



Norwegian University of
Science and Technology

PROJECT THESIS

**Hydropower reservoir management using multi-factor price
model and correlation between price and local inflow**

AUTHORS:

Joakim Dimoski & Sveinung Nersten

SUPERVISOR:

Stein-Erik Fleten

DEPARTMENT OF INDUSTRIAL ECONOMICS AND TECHNOLOGY MANAGEMENT (IØT)
NORWEGIAN UNIVERSITY OF SCIENCE AND TECHNOLOGY
TRONDHEIM
FALL 2017

Preface

This report is written as a specialization project within the field of Financial Engineering at the Department of Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU).

We would like to acknowledge and express our gratitude to the persons who have made the completion of this project thesis possible. First and foremost we are grateful to our supervisor Professor Stein-Erik Fleten, for his enthusiasm, assistance and helpful discussions throughout the entire semester. Further, we wish to thank Dr. Nils Löhndorf for his valuable help in model implementation and for providing us with QUASAR, a general-purpose solver for large-scale stochastic optimization. Next, we want to thank Gunnar Aronsen at TrønderEnergi, for providing us with data and for sharing his extensive knowledge within the field of hydropower optimization. At last, we would like to thank our families for their continuous support.

Trondheim, December 19, 2017

Joakim Dimoski

Sveinung Nersten

Abstract

Hydropower producers with reservoir capacity have a special challenge when it comes to weighing the short-term profit from selling power in the day-ahead spot market against waiting for better electricity prices. In this paper, we propose a medium-term scheduling model for a price-taking hydropower producer, using a horizon of two years. The paper contains several contributions to the field of reservoir management. First, we use the price of forward contracts to forecast future spot prices, and use multiple factors to describe movements in price. Second, we include a short-term correlation between movements in electricity price and local inflow. Third, we compare the performance of our scheduling model to a model in which price and local inflow are assumed to be independent and a model in which price movements are described using only one factor. We quantify the loss in expected revenues of using the latter two models compared to ours when price movements are in fact driven by multiple factors and correlated with local inflow. In both cases, we find the loss to be approximately 2-3 %. Lastly, we propose a new method for hedging hydropower producers against the risk associated with price uncertainty, introducing the delta of a hydropower plant.

Contents

1	Introduction	1
2	The environment of a hydropower production planner	4
2.1	Nordic electricity market	4
2.2	Hydropower optimization model	6
3	Price and inflow dynamics	9
3.1	Overview of models for electricity spot price movements	9
3.2	Constructing forward curves and volatility functions	12
3.3	Inflow model	19
3.4	Price and inflow lattice	20
3.5	Optimization algorithm used on lattice	24
4	Case: Sjøa hydropower plant	27
4.1	Sjøa hydropower plant	27
4.2	Revised optimization model	28
5	Results	31
5.1	Realized forward curves and corresponding volatility curves	31
5.2	Inflow model parameters	41
5.3	Monte Carlo simulations and lattice construction	44
5.4	Different starting dates and expected incremental water values	46
5.5	Sensitivity analysis using different correlation values and number of volatility functions	48
5.6	Backtesting the model policy with realized price and inflow data	51
5.7	Comparison of results using different volatility functions	53
5.8	Delta hedging	54
5.9	Further work	55
6	Conclusion	58
	References	62
	Nomenclature	65

1 Introduction

The decision problem of hydropower producers, which seek to dispatch the water in their reservoir optimally, has existed for many years. Massé (1946) argues that deterministic models are not good enough to solve such problems, as they do not incorporate the flexibility a production planner has when it comes to the timing of production. While deterministic models can provide the production planner with an optimal decision for each time step (considering discrete time decisions), a flexible model can obtain optimal decision policies based on the current state of the world and its uncertain future. Even though the text is old, the models and discussions presented by Massé (1946) are relevant for how reservoir management is done today.

Historically, power markets have been highly regulated with strong central planning. This is to ensure the reliability of the power distribution system, as a failure to match production and demand can result in a (partial) system break down. Therefore, as explained in Wallace and Fleten (2003), multiple publications on optimization models for reservoir management focus on central planning in hydrothermal systems and are mainly based on uncertainty in demand and area inflow. Such models aim to minimize total costs related to production, distribution and potential breakdowns. An example of this can be found in Rebennack (2015), which seeks to optimize production in the entire Panama power system. In recent time, some power markets, e.g., the Nordic, have experienced deregulation on the production and retailing side, allowing anyone with a power production plant to participate in the market. As Fosso et al. (1999) discuss, the decision problem in such markets becomes altered, as the focus is changed from minimizing system costs to a situation where all power producers seek to maximize their profits.

In this paper, we aim to create a scheduling model for a price-taking hydropower production planner that participates in such a deregulated market. The planner operates a plant that is assumed to be sufficiently small so that the decisions of the production planner do not affect the market as a whole. We also assume that the production planner only participates in the spot market. Further, the model contains two stochastic variables (spot price and inflow) and is meant for medium-term planning using a time horizon of two years and weekly granularity. This is in compliance with multiple current models for medium-term reservoir management, as described in Iladis et al. (2008), Wolfgang et al. (2009) and Abgottspon and Andersson (2014). As opposed to existing models for reservoir management, which often use single-factor processes to describe movements in price, we use multiple factors to capture the dynamics of the prices better. Also, we include a correlation coefficient between changes in price and local inflow, thereby treating them as dependent variables. In Section 5.5, we quantify the loss in expected revenues a single hydropower producer will experience if their scheduling model treats price and local inflow as independent when they are in reality correlated, or if the model describes price movements with one factor when they are in reality described by multiple factors. We also suggest a new method for risk hedging in hydropower production, treating the value of the production as an exotic option and calculating its delta.

To obtain optimal decision policies in each discrete state for the production planner, one can use stochastic dynamic programming. The term was first introduced by Bellman (1957), and it combines the problems of stochastic and dynamic programming. As the words suggest, stochastic programming revolves around problems with uncertainty in future state variables, while dynamic programming is a method of solving a large, complex problem by splitting it into smaller subproblems. In the case of hydropower, this is done by breaking down the overall problem of maximizing long-term profits into maximizing the profits earned in each state based on the decisions in that particular state and the expected future state space. Mathematically, Bellman (1957) summarizes this in the Bellman equation. Further, Pereira and Pinto (1991)

introduce an algorithm for stochastic dynamic programming based on a stochastic extension of Benders decomposition, a solution approach known as stochastic dual dynamic programming (SDDP). Rebennack (2015) extends the method of SDDP so that it can also be used on scenario trees, and Löhndorf and Wozabal (2017) introduce a framework that integrates SDDP with scenario lattices, referred to as approximate dual dynamic programming (ADDP). Scenario lattices are recombining scenario trees, and are useful for representing all states of the world and their transition probabilities when working with Markov decision processes (MDP). Multiple authors, e.g., Lamond and Boukhtouta (1996), show that it is reasonable to treat hydropower reservoir management problems as MDPs, an approach we adopt in this paper.

In order to model the problem as an MDP, all state variables must follow Markov processes. We show that it is reasonable to model movements in both price and inflow this way, and therefore use a scenario lattice to discretize all future states and transition probabilities. To construct the lattice, we use the method presented by Löhndorf and Wozabal (2017). We also use their method for solving stochastic dynamic programs, ADDP, to obtain all optimal decision policies. Since we use weekly granularity, we must construct a lattice with weekly time steps.

As previously mentioned, our model incorporates two stochastic state variables; spot price and inflow. EOPS (SINTEF (2017a)), which is one of the most common commercial software for medium-term reservoir management for local systems or single plants in the Nordic countries, uses spot price scenarios generated using EMPS (SINTEF (2017b)). EMPS is a fundamental model which, among others, can forecast spot prices by using historical scenarios of stochastic variables like inflow and demand in large, international power systems (Wolfgang et al. (2009)). Mo et al. (2001) show that there is a high correlation between the price of one week and the price of the previous week forecasted using EMPS. Therefore, until 2000, EOPS used an AR(1) process (a single-factor model) to describe the price movements found by EMPS, illustrated in Flatabø et al. (1998). As shown in Mo et al. (2000), EOPS also uses EMPS to generate price scenarios today, but the prices are now organized in a Markov chain using the scenarios directly instead of expressing them with an autoregressive model.

To distinguish ourselves from EOPS, we choose to model current and future spot prices with a multi-factor model based on the price of forward contracts traded in the market. Clewlow and Strickland (2000), Koekebakker and Ollmar (2005) and Bjerksund et al. (2008) argue that one-factor models such as AR(1) are unrealistic for accurately modeling forward and spot price movements. Instead, they propose using multi-factor models, which according to them give a much more realistic representation of the dynamics behind price movements. There exist several ways one can obtain the coefficients of such models. Koekebakker and Ollmar (2005) propose that they can be found empirically by first constructing smooth forward curves for many consecutive trading days and then by calculating the daily deviations between the curves. In this paper, we adopt a modified version of this approach and construct forward curves using three different methods, obtaining three sets of model coefficients. The methods used are presented in Benth et al. (2007), Fleten and Lemming (2003) and Alexander (2008a). A comparison of the first two methods can be found in Kiesel et al. (2017).

The other stochastic variable considered in our model is inflow. When determining the characteristics of the inflow, there are several questions that must be answered - whether the system is a local or a regional system consisting of a number of power plants, if there is a seasonal pattern to the inflow, if there are rain periods or snow melting periods, and the choice of temporal resolution of the inflow measurements. For inflow, there is often, depending on the time resolution, a significant degree of autocorrelation from one period to the next. E.g., after a period of precipitation or snow melting, one is likely to experience consecutive days and weeks of increased

inflow. A significant degree of autocorrelation favors autoregressive models. An example of an autoregressive model for inflow is the AR(1) model proposed by Rebennack (2015). In EOPS, the inflow for a local system is modeled using an ARIMA(1,1) process. At last, we have the geometric periodic autoregressive (GPAR) model presented by Shapiro et al. (2013), in which the deviations of the log inflows from their periodic mean are represented as an AR(1) process. In this paper, we adopt the GPAR model and use it to describe local inflow.

For hydropower dominated systems, multiple papers show that there exists a general a negative correlation between inflow and the electricity price, e.g., Mo et al. (2001). Naturally, when reservoir levels are low, prices increase as a result of lower supply. The nature and strength of this correlation will depend on several factors. Among these is the choice of time resolution, and whether we look at local or system-level inflow. All else equal, one will expect the strength of the relationship to be stronger for a coarse granularity of time (e.g., quarterly or yearly data), as the impact on the supply will be more substantial for inflow aggregated over a longer time.

The inflow-price relationship is in varying degree taken into account in the literature and commercial software. Kolsrud and Prokosch (2010) found a relationship between the spot price, the overall aggregated reservoir level in a given geographical area and the local reservoir level of a single plant. Further, EMPS, which is used to find spot price scenarios to be used in EOPS, finds spot price as a function of aggregated, regional inflow. On the contrary, Fosso et al. (1999) show that EOPS treats movements in local inflow and price as independent variables. Intuitively, we would expect the correlation between price and inflow to be stronger on an aggregated national level, than between the local inflow of a particular power plant and the system price. However, we do not expect local inflow and price to be independent. This is because the local inflow can be heavily correlated with the aggregated national inflow, as Boger and Vestbøstad (2016) and Seim and Thorsnes (2007) found. Therefore, we include a correlation between movements in local inflow and the price of forward contracts in our final model.

Lastly, inspired by Wallace and Fleten (2003) and Berger et al. (2016), we have looked at a method for how a hydropower production planner can hedge their position by trading in the forwards and options market. Using principles from financial theory, we introduce a way to calculate the delta of a hydropower plant with the forward curve as the underlying. We further discuss how this delta can be interpreted, and possible strategies for how a company operating multiple plants can delta-hedge their position, or at least reduce their vulnerability to fluctuations in the underlying forward curve.

In Section 2, we provide the reader with a short introduction of the most relevant aspects of the Nordic electricity market and present a general optimization model for reservoir management. In Section 3, we present the relevant theoretical background, including the stochastic models for price and inflow and how their coefficients are estimated. We also present the algorithm for discretizing all states into a scenario lattice and give a short overview of the optimization algorithm used to obtain all optimal decision policies. Further, in Section 4, we introduce a two-reservoir hydropower plant in Norway on which we have tested our model. Additionally, we also revise the general optimization model to fit this specific case. In Section 5, we present some empirical results, both regarding the obtained model coefficients and from running the reservoir management model itself. We test the impact of using a multi-factor model for price movements and a short-term correlation between price and inflow. In addition, we also backtest the model against historical realizations of price and inflow, discuss some hedging strategies and propose some further work that should be conducted with the model or in the field. Conclusions are made in Section 6.

2 The environment of a hydropower production planner

In this section, we give a short introduction of the main features of the Nordic electricity market. We especially emphasize the different products traded on the different physical and financial markets, and what role a hydropower producer takes in these markets. We also present a general optimization model for hydropower reservoir management, incorporating the most important constraints for a general case with multiple interconnected reservoirs and turbines.

2.1 Nordic electricity market

The material presented in this section is based mainly on NordPool (2017) and NASDAQ (2017). The companies that make up the Nordic electrical utility industry fall into one out of three categories - producers, grid companies, and retailers. Roughly speaking, the task of the producers is to produce electrical energy. The task of the retailers is to buy energy from the producers and sell it to the end users, which can both be households and business consumers. The grid companies connect the producers and the end users by transmitting and distributing the energy in the power grid. Since the focus of this paper is hydropower production, the emphasis will be put on the former category. While the grid companies are highly regulated monopolists, the markets for producers and retailers have been deregulated since 1991.

The primary trading place for a hydropower producer is the Nord Pool day-ahead market. The day-ahead market is an auction where power is traded for delivery each hour the next day. Despite the fact that it is not technically a spot market, as this would imply immediate delivery of power, it is referred to as the elspot or the electricity spot market. As shown in Figure 1, the day-ahead market is divided into several bidding areas with separate spot prices known as *area prices*. The area price is decided by the demand-supply equilibrium established in that particular area. Nord Pool also calculates the *system price*, which is an unconstrained market clearing reference price for the entire Nord Pool exchange. All orders from Nordic and Baltic regions are included in the system price calculation. Nord Pool also operates other markets in which the electricity producer can participate - the intraday market and the regulating power market.




Figure 1: Electricity price areas in the Nordic and Baltic countries. Source: NordPool (2017)

A hydropower producer can also participate in the trading of forward and futures contracts. The electricity spot price is subject to fluctuations, and hedging of the spot price risk is an important motivation for participating in the forward market. Additionally, some participants want to engage in speculative trading or market-making. The forward market also serves another purpose - it gives information about how the market participants are valuing the delivery of power in the future. Thus, the forward market can be used to determine the market expectation for the spot price, which is useful information for a profit-maximizing electricity producer.

Nordic power futures are traded on NASDAQ OMX Commodities Europe (hereby NASDAQ Commodities), the financial energy market for the Nordic countries. The market is purely *financial* in the sense that no physical energy is exchanged - only cash. This is in contrast to the elspot, which is a *physical* market. Unlike forward and futures contracts traded in equity markets, whose underlying is typically a single unit of an asset, the power futures traded at NASDAQ Commodities are for specified delivery periods (e.g. for a given day, week or month). Thus, an actor that goes short in a power futures contract is obliged to deliver a specified amount of power over the entire delivery period, earning the price specified by the contract. There are also options traded on NASDAQ Commodities, including both Asian-style options with the system price as the underlying instrument and options whose underlying is the price of Nordic power futures contracts with, e.g., quarterly and yearly delivery periods. The underlying price (reference price) for the Nordic power futures is the Nordic system price, discussed above. Further, the payoff of a long position of one MWh in a Nordic power futures contract is the average system price in the delivery period less the agreed-upon forward price. Generally, no trading occurs after the delivery period has begun. Both regular futures and deferred settlement (DS) futures are traded on NASDAQ Commodities. They have identical characteristics, apart from the settlement structure. DS futures have daily cash settlement during the delivery period, but not before. For regular futures, daily cash settlement also takes place throughout the trading period leading up to the delivery period. In general, the same product types are offered in both futures and DS futures format. The different product types are shown in Table 1.

Table 1: Nordic power futures traded at NASDAQ Commodities. The codes apply for DS futures. Regular Nordic power futures have the product code FUTBL in addition to the DS futures code.

Code	Length of delivery period	Trading begins
ENOWww-yy	Week	6 weeks ahead
ENOMmmm-yy	Month	7 months ahead
ENOQq-yy	Quarter	8 quarters ahead
ENOYR-yy	Year	10 years ahead

An explanation on the difference between a forward contract and a futures contract is due. McDonald (2014) defines a forward contract as an agreement that sets today the terms - including price and quantity - at which you buy or sell an asset or commodity at a specific time in the future. Further, he defines a futures contract as an agreement that is similar to a forward contract except that the buyer and the seller post margin, and the contract is marked-to-market periodically. Typically, futures contracts are exchange-traded while the forward contracts are traded over-the-counter. However, DS futures have historically been called forward contracts despite that they are exchange-traded (Koekebakker and Ollmar, 2005). Notice that the contracts traded on NASDAQ Commodities are different from a standard forward in the sense that they have a delivery period instead of a specific maturity time. Thus, one could argue that a more precise term would be to call them electricity swaps and that neither forwards or futures are the correct term. In this paper, we choose to treat forward, futures, DS futures and swap contracts as equal irrespective of their difference in settlement structure. This is in

line with existing literature such as Clewlow and Strickland (2000), Benth et al. (2008) and Boger and Vestbøstad (2016). In this thesis, we, therefore, use the term *forward* to refer to all contracts representing an agreement to sell an asset or commodity at a specific time or time interval in the future.

2.2 Hydropower optimization model

We now consider the problem faced by a price-taking hydropower production planner with multiple, interconnected reservoirs that participates in a deregulated market similar to the Nordic one. Based on a broad set of endogenous and exogenous variables such as reservoir level, inflow and spot price, they must decide how much water they should use for power production in a given period and how much they should store for future production. The production planner is limited by multiple constraints, e.g., on reservoir volume and turbine capacity, and his primary concern is how they can utilize their water to maximize the expected present value of all discounted future cash flows. Assuming that the planner only participates in the day-ahead market, the price they earn for produced power at time t will always be given by the corresponding time spot price.

The problem faced by the production planner is a stochastic decision problem, meaning that decisions must be made in light of uncertainty about future states of their environment. For each time step, there are two stochastic, exogenous variables that affect the decisions of the production planner; spot price P_t and inflow $I_{b,t}$ into all reservoirs $b = [1, \dots, B]$. For convenience, we denote $\hat{I}_t = \{I_{b,t} : b = [1, \dots, B]\}$ as a set of all inflows to all reservoirs at time t . In a complete market, all random future spot prices (and inflows) would be defined on a measure space (Ω, \mathcal{F}) and there would exist a risk-neutral, martingale measure \mathbb{Q} that represents the risk-neutral probabilities of all future states for spot price and inflow. Harrison and Pliska (1981) define a complete market as a market where the price of all securities is attainable, and there exists only one single price for each security. Thus, a complete market is also arbitrage free. If the market is incomplete, multiple such measures \mathbb{Q} could exist, but only if the market is also arbitrage-free and we can span all underlying assets. If this is not the case, one must use the measure \mathbb{P} , representing the real-world probability of all future states. In their paper, Löhndorf and Wozabal (2017) assume an incomplete market, and they do, among others, allow for non-zero expected gains from trading in the forward market. On the contrary, Bjerksund et al. (2008) assume a complete market with no riskless arbitrage opportunities. We adopt the latter approach in this model and consistently assume that the electricity market is complete throughout this paper. Using this and denoting π_t as a decision policy at time t providing a cash flow of $CF_t = CF_t(P_t, \hat{I}_t, \pi_t)$ and r as an appropriate continuously compounded discount rate, the expected discounted cash flows over a time horizon \hat{T} are given by.

$$\max_{\pi_t} \mathbb{E}_{\mathbb{Q}} \left(\sum_{t=1}^{\hat{T}} CF_t(P_t, \hat{I}_t, \pi_t) \exp(-rt) \right) \quad (1)$$

Like Lamond and Boukhtouta (1996), we model the reservoir management problem as a Markov decision process (MDP). Multiple formal definitions of such decision processes exist, but one of its most important traits is that the exogenous variables of the model must follow a Markov process. Given a current state of the world, the next state value of a variable following a Markov process is only dependent on its current state value, and not its entire history. Similarly, in an MDP, all decisions are made based on the current state of the world and its future expected states, irrelevant of all past states. Although Bellman (1954) does not use the term MDP explicitly, he introduces multiple concepts for dynamic programming (DP) that are also applied

in MDPs. As in DP, the objective of MDPs is to obtain optimal decision policies (π_t) for all current and future states of the world. These policies are based on the current state and all potential future states whose probabilities are incorporated in the measure \mathbb{Q} . When choosing their current state decisions, the production planner should do so to maximize the value of all future cash flows, denoted by V_t . We formulate this using the Bellman equation, first introduced by Bellman (1957). For the reservoir management problem, the objective function at time t will be given by

$$V_t(P_t, \hat{I}_t, \pi_t) = \max_{\pi_t} CF_t(P_t, \pi_t) + \exp(-r)\mathbb{E}[V_{t+1}(P_{t+1}, \hat{I}_{t+1}, \pi_{t+1}|P_t, \hat{I}_t, \pi_t)] \quad (2)$$

(2) is a recursive formula, meaning that the time t value of all future cash flows V_t is a function of the immediate cash flows CF_t and the expected next step value V_{t+1} . The possible values of V_{t+1} are, however, dependent on the current time decisions, indicating the importance of choosing π_t such that it does maximize not only the current cash flow, but also all expected future cash flows.

