimal Consumption and Portfolio Policies When Asset Prices Follow a diffusion Process, Journal of Economic Theory, 49,33-83

Cuoco, D. (1997) Optimal Consumption and Equilibrium Prices with Portfolio Constrains and Stochastic Income, Journal of Economic Theory 72, 33-73.

Cvitanic, J., Schachermayer W., and Wang H. (1997) Utility Maximization in Incomplete Markets with Random Endowment, Finance and Stochastics 5, 259-272.

Duffie, D. (2001) Dynamic Asset Pricing Theory, 3rd edition, Princeton University Press.

Faber, B.A., Stedinger J.R. (2001) Reservoir optimization using sampling SDP with ensemble streamflow prediction (ESP) forecasts, Journal of Hydrology 249, pp. 113-133

Fleten, S. E., Lemming J. (2003) Constructing forward price curves in electricity markets, Energy Economics 25, pp. 409-424

Fleten, S. E., Wallace S. W., Ziemba, T. (2002) Hedging electricity portfolios via stochastic programming, Decision Making Under Uncertainty: Energy and Power, pp. 71-93, Claude Greengard and Andrzei Ruszczynski (eds), Springer-Verlag, New York

Fosso, O. B., Gjelsvik, A., Haugstad, A., Mo, B., Wangensteen, I. (1999) Generation scheduling in a deregulated system. The Norwegian case. IEEE Transactions on Power Systems 14 (1), 75-80

Gjelsvik, A., Røtting, T. A., Røynstrand, J. (1992) Long-term scheduling of hydro-thermal power systems, Broch, E., Lysne, D. K. (Eds), Proceedings of the 2nd International Conference on Hydropower Development (Hydropower '92), Lillehammer, Norway. Balkema, Rotterdam, pp.539-546

Gjelsvik, A., Wallace S. W. (1996) Methods for stochastic medium-term scheduling in hydro-dominated power systems. EFI TR A 4438, Norwegian Electric Power Research Institute, Trondheim

Gröwe-Kuska, N., Kiwiel, K.C., Nowak, M.P., Römisch W., Wegner, I. (2002) Power management in a hydro-thermal system under uncertainty by Lagrangian relaxation, Decision Making under Uncertainty: Energy and Power (C. Greengard, A. Ruszczynski eds.), IMA Volumes in Mathematics and its Applications Vol. 128, Springer-Verlag, New York 39-70

International Energy Agency (IEA) (2003) IEA statistics 2003, Renewables Information, Paris, France

Jaillet, P., Ronn E., Tompaidis S. (2001) Valuation of Commodity based Swing Options, Working paper, UT Austin

Kall, P., Wallace S.W. (1997) Stochastic Programming, Wiley, Chichester

Keppo, J. (2002) Optimality with Hydropower System, IEEE Transactions on Power Systems 3, 583-589

Keppo, J. (2004) Pricing of Electricity swing options, forthcoming in the Journal of Derivatives

Karatzas, I, Lehocky J., Shreve S. (1987) Optimal Portfolio and Consumption Decisions for a “Small Investor” on a finite Horizon, SIAM Journal of Control and Optimization,25 , 1157-1186

Karatzas, I., J. Lehoczky, S. Shreve, and G.-L. Xu (1991) Martingale and Duality Methods for Utility Maximization in an Incomplete Market, Siam Journal of Control and Optimization 29, 702-730.

Koekebakker, S., F. Ollmar (2001) Forward Curve Dynamics in the Nordic Electricity Market, Preprint, Norwegian School of Economics and Business Administration.

Merton, R. (1969) Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case, Review of Economics and Statistics, 51, 247-257

Merton, R. (1971) Optimum Consumption and Portfolio Rules in a Continuous Time Model, Journal of Economic Theory, 3, 373-413

Mo B., Gjelsvik A., Grundt A. (2001) Integrated Risk Management of Hydro Power Scheduling and Contract Management, IEEE Transactions on Power Systems 16, 2,216-221

Pennanen T., King A. (2004) Arbitrage pricing of American contingent claims in incomplete markets- a convex optimization approach, Working Paper, Helsinki School of Economics

Pereira M.V.F. (1989) Optimal stochastic operations scheduling of large hydroelectric systems, International Journal of Electrical Power & Energy Systems 11, pp. 161-169

Pereira M.V.F., Pinto L. M. V. G. (1991) Multi-stage stochastic optimization applied to energy planning, Mathematical Programming 52(2), 359-375

Røtting, T. A., Gjelsvik, A. (1992) Stochastic dual dynamic programming for seasonal scheduling in the Norwegian power system, IEEE Transactions on Power Systems, Vol. 7 no. 1, pp.273-279

Salinger, D.H., Rockafellar R. T. (2003) Dynamic splitting: an algorithm for deterministic and stochastic multiperiod optimization, Working Paper

Schwartz, E.S. (1997) The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging, The Journal of Finance, 52 (3), pp. 923-973

Tejada-Guibert, Johnson, S. A., Stedinger, J. R. (1995) The value of hydrologic information in stochastic dynamic programming models of a multireservoir systems, Water Resources Research 31, 2571-2579

Thompson, A.C. (1995) Valuation of Path-Dependent Contingent Claims with Multiple Exercise Decisions over Time: The case of Take-or-Pay, Journal of Financial and Quantitative Analysis, 30(2), 271-293

Wallace, S. W. and Fleten S. E. (2003) Stochastic programming models in energy, Stochastic programming, vol. 10 of Handbooks in Operations Research and Management Science, A. Ruszczynski and A. Shapiro(eds), North-Holland, pp. 637-677

Vehviläinen, I. and J. Keppo (2003), Managing Electricity Market Price Risk, European Journal of Operational Research, 145, 136-147.

Appendix A

Due to the applied nature of the paper we omit the formal proof and give just an intuitive proof. As the total cash flow in (7) is linear function of the production strategy $u(t)$ for all $t \in [0,T]$ the optimal discharge strategy is a bang-bang control. A bang-bang control can be characterized with a threshold above which production is optimal and below which waiting is optimal. Thus, the production strategy omitting the spillage can be written as

$$u\left(K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right) = I\left\{S(t) \geq K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right\}\left(\bar{u} \wedge \left(x(t)+\nu(t)\right)\right), \qquad \text{(A1)}$$

where $K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)$ is the production threshold. The production threshold is a function of all the stochastic variables, i.e. forward curve $S(t,\cdot):(t,T] \to R_+$, inflow estimate $\nu(t,\cdot):(t,T] \to R_+$, time t, and water level $x(t)$. The water level is a state variable summarizing the past realizations of discharge strategy and inflow. The time dependence follows from the finite planning horizon. As mentioned earlier, the spillage is about to happen when the water level is about to rise above the size of the reservoir, i.e. when $x(t)+\nu(t) \geq \bar{x}$. In this case, the optimal production strategy is to prevent spillage by decreasing the water level to the upper bound of the reservoir or to discharge as much as possible, i.e. $u = \bar{u} \wedge \left(x(t)+d\nu(t)-\bar{x}\right)$. Combining this with production strategy in (A1) gives

$$u\left(K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right)=I\left\{S(t)\geq K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right)\right\}\left(\overline{u}\wedge\left(x(t)+\nu(t)\right)\right)$$
$$+I\left\{S(t)<K_S\left(t,x(t),S(t,\cdot),\nu(t,\cdot)\right),x(t)+\nu(t)\geq\overline{x}\right\}\left(\overline{u}\wedge\left(x(t)+\nu(t)-\overline{x}\right)\right)\ \ \forall t\in\left(0,\tau\right). \qquad \text{(A2)}$$

Q.E.D.