NTNU
Norwegian University of
Science and Technology

Dynamic Hedging for a Norwegian Hydropower Producer

Electricity prices, inflow and currency risk

Joakim Dimoski
Sveinung Nersten

Industrial Economics and Technology Management
Submission date: June 2018
Supervisor: Stein-Erik Fleten, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

Problem statement

We aim to create a dynamic hedging model for a Norwegian hydropower producer that seeks to reduce their exposure to risk factors such as price, inflow and currency exchange rate. We are interested in obtaining the relationship between the currency risk and the other risk factors, and we want to quantify the risk-reduction effect of including currency derivatives in a firm's hedging strategy. Further, we want to investigate how a dynamic hedging model performs compared to a heuristical approach used by many Norwegian hydropower companies.

Preface

This master's thesis was written during the spring of 2018 at the Norwegian University of Science and Technology (NTNU), Department of Industrial Economics and Technology Management, within the field of Financial Engineering.

We would like to acknowledge and express our gratitude to the persons who have made the completion of this master's 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 year. Further, we wish to thank Dr. Nils Löhndorf for his valuable input and assistance 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 and hedging. At last, we would like to thank our families for their continuous support.

Trondheim, June 8, 2018

Joakim Dimoski

Sveinung Nersten

Abstract

This master's thesis consists of two articles. In the first article, we present a global, dynamic model for hedging of hydropower production. We include stochastic processes for spot and futures prices, reservoir inflow and currency exchange rate. We present a sequential approach, obtaining optimal production and hedging decisions separately. To manage the risk, we allow for trading in currency forwards and power futures contracts. Risk preferences are modelled using conditional risk mapping. The risk-reduction effect of currency hedging is found to be moderate - currency hedging increases the 5% CVaR of the terminal discounted cash flows by 2.4%. We also find that including monthly power futures in the hedging strategy allows for precision hedging that can contribute to substantial reductions in risk. In the second paper, we propose a medium-term scheduling model for hydropower production. We use a multi-factor price process in which the price of futures contracts is used to forecast future spot prices. Further, we include a short-term correlation between prices and local inflow. Our main contribution is a comparison of 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 the case where price movements are in fact driven by multiple factors and correlated with local inflow. In both situations, we find the loss to be approximately 2-3 %.

Summary in Norwegian

Denne masteroppgaven består av to artikler. I den første artikkelen presenterer vi en global, dynamisk modell for sikring av vannkraftproduksjon. Vi inkluderer stokastiske prosesser for spot- og futurespriser, lokalt tilsig og valutakurs. Vi foreslår en sekvensiell tilnærming, der optimale produksjons- og sikringsbeslutninger blir funnet separat. For å redusere risikoen tillater vi handel i forwardkontrakter for valuta og futureskontrakter for kraft. Risikopreferanser modelleres ved bruk av *conditional risk mapping*. Risikoreduksjonseffekten av å inkludere valutasikring viser seg å være moderat – valutasikring øker 5 % CVaR av de totale diskonterte kontantstrømmene med 2.4 %. Vi finner også ut at bruken av månedlige futureskontrakter for kraft muliggjør presisjonssikring som kan bidra til en betydelig risikoreduksjon. I den andre artikkelen presenterer vi en mellomlangsiktig produksjonsplanleggingsmodell for vannkraft. Vi bruker en flerfaktorprosess til å modellere endringer i pris, der prisen til futureskontrakter brukes til å forutse fremtidige spotpriser. Vi inkluderer også en kortsiktig korrelasjon mellom pris og lokalt tilsig. Vårt hovedbidrag er en sammenligning av ytelsen til vår foreslåtte modell og ytelsen til en modell der pris og lokalt tilsig er antatt å være uavhengige, og ytelsen til en modell der prisbevegelser bare er drevet av én risikofaktor. Vi regner ut reduksjonen i forventede neddiskonterte kontantstrømmer av å bruke de to sistnevnte modellene gitt at prisbevegelser faktisk er drevet av flere faktorer eller faktisk er korrelert med lokalt tilsig. I begge situasjonene er reduksjonen på ca. 2-3 %.

Contents

Introduction

A hydropower producer with reservoir capacity is faced with multiple decision problems. First, they must dispatch the water in their reservoirs optimally so that they can maximize their revenues from production and avoid costly spillage. Second, they are exposed to numerous sources of uncertainty that affect their cash flows. These risk factors must be identified, analyzed and responded to. In this master's thesis, we aim to investigate both of these decision problems, and propose different approaches for how a hydropower company can 1) obtain optimal production policies and 2) manage the market risk associated with their production revenues.

This master's thesis consists of two academic articles, in which the first article builds on the findings and methods of the second. The first article, *Dynamic hedging for a Norwegian hydropower producer: Electricity prices, inflow and currency risk*, is the main contribution of this thesis. It proposes a dynamic model for hydropower risk management. Uncertainty in spot and futures prices, reservoir inflow, and currency exchange rate is considered. The paper uses a modified version of the scheduling model presented in the second paper to obtain optimal production decisions. The cash flows from production are then subject to hedging using financial instruments. A case study is conducted, using data from a Norwegian hydropower plant. The article has been written with the intention of being published in an academic journal, either in its current form or after additional adjustments. The second article, *Hydropower reservoir management using multi-factor price model and correlation between price and local inflow*, has been accepted at the 41st IAEE International Conference held in June 2018, where the authors will present its content. Subsequently, it will be published in the conference proceedings database of IAEE. The paper is based on the project thesis written in the fall of 2017, subject to corrections, improvements, and restructuring. Primarily, the paper proposes a multistage stochastic program for hydropower reservoir management, which has been tested using data from the aforementioned plant.