In hydropower production, the cash flows earned by the production planner equal the product of price and the amount of produced energy. When ignoring turbine and generator start-up costs, which is quite common in other papers discussing hydropower reservoir management (e.g., Wallace and Fleten (2003)), expected cash flows can be set equal to expected revenues. For a hydropower system consisting of B interconnected reservoirs, we denote by $w_{bi,t}$ the amount of water in [m^3] nominated for production in a turbine connecting reservoir b and reservoir i . In case the nominated water flows into an outlet (e.g., a river, lake or fjord), we set $i = O$. Further, ς and ϖ are the number of seconds and hours, respectively, the plant's turbines are running per week. Given that the plant produces at constant rate, we can define $q_{bi,t} = w_{bi,t}/\varsigma$ as the water discharge in [m^3/s] from reservoir b flowing into reservoir i at time t . Further, we let $H_{b,t}$ denote the head elevation in reservoir b , and $\eta_{bi,t}$ the efficiency rate of a turbine connecting reservoir b and i . In reality, these are typically functions of multiple decision variables, e.g., reservoir volume and water discharge. While we do not define these functions now, we discuss their form further in Section 4.2. Lastly, given a water density ρ and gravitational acceleration G , the cash flow CF_t at time t can be written as

$$CF_t = P_t \cdot \rho \cdot G \cdot \varpi \cdot \sum_{b \in B} \left[H_{b,t} \cdot \sum_{i \in B, i \neq b} q_{bi,t} \cdot \eta_{bi,t} \right] \quad \text{for } t = [1, \dots, \hat{T}] \quad (3)$$

We denote by $l_{b,t}$ the water level in [m^3] in reservoir b at time t . Further, $u_{bi,t}$ is the amount of water flowing from reservoir b to reservoir i outside a turbine, that is, either through a regulated channel or as spillage. Like above, we set $i = O$ if the water flows into an outlet. The general volume balance of all reservoirs will then be given by

$$l_{b,t} = l_{b,t-1} - \sum_{i \in B, i \neq b} [w_{bi,t} + u_{bi,t}] + I_{b,t} + \sum_{i \in B, i \neq b} u_{ib,t} \quad \text{for } t = [0, \dots, \hat{T}], \quad b = [1, \dots, B] \quad (4)$$

Further, the problem faces multiple restrictions. All reservoirs are subject to a maximum and minimum level of water, denoted by $\underline{l}_{b,t}$ and $\overline{l}_{b,t}$. These limits can be based on physical constraints such as reservoir geometry and dam robustness, but also on government regulations, some of

which may be seasonal. There are also restrictions in the turbines, stating the maximum allowed water discharge $\overline{q_{bi}}$ that they can handle. In case there exists no turbine at reservoir b whose water flows into reservoir i , $\overline{q_{bi}}$ will logically be 0. Lastly, due to infrastructural reasons (e.g., too small channels or insufficiently robust spillways), there might be a maximum constraint on the allowed amount of water flowing from reservoir b to i , $\overline{u_{bi,t}}$. If no water can flow from reservoir b to i , either due to the lack of physical connections or the effects of gravity (the head elevation of reservoir i is higher than that of reservoir b), $\overline{u_{bi,t}}$ will logically be 0. All these constraints can be summarized in (5)-(8).

$$l_{b,t} \leq \overline{l_{b,t}} \quad \text{for } t = [1, \dots, \hat{T}], \quad b = [1, \dots, B] \quad (5)$$

$$l_{b,t} \geq \underline{l_{b,t}} \quad \text{for } t = [1, \dots, \hat{T}], \quad b = [1, \dots, B] \quad (6)$$

$$q_{bi,t} \leq \overline{q_{bi}} \quad \text{for } t = [1, \dots, \hat{T}], \quad b = [1, \dots, B], \quad i \in B, i \neq b \quad (7)$$

$$u_{bi,t} \leq \overline{u_{bi,t}} \quad \text{for } t = [1, \dots, \hat{T}], \quad b = [1, \dots, B], \quad i \in B, i \neq b \quad (8)$$

By combining all expressions and restrictions, our dynamic program can be summarized as solving the following problem at time t

$$\begin{aligned} & \max && V_t(P_t, \hat{I}_t, \pi_t) \\ & \text{subject to} && (4), (5), (6), (7), (8) \end{aligned}$$

3 Price and inflow dynamics

A production planner that only participates in the spot market faces two uncertain state variables; spot price and inflow. In the following sections, we will demonstrate how these variables can be modeled as stochastic processes. As mentioned in the previous section, we want to model our optimization problem as a Markov decision process (MDP). To do this, both price and inflow must be modeled as Markov processes. In literature, most processes used to describe price movements are Markov processes, but we show that it is reasonable to model inflow movements as a Markov process as well. We begin this section by presenting processes used to model spot price and inflow movements and how their parameters can be found. For the chosen spot price model, which is in fact a forward price model, we present three different methods to construct smooth forward curves, which are then used to obtain the model coefficients. Further, we present how all future states of price and inflow are discretized using a scenario lattice, and we also present the optimization algorithm used to obtain optimal decision policies for each state.

3.1 Overview of models for electricity spot price movements

In literature, multiple stochastic processes for spot price movements have been proposed, all of which have associated advantages and drawbacks. Here, we introduce some of these processes, discuss their applications and decide which of them to use in our final model. For more extensive details on all processes and their applications in commodity markets, the reader should consult Clewlow and Strickland (2000). For a process to appropriately describe the changes in electricity spot price, it should incorporate the most important features and dynamics driving movements in the price. Johnson and Barz (1999) present multiple properties of electricity spot prices, and a selection of these include:

- **Mean reversion:** Although the electricity spot price has a stochastic nature, it tends to revert towards a mean value if the spot price deviates significantly from it.
- **Seasonality effects:** Spot prices tend to show periodic fluctuations over the course of a day, week and year. Seasonality effects are also common in other parameters such as volatility.
- **Price-dependent volatility:** Although volatility shows seasonal variation, it also tends to be dependent on the spot price itself. For periods of higher spot prices, the volatility of spot price returns tends to be larger than for time periods with lower prices.
- **Price spikes and jumps:** Electricity spot prices sometimes experience sudden price spikes or jumps, both in positive and negative directions.

A simple one-factor model used to model price movements both in equity and commodity markets is the Geometric Brownian motion (GBM). The model can, however, not incorporate any of the points mentioned above except for the one regarding price-dependent volatility, meaning it will perform poorly when the objective is to make accurate predictions of future spot prices. We can resolve the issue of mean reversion by adjusting the GBM into a mean-reverting model, represented by (9). This process is often referred to as an Ornstein–Uhlenbeck process.

$$dP_t = \alpha_t(\ln \bar{P}_t - \ln P_t)P_t dt + \sigma_t P_t dZ_t \quad (9)$$

Here, dt is an infinitesimally short time interval, and dP_t denotes a change in spot price during dt . The spot price P_t reverts towards a long-term level \bar{P}_t with the (possibly) time-dependent mean reversion rate α_t . For higher values of α_t , the spot price reverts faster, and vice versa. σ_t is the spot price volatility, and dZ_t is an increment in a Wiener process (also known as Brownian motion) during dt . In discrete time, dZ_t is drawn from a normal distribution with zero mean and a volatility of \sqrt{dt} , that is, $dZ_t \sim N(0, dt)$. dZ_t is often referred to as a source of risk or a *source of uncertainty*. In order to incorporate jumps, one can extend the mean-reverting model into a two-factor model given by

$$dP_t = \alpha_t(\ln \bar{P}_t - \ln P_t)P_t dt + \sigma_t P_t dZ_t + \kappa_t P_t dq_t \quad (10)$$

In this model, dq_t is the increment of a Poisson process, also known as a jump process, during dt . It can take one of two values: $dq_t = 0$ with probability $1 - \theta_t dt$ and $dq_t = 1$ with probability $\theta_t dt$. Here, θ_t is the jump frequency, and κ_t is the proportional jump size. κ_t is random, and is drawn from the distribution

$$\ln(1 + \kappa_t) \sim N(\ln(1 + \bar{\kappa}_t) - \frac{1}{2}\gamma_t^2, \gamma_t^2) \quad (11)$$

where $\bar{\kappa}_t$ is the mean jump size and γ_t the jump volatility. In order to adjust the process so that it also incorporates spikes, the value of the mean reversion rate α_t can be modified such that it takes higher values right after a jump making the price revert faster towards its mean, as proposed by Mayer et al. (2015). Alternatively, one could modify the Poisson process such that it makes a jump of equal magnitude in the opposite direction in the timestep dt after it has made a jump, making the price return to its pre-jump value.

Although the process described by (10) can capture many of the typical spot price characteristics, it does not capture seasonal spot price fluctuations. Lucia and Schwartz (2002) solve this by adding a deterministic, underlying seasonality function, whose values are based on empirical prices. Thereby, they obtain a model that incorporates all of the four most important findings in Johnson and Barz (1999). A natural question that arises is how one should find all the model coefficients. In order to set the long-term mean spot price \bar{P}_t , we need an expectation about which value the spot price is reverting towards. Finding the value of the mean reversion rate α_t for periods with less or more jumps might also prove to be a difficult task, in addition to setting correct, possibly time-dependent values for the coefficients of the jump process. One can also question whether it is correct to use a deterministic underlying seasonal function. Is it reasonable to expect that the average empirical seasonal spot price is a precise forecast for future spot prices? If no, then how should the function be chosen?

It may be difficult to find the coefficients of the process described by (10). However, the main disadvantage of (10) is that it is a two-factor model. Bjerksund et al. (2008) claim that these models are unrealistic, as they cannot capture the entire dynamics behind spot price movements. They find that six factors are needed to explain the variation in price movements in the UK gas market, while Koekebakker and Ollmar (2005) find that ten factors were needed to explain the variation in the Nord Pool electricity market (1995-2000). Consequently, we reject the processes mentioned above and try to obtain a similar, multi-factor model for describing spot price movements.

Next, we look at the possibility of using movements in the price of forward contracts to predict future spot prices, as suggested by Clewlow and Strickland (2000). At time t , $F_{t,T}$ is the price of

a forward contract with maturity (or delivery) at time T . For a forward contract with immediate delivery ($T = t$), the price of that contract is simply the current time spot price, that is $P_t = F_{t,t}$. Thus, a stochastic model for the price development of forward contracts with different times to maturity can be used to model future spot prices. In a liquid power market, the available future and forward contracts traded at time t should represent the current time risk-adjusted market expectations for future spot prices, meaning that the spot prices projected by the model will incorporate these expectations. A further advantage is that the price of all forward contracts traded in the market include the seasonality of electricity prices, implying that the model does not need any deterministic function to account for seasonality.

We start with a simple single-factor model for forward price movements presented by Heath et al. (1992). It is given by

$$\frac{dF_{t,T}}{F_{t,T}} = \sigma_{t,T} dZ_t \quad (12)$$

Here, $\sigma_{t,T}$ denotes the time t volatility of a forward contract with maturity at time T , and dZ_t is a source of uncertainty. The observant reader has probably noticed that (12) contains no drift term, but only a stochastic part. Because future and forward contracts have zero initial investment, their expected return in the risk-neutral world is zero. No drift term complies with our assumption that the market is complete, as a non-zero drift term would mean that the price process would not be martingale (risk-neutral). In that case, there would not exist a measure \mathbb{Q} representing the risk-neutral probabilities. Instead, we would have to use the real-life measure, \mathbb{P} . In their paper, Löhndorf and Wozabal (2017) use a non-zero drift to describe the movement of forward prices, which is one of the reasons why they cannot use a measure \mathbb{Q} .

As mentioned above, a single-factor model cannot capture the entire dynamics driving forward price movements. Clewlow and Strickland (2000) show that the process described in (12) can be extended into a multi-factor model with N sources of uncertainty $dZ_{i,t}$, given by

$$\frac{dF_{t,T}}{F_{t,T}} = \sum_{i=1}^N \sigma_{i,t,T} dZ_{i,t} \quad (13)$$

Here, $\sigma_{i,t,T}$ is the i th volatility function, and together, $\sigma_{i,t,T}$ for $i = [1, \dots, N]$ explain the dynamics driving the time t movement of a forward contract with maturity at time T . The i th volatility function is associated with the i th source of uncertainty, $dZ_{i,t}$. There exist multiple papers on how $\sigma_{i,t,T}$ should be found. Bjerksund et al. (2000) propose a three-factor model where $\sigma_{i,t,T}$ for $i = [1, 2, 3]$ are given by three theoretical expressions. Koekebakker and Ollmar (2005) try to use this model for the Nord Pool electricity market, but find that it is inconsistent with empirical findings. Like Clewlow and Strickland (2000), they therefore propose constructing *volatility functions* on the form $\sigma_{i,t,T} = \Psi_i(T - t)$ based on empirical returns data. That is, they let the volatility coefficients be a function of time to maturity $T - t$. They use principal component analysis (PCA) to obtain multiple volatility functions, and choose the number of factors N such that the amount of cumulative explained variance by all the functions exceeds some threshold.

In our final model, we choose to adopt the multi-factor process described by (13) and estimate the volatility functions the same way as Clewlow and Strickland (2000) and Koekebakker and Ollmar (2005). There are multiple reasons for choosing this model above the spot price process

in (10). Firstly, by using the forward contracts traded in the market, all predicted spot prices are a function of their expected price. Further, the forward contracts incorporate yearly price seasonality, making the process satisfy the second finding of Johnson and Barz (1999). Including multiple factors also allows the model to better explain the real dynamics driving forward price movements, and thereby make better price predictions. Our model could also have been extended to incorporate spikes and jumps using a Poisson process, but for our model granularity we do not expect average spot prices to exhibit movements larger than those driven by the volatility functions and the Wiener processes $Z_{i,t}$, and therefore do not add a Poisson process.

3.2 Constructing forward curves and volatility functions

As mentioned in Section 3.1, we want to model future spot prices using a forward price model. Since the volatility functions $\sigma_{i,t,T} = \Psi_i(T-t)$ are to be estimated empirically, we must construct a sufficiently large dataset of daily returns for multiple types of forward contracts $m = [1, \dots, M]$ with time to maturity $\tau_m = T - t$. Using τ as time to maturity, the volatility functions are now denoted $\sigma_{i,\tau}$. In order to calculate these returns series, Koekebakker and Ollmar (2005) propose constructing multiple forward curves for a large set of historical trading days and then calculate daily returns as the deviations between two consecutive curves.

As explained earlier, forward contracts traded in the market have delivery periods stretching over longer time periods. A forward curve is a curve that aims to explain the expected forward price for delivery in each hour/day/week in a time interval (t_b, t_e) based on all contracts available in the market whose delivery periods span the interval. A forward curve constructed on the date t_s is denoted $f(t_s)$. $f(t_s, t)$ where $t > t_s$ denotes the value of that curve for time t , and intends to represent the price of a fictional forward contract with delivery exactly at time t . There exist multiple methods for constructing forward curves, all of which are more or less appropriate for different purposes of usage. In the following sections we explain three different methods for constructing forward curves, presented by Benth et al. (2007), Fleten and Lemming (2003) and Alexander (2008a). We then show how a set of forward curves $f(t_s, t)$ constructed on multiple dates t_s can be used to obtain a set of daily returns series and to find the desired volatility functions $\sigma_{i,\tau}$.

3.2.1 Constructing smooth forward curves using method of Benth, Koekebakker and Ollmar

In their paper, Benth et al. (2007) propose using multiple splines to construct a continuous forward curve. Since the forward price can then be expressed as a continuous function of time, the time t value of a forward curve $f(t_s, t)$ constructed using this method represents a forward contract with delivery at exactly that time. We denote

$$\mathbb{S} = \{(t_{b,1}, t_{e,1}), \dots, (t_{b,M}, t_{e,M})\} \quad (14)$$

as a set of delivery periods for M observable forward contracts. Using these contracts, we can construct a forward curve starting at $t_b = t_{b,1}$ and ending at $t_e = t_{e,M}$. Note that $t_b \geq t_s$ must not necessarily equal t_s . An example of this is when the first contract used to construct a curve is a weekly contract, indicating that t_b will denote the starting date of the upcoming week. Mathematically, the method defines the value of the forward curve $f(t_s, t)$ as a function of time

t as

$$f(t_s, t) = U(t) + \varepsilon(t) \quad (15)$$

where $U(t)$ is an optional underlying seasonal function and $\varepsilon(t)$ is the adjustment function whose coefficients are what the method seeks to find. The function $\varepsilon(t)$ is constructed using multiple splines for different intervals of time. In order to construct these splines, the time intervals defined by (14) must be split such that we get a set of D non-overlapping intervals divided by the dates $\{t_0, t_1, \dots, t_D\}$. Figure 2 illustrates the logic of obtaining all dates t_i . The dates can be found by taking all start and end dates from \mathbb{S} , sort them in ascending order and then remove all duplicated dates.




Figure 2: Splitting overlapping contracts

Having found the intervals, the adjustment function $\varepsilon(u)$ is defined by (16). Benth et al. (2007) argue that the polynomial function of each spline should be of order four, as it is the most appropriate class for their purposes. The intuition in their method can, however, be extended to any class of polynomial functions or other function types.

$$\varepsilon(u) = \left\{ \begin{array}{ll} a_1 u^4 + b_1 u^3 + c_1 u^2 + d_1 u + e_1, & u \in [t_0, t_1], \\ a_2 u^4 + b_2 u^3 + c_2 u^2 + d_2 u + e_2, & u \in [t_1, t_2], \\ \vdots & \\ a_D u^4 + b_D u^3 + c_D u^2 + d_D u + e_D, & u \in [t_{D-1}, t_D]. \end{array} \right\} \quad (16)$$

Further, all parameters are collected in the vector $\mathbf{x} = [a_1 b_1 c_1 d_1 e_1 \dots a_D b_D c_D d_D e_D]'$. In order to obtain the parameter values that give a smoothest possible forward curve, the objective is to minimize the expression

$$\int_{t_b}^{t_e} [\varepsilon''(u; \mathbf{x})]^2 du \quad (17)$$

subject to a set of constraints regarding the shape of the curve. Firstly, we want the forward curve to be smooth and continuous in all points. We must, therefore, ensure that not only the forward curve itself but also its first and second derivative with respect to time are continuous at all time. The constraints (18)-(20) ensure that this is the case for all time points $t_j = [t_1, \dots, t_{D-1}]$ where two spline functions meet. Also, constraint (21) sets the first derivative of the forward curve at $t = t_D$ equal to zero.

$$(a_{j+1} - a_j)u_j^4 + (b_{j+1} - b_j)u_j^3 + (c_{j+1} - c_j)u_j^2 + (d_{j+1} - d)u_j + e_{j+1} - e_j = 0 \quad (18)$$

$$4(a_{j+1} - a_j)u_j^3 + 3(b_{j+1} - b_j)u_j^2 + 2(c_{j+1} - c_j)u_j + d_{j+1} - d_j = 0 \quad (19)$$

$$12(a_{j+1} - a_j)u_j^2 + 6(b_{j+1} - b_j)u_j + 2(c_{j+1} - c_j) = 0 \quad (20)$$

$$4a_D u_D^3 + 3b_D u_D^2 + 2c_D u_D + d_D = 0 \quad (21)$$

Lastly, constraint (22) ensures that the value of the forward curve in an interval $(t_{b,i}, t_{e,i})$ is consistent with the price of a forward contract $F_{t, t_{b,i}, t_{e,i}}$ with delivery period over the same interval.

$$F_{t, t_{b,i}, t_{e,i}} = \int_{t_{b,i}}^{t_{e,i}} \zeta(u, t_{b,i}, t_{e,i})(\varepsilon(u) + \Lambda(u)) du \quad (22)$$

For reasonable levels of interest rate, Lucia and Schwartz (2002) show that $\zeta(u, t_x, t_y) = 1/(t_y - t_x)$ is a good approximation. (22) can therefore be interpreted such that the average price of the forward curve over an interval $(t_{b,i}, t_{e,i})$ must equal the price of a forward contract $F_{t, t_{b,i}, t_{e,i}}$.

Using Lagrange multipliers, the problem can be simplified to solving the matrix equation given by (23). This must be done numerically, e.g., by using QR factorization.

$$\begin{bmatrix} 2\mathbf{H} & \mathbf{A}' \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\delta} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{b} \end{bmatrix} \quad (23)$$

Here, we first organize all constraints (18)-(22) in matrix form as $\mathbf{A}\mathbf{x} = \mathbf{b}$ where \mathbf{x} is the parameter vector defined earlier. Totally, there are $3D + M - 2$ constraints and $5D$ coefficients to be obtained, making \mathbf{A} a $(3D + M - 1) \times 5D$ matrix and \mathbf{b} a $(3D + M - 1) \times 1$ matrix. Further, $\boldsymbol{\delta} = [\delta_1, \dots, \delta_{3D+M-2}]'$ is the Lagrange multiplier coefficient, and \mathbf{H} is a $5D \times 5D$ matrix defined by (24) and (25)

$$\mathbf{H} = \begin{bmatrix} h_1 & & 0 \\ & \ddots & \\ 0 & & h_n \end{bmatrix}, h_j = \begin{bmatrix} \frac{144}{5}\Delta_j^5 & 18\Delta_j^4 & 8\Delta_j^3 & 0 & 0 \\ 18\Delta_j^4 & 12\Delta_j^3 & 6\Delta_j^2 & 0 & 0 \\ 8\Delta_j^3 & 6\Delta_j^2 & 4\Delta_j^1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (24)$$

$$\Delta_j^k = t_j^k - t_{j-1}^k \quad (25)$$

where t_j with $j = [0, \dots, D]$ are the time points separating the D non-overlapping time intervals.

3.2.2 Constructing forward curves using method of Fleten and Lemming

Fleten and Lemming (2003) propose a method to construct forward curves with different levels of smoothness. Unlike Benth et al. (2007), they model their curve discretely by finding one unique price for a set of constant time steps. We still let $f(t_s)$ be a forward curve constructed on an arbitrary date t_s , and assume that the curve is constructed using M contracts with delivery intervals given as in (14). Since the model aims to find the value of the forward curve in discrete time steps, $f(t_s)$ can be represented as a vector $f(t_s) = [f(t_s, t_b), f(t_s, t_b + 1), \dots, f(t_s, t_e)]'$ where $f(t_s, t)$ is the value of the forward curve at time $t \in [t_b, t_e]$. The dimensions of the vector $f(t_s)$ is $C \times 1$, where C denotes the number of discrete prices contained by the curve. Further, we let $F_{t_s, t_b, j, t_b, j}^{ask}$ and $F_{t_s, t_b, j, t_b, j}^{bid}$ denote the ask and bid price of forward contract $j = [1, \dots, M]$, where $(t_{b, j}, t_{b, j})$ denotes its delivery period. $U(t) = [U(t_b), U(t_b + 1), \dots, U(t_e)]'$ denotes an underlying deterministic price function in vector form, and r is the model discount rate. Lastly, we denote $\lambda \in \langle 0, \infty \rangle$ as the smoothness parameter. For high values of λ , the method will construct forward curves with maximum smoothness, whereas lower values create curves with smaller smoothness and larger price jumps. Given all these parameters, the forward curve $f(t_s, t)$ is constructed by solving the minimization problem

$$\begin{aligned} & \underset{f(t_s, t)}{\text{minimize}} && \sum_{t=t_b}^{t_e} (f(t_s, t) - U(t))^2 + \lambda \sum_{t=t_b+1}^{t_e-1} (f(t_s, t-1) - 2f(t_s, t) + f(t_s, t+1))^2 \\ & \text{subject to} && F_{t_s, t_b, j, t_b, j}^{bid} \leq \frac{1}{\sum_{t=t_b, j}^{t_{b, j}} \exp(-rt)} \sum_{t=t_b, j}^{t_{b, j}} \exp(-rt) f(t_s, t) \leq F_{t_s, t_b, j, t_b, j}^{ask} \quad \text{for } j \in [1, \dots, M] \end{aligned}$$

Since we assume a complete market and no drift, we set $r = 0$. We also set $U(t) = 0$ to avoid deterministic elements in our model. Further, if we only consider market closing prices, we can set $F_{t_s, t_b, j, t_b, j}^{bid} = F_{t_s, t_b, j, t_b, j}^{ask} = F_{t_s, t_b, j, t_b, j}$. This simplifies the problem to

$$\begin{aligned} & \underset{f(t_s, t)}{\text{minimize}} && \sum_{t=t_b}^{t_e} f(t_s, t)^2 + \lambda \sum_{t=t_b+1}^{t_e-1} (f(t_s, t-1) - 2f(t_s, t) + f(t_s, t+1))^2 \\ & \text{subject to} && F_{t_s, t_b, j, t_b, j} = \frac{1}{t_{b, j} - t_{b, j}} \sum_{t=t_b, j}^{t_{b, j}} f(t_s, t) \quad \text{for } j \in [1, \dots, M] \end{aligned}$$

Using the method of Lagrange multipliers, the problem can be reformulated to solving the matrix equation given by (26), where $\boldsymbol{\delta} = [\delta_1, \dots, \delta_M]'$ is the vector of Lagrange multipliers and $f(t_s)$ is the forward curve on vector form.

$$\begin{bmatrix} 2\mathbf{H} & \mathbf{A}' \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \cdot \begin{bmatrix} f(t_s) \\ \boldsymbol{\delta} \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{F}_{t_s} \end{bmatrix} \quad (26)$$

In 26, \mathbf{A} is an $M \times C$ matrix whose elements $A_{j, t}$ can take the values $A_{j, t} = 1$ if time t is part of the delivery period of the j th forward contract and $A_{j, t} = 0$ otherwise. $\mathbf{F}_{t_s} = [F_{t_s, t_b, 1, t_e, 1}, \dots, F_{t_s, t_b, M, t_e, M}]'$ is a vector containing the price of all M forward contracts traded on

the market. The matrix \mathbf{H} is $C \times C$ and given by (27).

$$\mathbf{H} = \begin{bmatrix} 1 + \lambda & -2\lambda & \lambda & 0 & 0 & 0 & 0 & \dots & 0 \\ -2\lambda & 1 + 5\lambda & -4\lambda & \lambda & 0 & 0 & 0 & \dots & 0 \\ \lambda & -4\lambda & 1 + 6\lambda & -4\lambda & \lambda & 0 & 0 & \dots & 0 \\ 0 & \lambda & -4\lambda & 1 + 6\lambda & -4\lambda & \lambda & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \lambda & -4\lambda & 1 + 6\lambda & -4\lambda & \lambda & 0 \\ 0 & \dots & 0 & 0 & \lambda & -4\lambda & 1 + 6\lambda & -4\lambda & \lambda \\ 0 & \dots & 0 & 0 & 0 & \lambda & -4\lambda & 1 + 5\lambda & -2\lambda \\ 0 & \dots & 0 & 0 & 0 & 0 & \lambda & -2\lambda & 1 + \lambda \end{bmatrix} \quad (27)$$

3.2.3 Constructing forward curves by linear interpolation

Alexander (2008a) presents a different approach for constructing forward curves. It involves creating what she calls *constant maturity futures* by interpolating between the prices of adjacent forward contracts traded in the market. We modify this method so that it can be used to create forward curves. Similar to the method of Fleten and Lemming (2003) the forward curve is made up of values at discrete, predefined time steps. Note that since the forward curve is found using linear interpolation, it will generally not be smooth at all points.

There is a special challenge in applying the approach described by Alexander (2008a) to electricity forwards, as it has no obvious way of handling forwards with a delivery period instead of delivery in a specific point in time. In this method, the value of the forward curve $f(t_s, t)$ constructed on t_s for delivery time t is defined as the value of a forward contract whose delivery period starts at t . The length of the delivery period is, however, given by the delivery period of the two contracts used to find this value of the curve. This is slightly different from the previous two methods, in which we used all available contracts to construct a smooth forward curve where the value of the curve for a given delivery time denoted the price of a forward contract with delivery on that particular point.

In this method, each element of the forward curve is calculated by linear interpolation. Say that we want to calculate the price of a forward contract $F_{t,T}$ at time t with delivery at time T . Intuitively, it can be found by weighting the market prices of two tradable forward contracts such that their weighted average delivery time is equal to T . We extend this logic and use it to find every entry of a forward curve $f(t_s)$ whose underlying contracts have delivery periods as described in (14).

As before, we let t_s be the date for which a forward curve is constructed and let t be the delivery time of the forward curve element we want to find. Further, let $F_{t_s, t_{b,i}, t_{e,i}}$ be the market price of a forward contract with delivery period $(t_{b,i}, t_{e,i})$ where $t_{b,i} \leq t$. Among the contracts with the beginning of the delivery period *earlier* than t , $F_{t_s, t_{b,i}, t_{e,i}}$ is the contract having the beginning of the delivery period $t_{b,i}$ closest in time to t . Let $F_{t_s, t_{b,j}, t_{e,j}}$ be the market price of a future contract with delivery period $(t_{b,j}, t_{e,j})$ where $t_{b,j} \geq t$. Among the contracts with the beginning of the delivery period *later* than t , $F_{t_s, t_{b,j}, t_{e,j}}$ is the contract having the beginning of the delivery period $t_{b,j}$ closest in time to t . In summary, we have that

$$t_{b,i} \leq t \leq t_{b,j} \quad (28)$$

By linear interpolation, the forward curve element ($f(t_s, t)$) located in the interval $(t_{b,i}, t_{b,j})$ is

given by (29).

$$f(t_s, t) = F_{t_s, t_{b,i}, t_{e,i}} + \frac{t - t_{b,i}}{t_{b,j} - t_{b,i}} \cdot (F_{t_s, t_{b,j}, t_{e,j}} - F_{t_s, t_{b,i}, t_{e,i}}) \quad (29)$$

Further, we want to interpolate between forward contracts of the same delivery period length, meaning that the part of the curve spanning the interval $(t_{b,1}, t_{b,j})$ must be constructed using contracts where

$$t_{e,i} - t_{b,i} = t_{e,j} - t_{b,j} \quad (30)$$

Denote R the number of different contract types (e.g., contracts with weekly, monthly, quarterly and yearly delivery periods). Since we use all contract types available we will have R different forward curves for each trading day. To create a complete forward curve with one unique value $f(t_s, t)$ for each value of t , portions of these forward curves are used for different intervals of time. In the near end of the curve, one should use contracts with shorter delivery periods (e.g., weekly) to construct the values of the curve. When t is increased, $f(t_s, t)$ will eventually have to be constructed using contracts with a longer delivery period (e.g., monthly). This is due to the nature of electricity forward markets, where the delivery period of the contracts traded in the market generally increases for larger times to maturity. Therefore, one must use two forward contracts with larger delivery periods to comply with restriction (28). For even larger values of t , the curve is constructed using a contract type with an even longer delivery period, and so on. Apart from using contracts with the shortest possible delivery period, a heuristic to decide which contract type is to be used in the complete forward curve is to choose the contract with the most historical price observations for a given time to delivery.

3.2.4 Constructing returns series from forward curves

Having constructed smooth forward curves for multiple consecutive days, we can now create a returns dataset that can be used to find the volatility function. We apply a modified version of the method used by Koekebakker and Ollmar (2005), as we choose to calculate continuously compounded logarithmic returns rather than discrete compounded returns. We do this because it allows us to calculate returns over longer time periods by addition, thereby simplifying many calculations. This approach is also used by Bjerksund et al. (2008). Since we want volatility functions on the form $\sigma_{i,t,T} = \sigma_{i,\tau} = \Psi_i(T-t)$, we must create returns series for a set of contracts with equal time to maturity $\tau = T - t$. $f(t_j)$ still denotes a forward curve constructed at date t_j , and $f(t_j, t_j + \tau_a)$ denotes the value of this curve for a delivery date $T_a = t_j + \tau_a$. We then use (31) to calculate daily returns at time t_j for contracts with time to maturity τ_a .

$$x_{j,a} = \ln(f(t_j, t_j + \tau_a)) - \ln(f(t_{j-1}, t_j + \tau_a)) \quad (31)$$

Here, $j = [2, \dots, J]$ and $a = [1, \dots, A]$, where J is the number of forward curves and A is the number of maturity dates for which we want to construct a dataset. Next, we try to illustrate the use of the formula. For a forward curve generated on a Wednesday and a time to maturity $\tau_a = 5$ days, t_j will represent that Wednesday and $t_j + \tau_a$ the upcoming Monday. The price of a forward contract with these properties is given by the forward curve, $f(t_j, t_j + \tau_a)$. When we calculate the daily price change of this particular contract, we also need to obtain $f(t_{j-1}, t_j + \tau_a)$, which is the value of the contract given by the forward curve constructed on the previous date

(Tuesday) with maturity on the same upcoming Monday. If we call the forward curve constructed on the Wednesday for *Wednesdaycurve* and on Tuesday for *Tuesdaycurve*, the described return is given by $\ln(\text{Wednesdaycurve}(5)) - \ln(\text{Tuesdaycurve}(6))$. If t_j is on a Monday, t_{j-1} will be the prior week Friday. Though some may argue that the return between Friday and Monday should not be treated as a daily return, it is, in fact, a return between two trading days, and we consider the number of days per year equal to the number of trading days. Further, Bjerksund et al. (2008) show that they do not see any major differences in their returns series when they treat such returns as daily returns, and neither do we.

The returns series matrix calculated using $J + 1$ forward curves (meaning we can find J returns) and A different time to maturities is then given by

$$\mathbf{X}_{\mathbf{J} \times \mathbf{A}} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,A} \\ x_{2,1} & x_{2,2} & \dots & x_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ x_{J,1} & x_{J,2} & \dots & x_{J,A} \end{bmatrix} \quad (32)$$

3.2.5 Finding volatility functions using PCA

The description in this section is based on material from Alexander (2008b) and Koekebakker and Ollmar (2005). After constructing the returns dataset $\mathbf{X}_{\mathbf{J} \times \mathbf{A}}$, we must use principal component analysis (PCA) to find the desired N volatility functions. PCA is an orthogonalization technique used to reduce a dataset consisting of highly correlated variables into a set of non-correlated factors (principal components) that explain the total variation in the data. Since $\mathbf{X}_{\mathbf{J} \times \mathbf{A}}$ has the dimension $J \times A$, its covariance matrix \mathbf{V} should have the dimension $A \times A$. The principal components of \mathbf{V} are the columns of \mathbf{P} given by

$$\mathbf{P} = \mathbf{X}_{\mathbf{J} \times \mathbf{A}} \mathbf{W} \quad (33)$$

Here, \mathbf{W} is a matrix whose columns are the eigenvectors \mathbf{w}_i of \mathbf{V} sorted in descending order based on their corresponding eigenvalue Λ_i . By definition, the following relationship exists between the covariance matrix \mathbf{V} and its eigenvalues and eigenvectors:

$$\mathbf{V} = \mathbf{W} \mathbf{\Lambda} \mathbf{W}' \quad (34)$$

Here, $\mathbf{\Lambda}$ is an $A \times A$ matrix whose non-diagonal elements are zero and diagonal elements are the eigenvalues Λ_i in descending order. By using this, the covariance matrix of the principal components is given by (35).

$$\text{Var}(\mathbf{P}) = \mathbf{W}' \mathbf{V} \mathbf{W} = \mathbf{W}' \mathbf{W} \mathbf{\Lambda} \mathbf{W}' \mathbf{W} = \mathbf{\Lambda} \quad (35)$$

(35) illustrates that the variance of the i th principal component is the eigenvalue Λ_i and that the principal components are uncorrelated. As explained above the principal components explain the entire variation in $\mathbf{X}_{\mathbf{J} \times \mathbf{A}}$. The proportion of total variance explained by the a th component

is given by $\Lambda_a / \sum_{i=1}^A \Lambda_i$, letting us calculate the cumulative proportion of variance explained by the first N components as

$$\frac{\sum_{i=1}^N \Lambda_i}{\sum_{i=1}^A \Lambda_i} \quad (36)$$

Typically, one would choose the number of principal components N to describe $\mathbf{X}_{\mathbf{J} \times \mathbf{A}}$ such that the proportion of explained variance is around 90%-95%. Clewlow and Strickland (2000) show that only two components are needed to explain 96.8% of total variation of NYMEX crude oil futures contracts, whereas Koekebakker and Ollmar (2005) needed more than ten components to explain the same proportion for Nordic electricity forwards in the period 1995-2000. Further, in order to find the empirical volatility functions σ_{i,τ_a} , we use

$$\sigma_{i,\tau_a} = \sqrt{\lambda_i} w_{ai} \quad (37)$$

where w_{ai} is the a th element of the i th eigenvector of \mathbf{V} , or equivalently, the a th element of the i th column of \mathbf{W} .

Lastly, we give an interpretation of the first three principal components. In a highly correlated system, the first principal component (\mathbf{p}_1) is said to capture a common trend in the returns series. Therefore, it is sometimes called the *trend* or shift factor. That is, if the first principal component changes in one direction while the other components are fixed, the entire forward curve shifts in the same direction. Thus, the volatility function associated with the first principle component will often have a similar shape as the overall volatility function, but of smaller magnitude. If the system has no natural ordering, then the second and higher order principal components have no intuitive interpretation. However, if the system is ordered, such as a set of returns on forwards of different maturities, a change in the second principal component (\mathbf{p}_2) will cause the long and short maturity end of the forward curve to move in opposite directions. Thus, it is called the *tilt* factor. On the contrary, the third component (\mathbf{p}_3) is often referred to as the *bending* factor, as a change in this component will make the middle part of the forward curve move in the opposite direction of the long and short ends. A more extensive explanation of these concepts can be found in Clewlow and Strickland (2000).

3.3 Inflow model

The inflow model is based on the geometric periodic autoregressive (GPAR) model presented by Shapiro et al. (2013). The authors found that a first-order periodic autoregressive model of the log-inflows provides a good description of the dataset, which contained inflow observations from the Brazilian hydropower system. They found that the distribution of inflow observations I_t was highly right-skewed. Therefore, they worked with $Y_t = \ln(I_t)$ to obtain a distribution with less skew.

Let $\hat{\mu}_t$, $t = 1, \dots, 52$ be the weekly averages of Y_t and $Z_t = Y_t - \hat{\mu}_t$ be the corresponding deviations. Shapiro et al. (2013) found that Z_t could be described by an AR(1) process. (38) shows how the deviations of the log inflows from their mean can be described as an 1-lag autoregressive process.

$$Z_t = \phi_0 + \phi_1 Z_{t-1} + \epsilon_t \quad (38)$$

Here, ϕ_0 and ϕ_1 are parameters of the model, and ϵ_t is the error term representing the difference between the observed and predicted value. To be able to model the inflow as a stochastic process, we assume that the error terms are distributed $\epsilon_t \sim N(0, \sigma_{\text{INF}}^2)$, where σ_{INF} is the standard deviation of the error terms. The parameters ϕ_0 and ϕ_1 are estimated by ordinary least squares regression. Because Z_t observations are themselves deviations, ϕ_0 is highly insignificant. We set $\phi_0 = 0$ and use $\phi_1 = \phi$ from this point on. Furthermore, we find the log inflow Y_t by substituting for Z_t and Z_{t-1} .

$$Y_t = Z_t + \hat{\mu}_t = \phi Z_{t-1} + \epsilon_t + \hat{\mu}_t = \phi(Y_{t-1} - \hat{\mu}_{t-1}) + \epsilon_t + \hat{\mu}_t \quad (39)$$

The inflow I_t can be expressed as a function of Y_t . We insert the obtained expression of Y_t into $I_t = \exp(Y_t)$, and get

$$I_t = \exp(Y_t) = \exp(\phi Y_{t-1}) \exp(-\phi \hat{\mu}_{t-1} + \epsilon_t + \hat{\mu}_t) \quad (40)$$

By rewriting, we obtain the inflow process described by (42).

$$I_t = \exp(\phi \ln I_{t-1}) \exp(\hat{\mu}_t - \phi \hat{\mu}_{t-1} + \epsilon_t) \quad (41)$$

$$I_t = \exp(\epsilon_t) \exp(\hat{\mu}_t - \phi \hat{\mu}_{t-1}) I_{t-1}^\phi \quad (42)$$

We further allow the error term standard deviation σ_{INF} and the coefficient ϕ to be time-dependent. The final inflow model can then be expressed as

$$I_t = \exp(\epsilon_t) \exp(\hat{\mu}_t - \phi_t \hat{\mu}_{t-1}) I_{t-1}^{\phi_t} \quad (43)$$

where t is the week number and ϵ_t now follows the distribution $\epsilon_t \sim N(0, \sigma_{\text{INF},t}^2)$. Since inflow I_t is a function of its first lag only, future values of inflow are only dependent on their current value and not the entire history. Thus, inflow follows a Markov process. This means that both inflow and price follow Markov processes, which was one of the prerequisites for representing our decision problem as a Markov decision process.

3.4 Price and inflow lattice

Having established a stochastic optimization model and separate models for spot price and inflow development, we must now find a method for solving the given MDP. As mentioned previously, the objective in an MDP is to obtain an optimal decision policy for all states of the world. We choose a discrete representation of the decision problem, meaning that we need to discretize all future states of the world for given intervals of time. Afterwards, we organize the discrete states using a scenario lattice. In the following sections, we describe the nature of scenario lattices and a method for discretizing states and transition probabilities into a lattice.

3.4.1 Lattice description

A lattice is a recombining tree with a finite number of states (represented by nodes) for each time step, where all nodes for a time step t are possible successor nodes of a node at $t - 1$. Each node represents a possible state of one or multiple stochastic variables. In our case, this implies that a lattice node represents a state of both inflow and spot price. Possible state transitions are illustrated by connecting arcs, and each arc has a probability weight denoting the probability of a transition between the connected states. If our problem had not been an MDP, meaning that we could not find one single decision policy for each state independent of its path, we could not have used a scenario lattice. Instead, we would be forced to use a scenario tree. In a scenario tree, each node can only have one predecessor node, meaning that each node represents a single state AND the history of previous states. This implies that more nodes are needed to represent all possible states, as multiple nodes for a given time step will most likely represent very similar or equal states. Unlike a scenario tree, a scenario lattice allows each state to have multiple predecessor states. This means that fewer nodes can be used to represent all possible states, and more importantly, that fewer nodes can be used to represent multiple scenarios. We denote N_t as the number of nodes at time t . Further, $\overline{S}_{tn} = \{P_{tn}, I_{tn}\}$ denotes the n th state (or node) at time step t , where P_{tn} is the state spot price and $I_{tn} = \{I_{b,tn} : b = [1, \dots, B]\}$ is a set of inflows into all B reservoirs for the same state. We also let $n \in [N_t]$, where N_t is the total number of states at time t . If $\overline{S}_t = \{\overline{S}_{tn} : n \in [N_t]\}$, that is, \overline{S}_t is a set of all states at time t , the set of possible scenarios for a lattice with time horizon \hat{T} where all arcs have positive probability weights is given by

$$\overline{S}_1 \times \overline{S}_2 \times \dots \times \overline{S}_{\hat{T}-1} \times \overline{S}_{\hat{T}} \quad (44)$$

Figure 3 illustrates the difference between a scenario tree (left) and scenario lattice (right). If we wanted to expand the scenario tree with more time steps, the number of extra nodes needed would grow exponentially for each step. For the lattice, we can ourselves choose the number of nodes we add, and it is often common to operate with a fixed number of nodes per stage, as in Löhndorf and Wozabal (2017). Therefore, a lattice is less computationally expensive for the representation of all possible future states than a scenario tree.




Figure 3: Scenario tree (left) and scenario lattice (right) with five time steps. The scenario tree contains 31 nodes representing 16 scenarios, whereas the lattice contains 15 nodes representing 120 scenarios. Source: Löhndorf and Wozabal (2017)

3.4.2 Lattice quantization

In order to construct our lattice, we use the two-step method proposed by Löhndorf and Wozabal (2017). The first step involves optimally locating all nodes and the second step is to find the

correct transition probabilities. For a time stage t with N_t nodes, the location of all nodes is found by minimizing the Wasserstein distance between the nodes and a set of simulated observations of price and inflow. Consider K Monte-Carlo simulations (\hat{S}^k) where \hat{S}_t^k is the time t state of simulation k where $k \in [K]$. Then, the Wasserstein distance between a simulated state S_t^k and its closest node \overline{S}_{tn} is given by the vector distance $\|\hat{S}_t^k - \overline{S}_{tn}\|$. It is, however, difficult to find the optimal nodes \overline{S}_{tn} analytically. We, therefore, use a method based on stochastic gradient descent first proposed by Bally and Pages (2003) and later redeveloped by Löhndorf and Wozabal (2017). We define a set of stepsizes $\beta = \{\beta_k : k \in [K]\}$. Further, we define \overline{S}_{tn}^k as

$$\overline{S}_{tn}^k = \left\{ \begin{array}{ll} \overline{S}_{tn}^{k-1} + \beta_k(\hat{S}_t^k - \overline{S}_{tn}^{k-1}) & \text{if } n = \arg \min_m \{\|\hat{S}_t^k - \overline{S}_{tm}^{k-1}\|, m \in [N_t]\} \\ \overline{S}_{tn}^{k-1} & \text{otherwise} \end{array} \right\} \quad (45)$$

where $\overline{S}_{tn}^0 \equiv 0$, $n \in [N_t]$, $t = 2, 3, \dots, \hat{T}$ and $k \in [K]$. Then, the value of all nodes \overline{S}_{tn} can be found by setting $\overline{S}_{tn} \equiv \overline{S}_{tn}^K$. In order to ensure that the resulting nodes are actually local minimizers of the Wasserstein distance, β_k must be defined such that $\sum_{k=1}^{\infty} \beta_k = \infty$ and $\sum_{k=1}^{\infty} \beta_k^2 < \infty$.

Having found all lattice nodes, we must now find all transition probabilities between them. We denote p_{tnm} as the conditional probability of a node transition from \overline{F}_{tn} to $\overline{F}_{t+1,m}$. Bally and Pages (2003) propose estimating the transition probabilities by (46), a method which we choose to adopt.

$$p_{tnm} = \frac{\sum_{k=1}^K \mathbb{I}_{\Gamma_{tn}}(\hat{S}_t^k) \mathbb{I}_{\Gamma_{tm}}(\hat{S}_{t+1}^k)}{\sum_{k=1}^K \mathbb{I}_{\Gamma_{tn}}(\hat{S}_t^k)}. n \in [N_t], m \in [N_{t+1}], t \in [T-1] \quad (46)$$

Here, $\mathbb{I}_A(x)$ is an indicator function taking the value 1 if x is a part of the set A and 0 otherwise. Γ_{tn} is the Voronoi decomposition of the nodes \overline{S}_{tn} . In other words, Γ_{tn} is the set of simulated states \hat{S}_t^k whose nearest node defined by the Wasserstein distance is \overline{S}_{tn} . Mathematically, Γ_{tn} can be written as

$$\Gamma_{tn} = \left\{ \hat{S}_t^k : n = \arg \min_m \{\|\hat{S}_t^k - \overline{S}_{tm}\|, m \in [N_t]\} \right\} \quad (47)$$

Intuitively, (46) implicates that the transition probability p_{tnm} is given by the number of simulated paths whose time t and $t+1$ states lie closest to the nodes \overline{S}_{tn} and $\overline{S}_{t+1,m}$, respectively, divided by the total number of paths going from node \overline{S}_{tn} to any time $t+1$ node.

3.4.3 Monte Carlo simulations

We have now established a method for constructing a lattice. As mentioned, a set of simulated state paths are needed to do this. In this section, we describe how Monte Carlo simulations are used to construct correlated time series for our two state variables; spot price and inflow. Recall that movements of a forward contract $F_{t,T}$ are given by

$$\frac{dF_{t,T}}{F_{t,T}} = \sum_{i=1}^N \sigma_{i,\tau} dZ_{i,t} \quad (48)$$

where $dZ_{i,t}$ for $i = [1, \dots, N]$ are uncorrelated increments of a Wiener process and $\sigma_{i,\tau}$ is a volatility function found empirically using PCA. Since the principal components used to construct the volatility functions are orthogonal, the movements described by each volatility function must be independent. Therefore, $dZ_{i,t}$ and $dZ_{j,t}$ for $i \neq j$ are independent for all values of t . Using Ito's lemma, (48) can be represented in logarithmic form as

$$d \ln F_{t,T} = -\frac{1}{2} \sum_{i=1}^N \sigma_{i,\tau}^2 dt + \sum_{i=1}^N \sigma_{i,\tau} dZ_{i,t} \quad (49)$$

For small time changes Δt , (49) can be discretized into

$$\ln F_{t+\Delta t,T} - \ln F_{t,T} = -\frac{1}{2} \sum_{i=1}^N \sigma_{i,\tau}^2 \Delta t + \sum_{i=1}^N \sigma_{i,\tau} \Delta Z_{i,t} \quad (50)$$

Here, $\Delta Z_{i,t}$ can be rewritten as $\Delta Z_{i,t} = \sqrt{\Delta t} \epsilon_{i,t}$, where $\epsilon_{i,t}$ is a random draw from a distribution $\epsilon_{i,t} \sim N(0, 1)$. Using this, the price of a forward contract $F_{t,T}$ with time to delivery $T - t$ is given by

$$F_{t,T} = F_{t-\Delta t,T} \cdot \exp \left(-\frac{1}{2} \sum_{i=1}^N \sigma_{i,\tau+\Delta t}^2 \Delta t + \sum_{i=1}^N \sigma_{i,T-t+\Delta t} \sqrt{\Delta t} \epsilon_{i,t} \right) \quad (51)$$

We try to get an intuitive interpretation of (51). If we set $\Delta t = 1$ week, $F_{t,t+n\Delta t}$ for $n = [1, \dots, T]$ is the time t price of a forward contract with n weeks to maturity and delivery period Δt . Next week, the price of the forward contract will evolve to $F_{t+\Delta t,t+n\Delta t}$ following the process above, and become the price of a forward contract with $n - 1$ weeks to maturity. If we set $n = 1$, the price of a forward contract $F_{t,t+\Delta t}$ will evolve into $F_{t+\Delta t,t+\Delta t}$, that is, the next week spot price $P_{t+\Delta t}$. (51) can therefore be rewritten into an expression of the time t spot price P_t as a function of $F_{t-\Delta t,t}$, given by

$$P_t = F_{t,t} = F_{t-\Delta t,t} \cdot \exp \left(-\frac{1}{2} \sum_{i=1}^N \sigma_{i,\Delta t}^2 \Delta t + \sum_{i=1}^N \sigma_{i,\Delta t} \sqrt{\Delta t} \epsilon_{i,t} \right) \quad (52)$$

Further, we recall that the time t inflow I_t is given by (43). We modify it by rewriting $\epsilon_t \sim N(0, \sigma_{\text{INF},t}^2)$ into $\epsilon_{\text{INF},t} \cdot \sigma_{\text{INF},t}$ where $\epsilon_{\text{INF},t} \sim N(0, 1)$ to get

$$I_t = \exp(\epsilon_{\text{INF},t} \cdot \sigma_{\text{INF},t} + \hat{\mu}_t - \phi_t \hat{\mu}_{t-1}) I_{t-1}^\phi \quad (53)$$

Thus, the Monte Carlo simulated times series used to construct the lattice will consist of parallel paths of (53) for $t \in [1, \dots, \hat{T}]$ and (51) for $t \in [1, \dots, \hat{T}]$ and $T \in [t, \dots, \hat{T}]$. If the number of discrete time steps used in the model is \hat{T} , the first time step of the simulation must include the starting period spot price and $\hat{T} - 1$ forward contracts. For each consecutive time step, the contract from the previous step with the shortest time to maturity becomes the new spot price, thereby reducing the number of contracts in the set. Consequently, the simulation will only contain one

price for the last time step, the spot price $P_{\hat{T}}$. Though we must perform Monte Carlo simulations for many forward prices with maturity dates up to \hat{T} , only the spot price $P_t = F_{t,t}$ is included in the lattice nodes \overline{S}_{tn} .

Further, we want our model to include a correlation between movements in price and reservoir inflow. We implement this by adding a correlation between the normalized increment of the Wiener process associated with the first volatility function ($\epsilon_{1,t}$) and the inflow function residual ($\epsilon_{\text{INF},t}$). In order to do this, we use Cholesky decomposition. For two random draws $\epsilon_{a,t}$ and $\epsilon_{b,t}$ with distributions $\epsilon_{a,t}, \epsilon_{b,t} \sim N(0, 1)$, the two residuals $\epsilon_{1,t}$ and $\epsilon_{\text{INF},t}$ with correlation ρ can be found by the following expressions

$$\epsilon_{1,t} = \epsilon_{a,t}, \quad \epsilon_{\text{INF},t} = \rho\epsilon_{a,t} + \epsilon_{b,t}\sqrt{1 - \rho^2} \quad (54)$$

In McDonald (2014), it is shown that $\epsilon_{1,t}$ and $\epsilon_{\text{INF},t}$ will in fact have a correlation ρ if they are defined as in (54).

3.5 Optimization algorithm used on lattice

To conclude this section, we give an overview of how the optimization algorithm used on the price and inflow lattice works. In principle, obtaining optimal decision policies for a Markov decision process should be possible using traditional dynamic programming (DP) as introduced by Bellman (1957). In DP, the optimal decision policy for each state can be found using the Bellman equation, which we first introduced in Section 2.2. We rewrite it so that its notation is coherent with the one we use to represent all lattice states, giving us

$$V_t(\overline{S}_{tn}, \pi_{tn}) = \max_{\pi_{tn}} CF_t(\overline{S}_{tn}, \pi_{tn}) + \exp(-r)\mathbb{E}[V_{t+1}(S_{t+1}, \pi_{t+1}|\overline{S}_{tn}, \pi_{tn})] \quad (55)$$

Here, π_{tn} is the optimal decision policy of the n th state at time t . The objective of the production planner is to maximize the value of their future cash flows (also known as the value function), V_t , subject to their current state and decision policy. All future cash flows are represented by the current time cash flow CF_t plus the expected discounted next state value function. To find the optimal decision policy of each state, we must start in the end nodes. There, one assumes that the world has no future states, and the value function is given by the present time cash flows. Thus, the optimal decision policies in the ultimate states is the one maximizing these cash flows. Having found these policies and the corresponding value functions, we can move one step back in the lattice. There, one can use (55) to find the optimal decision policies and value functions for all states, using the next state value functions found in the ultimate states. By repeating this process for all time steps, one can obtain the optimal decision policy of all values of the lattice.

A common problem with dynamic programming is the "curse of dimensionality". It has been addressed by multiple authors, more recently by Mes and Rivera (2017). In their paper, they list three "curses". The most critical in our case is that the decision space can become too large to find the optimal decision for all states within a reasonable amount of time. We must, therefore, use a method that resolves this issue by obtaining decision policies that are approximately optimal. Multiple such methods are proposed in the literature, and they are often referred to as approximate dynamic programming. A method that has been widely used to manage hydropower reservoirs is stochastic dual dynamic programming (SDDP), first introduced by

Pereira and Pinto (1991). Löhndorf and Wozabal (2017) extend the method of SDDP so that it can also be used for scenario lattices, calling it ADDP. When using SDDP and ADDP, one of the main simplifications is that the value function V_t is approximated to be a piece-wise linear, concave function of all resource variables (e.g., reservoir levels). The shape of the value function is illustrated in Figure 4. We adopt the method of ADDP and give an overview of how ADDP can be applied to find the optimal decision policies π_{tn} for each state in a lattice. In order to do so, one must use the following algorithm for $\xi \in [\Xi]$ iterations:

1. Perform a *forward pass* by drawing a sequence of states $(\hat{S}_{tn}^\xi)_{t=1}^{\hat{T}}$ through the lattice based on the state transition probabilities. For each state, find an optimal decision policy π_{tn}^ξ such that it maximizes the approximate post-decision value function denoted \hat{V}_t^ξ , which is defined as in the Bellman equation, for each drawn state.
2. Perform a *backward pass* where the approximated value function is updated relative to the sampled sequence of states and all state decision policies. In practice, this is done by adding new hyperplanes (linear constraints) to the sets of supporting hyperplanes that define the approximate post-decision value function. These linear constraints are illustrated in Figure 4.
3. If $\xi \leq \Xi$, return to step 1. If not, terminate and set $V_t = \hat{V}_t^\Xi$ and $\pi_{tn} = \pi_{tn}^\Xi$ for all states.

Here, we use that $\hat{V}_t^0 = 0$. For a more detailed description of the algorithm, the reader should consult Löhndorf and Wozabal (2017). Compared to traditional DP, ADDP can significantly reduce the computational time needed to find all the decision policies by choosing them so that they are close to optimal. To further speed up the process, we also use the ε -approximation introduced by Löhndorf et al. (2013). It says that the hyperplanes made during the backward pass should be rejected if they do not improve the approximate post-decision value function by more than $\varepsilon > 0$. This will result in a looser approximation, but it decreases the problem size and thereby the computational efforts needed to obtain the approximate optimal policies.




Figure 4: Illustration of how the linear constraints C_1 , C_2 and C_3 construct the approximate value function as a function of the water level in reservoir b .

One of the outputs from the model is the optimal current state decision policy π_1 . This policy tells us what decision the production planner should make right now to maximize their expected, long-term revenues. However, the model can also be used to calculate the expected revenues or

the mean of the optimal decisions in all future time steps. In order to obtain these values, one must simulate multiple state paths through the lattice. In all simulations, the transition from one state to another is drawn based on the transition probabilities of going from that particular state to any other states. Since some of the decision variables (e.g., water level $l_{b,t}$) are dependent on their previous state value, no single solution can be found for all decision variables in all states. Instead, we find their possible value in each time step based on the decisions made in this time step for all simulated paths, and then set their expected value equal to the mean of the values found in the different simulations.

4 Case: Sjøa hydropower plant

In this section, we present the hydropower plant for which we have developed the initial version of our model. Also, we present a revised version of the optimization model given the current case and some necessary simplifications.

4.1 Sjøa hydropower plant

To be able to test our model, we have received empirical data from the Sjøa hydropower plant, a plant owned and operated by the integrated electric utility company TrønderEnergi. Apart from sharing the relevant characteristics of the plant, TrønderEnergi has also provided us with historical time series for inflow and production. The plant is mid-sized both in terms of regulating capacity and power capacity, and it is located in the NO3 area in Central Norway. It consists of two reservoirs - Vasslivatn and Sjøvatn, and one Francis turbine. The discharge from the Sjøvatn reservoir to the Vasslivatn reservoir is controllable. In Table 3, we have listed the physical boundaries of both reservoirs. There is also a special summertime restriction that applies for Sjøvatn, which is set by local authorities. This restriction and its duration are also listed in Table 3. The outlet of the hydropower plant is in Hemnefjorden, which has an average head of -1 MASL (meters above sea level).




Figure 5: The Sjøa hydropower plant and the reservoir capacities. The elevation of 273.1m is the production-weighted average head difference between Vasslivatn and Hemnefjorden.

Table 2: Characteristics of the Sjøa hydropower plant

	Value	Unit
Maximum power capacity	36	MW
Mean yearly production	191.3	GWh
Avg. yearly inflow, total	311	mill m^3
Average inflow to Sjøvatn	60.5	% of total
Average inflow directly to Vasslivatn	39.5	% of total
Energy coefficient	0.6748	kWh/ m^3
Turbine capacity	17	m^3/s

The energy coefficient takes into account all sources of energy loss in the system, including head loss, turbine losses, generator losses and transformer losses. It is calculated using the production-weighted average head elevation (273.1 m) and production-weighted average discharge to the turbine.

Table 3: Water level constraints for Soa. All water levels are denoted in meters above sea level

Reservoir	Restriction type	Min [MASL]	Max [MASL]
Vasslivatn	Physical	260.00	279.83
Sovatn	Physical	275.00	279.83
Sovatn	Regulatory (May 25 - Oct. 15)	278.33	279.83

4.2 Revised optimization model

In order to construct the lattice mentioned earlier and perform ADDP on Soa hydropower plant, we have used the Java library QUASAR. QUASAR is a general-purpose solver for stochastic optimization developed by Dr. Nils Lohndorf and provided by Quantego (Quantego (2017)). Though the software is highly sophisticated and can be used to solve large-scale problems, it is currently only applicable for linear problems. This forces us to make some adjustments to our optimization model with regards to formulating a linear objective function and linear constraints.

We start with the general model formulation from chapter 2.2 and specify the parameters of our case. The number of reservoirs is $B = 2$, and we let $l_{1,t}$ and $l_{2,t}$ denote the water levels in Vasslivatn and Sovatn respectively. Also, $I_{1,t}$ and $I_{2,t}$ denote the inflows into each reservoir. Since the system only contains one turbine which connects Vasslivatn to the outlet of Hemnefjorden, we denote the amount of water nominated for production as w_t and its discharge as q_t . The amount of water flowing from Sovatn to Vasslivatn (denoted $u_{12,t}$ using logic in Section 2.2) is now denoted by u_t , and the amount of spilled water flowing from Vasslivatn to Hemnefjorden is (previously denoted $u_{1O,t}$) is denoted by s_t . With these new notations, the general volume balance constraint from (4) can be written as (56) and (57) for Vasslivatn and Sovatn, respectively.

$$l_{1,t} = l_{1,t-1} - w_t - s_t + I_{1,t} + u_t \quad \text{for } t = [1, \dots, \hat{T}] \quad (56)$$

$$l_{2,t} = l_{2,t-1} + I_{2,t} - u_t \quad \text{for } t = [1, \dots, \hat{T}] \quad (57)$$

Neither u_t nor s_t are restricted by an upper bound $\overline{u_{ij,t}}$, so this constraint is omitted from the model. Further, the head elevation of Vasslivatn is denoted by H_t and the combined turbine and generator efficiency rate by η_t . Due to the linearity constraint of the optimization software, the cash flow expression CF_t defined by (3) can only be a function of one decision variable. Typically, head elevation H_t is a function of volume level $l_{1,t}$ and turbine efficiency a function of water flow q_t . Generator efficiency is typically a function of both. Multiplying these three parameters with the water flow q_t , as is the case in the cash flow expression, would give a non-linear expression. Thus, we must simplify the model by using constant head elevation and efficiency rate, a simplification made in multiple similar models for reservoir management, e.g., EOPS (SINTEF (2017a)). Madani and Lund (2009) also use a fixed head in their model and argue that this is a reasonable assumption for high-elevation hydropower systems. There is no formal definition of high-elevation plants, but they typically have a head elevation above 250-300

meters. As the head elevation of Sjøa is within this interval, it is not highly unreasonable to argue for using a constant head. Also, if the head is chosen as the centre of gravity for the reservoir (i.e., about 270 MASL, indicating an elevation of 271 meters between the reservoir and the outlet), the deviations between realized power and approximated power will be in the range $[-3.7\%, 3.7\%]$. We believe this is acceptable, considering the granularity of our model.

Since we must use constant values for head elevation and efficiency rate, the objective function of the optimization problem suddenly consists of many constants whose product is the energy coefficient. By definition, the energy coefficient is the average amount of energy a hydropower plant can produce by using one cubic meter of water. In the objective function, we, therefore, make the simplification $\rho GH\eta\varpi/\varsigma = EC$, where EC denotes the energy coefficient. As mentioned above, the mean energy coefficient has been provided to us by TrønderEnergi. It is the same as the mean energy coefficient used in EOPS, the software TrønderEnergi is currently using for medium-term reservoir management.

Further, we only have available data on the aggregated inflow into both reservoirs, forcing us to model inflow as a single stochastic variable $I_t = I_{1,t} + I_{2,t}$. In order to obtain $I_{1,t}$ and $I_{2,t}$, we have used the historical inflow split given in Table 2. We let $\iota = 0.395$ denote the historical amount of inflow flowing into Vasslivatn, and thereby set $I_{1,t} = \iota I_t$ and $I_{2,t} = (1 - \iota)I_t$. Also, since the water level in Sjøvatn is subject to a minimum restriction during the summer $\overline{l_{2,t}} > 0$, we must include a dummy variable $l_{2,t}^{DUM}$ in our model to account for cases where this constraint cannot be held. Since we do not know the exact cost of violating the constraint, we add a sufficiently large cost Υ associated with the dummy variable to the value function such that its value is kept at a minimum. By combining all the mentioned simplifications and adjustments, our optimization problem at time t is reduced to

$$\begin{aligned}
\max \quad & V_t = P_t \cdot EC \cdot w_t - \Upsilon \cdot l_{2,t}^{DUM} + \exp(-r)\mathbb{E}[V_{t+1}|P_t, I_t, \pi_t] \\
\text{subject to} \quad & l_{1,t} = l_{1,t-1} - w_t - s_t + \iota I_t + u_t \\
& l_{2,t} = l_{2,t-1} + (1 - \iota)I_t - u_t \\
& l_{1,t} \leq \overline{l_{1,t}} \\
& l_{2,t} \leq \overline{l_{2,t}} \\
& l_{1,t} \geq \underline{l_{1,t}} \\
& l_{2,t} \geq \underline{l_{2,t}} \\
& q_t \leq \overline{q}
\end{aligned}$$

where $\pi_t = \{w_t, l_{1,t}, l_{2,t}, u_t, s_t, l_{2,t}^{DUM}\}$. All coefficients and constant parameter values are given in Table 4. We recall that water discharge is defined as $q_t = w_t/\varsigma$ where ς is the number of seconds of production per week. The larger we choose ς , the larger becomes the maximum limit for w_t , water nominated for production at time t . In cases of large inflows, low values for ς will only result in larger amounts of spilled water, indicating that ς should be set as large as possible. Also, since efficiency rate is not modeled as a function of water discharge q_t , the choice of ς will be irrelevant for the value function in all stages where spillage is of no concern. We, therefore, set $\varsigma = 604800s$, which is the total number of seconds in one week. We also let the risk-free rate r be equal to the Norwegian Interbank Offered Rate (NIBOR). To get comparable results between runs for different days, we chose to use a constant value of r (NIBOR for 6-month maturity debt on January 7, 2013) in our model. Optimally, we would have used an estimate of the two-year maturity risk-free rate, but six months was the longest maturity available.

Table 4: Model coefficients and constants

Coefficient/Parameter	Value	Unit	Dates
$\overline{l_{1,t}}$	44.5	Mm^3	$t = [1, \dots, \hat{T}]$
$\overline{l_{2,t}}$	22.5	Mm^3	$t = [1, \dots, \hat{T}]$
$l_{1,t}$	0	Mm^3	$t = [1, \dots, \hat{T}]$
$\underline{l_{2,t}}$	0	Mm^3	$t = [\text{October 16}, \dots, \text{May 24}]$
$\overline{l_{2,t}}$	15.05	Mm^3	$t = [\text{May 25}, \dots, \text{October 15}]$
\overline{EC}	0.6747	kWh/m^3	$t = [1, \dots, \hat{T}]$
ς	604800	s	$t = [1, \dots, \hat{T}]$
\overline{q}	17	m^3/s	$t = [1, \dots, \hat{T}]$
r	0.0198	–	$t = [1, \dots, \hat{T}]$

5 Results

In this section of our thesis, we show the results obtained by running the model. In the first part, we show and discuss the forward curves and corresponding volatility functions found using the methods of Benth et al. (2007), Fleten and Lemming (2003) and Alexander (2008b). Then, we present the parameters of the inflow model and show how we construct the scenario lattice using correlated Monte Carlo simulations. Furthermore, we present the most important results and decisions obtained when running the model on five different dates. Next, we plot the expected marginal water values as a function of reservoir level and time. Then, we perform a set of sensitivity analyzes with regards to different values of the correlation coefficient ρ and the number of factors N used in the forward price model. We also perform a backtest of our model compared to historical operations and discuss some simple hedging strategies. Except for the algorithms used to construct a lattice from a set of Monte Carlo simulations and to obtain optimal decision policies in a lattice using ADDP, all lines of code have been written by us using MATLAB and R.

5.1 Realized forward curves and corresponding volatility curves

By first constructing forward curves according to the three methods described in Section 3, we have obtained three different sets of volatility functions that describe forward price movements. Remember that each volatility function is associated with an independent uncertainty factor. The volatility function determines by how much, and in which direction the random shock associated with the uncertainty factor moves each point of the forward curve. The number of factors, and thus volatility functions, needed to explain the entire forward curve dynamics will depend on the return series that is used when finding the volatility functions by principal component analysis. As we use weekly granularity and a time horizon of 105 weeks, the volatility functions $\sigma_{i,\tau} = \Psi_i(\tau)$ must be constructed for the same granularity and length. That is, we must find the volatility functions $\sigma_{i,\tau}$ for all $i = [1, \dots, N]$ and time to maturity given by $\tau = [1, 2, \dots, 104]$ weeks.

In all three methods, the forward curves were constructed using closing prices of DS futures contracts traded NASDAQ Commodities. These contracts are listed in Table 5. The dataset includes closing prices for all trading days between April 28, 2011 and June 30, 2017. As discussed in Section 2.1, futures contracts and DS futures contracts are treated as equal despite their difference in settlement structure. Hence, our results are not to be understood as specific for DS futures contracts.

Table 5: Electricity forward contracts traded on NASDAQ Commodities

Code	Length of delivery period	Contracts used
ENOW	Week	1, 2, 3, 4, 5, 6 weeks ahead
ENOM	Month	1, 2, 3, 4, 5 and 6 months ahead
ENOQ	Quarter	1, 2, 3, 4, 5, 6, 7 and 8 quarters ahead
ENOYR	Year	1, 2 and 3 years ahead

5.1.1 Forward and volatility curves found using the method of Benth, Koekebakker and Ollmar

Based on data for all trading days between June 29, 2015, and June 30, 2017, we have constructed 507 forward curves using the method of Benth et al. (2007). In Figure 6, the forward curve obtained for June 29, 2015, is illustrated together with its underlying contracts. Further, Figure

7 illustrates multiple forward contracts constructed for all trading days between June 29 and September 18, 2015. While both figures show that the constructed forward curves tend to be smooth and move up and down together, they reveal three weaknesses with the method:

1. **Discontinuity:** Though it is included as a requirement, many forward curves are not strictly continuous. This can be seen for $t = 1.25$ years in Figure 6, where the curve makes a jump in the downward direction.
2. **Obviously wrong prices:** One of the requirements stated in the model says that if there exists a forward contract with delivery period (t_b, t_e) whose price is F_{t_b, t_e} , the average value of the forward curve for that interval must equal F_{t_b, t_e} . In some cases, the forward curve obviously does not comply with this requirement as it lies either entirely above or below F_{t_b, t_e} for all values in the interval. This is also visible in Figure 6, where the curve lies below the price of the fifth quarterly contract (illustrated by the green line ending at $t = 0.5$) for the entire duration of the contract.
3. **Large oscillations:** Some forward curves oscillate unreasonably much, giving unrealistic forward prices. These oscillations tend to occur during the first six-seven weeks of the forward curves, which is where weekly and monthly contracts overlap. For this interval of time, complying with the requirement mentioned in point 2 is more challenging, and the algorithm solves this by creating an oscillating curve. The possible magnitude of these oscillations is visible in Figure 7, though there were also cases where the curves oscillated to prices below 0 EUR/MWh or above 1000 EUR/MWh.




Figure 6: Forward curve constructed on June 29, 2015. The horizontal lines illustrate the price and duration of the forward contracts used to construct the curve

Even if the first two weaknesses are not negligible, the third one is the most critical for our purpose. Our objective by constructing forward curves was to use them to obtain a dataset of returns series. Applying the methods of Section 3.2.4 would give us a set consisting of many daily returns taking unrealistically high or low values. We try to resolve this issue by removing all forward curves that oscillate to values beyond certain limits. The lowest and highest observed forward contract prices in our dataset are respectively 8.73 EUR/MWh and 46.91 EUR/MWh. Thus, we set the filtering limits 2 EUR/MWh below and above these values, respectively. That is, we remove all forward curves with values below 6.73 EUR/MWh or above 48.91 EUR/MWh. This requires us to remove 42 curves, 8.3% of the original 507. Doing this, we obtain the volatility functions illustrated in Figure 8.

Forward curves constructed between June 29, 2015 and September 18, 2015




Figure 7: Forward curve constructed for 60 trading days between June 29, 2015 and September 18, 2015. The X-axis denotes time t in daily granularity and the Y-axis denotes the date for which the forward curve is constructed. For both axes, $t = 1$ for June 29, 2015




Figure 8: Volatility curves for $N = 6$ using method of Benth et al. (2007). 'Overall' denotes the overall volatility function, and $fn\ i$ denote the volatility functions given by principal component i , that is, $\sigma_{i,\tau}$

In Figure 8, the overall volatility function can be understood as the volatility of price returns of forward contracts with time to maturity τ . It is equivalent to $\sigma_{t,\tau}$, which was introduced in the single-factor model in (12). Since it represents the actual volatility of forward price returns, it will always be positive. Also, the overall volatility function can be found by adding up the volatility functions associated with all 104 principle components. These functions must, however, not be interpreted the same way as the overall volatility function, as they do not represent the volatility in price movements of a single asset. Therefore, they can also take negative values, as opposed to the overall volatility function.

One would expect the overall volatility function to be strictly decreasing for ascending values of τ , as forward prices tend to change more the closer they come to maturity. This is called the *Samuelson effect*, discussed by Jaeck and Lautier (2016) and originally proposed by Samuelson

(1965). The reasoning behind this phenomenon is that an information shock that affects the short-term price has an effect on the succeeding prices that decreases as the time to maturity increases. Weather forecasts are an example of information that one would expect to have short-term effects only on the electricity price.

By visual inspection, the volatility functions seem quite unreasonable. Firstly, one would expect the overall volatility function to be strictly decreasing, due to the Samuelsen effect. In our case, the maximum volatility is recorded for contracts with seven weeks to maturity, that is, in the region where the forward curves oscillate the most. Secondly, Koekebakker and Ollmar (2005), who also investigated the Nordic electricity forward market, record no annualized volatility above 80%. In our case, we get values going above 200%. Further, Table 6 shows the cumulative proportion of explained variance when using different numbers of principal components and corresponding volatility functions. In order to reach a level of 95% explained variation, which Koekebakker and Ollmar (2005) use as a target in their paper, 19 components (or factors) are needed. Koekebakker and Ollmar (2005) only needed 11 factors to obtain the same explanation.

Table 6: Proportion of explained variance for different numbers of explanatory factors

Number of factors N	1	2	3	4	5	...	18	19
Explained variance	0.2561	0.4004	0.4974	0.5850	0.6474	...	0.9447	0.9520




Figure 9: Correlation matrix obtained using method of Benth et al. (2007). The column and row names are both the number of weeks to delivery

If we look at the correlation matrix (Figure 9) of our dataset, which serves as the basis for constructing the volatility functions, we see why the functions are of poor quality. Despite that one would expect all correlations to be positive (in the market, the prices of all forward contracts tend to move in the same direction between two consecutive trading days), some

forward contracts are negatively correlated in our case. Further, Alexander (2008a) presents two different models for correlation of the term structure of interest rates, both of which display an almost monotonic decrease in correlation for increasing contract maturity spreads. We would expect this to be true for a term structure of forwards as well, but as one can see, forward contracts with larger maturity spreads are sometimes significantly more correlated than contracts with lower spreads. Consequently, we reject the volatility functions obtained by this method and move on to the other two methods.

5.1.2 Forward and volatility curves found using the method of Fleten and Lemming

Using the same dataset of forward prices as in the previous method, we have also constructed 507 forward curves with the method of Fleten and Lemming (2003). We have used daily time steps, meaning that the forward curves are constructed with one value for each day in the interval (t_b, t_e) . In Figure 10, we have displayed four different forward curves found for the trading day of January 7, 2013, using different values of λ , the smoothing parameter. Here, we have removed the two contracts ENOQ-1 (one quarter ahead) and ENOYR-1 (one year ahead), as their delivery periods are entirely spanned by other contracts with shorter delivery periods. We see that for larger values of λ , the curves become smoother. Also, high λ values are needed to simulate seasonality in the long end of the curve, as the less smooth curves become straight lines when their only underlying contracts are yearly ones (ENOYR).

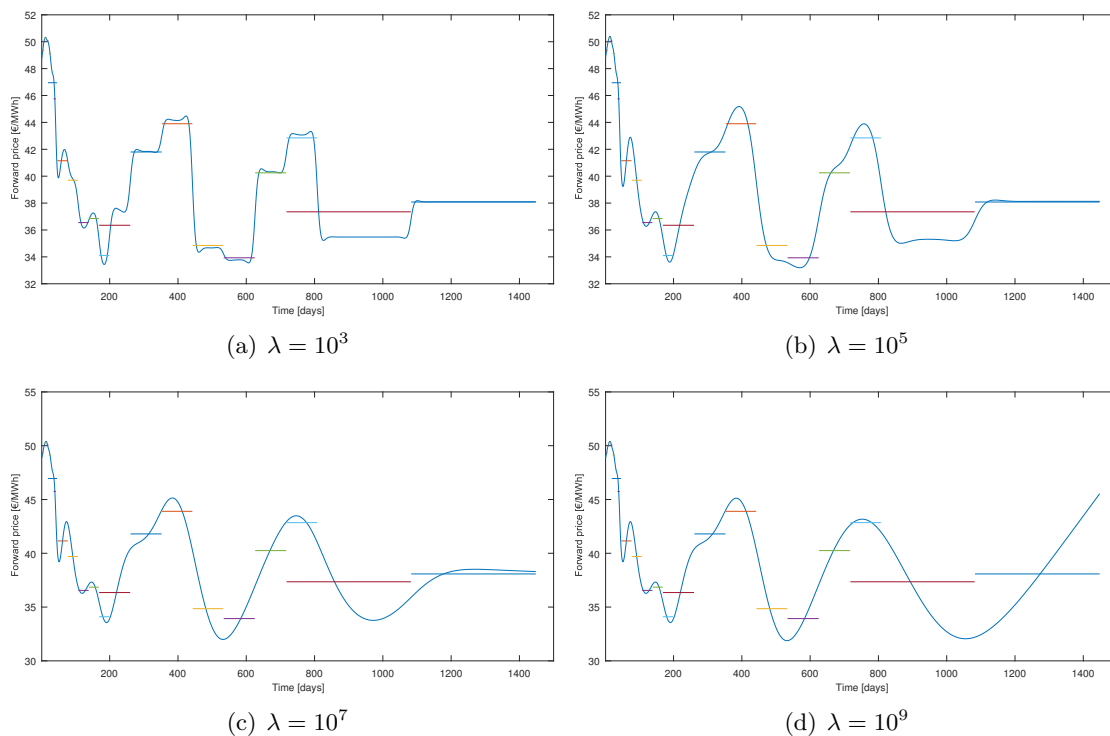


Figure 10: Forward curves obtained for January 7, 2013, using method of Fleten and Lemming (2003) with different values of λ .

By visual inspection, we can see that the first two weaknesses found in the method of Benth et al. (2007) are not an issue when using the method of Fleten and Lemming (2003). Since their method only calculates the value of the forward curve for each day and not as a function of time, the curve will logically appear to be strictly continuous. Additionally, the resulting curves also seem to comply with the requirement that the average value of the curve in a region must equal

the price of a contract covering the same region. We still experience problems with oscillating curves, but the fact that we can use different values of λ resolves this problem to some extent. For all values of λ , we can create a set of forward curves and a corresponding set of volatility functions. Furthermore, we can choose λ such that we minimize oscillations and thereby obtain more plausible volatility curves. In Figure 11, we have displayed four sets of volatility functions constructed using the same values of λ as in Figure 10. As in the previous method, we remove all forward curves that oscillate out of the price region [6.73, 48.91] EUR/MWh from the dataset.




Figure 11: Volatility functions obtained using method of Fleten and Lemming (2003) with different values of λ . 'Overall' denotes the overall volatility function, and fn i denote the volatility functions given by principal component i

Once again, we observe that the volatility functions take very high values for contracts with 7-8 weeks to maturity. In order to lower these oscillations, we try to remove the contract ENOM-1 (one month ahead) when constructing the underlying forward curves as well. The method of Benth et al. (2007) did not allow this, as it would imply that some intervals of the curve were not covered by any forward contracts, but this is possible when using the method of Fleten and Lemming (2003). Doing this, we obtain the volatility functions displayed in Figure 12. Their proportion of explained variance is displayed in Table 7.

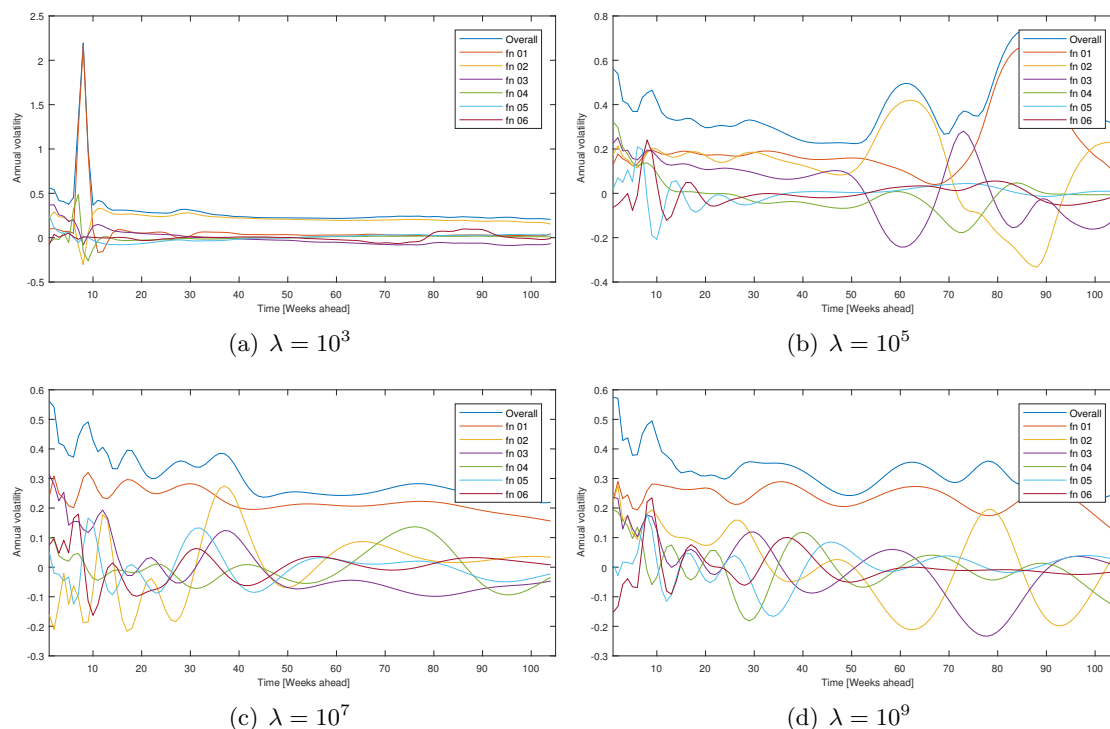


Figure 12: Volatility functions obtained using method of Fleten and Lemming (2003) with different values of λ and without the month ahead ENOM contract. 'Overall' denotes the overall volatility function, and fn i denote the volatility functions given by principal component i

Table 7: Proportion of explained variance for different numbers of explanatory factors using method of Fleten and Lemming (2003). The bottom line indicates the number of factors needed to reach a level of 95% explained variance.

Number of factors	$\lambda = 10^3$	$\lambda = 10^5$	$\lambda = 10^7$	$\lambda = 10^9$
1	0.4918	0.4463	0.5368	0.5179
2	0.8229	0.7167	0.6564	0.6620
3	0.8794	0.8357	0.7483	0.7555
4	0.9116	0.8755	0.7883	0.8049
5	0.9256	0.8918	0.8210	0.8357
6	0.9355	0.9069	0.8490	0.8633
7	0.9445	0.9206	0.8761	0.8876
# of needed factors	8	10	12	11

By visual observation and using the same arguments as for the method of Benth et al. (2007), we see that the volatility functions calculated for $\lambda = 10^3$ and $\lambda = 10^5$ seem implausible. The curves obtained for $\lambda = 10^7$ and $\lambda = 10^9$ seem more plausible, as the overall curve shows a more general decreasing trend for contracts with longer time to maturity. In fact, the obtained overall curves have quite similar shapes as the one Bjerksund et al. (2008) find in their paper. The number of factors needed to explain 95% of the forward price variance is also quite reasonable, as it is similar to the one obtained by Benth et al. (2007). However, they both show large volatilities for contracts with 6-7 weeks to maturity, indicating that we have not fully managed to damp all unreasonable oscillations. We try to analyze how much sense the volatility functions make by looking at their correlation matrices. Unlike the case with the method of Benth et al. (2007), we experience no negative correlations. We do, however, see that multiple contracts with larger maturity spreads are more correlated than contracts with significantly shorter spreads, which

by the logic of Alexander (2008a) should not be the case. Hence, we conclude that the method of Fleten and Lemming (2003) provides us with more reasonable volatility curves than that of Benth et al. (2007), but still not as good results as one would expect.

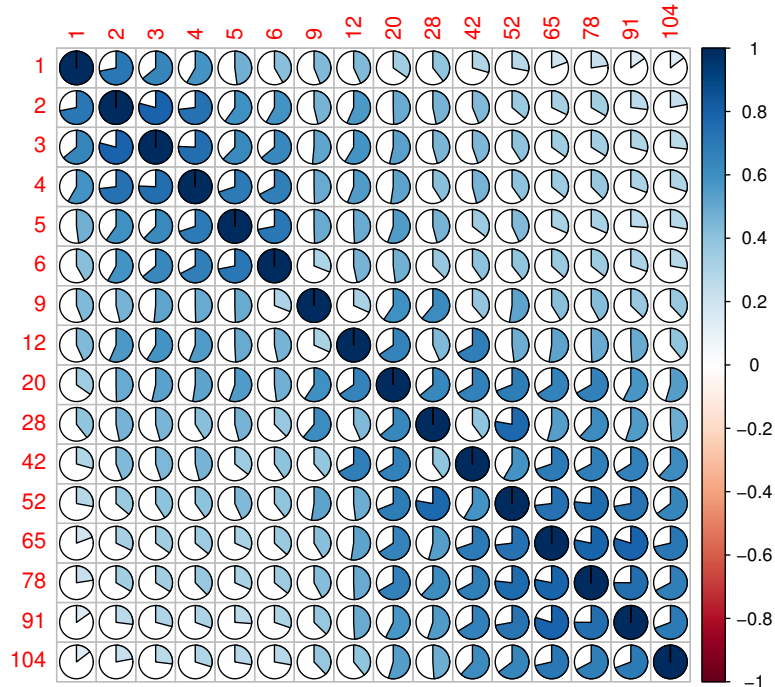


Figure 13: Correlation matrix obtained using method of Fleten and Lemming (2003) with $\lambda = 10^7$. The column and row names are both the number of weeks to delivery

5.1.3 Forward and volatility curves found using linear interpolation

Historical forward curves were found by interpolating between forward prices as described by Alexander (2008a). Out of the three methods investigated, this was arguably the least computationally complicated method. We will see that although the forward curves calculated using this method are not the most realistic, the resulting return series reflect the nature of the forward market better than the return series obtained using the methods of Benth et al. (2007) and Fleten and Lemming (2003).

For this method, a larger dataset was used than in the two previous methods. The dataset included forward prices for all trading days between April 28, 2011, to December 30, 2016, resulting in 1450 forward curves. Figure 14 shows the forward curve found for January 7, 2013.

Remember that initially, for every trading day, one forward curve was created for each contract type associated with a particular delivery period length. That resulted in four forward curves (week, month, quarter and year) for each trading day. These forward curves cover different time intervals. In order to obtain a forward curve spanning the entire planning horizon (1 to 104 weeks ahead), it was necessary to use parts of all four forward curves. The forward curves overlap for some time intervals. To decide which contract types to use in the complete forward curve for these overlapping regions, a heuristic was applied: Use whichever contract that has the most price observations in the time series. As can be seen in Table 8, this resulted in weekly contracts being used in the short end of the curve, monthly contracts in the mid-short part of the curve, quarterly contracts in the mid-long part of the curve, and yearly contracts in the




Figure 14: Forward curve created using linear interpolation

long end of the curve. After the complete forward curves were created, a time series of daily log returns was calculated for each relevant time to delivery, resulting in a 1449×104 matrix.

Table 8: Forward contract type used to create forward curves in given intervals of time

Weeks ahead	1 - 4	5 - 21	22 - 91	92 - 104
Contract type used	Week	Month	Quarter	Year

It is clear that the forward curve found using this method will be discontinuous in the points where we switch from one contract type to another, as illustrated in Figure 14. This will, however, not result in volatility spikes or discontinuous volatility functions, as all daily returns are calculated using the parts of the curves associated with the same contract type.

Using the time series of returns, we could estimate an overall volatility curve for the term structure of forward prices, as well as the volatility functions associated with the principal components. As can be seen from the dark blue curve in Figure 15, the overall volatility is *monotonically decreasing*, as we would expect in a term structure of forward prices. While this is not a requirement for the validity of an empirically estimated volatility curve, we would certainly expect the volatility of a forward contract to increase as its maturity approaches, as discussed in Section 5.1.1. The overall volatility curve also does not feature the volatility spike at 7-8 weeks to delivery, which was prominent in the overall volatility curves obtained by the method of Benth et al. (2007) and Fleten and Lemming (2003). That favors this method of creating forward curves.




Figure 15: Volatility functions found by using method of linear interpolation. 'Overall' denotes the overall volatility curve, and fn i denote the volatility functions given by principal component i

The return series resulting from this method show a higher degree of inter-correlation than the return series obtained using the two previous methods. This can be seen from the correlation matrix that is shown in Figure 16. A high degree of inter-correlation is in accordance with

our experience, which is that forward electricity prices more often than not move in the same direction. Further, the correlation matrix shows that there is a clear decreasing trend in the correlation between contracts with larger maturity spreads. This is in contrast to the correlation matrix attained by the method of Fleten and Lemming (2003).

The high degree of inter-correlation is also demonstrated by the explanatory power of the first principal component, which explains 73.3% of the variance in the dataset. This complies with what was described in Section 3.2.5, where we explained that the first principal component represents vertical movements in the entire forward curve. As opposed to in the previous methods, only six principal components are needed to explain 95% of the cumulative variance, as shown in Table 9.

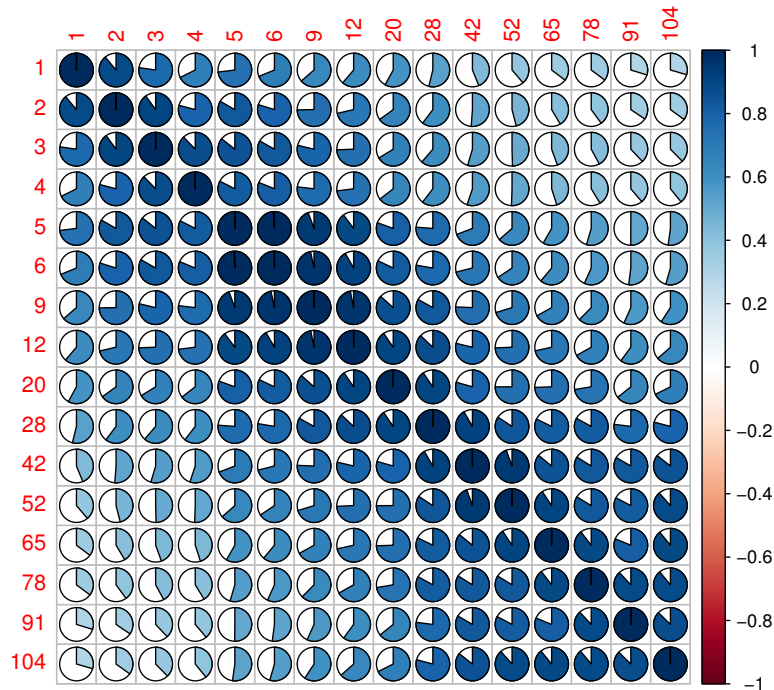


Figure 16: Correlation matrix associated with returns found by the method of linear interpolation. The column and row names are both the number of weeks until the beginning of the delivery period

Table 9: Proportion of explained variance for different numbers of explanatory factors

Number of factors n	1	2	3	4	5	6
Explained variance	0.7338	0.8755	0.9153	0.9348	0.9482	0.9604

5.1.4 Evaluating the methods

When evaluating the methods, it is interesting to discuss why both methods involving creating smooth forward curves to find the volatility functions perform so poorly. Koekebakker and Ollmar (2005) used this method to analyze the Nordic electricity market as well, obtaining forward curves with less or no oscillations. The same is the case for Bjerksund et al. (2008), which analyzed the UK gas market. We do, however, note an important difference in the underlying contracts used in both papers and our thesis. In Koekebakker and Ollmar (2005), price data from 1995 to 2000 is used, a period when the types of forward contracts traded

in the Nordic market where different from now. While we constructed forward curves using weekly, monthly, quarterly and yearly contracts, they used daily, weekly, block and seasonal contracts. The block and seasonal contracts, which are no longer traded in the Nordic market, had delivery periods of four weeks and four months, respectively. This means that Koekebakker and Ollmar (2005) used no overlapping contracts for the near end of the curve, as there are no overlapping weekly and monthly contracts in their dataset. Thereby, they most likely avoid the large oscillations that we experience in this region. There will, of course, be some overlap between block contracts and seasonal contracts, but these contracts have such long delivery periods that an overlap should not result in an oscillating curve. Bjerksund et al. (2008) also use non-overlapping contracts, but they construct their forward curves using the software Elviz Front Manager. Unfortunately, we did not have access to such software, forcing us to implement the algorithms for forward curve construction ourselves. We might have obtained better forward curves using such software. However, we find it reasonable to conclude that the algorithms of Benth et al. (2007) and Fleten and Lemming (2003) perform poorly when the objective is to construct empirical volatility functions in the Nordic electricity market. For commodity markets with overlapping curves in the near end, it is most likely necessary to use more sophisticated algorithms for forward curve construction.

By using the method of Alexander (2008a), the price will be a linear combination of two adjacent contracts which are traded in the market. In comparison, there are much more complicated relationships between the forward curve and the real market prices in the methods of Fleten and Lemming (2003) and Benth et al. (2007). This can, in turn, cause unexpected changes in the forward curves from one day to another that do not necessarily represent the price changes in the market realistically. A criticism that could be directed at *all* forward curve models is that their values $f(t_s, t)$ represent the prices of synthetic products that are not actually traded in the market. This will, in turn, cause the historical returns series to be based on synthetic prices, and not directly on the actual prices observed in the market. However, it is clear that creating forward curves is necessary for several reasons. In our case, we need volatility functions $\sigma_{i,\tau}$ that can describe price movements of forward contracts with time to maturity $\tau = [1, \dots, 104]$ weeks. Thus, we need a dataset containing the price of these 104 contracts. Since many of these contracts do not exist in the market, they must be constructed synthetically, which is where we can make use of forward curves.

Due to all findings, we conclude that we should use the volatility functions obtained using the method of linear interpolation, presented by Alexander (2008a). Its overall volatility function is strictly decreasing, and the correlation matrix shows that there is a clear decreasing trend in the correlation between contracts with larger maturity spreads. It also needs fewer factors to explain a larger proportion of the variance in forward price movements, making it more computationally efficient. In our final model, we, therefore, use the volatility functions associated with the first 6 principal components of the returns series. We do, however, test what lattices and decision policies we would get if we had used the volatility functions obtained using the methods of Benth et al. (2007) and Fleten and Lemming (2003) to see how these compare with the one obtained using the volatility functions from the chosen method. This analysis is performed in Section 5.7.

5.2 Inflow model parameters

We needed to fit the geometric periodic autoregressive (GPAR) model suggested by Shapiro et al. (2013) to the inflow data for the Soa hydropower plant. The dataset consists of daily inflow observations for each day between January 1, 1958, and December 31, 2016. According to TronderEnergi, the data set has been constructed by combining observations from two different sources. The observations from the most accurate source are found by measuring the change in

water level at the reservoirs and finding the inflow by adjusting for water used in production and spilled water. For days without available production data, the inflow is calculated by measuring the water level in the rivers in the catchment area of the hydropower plant.

Remember that the inflow is given by

$$I_t = \exp(\epsilon_t) \exp(\hat{\mu}_t - \phi_t \hat{\mu}_{t-1}) I_{t-1}^{\phi_t} \quad (43)$$

Here,

- I_t is the inflow in week t
- $\hat{\mu}_t$ is the mean log inflow in week $t = 1, \dots, 52$
- ϕ_t is the coefficient in the autoregressive process in week $t = 1, \dots, 52$
- $\epsilon_t \sim N(0, \sigma_{\text{INF},t}^2)$ is the error term representing the difference between the observed and predicted value in the autoregressive process
- $\sigma_{\text{INF},t}$ is the standard deviation of the error terms in week $t = 1, \dots, 52$

The inflow model is *geometric* in the sense that logarithmic inflows are used. This is due to the skewness of the inflow distribution, which can be seen in Figure 17. Similar to the inflow data used by Shapiro et al. (2013), the distribution of weekly inflow observations for the Sjøa hydropower plant is highly right-skewed. To obtain a distribution with less skew, the inflow data is log-transformed.




Figure 17: Distribution of inflow and log inflow observations

The inflow model is *autoregressive* because Z_t , the deviation of the log inflows from their mean, is represented as an AR(1) process. Recall that the AR(1) process is given by

$$Z_t = \phi_t Z_{t-1} + \epsilon_t \quad (58)$$

The suitability of a 1-lag model can be determined by investigating the partial autocorrelation of the historical data for Z_t . Partial autocorrelation is the correlation for a time series with its own lagged variables, but removing the correlation effects of the values of the time series at all shorter lags. Figure 18 shows the partial autocorrelation of the Z_t time series. Similar to the findings of Shapiro et al. (2013), our dataset showed a high value at lag 1 and insignificant values for larger lags, indicating that it is sufficient to include one lag only in the autoregressive model.




Figure 18: Partial autocorrelation of the Z_t time series

The inflow model is *periodic* in the sense that it accounts for seasonality - both in terms the expected weekly log inflow $\hat{\mu}_t$, the strength of the autoregressive coefficient (ϕ_t) and the standard deviations of the error terms ($\sigma_{\text{INF},t}$). Figure 19 shows the seasonal pattern in the inflows. Specifically, there is an inflow peak during the spring due to snow melting, and there are higher inflow levels in the fall due to high precipitation levels in September, October, and November.




Figure 19: Average inflow for a certain week of the year

As mentioned above, the autoregressive coefficient ϕ_t and the standard deviation of the error terms $\sigma_{\text{INF},t}$ are both time-varying. For each week $W = [1, \dots, 52]$, empirical estimates of $\sigma_{\text{INF},t}$ and ϕ_t were found. These are shown Figure 20 and Figure 21. In an attempt to reduce the statistical noise in the parameter estimates, the values of $\sigma_{\text{INF},t}$ and ϕ_t were set to the 5 weeks centered moving average of the empirical estimates. By looking at Figure 20 and Figure 21, it is clear that both $\sigma_{\text{INF},t}$ and ϕ_t are considerably smaller during the snow melting period in the spring.




Figure 20: Time-varying standard deviation of the error terms $\sigma_{\text{INF},t}$. The blue line is the empirical estimates of $\sigma_{\text{INF},t}$, and the orange line is the centered moving average of the empirical estimates, using a span of 5 observations. The moving average is used in the inflow model.




Figure 21: Time-varying autoregressive coefficient ϕ_t . The blue line is the empirical estimates of ϕ_t , and the orange line is the centered moving average of the empirical estimates, using a span of 5 observations. The moving average is used in the inflow model.

5.3 Monte Carlo simulations and lattice construction

In order to create a lattice, we had to run multiple parallel Monte Carlo simulations of water and inflow. As we propose a correlated model, we start by explaining how we calculated the correlation between the residuals of the inflow model and the increment of the Wiener process associated with the first volatility function. Mathematically, this was done by estimating the historical correlation coefficient between the normalized error term of the inflow model and the normalized first principal component (\mathbf{p}_1).

The error term in the inflow model is the difference between the predicted and realized log-inflow. To be able to find a correlation with the weekly inflow data, \mathbf{p}_1 had to be transformed into a weekly resolution as well. Similar to how one would transform daily log returns to weekly log returns, the historical \mathbf{p}_1 observations were aggregated from daily to weekly observations by simple addition.

The resulting Pearson correlation coefficient was found to be -0.1765, based on a time series of 248 weekly observations from April 28, 2011, to December 30, 2016. The 95% confidence interval was [-0.28, -0.06]. This suggests that there has been a weak offsetting effect between weekly inflow deviations and vertical movements in the entire forward curve, historically. A larger dataset would most likely decrease the confidence interval, but this was not possible to obtain in this case.

We move on to the Monte Carlo simulations. In Section 3.4.3, we stated that the starting values of the price simulations included one current time spot price and the price of $\hat{T} - 1$ forward contracts. Since our model uses weekly granularity and a horizon of $\hat{T} = 105$ weeks, this requires 104 weekly forward contracts with time to delivery $\tau = [1, \dots, 104]$. However, only six weekly contracts are traded at NASDAQ Commodities, meaning that we must construct 99 synthetic weekly contracts. This is done by constructing a forward curve using the method of Fleten and Lemming (2003) and then discretizing it into 104 weekly prices. We recall that using the method of Fleten and Lemming (2003) did not provide us with plausible volatility functions. It does, however, produce forward curves of sufficient quality when we manage to damp the oscillations in the near end of the curve, both by removing the first monthly contract and by adjusting the value of the smoothing parameter λ . We still use the volatility functions obtained when we constructed forward curves by the method of Alexander (2008a), but since these forward curves were neither smooth nor continuous, they cannot be used to make synthetic weekly contracts.

Mathematically, the forward price of a contract with delivery in a given week $W = [2, \dots, 105]$ is calculated using the average value of the forward curve within the time interval of that particular week. For the weeks $W = [2, \dots, 7]$, the weekly average value of the forward curve will be the price of the six weekly forward contracts sold in the market. This is due to the single restriction in the method of Fleten and Lemming (2003) for construct forward curves, which stated that the average value of the curve within an interval $(t_{b,j}, t_{e,j})$ had to equal the price of a forward contract with delivery in the same interval. Also, if the model is run on a Monday, the starting week spot price is set equal to the price of the one week ahead weekly contract from the last trading day. Typically, this will be the previous week Friday.

It is important to note that while the spot prices in the Nordic electricity market are area specific, the price of forward contracts is the same for the entire Nordic and Baltic region. Thus, the spot price forecasted by our model is actually the system spot price and not the NO3 area spot price, the price Sjøa hydropower plant receives for their production. In this paper, have not tried to model the relationship between the system price and the NO3 price. We do, however, see that the two prices are quite similar to each other, and believe that using the system price instead of the area price is an acceptable approximation considering the granularity of our model and the scope of this project.

Further, we have used 380.000 Monte Carlo simulations to construct the lattice. Ideally, we should have used more (Löhndorf and Wozabal (2017) use 10^6 simulations), but this was not possible due to computational limits. Each lattice consists of 100 nodes for all time steps except the starting one, giving a total count of 10401 nodes. When obtaining the value of each node \bar{S}_{tn} using (45), we defined the stepsize β_k as

$$\beta_k = \frac{100}{k + 1000}, k \in [K] \quad (59)$$

where K is the number of Monte Carlo simulations. Figure 22(a) displays the spot price lattice with starting date January 7, 2013, while Figure 22(b) displays the inflow lattice. Since the lattice nodes are found by minimizing the Wasserstein distance, we have scaled the inflow values down with a factor of 10^5 such that their magnitudes are closer to those of the spot prices. In the same figure, we also plot the first ten stages of both lattices to illustrate their shape better.

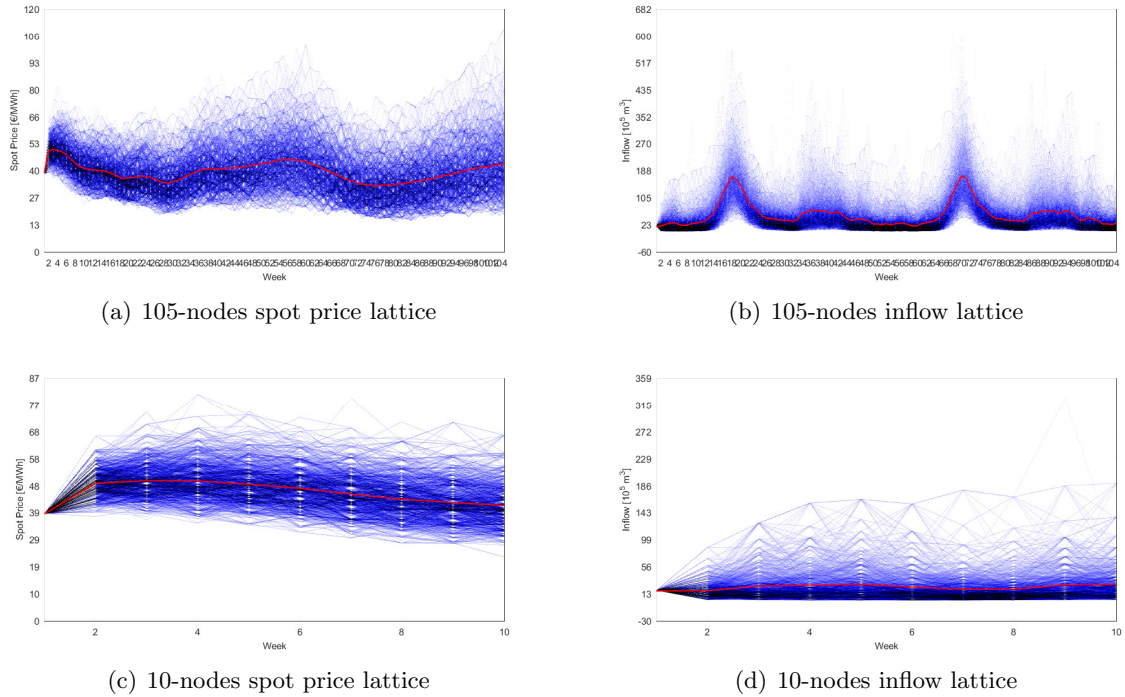


Figure 22: Spot price and inflow lattices constructed on with data for January 7, 2013. The Y-axis for all plots denotes the time steps (weeks), while the X-axis of the price lattice is denotes the spot price in EUR/MWh. For the inflow lattice, the Y-axis is denoted in $10^5 m^3$. The red lines in the figures represent the mean values.

5.4 Different starting dates and expected incremental water values

Now, we run the model for five different starting dates with different underlying forward curves and historical starting values for the reservoir level. One of the key figures we are interested in is the *expected discounted revenues* for the planning horizon. This is the value of production during the next two years, assuming negligible variable costs. The revenues are discounted using the risk-free rate, as we are using risk-free probabilities. Furthermore, a key figure is the expected discounted revenues per produced unit of electricity, which we will call *expected discounted revenues per production*. This figure is denominated in EUR/MWh, and it allows us to compare the performance of policies without differences in total inflow affecting the results. For intuition, this figure can be thought of as the average price at which the hydropower producer sells their power. However, this will be inaccurate in this case, since the average price should be calculated using undiscounted revenues.

Table 10 shows the expected discounted revenues for the upcoming 105 weeks, in addition to the expected discounted revenues per production. These results are based on the revenues obtained by 50.000 simulated paths through the lattice. The number (50.000) is chosen because it enables the first three digits of all mean values to converge, while simultaneously keeping the computation time at an acceptable level. In general, the summer restriction at Søvattnet is violated in 2% of the simulations. Since we have included a substantial cost for such violations in our model, the total revenues earned in these simulations are significantly lower than the other ones, often negative. Hence, we remove these simulations from the set used to calculate the expected discounted revenues and the expected discounted revenues per production.

In Table 10, we also include one of the most important immediate result for the production planner; the value of w_1 . We recall that w_t is the amount of water nominated for production at time step t . Based on all possible future states and their corresponding probabilities, w_1 tells

the production planner how much water they should nominate for production in the current week in order to maximize their expected discounted revenues over the next 105 weeks. We also include the average water dispatch $q_1 = w_1/\varsigma$, where $\varsigma = 604800s$ is the number of seconds per week.

Table 10: Expected discounted revenues (EDR), expected discounted revenues per production (EDR/Prod.), the amount of water nominated for production w_1 , and average water discharge q_1 in week one for five different starting dates. The starting values for reservoir levels are set according to their historical values.

Parameter	Unit	Jan 7-13	Apr 8-13	Jul 8-13	Oct 7-13	Jan 6-14
EDR	[M EUR]	15.93	15.95	14.36	15.26	13.83
EDR/Prod.	[EUR/MWh]	40.12	41.24	35.35	38.36	33.21
w_1	[M m^3]	0	5.21	10.28	0	0
q_1	[m^3/s]	0	8.62	17.00	0	0

It is somewhat surprising that the model suggests no production on multiple starting weeks, especially those of January 7, 2013, and January 6, 2014. However, this is because the input forward curve suggests that the spot prices will be higher in the upcoming weeks, making it optimal to wait. It is also interesting to see that the model expects almost identical revenues in the upcoming two year period both on January 7, 2013, and April 8, 2013, despite the difference in expected discounted revenues per production. This is because less inflow is expected in a two year period from April 8 than January 7, meaning that higher expected prices are necessary to get the same expected discounted revenues.

An important question that arises is how one should handle the end level of the reservoir. Our model does not have any end level restrictions. Thus, the optimal policy in the last time step is to empty the reservoirs, or at least empty them as much as possible. In some of the above simulations, e.g., the ones starting and ending in January, this would probably be a poor decision since one would normally expect high prices in the upcoming periods. Emptying the reservoirs in the last time step will result in expected discounted revenues that are slightly larger than what one would achieve in reality. Therefore, if the model had considered a longer time perspective, the expected two-year revenues at January 7, 2013, should be lower than the ones obtained in our model. The issue with introducing an end level restriction is that the problem might become unfeasible if we try to restrict the end levels for reservoir volume. In order to avoid this, we have therefore chosen not to restrict the decision policies of the ultimate time step. Nevertheless, this should not affect the optimal immediate decision policy π_1 , which is the most interesting one for the production planner, in addition to most decision policies π_{tn} when t is substantially smaller than 2 years. Note that for models in which the time horizon ends during the spring when inflows are typically at their maximum, it is reasonable to allow the model to empty as much of the reservoir as possible.

Further, using the simulation of January 7, 2013, we have calculated the plant's expected incremental water value. The incremental water value is the expected marginal value of the energy stored in the reservoir, i.e., the value of one more unit of water, measured in MWh. In Figure 23, we have plotted the incremental water value as a function of aggregated reservoir level and time. In order to construct the surface, we have used a modified version of the method presented by Tipping et al. (2005). After drawing 50.000 simulated paths, we identified the prices for which it was decided to produce in each week. Then, for each week, based the spot price and the cumulative water volume in both reservoirs for the identified simulations, we grouped the observations in bins for different reservoir levels. The expected marginal water value for a given reservoir level was then calculated as the smallest price of that bin. That price can interpreted

as a critical price for which it is barely optimal to produce rather than wait.




Figure 23: Incremental water values for next 52 week based on simulation starting on January 7, 2013

Looking at the plot, we see that it incorporates most traits one would expect for a water value surface as described in, e.g., Boger and Vestbøstad (2016). For higher reservoir levels, the incremental water value is decreased and vice versa. This makes sense as the incremental water value for a full reservoir should be 0, since this water cannot be utilized and must be spilled. Also, the plot shows a seasonal pattern, indicating that the incremental value of water is larger during winter, when prices are higher, than during summer. We do, however, see a general decrease in water value between week 40 and 41. This is due to the reservoir restriction in Søvnet being lifted, increasing the amount of available water for the production planner and thereby reducing the value of one extra water unit. Overall, the water value plot complies with our expectations, thereby strengthening the validity of our model.

5.5 Sensitivity analysis using different correlation values and number of volatility functions

In order to further test the performance and validity of our model, we have performed a set of sensitivity analyses. Our model incorporates two features that are new to reservoir management models, namely a correlation between movements in price and local inflow and a multi-factor model to describe price movements. It is interesting to test the effect of introducing these features, as it can tell us how models that assume no correlation or that only use one factor to describe price movements perform compared to ours. For both the correlation coefficient and number of factors, we have calculated the expected revenues using different values of ρ and N . More importantly, we have also tested how decision policies obtained using $\rho = 0$ and $N = 1$ perform when inflow and price movements are in fact correlated or driven by multiple factors.

As mentioned earlier, our model incorporates a correlation coefficient between the error term of the inflow model and the increment of the Wiener process associated with the first principal

component. In order to test the effect of introducing such a correlation, we have first made three lattices with different correlation values $\rho = [0, -0.1765, -0.353]$. We then compare the simulated expected discounted revenues and average reservoir level curves for all three lattices and corresponding decision policies to see how much they deviate. Once again, we have used 50.000 simulated paths, and for all three runs, the starting date is January 7, 2013. The expected return of all three runs is displayed in Table 11 and mean optimal reservoir curves for Vasslivatn are displayed in Figure 24.

Table 11: Expected discounted revenues (EDR) obtained using three different correlation coefficients ρ .

Correlation coefficient	[-]	$\rho = 0$	$\rho = -0.1765$	$\rho = -0.353$
EDR	[M EUR]	16.10	15.93	15.83

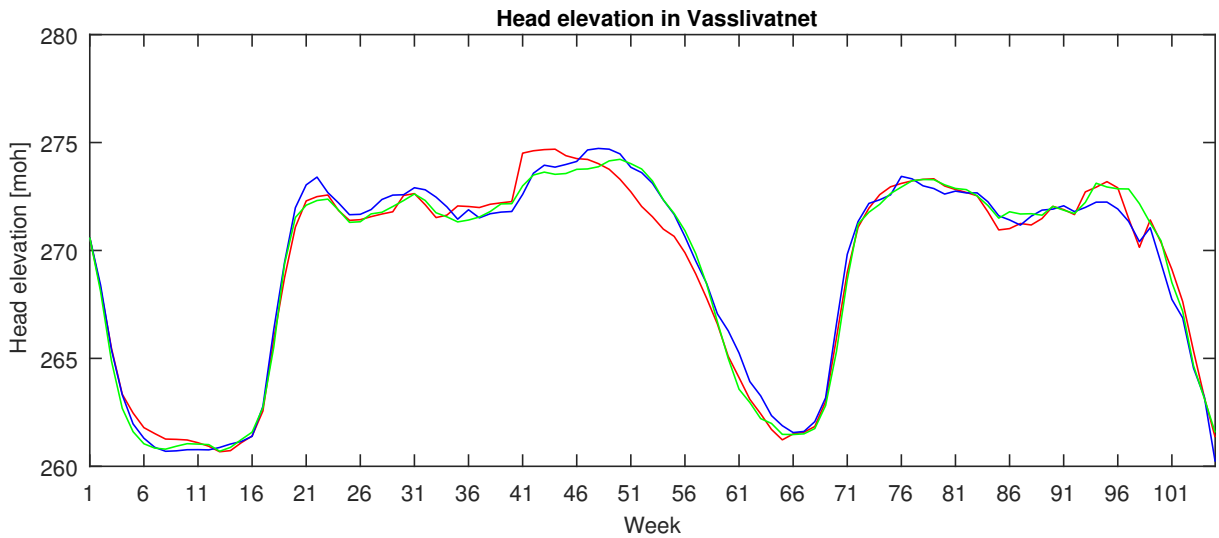


Figure 24: Average reservoir curves for Vasslivatnet found using three different values of the correlation coefficient ρ . The red curve denotes $\rho = 0$, the blue curve $\rho = -0.1765$ and the green curve $\rho = -0.353$.

By looking at the expected discounted revenues, we see that the correlation coefficient ρ does affect the result. The higher we choose the correlation coefficient, the larger are the expected discounted revenues. The reservoir elevation curves have similar shapes, but their differences still imply that the optimal policies are different in all three cases. We made multiple retests to verify the results, obtaining expected discounted revenues quite close to those in Table 11. There are some potential sources of error in the calculations. For example, it could be that 380.000 Monte Carlo simulations are not sufficient to represent all future states. The differences in expected discounted revenues are, however, large enough to indicate that the correlation coefficient undoubtedly affects the decision policy, implying that it must be estimated correctly.

It is also interesting to test how a decision policy created using the correlation $\rho = 0$ performs when we use it in a stochastic process where $\rho \neq 0$. We test this by first creating a lattice and obtaining the optimal decision policies for each node using $\rho = 0$. Instead of drawing simulated lattice paths based on the risk-neutral probabilities provided when $\rho = 0$, we draw paths corresponding to a stochastic price and inflow process where $\rho \neq 0$. Then, to compare the policies, we look at the difference between the expected discounted revenues obtained using policies with $\rho = 0$ and policies with $\rho \neq 0$. In Table 12, we present the expected discounted revenues obtained when the stochastic processes in reality have a correlation $\rho = -0.1765$ and $\rho = -0.353$. As the results indicate, if the real correlation is $\rho = -0.1765$, the policies will provide expected discounted revenues that are 2.5% lower than if the policies incorporated this

correlation. For $\rho = -0.353$, the expected discounted revenues become 3.1% lower. Although these differences might seem small, they show that the producer at Sjøa can miss out on discounted revenues of multiple 100.000 EUR yearly if they misspecify the correlation coefficient. Therefore, we find it reasonable to conclude that the choice of the correlation coefficient does have an effect on the model performance, and should be considered by the production planner in their model.

Table 12: Expected discounted revenues (EDR) calculated when using a policy in which $\rho = 0$, but where the real stochastic process has $\rho = [-0.1765, -0.353]$. The bottom row indicates the difference between the expected discounted revenues obtained using these policies versus the expected discounted revenues obtained using a policy with the same ρ as in the stochastic process, as shown in Table 11.

Correlation coefficient	[-]	$\rho = -0.1765$	$\rho = -0.353$
EDR	[M EUR]	15.54	15.40
Performance difference	[-]	-2.5%	-3.1%

Further, up until now, we have used a model with six factors to describe the movements of a forward contract. The number was chosen such that the proportion of explained variance would be larger than 95%, a threshold value used by Koekebakker and Ollmar (2005). However, Bjerksund et al. (2008) claim that a proportion of 90% is sufficient, while Clewlow and Strickland (2000) choose the number of factors such that the proportion becomes 98.4%. Therefore, we perform a sensitivity analysis of our model using a different number of factors. We investigate four different numbers of factors: 1 (that is, we use the overall volatility function), 3 (91.53% explanation), 6 (96.04% explanation) and 10 (98.44% explanation). The obtained expected discounted revenues are shown in Table 13, and the mean optimal reservoir levels in Figure 25.

Table 13: Expected discounted revenues (EDR) obtained using different number of factors N to describe the underlying price process.

Number of factors	[-]	$N = 1$	$N = 3$	$N = 6$	$N = 10$
EDR	[M EUR]	15.80	15.86	15.93	16.00




Figure 25: Average reservoir curves for Vasslivatnet obtained using different number of factors in the forward price model. The red curve is from a run with $N = 1$ factors, the dark blue one for a run with $N = 3$, the red curve for $N = 6$ and the light blue curve for $N = 10$.

The results in Table 13 show that there is an increasing trend in expected discounted revenues

when we use more factors to describe the price process. This should make sense, as more factors can result in larger price fluctuations, thereby resulting in a lattice with a larger difference between the highest and lowest possible price at a time step. The optimal policies utilize the higher prices in the lattices with more factors, and the model thereby forecasts larger expected discounted revenues.

As for the case with different values of ρ , it might be more interesting to test how the policies obtained using one-factor price model perform when the price process can, in reality, be described using $N = [3, 6, 10]$ factors. We, therefore, redo the steps explained above for the case of different numbers of factors N instead of correlation coefficient ρ . The expected discounted revenues are displayed in Table 14. By looking at the numbers, we see that a policy created using a one-factor price model will underperform by approximately 2% when the price process is in fact driven by multiple factors. Similar to the case for different values of ρ , this can result in a decrease in revenues of multiple 100.000 EUR yearly for a hydropower plant, underlining the importance of using a price process that is as correct as possible when modeling reservoir management.

Table 14: Expected discounted revenues (EDR) calculated when using a policy where the number of factors is $N = 1$, but where the real stochastic price process is described by $N = [3, 6, 10]$. The bottom row indicates the difference between the expected discounted revenues using these policies versus the expected discounted revenues obtained using a policy with the same number of factors N as in the stochastic process, as shown in Table 13.

Number of factors N	[-]	$N = 3$	$N = 6$	$N = 10$
EDR	[M EUR]	15.52	15.61	15.63
Performance difference	[-]	-2.1%	-2.0%	-2.3%

5.6 Backtesting the model policy with realized price and inflow data

Another interesting analysis to perform is a backtest. When backtesting, we have collected the realized weekly inflows and average area spot prices over the entire simulation horizon. Then, we apply the model policy to the realized history of price and inflow and get all decisions that our model would have made for the given history of inflow and price. Using this, we can compare how our model performs compared to the existing strategies of the hydropower production planner. Using January 7, 2013, as our starting date, we have found the realized weekly inflows and spot prices over the next 105 weeks. Next, we have found the total revenues earned using the model policy, and compare this with the actual income earned by the power plant in the same time interval. In reality, the Sjøa power plant generated discounted revenues of **10.89 million EUR** between January 7, 2013, and January 11, 2015, by trading in the spot market. By applying the policy obtained by our model, the plant would have had discounted revenues of **11.69 million EUR**, meaning that using our model could have provided the production planner with approximately 400.000 EUR in extra yearly revenues. To explain this difference, we look at the modeled and realized head elevation curves for both reservoirs in the corresponding period. These are interesting to compare, as they show whether the model policy agrees or disagrees with the realized strategy. In Figure 26 we have plotted the modeled and realized head curves for both reservoirs over the simulation period.

By visual inspection, we see that our model empties both reservoirs in the last time step, providing it with some additional revenues compared to the historical operations. Thus, it might be more accurate to compare the revenues obtained during the first 52 weeks - that is, between January 7, 2013, and January 5, 2014, instead. In this period, the plant earned discounted revenues of **6.22 million EUR**. On the contrary, using the policies from our model, the discounted revenues provided to the plant would be **6.34 million EUR**, meaning that our




Figure 26: Modelled (red curve) and realized (blue curve) reservoir curves for Vasslivatnet and Søvavatnet between January 7, 2013 and January 11, 2015.

model performs well compared to reality also without emptying the reservoir in the last time step. Our model did, however, use more water in these 52 weeks, so we should also compare the expected discounted revenues per production as well. In reality, the plant had discounted revenues per production of **39.00 EUR/MWh** between January 7, 2013, and January 5, 2014, while our model had discounted revenues per production of **38.92 EUR/MWh**. This indicates that our model performed almost as good as the realized production in this period.

Further inspection of the head curves shows that our model is less risk-averse than the real-life production planner. One example of this can be seen by looking at the figure for Vasslivatnet around week 36, that is, in the middle of September 2013. Here, spot prices were quite high, so both our model and the real-life operation planner chose to nominate quite much water for production. However, since spot prices tend to be higher during winter, it is risky to empty the reservoirs in September. Therefore, the real-life production planner chooses to nominate only half of the amount that the model nominates. The model is, however, expecting high inflows in the upcoming weeks, and therefore nominates quite much water before it fills up the reservoir around week 46. Another good example is around week 71, that is, one week before the summertime restriction on the reservoir level in Søvavatnet starts to apply. While the real-life production planner fills up Søvavatnet a few weeks ahead, the model expects sufficient inflows during the next week and decides to reduce the water level in the reservoir to approximately 1 meter below the summertime restriction. It does, however, still manage to fill up the reservoir and meet the constraint in time.

The aforementioned points should help our model perform better than the real-life production planner. Another factor helping our model is that it does not have to perform maintenance, which is an event that forces all operation to be temporarily suspended. However, the real-life production planner has an advantage that our model does not have. Since our model uses weekly granularity, it can only make production decisions on a weekly basis. We assume that our model sells the electricity at a price equal to the average price of that week. In real life, the production planner makes hourly decisions and can utilize the fluctuations of the electricity spot price both within a single day and within a week. They do also have access to the intraday market, allowing them to optimize their production further. The opportunity to optimize production on an hourly level should give the real-life production planner an advantage compared to our model. At last, the real-life production planner has access to short-term weather forecasts the our model does not. Despite the circumstances discussed above, our model still manages to achieve similar results.

5.7 Comparison of results using different volatility functions

In this subsection, we test what results we would get if we had used the volatility functions obtained using the methods of Benth et al. (2007) or Fleten and Lemming (2003) instead. For both cases, we chose the number of factors such that the proportion of explained variance is above 95%, choosing $N = 19$ when we used the method of Benth et al. (2007) and $N = 12$ for the method of Fleten and Lemming (2003). We have constructed a separate lattice using the volatility functions from both methods, and calculated the expected discounted revenues and the expected discounted revenues per production. The results are displayed in Table 15, where we have also included the results obtained using the method of Alexander (2008a).

Table 15: Expected discounted revenues (EDR) and expected discounted revenues per production (EDR/Prod.) using volatility functions based on forward curves constructed by the methods of Alexander (2008a), Fleten and Lemming (2003) and Benth et al. (2007).

Method	Unit	Alexander	Fleten and Lemming	Benth et al.
EDR	[M EUR]	15.93	16.03	19.67
EDR/Prod.	[EUR/MWh]	40.12	40.47	50.87

From Table 15, it is clear that the expected discounted revenues are largest when we use the volatility functions from the method of Benth et al. (2007). This is mainly because the volatility functions found using this method were of larger magnitudes, resulting in a lattice with a larger spread between the different state values of spot price. At one point, the lattice contains spot price scenarios of 400 EUR/MWh, which is highly unrealistic for weekly prices. Looking at Figure 27, which illustrates the first 20 weekly steps of the lattice constructed using the volatility functions from the method of Benth et al. (2007), we also notice another weakness. Recall the volatility functions plotted in Figure 8. These suggested that the volatility of forward price movements were at their largest for contracts with time to maturity of around seven weeks. A result of this can be seen in the lattice at the eighth week, where the price spread increases significantly from the last step. This odd behaviour indicates that the volatility functions obtained using this method are of poor quality.




Figure 27: Price lattice constructed with volatility functions from method of Benth et al.

On the contrary, the expected discounted revenues calculated using the method of Fleten and Lemming (2003) are quite similar to the ones used in our model. The difference in expected discounted revenues is similar to the one we obtained when we included four extra factors in Section 5.5. Also, the lattice in Figure 28 does not incorporate any unreasonable traits and is similar to the one obtained using the volatility functions from the method of Alexander (2008a). Based on these findings, it is, therefore, difficult to say anything about the validity of the volatility functions obtained using the method of Fleten and Lemming (2003). Nevertheless, we

pointed out in Section 5.1.2 that the correlation matrix featured some unexpected characteristics, favoring the usage of the volatility functions found using linear interpolation.




Figure 28: Price lattice constructed with volatility functions from method of Fleten and Lemming

5.8 Delta hedging

Lastly, we make some considerations about how our model can be used in the context of risk management. As described in Flatabø et al. (1998), hydropower producers face multiple types of risk, including but not limited to uncertainty in price, inflow, and demand. Since our model creates all future spot price scenarios using the price of forward contracts traded in the market and not from a fundamental model, it allows us to look at new ways hydropower producers can hedge themselves against price fluctuations. Inspired by financial theory, we introduce a new concept: the delta of a hydropower plant. For a financial option, the delta (Δ) denotes the change in the value of the option if the price of the underlying (typically a stock) increases by 1 EUR (or any other currency). The reason why investors are interested in calculating deltas is that they want their portfolio to be hedged against price fluctuations, meaning that they want to minimize the change in the value of their portfolio given a change in the price of the underlying asset. To achieve this, the investor should construct their portfolio such that $\Delta_{port} = \sum_{i \in port} \Delta_i = 0$, where Δ_i is the delta of the i th asset of the portfolio, thereby becoming delta hedged.

There is an element of optionality to hydropower production scheduling because the production planner must decide how much to produce in the current and all subsequent time steps. Thereby, the expected value of all future cash flows given optimal production scheduling can be considered as the value of an exotic option held by the production planner. We choose the underlying asset of these cash flows to be the forward curve. Using the analogy of financial options, we define the delta of this option as the change in expected future cash flows given a positive parallel shift of 1 EUR/MWh in the forward curve. Since a hydropower plant can potentially produce in all eternity, the value of this option can become quite large. In this case, we limit the value of the option to be the expected discounted cash flows during the next 105 weeks. In Figure 29, we have plotted a delta curve illustrating the delta of the hydropower plant as a function of its underlying forward curve. The deltas have been calculated by finding the expected discounted revenues for a set of parallel forward curves and then calculating the difference in expected discounted revenues between two forward curves with a constant spread of 1 EUR/MWh.

For financial options on a single unit of the underlying, Δ is always in the region $-1 \leq \Delta \leq 1$ since an option cannot be worth more than its underlying. In our case and the way we have defined our option, it can, take significantly higher values. By observing the delta curve, we see that it is strictly positive, which is reasonable since higher prices should give higher expected discounted revenues. It is, however, interesting to see that the curve has no smooth shape, but it oscillates. This might be because we had to reduce the number of Monte Carlo simulations




Figure 29: Delta curve for the option to produce in the period January 7, 2013 and January 11, 2015. F denotes the realized forward curve, and $F + i$ denotes a forward curve with a positive parallel shift of 1 EUR/MWh.

for each run to 35.000 and nodes per time step to 50 due to computational limits, giving us less accurate results. However, assuming that the shape of the curve is correct, it can be used by the company operating the plant to figure out how they can hedge their entire portfolio of operating plants and financial products. To become delta hedged, or at least reduce the portfolio delta subject to changes in the forward curve, they must take positions with a negative delta. For example, they can go short in forward contracts, go short in call options or go long in put options. As mentioned in Section 2.1, options with quarterly and yearly forward contracts as their underlying are traded at NASDAQ Commodities. A parallel shift in the forward curve should be equivalent to a shift in all of its underlying contracts, meaning that these options can be considered to have the same underlying as the power plant. The larger the Δ of the company's operating plants, the more extensive becomes the amount of forward and options trades required to become hedged.

In our case, if the current forward curve had been given by $F - 3$, less trade in the forward and options market would be needed to become delta hedged. However, if the forward curve had been given by $F - 4$, the delta of the plant would be larger, implying that more trade would have to be undertaken for the plant owner to become delta hedged. While it would maybe not be common for a hydropower plant to become completely delta hedged, as it would imply that almost all of its power would be hedged, the plant owner could set a maximum cap for their total delta Δ_{max} . Using Δ_{max} and the delta of the plant, the plant owner could then calculate the amount of trades needed to keep $\Delta_{port} \leq \Delta_{max}$.

5.9 Further work

Although our model provides us with multiple, interesting results, many aspects of it should either be improved or modified. In this subsection, we discuss some of these aspects and propose some work that can be conducted to develop the model further. We also propose some possible extensions of the model, and how it can be combined with other models to improve its applicability further.

Obtain more sophisticated forward curves and volatility functions

In this paper, we tested three different approaches for constructing forward curves and corresponding volatility functions. None of the two methods that involved constructing smooth curves provided us with particularly reasonable volatility functions, which we found a bit surprising. Although the method of Alexander (2008a) provided us with satisfactory volatility functions, the forward curves themselves were inadequate. Hence, we recommend that one should redo our calculations of volatility functions using more sophisticated methods for forward curve construction, e.g., with software like Elvis Front Manager. We also propose reducing the granularity in the method of Fleten and Lemming (2003) so that it, for example, calculates the spot price for every quarter of

a day instead of once a day, as this may contribute to damp oscillations in overlapping regions.

Use more sophisticated starting values of price and inflow

Currently, the starting period spot price corresponds the previous week price of an ENOW-1 (one week ahead) contract, while the starting period inflow is set equal to the empirical mean inflow of that current week. Considering present-day research on short-term price and weather forecasting, it should be possible to determine the starting values of price and inflow in a more sophisticated way, especially for inflow. If the current time value of inflow is chosen to be quite different from its actual value, it will affect all future inflow states in the lattice, resulting in biased expectations of future inflows.

Quality check the value of the correlation coefficient

In our paper, we calculate the short-term correlation between local inflow and price movements to be $\rho = -0.1765$ using Pearson product-moment correlation coefficient. We do, however, have a limited dataset, resulting in a 95% confidence interval of $[-0.28, -0.06]$. Considering the potential loss of revenues from constructing decision policies based on the assumption that $\rho = 0$, it is crucial to determine the correlation coefficient correctly to avoid the losses associated with setting it either too high or too low. Furthermore, note that there are other ways to measure correlation, for instance by using weighted correlation or rolling window correlation. Such analyses could be undertaken to test the stability of the correlation coefficient.

Obtain better relationship between system spot price and area price

As mentioned in the paper, the spot prices obtained in the model are in fact representations of the system price and not area prices. Thus, we suggest that the price process should be extended to incorporate the difference between system and area spot price, thereby providing more realistic forecasts of the relevant area price for a given hydropower plant.

Extend the model such that it can also handle non-linear expressions

A short-coming of our model is that we had to treat both head elevation and efficiency rate as constants. In a more realistic model, it should be possible to model them as non-linear functions of the decision variables. This would, however, not be possible using classical ADDP.

Extend the model to allow for trade in the forward market

By allowing the production planner to also participate in the forward market, they would receive more flexibility when it comes to making decisions. It is reasonable to assume that this flexibility has value for the production planner. However, since there is no drift in the price of forward contracts, their prices should all equal the expected future spot prices given risk neutral transition probabilities. Thus, a model incorporating trading in the forward market should not result in changes in expected revenues. Irrespective of the potential change in expected revenues, a model that allows for forward trading could provide the production planner with valuable decision support when it comes to optimally hedging their positions against price fluctuations. It should be relatively straightforward to extend our model to also allow for trading in the forward market. Since the model uses weekly forward contracts to obtain all spot prices, a contract type which is traded for the six nearest weeks at NASDAQ Commodities, we could extend all states in the lattice to also contain the price of the six nearest weekly forward contracts. The model constraints would also have to be altered to account for the possibility to trade in the forward market.

Combine the model with a short-term production scheduling model

Many medium-term models for reservoir management, e.g., EOPS, are used in combination with short-term production scheduling models. It would be interesting to test how our model

performs in combination with such a model, e.g., with SHOP (Short-term Hydro Operation Planning), a model developed by SINTEF (SINTEF (2017c)). In such a combined model, our model could be used to calculate water values and/or decide how much water that should be used for production in a particular week with the goal of maximizing revenues over the next two years. The short-term model would then use these quantities decide at what time within that week the production should be scheduled.

Include maintenance scheduling

Occasionally, hydropower plants must undergo maintenance that forces all operation to be temporarily suspended. To incorporate maintenance in the model, one could add a restriction to the model that sets the maximum level of water discharge to 0 during the weeks scheduled for maintenance. It should also be possible to extend the model such that it finds which particular weeks it is optimal to undergo maintenance. For example, this could be done by calculating the expected discounted revenues over a two-year horizon for multiple situations where maintenance is scheduled in different weeks, and then select the week for which the total revenues are maximized.

6 Conclusion

In this paper, we have developed a medium-term model for reservoir management. We have modeled the problem as a Markov decision process, forcing both inflow and price movements to follow Markov processes. Besides, we have used a multi-factor model to describe changes in forward and spot prices, and have constructed forward curves using three different methods to obtain the coefficients of the price model. We have incorporated a short-term correlation between the local inflow model and price model and used a scenario lattice to discretize all future states of price and inflow and their corresponding transition probabilities. Lastly, we have used approximate dual dynamic programming to obtain the optimal decision policies of all states.

When we constructed smooth forward curves and used them to find the volatility functions driving price movements, the three methods provided us with rather different results. The forward curves formed using the procedure of Benth et al. (2007) showed three weaknesses, and resulted in volatility functions that failed to represent the dynamics of the forward curve in a good way. The procedure of Fleten and Lemming (2003) resulted in more reasonable forward curves, but the volatility functions still showed some unexpected traits. The method of linear interpolation provided less sophisticated forward curves than the methods mentioned above, but the volatility functions incorporated all the features one would expect them to do according to previous research. However, since our final model needed volatility functions that could be used to explain movements of weekly forward contracts, one could also argue that the volatility functions found using the method of Alexander (2008a) are not entirely correct. That is because they were constructed based on movements in not only weekly, but also monthly, quarterly and yearly contracts. It is difficult to make any absolute conclusions about the results obtained using the method of Fleten and Lemming (2003) as we believe that one could have damped the oscillations of the curves by reducing the granularity of the model, thereby getting more plausible volatility functions. We do, however, find it reasonable to conclude that the method of Benth et al. (2007) performs poorly when the objective is to find volatility functions describing the price movements of forward contracts in the Nordic electricity market. We also conclude that the method of Benth et al. (2007) should be avoided when constructing smooth forward curves using the contracts available at NASDAQ Commodities. That is because it shows two weaknesses that are not present when, e.g., using the method of Fleten and Lemming (2003), namely discontinuity and obviously wrong values.

We further conclude that there exists a short-term correlation between the weekly residuals of the inflow model and the increment associated with the first volatility function of the forward curve movements. The 95% confidence interval of this correlation coefficient is $[-0.28, -0.06]$ for this particular plant. Our analyzes further show that disregarding such a correlation in a model for reservoir management can result in a reduction of expected revenues by 2.5% (that is, multiple 100.000 EUR yearly) if the correlation coefficient is in fact $\rho = -0.1765$. We thereby conclude that it is necessary to include the short-term correlation between price and inflow movements when constructing a lattice and the optimal decision policies, as it will provide the production planner with more realistic expectations about price and inflow and better decision policies.

Our analyzes also show that it is important to use multiple factors when describing price movements. Constructing a lattice and optimal decision policies using a one-factor model can result in a loss of expected revenues by approximately 2% when the price is in fact driven by three, six or ten factors. We thereby conclude that it is important to model price movements using multiple factors, as it can result in increased revenues for the production planner. When

it comes to the correct number of factors one should use, we find it hard to make an absolute conclusion. While we have mostly used a six-factor model explaining 96.04% of all variation, the ten-factor model tested in Section 5.5 explaining 98.44% provided us with 0.4% higher expected revenues. It could have been correct to use more factors and thereby obtain a higher proportion of explained variance, but at the same time, one should not seek to overfit the model. Also, using more factors results in longer computational time, as one must make more random draws of the residuals $\epsilon_{i,t} \sim N(0, 1)$ for $i = [1, \dots, N]$ when we increase the number of factors N .

Lastly, we conclude that the performance of our model is quite good, as using it would have provided revenues similar to those realized by the plant in the period between January 7, 2013, and January 5, 2015. These results are especially interesting since our model receives the average weekly price, while the real-life production planner can make decisions on an hourly basis, allowing them to produce during periods with higher prices within a particular week that are not available for our model.

References

- [1] Pierre Massé. *Les Reserves et la Regulation de l'Avenir dans la vie Economique*. Vol. I and II. Hermann, 1946.
- [2] Stein W. Wallace and Stein-Erik Fleten. “Stochastic Programming Models in Energy”. In: *Handbooks in OR and MS*. Ed. by Andrzej Piotr Ruszczyński and Alexander Shapiro. Vol. 10. Elsevier Science B.V., 2003. Chap. 10, pp. 637–677.
- [3] Steffen Rebennack. “Combining sampling-based and scenario-based nested Benders decomposition methods: application to stochastic dual dynamic programming”. In: *Mathematical Programming* 156 (2015), pp. 343–389.
- [4] Olav B. Fosso et al. “Generation scheduling in a deregulated system. The Norwegian case”. In: *IEEE Transactions on Power System* 14 (1 1999), pp. 75–81.
- [5] Niko A. Iladis et al. “Optimal operation of hydro-dominated multireservoir systems in deregulated electricity markets”. 2008. URL: https://www.researchgate.net/publication/228746176_Optimal_operation_of_hydro-dominated_multireservoir_systems_in_deregulated_electricity_markets.
- [6] Ove Wolfgang et al. “Hydro reservoir handling in Norway before and after deregulation”. In: *Energy* 34 (10 2009), pp. 1642–1651.
- [7] Hubert Abgottspon and Göran Andersson. “Medium-term optimization of pumped hydro storage with stochastic intrastage subproblems”. In: *Power Systems Computation Conference (PSCC)*. IEEE, 2014.
- [8] Richard Bellman. *Dynamic Programming*. Princeton Landmarks in Mathematics and Physics. Princeton University Press, 1957.
- [9] M.V.F. Pereira and L.M.V.G. Pinto. “Multi-stage stochastic optimization applied to energy planning”. In: *Mathematical Programming* 52 (1-3 1991), pp. 359–375.
- [10] Nils Löhndorf and David Wozabal. “Indifference pricing of natural gas storage contracts”. 2017. URL: http://www.optimization-online.org/DB_HTML/2017/02/5863.html.
- [11] Bernard F. Lamond and Abdeslem Boukhtouta. “Optimizing long-term hydro-power production using Markov decision processes”. In: *International Transactions in Operational Research* 3 (3-4 1996), pp. 223–241.
- [12] SINTEF. *EOPS - one area power-market simulator*. 2017. URL: <http://www.sintef.no/en/software/eops-one-area-power-market-simulator/>.
- [13] SINTEF. *EMPS - multi area power-market simulator - SINTEF*. 2017. URL: <https://www.sintef.no/en/software/emp-multi-area-power-market-simulator/>.
- [14] Birger Mo et al. “Optimisation of hydropower operation in a liberalised market with focus on price modelling”. In: *Porto Power Tech Conference*. IEEE, 2001.
- [15] Nils Flatabø et al. “Short-term and Medium-term Generation Scheduling in the Norwegian Hydro System under a Competitive Power Market Structure”. In: *EPSOM '98 Proceedings*. Switzerland: ETH Zürich, 1998.
- [16] Birger Mo et al. *Dokumentasjon av ny prismodell til bruk i Vansimtap, Plansddp og Beta*. Tech. rep. Trondheim: Sintef Energiforskning AS, Sept. 2000.
- [17] Les Clewlow and Chris Strickland. *Energy Derivatives: Pricing and Risk Management*. Tottenham Court Road, London, UK: Lacima Publications, 2000.
- [18] Steen Koekebakker and Fridthjof Ollmar. “Forward curve dynamics in the Nordic electricity market”. In: *Managerial Finance* 31 (6 2005), pp. 73–94.

- [19] Petter Bjerksund, Gunnar Stensland, and Frank Vagstad. “Gas Storage Valuation: Price Modelling v. Optimization Methods”. 2008. URL: <https://brage.bibsys.no/xmlui/bitstream/handle/11250/163949/2008.pdf?sequence=1>.
- [20] Fred Espen Benth, Steen Koekebakker, and Fridthjof Ollmar. “Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation”. In: *The Journal of Derivatives* 15 (1 2007), pp. 52–66.
- [21] Stein-Erik Fleten and Jacob Lemming. “Constructing forward price curves in electricity markets”. In: *Energy Economics* 25 (5 2003), pp. 409–424.
- [22] Carol Alexander. *Pricing, Hedging and Trading Financial Instruments*. Vol. III. Market Risk Analysis. The Atrium, Souther Gate, Chichester, West Sussex PO19 8SQ, England: John Wiley & Sons, Ltd, 2008.
- [23] Rüdiger Kiesel, Florentina Paraschiv, and Audun Sætherø. “On the Construction of Hourly Price Forward Curves for Electricity Prices”. Aug. 2017. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2845302.
- [24] Alexander Shapiro et al. “Risk neutral and risk averse Stochastic Dual Dynamic Programming method”. In: *European Journal of Operational Research* 224 (2 2013).
- [25] Camilla Westad Kolsrud and Martin Prokosch. “Reservoir Hydropower: The value of flexibility”. Master thesis. Norwegian University of Science and Technology, June 2010.
- [26] Maren Boger and Einar Midttun Vestbøstad. “Structural Estimation Analysis in Hydropower Scheduling”. Master thesis. Norwegian University of Science and Technology, June 2016.
- [27] Tor Olav Seim and Ole Ronny Thorsnes. “Analyzing the price- and inflow relationships in hydroelectric scheduling”. Master thesis. Norwegian University of Science and Technology, Dec. 2007.
- [28] Benjamin Berger et al. “Analytical Pricing and Hedging of Fixed and Indexed Energy Swing Contracts”. 2016.
- [29] NordPool. *Bidding areas — Nord Pool*. 2017. URL: <https://www.nordpoolgroup.com/the-power-market/Bidding-areas/>.
- [30] NASDAQ. *Nordic Power Products - Nasdaq*. 2017. URL: <http://www.nasdaqomx.com/commodities/markets/power/nordic-power>.
- [31] Robert L. McDonald. *Derivatives Markets*. 3rd ed. Pearson New International Edition. Edinburgh Gate, Harlow, Essex CM20 2JE: Pearson Education Limited, 2014.
- [32] Fred Espen Benth, Jurate Saltyte Benth, and Steen Koekebakker. *Stochastic modelling of electricity and related markets*. 5 Toh Tuck Link, Singapore 596224: World Scientific Publishing, 2008.
- [33] J. Michael Harrison and Stanley R. Pliska. “Martingales and stochastic integrals in the theory of continuous trading”. In: *Stochastic Processes and their Applications* 11 (3 1981), pp. 215–260.
- [34] Richard Bellman. “The theory of dynamic programming”. In: *Bulletin of the American Mathematical Society* 60 (6 1954), pp. 503–515.
- [35] Blake Johnson and Graydon Barz. “Selecting stochastic processes for modelling electricity prices”. In: *Energy Modelling and the Management of Uncertainty*. Ed. by Vincent Kaminski. Risk Books, 1999. Chap. 1.
- [36] Klaus Mayer, Thomas Schmid, and Florian Weber. “Modeling electricity spot prices: combining mean reversion, spikes, and stochastic volatility”. In: *The European Journal of Finance* 21 (4 2015), pp. 292–315.

- [37] Julio Lucia and Eduardo Schwartz. “Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange”. In: *Review of Derivatives Research* 5 (1 2002), pp. 5–50.
- [38] David Heath, Robert Jarrow, and Andrew Morton. “Bond Pricing and the Term Structure of Interest Rates A New Methodology for Contingent Claims Valuation”. In: *Econometrica* 60 (1 1992), pp. 77–105.
- [39] Petter Bjerksund, Heine Rasmussen, and Gunnar Stensland. “Valuation and Risk Management in the Norwegian Electricity Market”. 2000. URL: https://brage.bibsys.no/xmlui/bitstream/id/361246/Bjerksung_Valuation.pdf.
- [40] Carol Alexander. *Practical Financial Econometrics*. Vol. II. Market Risk Analysis. The Atrium, Souther Gate, Chichester, West Sussex PO19 8SQ, England: John Wiley & Sons, Ltd, 2008.
- [41] Vlad Bally and Gilles Pages. “A quantization algorithm for solving multidimensional discrete-time optimal stopping problems”. In: *Bernoulli* 9 (6 2003), pp. 1003–1049.
- [42] Martijn R.K. Mes and Arturo Perez Rivera. “Approximate Dynamic Programming by Practical Examples”. In: *Markov Decision Processes in Practice*. Ed. by R.J. Boucherie and N.M. van Dijk. Springer International Publishing, 2017. Chap. 3.
- [43] Nils Löhndorf, David Wozabal, and Stefan Minner. “Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming”. In: *Operations Research* 61 (4 2013), pp. 810–823.
- [44] Quantego. *QUASAR - The Next Generation Software for Stochastic Optimization*. 2017. URL: <http://www.quantego.com/>.
- [45] Kaveh Madani and Jay R. Lund. “Modeling California’s high-elevation hydropower systems in energy units”. In: *Water Resources Research* 45 (9 2009).
- [46] Edouard Jaeck and Delphine Lautier. “Volatility in electricity derivative markets: The Samuelson effect revisited”. In: *Energy Economics* 59 (Aug. 2016), pp. 300–313.
- [47] Paul A. Samuelson. “Proof That Properly Anticipated Prices Fluctuate Randomly”. In: *Industrial Management Review* 6 (2 1965), pp. 41–49.
- [48] James Tipping et al. *A Model for New Zealand Hydro Storage Levels and Spot Prices*. Presented at the EPOC Winter Workshop, Auckland, 2 September 2005. 2005. URL: http://www.mang.canterbury.ac.nz/people/jimbo/James_Tipping_EPOC_WW_2005.pdf.
- [49] SINTEF. *SHOP - SINTEF*. 2017. URL: <http://www.sintef.no/en/software/shop/>.

Nomenclature

Symbol	Parameter	Unit
(Ω, F)	Measure space	[-]
\mathbb{Q}	Measure representing risk neutral probabilities of all future states	[-]
\mathbb{P}	Measure representing real life probabilities of all future states	[-]
P_t	Spot price at time t	[EUR/MWh]
P_{tn}	Spot price in the n th discrete state at time t	[EUR/MWh]
I_t	Total inflow at time t	[m^3]
$I_{b,t}$	Inflow into reservoir b at time t	[m^3]
\hat{I}_t	Set of inflows into multiple reservoirs	
$I_{b,tn}$	Inflow into reservoir b in the n th discrete state at time t	[m^3]
I_{tn}	Set of inflows into all reservoirs in the n th discrete state at time t	
B	Total number of reservoirs in a connected hydropower system	[-]
\hat{T}	Time horizon	[-]
π_t	Optimal decision policy at time t	[-]
π_{tn}	Optimal decision policy in the n th discrete state at time t	[-]
CF_t	Cash flows at time t	[EUR]
r	Appropriate discount rate	[-]
$w_{bi,t}$	Water nominated for production in a turbine connecting reservoir b and i	[m^3]
ς	Number of seconds all turbines are running per week	[s]
ϖ	Number of hours all turbines are running per week	[h]
ϱ	Density of water	[kg/m^3]
G	Gravitational acceleration	[m/s^2]
$q_{bi,t}$	Water discharge through a turbine connecting reservoir b and i	[m^3/s]
$H_{b,t}$	Head elevation of reservoir b	[$MASL$]
$\eta_{bi,t}$	Combined turbine and generator efficiency rate	[-]
$l_{b,t}$	Water level in reservoir b	[m^3]
$u_{bi,t}$	Water flowing from reservoir b to reservoir i outside of turbines	[m^3]
O	An outlet of a hydropower system	[-]
dP_t	Small change in spot price at time t	[EUR/MWh]
α_t	Mean reversion rate	[-]
\bar{P}_t	long-term spot price level	[EUR/MWh]
σ_t	Volatility of spot price movements	[-]
dt	Small increment of time	[-]
dZ_t	Increment of a Wiener process during dt	[-]
dq_t	Increment of a Poisson process during dt	[-]
θ_t	Jump frequency	[-]
κ_t	Proportional jump size	[-]
γ_t	Jump volatility	[-]
$\bar{\kappa}_t$	Mean jump size	[-]
T	Maturity time of a forward contract	[-]
$F_{t,T}$	Price of a forward contract at time t with maturity at time T	[EUR/MWh]

$dF_{t,T}$	Small change in price of forward contract	[EUR/MWh]
$\sigma_{t,T}$	Volatility of price movements in a forward contract at time t with maturity at time T	[-]
$\sigma_{i,t,T}$	Volatility function/factor i of price movements in a forward contract at time t with maturity at time T	[-]
$\sigma_{i,\tau}$	Volatility function/factor i of price movements in a forward contract at time t with time to maturity τ	[-]
$\Psi_i(x)$	Denotes an arbitrary function of x	[-]
τ	Time to maturity/delivery of a forward contract	[-]
N	Number of factors in multi-factor model used to describe price movements of forward contracts	[-]
M	Number of forward contracts traded in the market/used to construct forward curves	[-]
$f(t_s)$	Forward curve constructed using all available contracts at time t_s	[-]
$f(t_s, t)$	Value of a forward curve constructed at time t_s for time $t > t_s$	[-]
(t_b, t_e)	Time interval spanned by forward contracts used to construct a forward curve	[-]
$(t_{b,j}, t_{e,j})$	Time interval of delivery of j th forward contract	[-]
S	Set of delivery intervals	[-]
$U(t)$	Underlying seasonal function of price	[EUR/MWh]
$\varepsilon(t)$	Adjustment function	[EUR/MWh]
D	Number of non-overlapping time intervals used to construct forward curves	[-]
a_i, b_i, c_i, d_i, e_i	Parameters of the adjustment function for time interval $i = [1, \dots, D]$	[-]
x	Vector of parameters	[-]
$F_{t,t_b,j,t_e,j}$	Closing price of a forward contract with delivery period $(t_{b,j}, t_{e,j})$	[EUR/MWh]
$F_{t,t_b,j,t_e,j}^{BID/ASK}$	Bid/ask price of a forward contract with delivery period $(t_{b,j}, t_{e,j})$	[EUR/MWh]
$\zeta(t, t_i, t_j)$	Function used to find adjustment function parameters	[-]
A and b	Matrices used to get all constraints on the form $\mathbf{Ax} = \mathbf{b}$	[-]
δ	Vector of Lagrange multipliers	[-]
H	Matrix used to solve optimization problem using method of Lagrangian Multipliers	[-]
λ	Smoothness parameter	[-]
C	Number of discrete prices contained in a forward curve constructed using the method of Fleten and Lemming (2003)	[-]
F _{t_s}	Vector of forward prices	[EUR/MWh]
J	Number of forward curves used to construct returns data series	[-]
R	Number of different delivery periods	[-]
A	Number of discrete times to maturity τ_a for which we obtain the value of the volatility functions $\sigma_{i,\tau}$	[-]
$x_{j,a}$	Daily logarithmic deviation/return between forward curve $f(t_j)$ and $f(t_{j-1})$ for time to maturity τ_a	[-]
X _{$J \times A$}	Returns dataset for J observations of A different times to maturity	[-]
V	Variance-covariance matrix	[-]

\mathbf{P}	Matrix of principal components	[-]
\mathbf{p}_i	Principal component i , i th column of \mathbf{P}	[-]
\mathbf{W}	Matrix of eigenvectors	[-]
\mathbf{w}_i	Eigenvector i , i th column of \mathbf{W}	[-]
w_{mi}	Element m of i th eigenvector. Not to be confused with $w_{bi,t}$, the amount of water nominated for production in a turbine connecting reservoir b and i at time t .	[-]
Λ	Diagonal matrix of eigenvalues Λ_i	[-]
Y_t	Logarithm of inflow I_t	$[\ln(m^3)]$
$\hat{\mu}_t$	Weekly average logarithmic inflow	$[\ln(m^3)]$
Z_t	Deviation between log inflow and log average	$[\ln(m^3)]$
ϵ_t	Residual/error term	[-]
ϕ_t	Inflow model coefficient	[-]
$\sigma_{\text{INF},t}$	Standard deviation of residual term in inflow model	[-]
\overline{S}_{tn}	n th state of a lattice at time t	[-]
N_t	Number of nodes in a lattice at time t	[-]
\overline{S}_t	Set of states at time t	[-]
K	Number of Monte Carlo simulations used to construct a lattice	[-]
$(\hat{S}^k)_{t=1}^{\hat{T}}$	k th Monte Carlo simulation	[-]
\hat{S}_t^k	Time t state of k th Monte Carlo simulation	[-]
Γ_{tn}	Voronoi decomposition of all nodes \overline{S}_{tn} for $n \in [N_t]$	[-]
β_k	Stepsize used to construct lattice nodes	[-]
\overline{S}_{tn}^k	Variable used to construct lattice states \overline{S}_{tn}	[-]
p_{tnm}	Conditional probability for a transition between node n and m at time t	[-]
$\mathbb{I}_A(x)$	Indicator function indicating whether x is part of the set A	[-]
V_t	Value of all future discounted cash flows, also referred to as the value function	[EUR]
$\mathbb{E}[V_{t+1}]$	Expected value of all future cash flows from next state and onwards	[EUR]
Ξ	Number of iterations used in ADDP algorithm	[-]
s_t	Spilled water	$[m^3]$
EC	Energy coefficient	$[kWh/m^3]$
ι	Split factor for inflow	[-]
W	Given week	[-]
Δ	Delta of an option	[-]