In the first article, we have three main contributions to the field of hydropower hedging. First, we include currency exchange risk and quantify the potential risk reduction obtainable by hedging it using derivatives. Second, we model the risk preferences of the hydropower firm using conditional convex risk mapping (Ruszczynski and Shapiro, 2006), more specifically the nested conditional value-at-risk (nested CVaR), initially proposed by Shapiro et al. (2013). According to Shapiro (2009) and Löhndorf and Wozabal (2017), risk measures like the nested CVaR are time consistent. To the authors' knowledge, conditional risk mapping has not been used in previous papers on dynamic hedging of electricity production. Third, we propose using variables to track confirmed future cash flows from a portfolio of forward positions in a multistage stochastic program. These variables enable us to accurately replicate the timing of the cash flows from financial contracts that are marked-to-market regularly, and they are necessary to incorporate the currency risk. Additionally, we investigate how the performance of a sequential approach compares to a simultaneous approach. In the sequential approach, production decisions and hedging decisions are made separately, while all decisions are taken simultaneously in an integrated model. Finally, we assess how the sequential approach

performs compared to heuristics used by industry practitioners.

In the hydropower scheduling paper, we use a multi-factor process to model spot price movements. This method is in contrast to existing models for hydropower scheduling, which often use single-factor processes. Further, we calculate the correlation coefficient between movements in inflow and the first factor of the price process, thereby treating inflow and price as dependent variables. We also quantify the loss in expected cash flows if price and local inflow are assumed to be independent when they are in fact correlated, and equivalently, the losses that occur when using a single factor price process if price movements are in fact described by multiple factors. At last, we perform a backtest to compare the decision policies of our model to the realized decisions and cash flows of the case plant.

The master's thesis is organized as follows. First, we present a section considering advantages of managing risk in the context of hydropower production, serving as motivation for writing a paper on hydropower risk management. Next comes the first article, *Dynamic hedging for a Norwegian hydropower producer: Electricity prices, inflow and currency risk*. It is followed by *Hydropower risk management using multi-factor price model and correlation between price and local inflow*. Afterward, we provide some discussions on both papers as a whole and propose some further work that may be conducted. Finally, we include an extensive appendix which provides depth to the methods, concepts, and results of both articles. The appendix also contains relevant work that is not included in the final versions of the articles.

There are minor differences between the nomenclature used in the first and second paper. The list of nomenclature found at the end of the hedging paper is meant for that paper only. Additionally, while the production model used in the first article is based on the one proposed in the production paper, some aspects are slightly different. This includes, among others, the dynamics of the price process and the model granularity.

The motivation behind hedging for a power producer

Risk aversion is usually the primary motivation for hedging, as hedging can reduce the variability of cash flows. The traditional hedging perspective is that of the risk-averse producer who uses financial markets to reduce the diversifiable risk of their profits (Anderson and Danthine, 1980). Norwegian power producers are predominantly publicly owned, according to the Norwegian Ministry of Petroleum and Energy (2017). Typically, the producers pay out large yearly dividends that contribute to financing public authorities (Sanda et al., 2013). Fluctuations in these dividends can have large consequences for the public owners. Thus, owner risk-aversion can be an important motivation for hedging among Norwegian hydropower producers.

Neoclassical economics and theory on perfect markets (e.g. Modigliani and Miller, 1958) state that hedging cannot add value to a firm, other than reducing their level of risk. However, evidence can be found in the literature that hedging can add value, beyond risk reduction. Smith and Stulz (1985) propose multiple types of non-linear costs that companies can reduce through risk management. These are, among others, related to tax function convexity and reduced default risk. The first cost is irrelevant for Norwegian companies, as they are allowed to carry forward losses to the next year, meaning that their tax function is linear. Sanda et al. (2013) argue that the reduced default risk is an irrelevant argument as well. They claim that most hydroelectric companies in Norway have a negligible risk of default because they are publicly owned and their variable costs related to production are low. However, we believe this to be too generalizing. The main drivers of default risk of a company is its leverage and the earning capacity, not whether the ownership is public or private.

Stulz (1996) argues that since hedging can reduce the default risk of a company, firms that hedge their profit can benefit from a more leveraged capital structure. In many countries, there is a tax advantage to debt financing due to deductable interest cost. Besides, increasing the debt ratio releases equity capital. The equity can be paid back to the shareholders, or it can be invested in promising or strategically important projects within the company. As stated in May and Neuhoff (2017), power systems based on renewable energy sources like solar and wind are associated with high upfront investment costs. Thus, hydropower companies that want to extend their power generation portfolio with new sources of energy can benefit from hedging to increase their total amount of available capital. Finally, the company can choose to keep its capital structure constant. Assuming that hedging contributes to reduced default risk, hedging should result in an improved credit rating and decreased costs of raising funds.

Stulz (1996) introduces the term *selective hedging*. His research shows that many companies let their market views influence their hedging practice, thereby taking speculative positions in their portfolios. Sanda et al. (2013) show that selective hedging is present in many of the hydropower companies they study, even documenting negative hedge ratios in some of the companies. Adam and Fernando (2006) separate hedging into two components, namely *predictive* and *selective* hedging. As opposed to selective or speculative hedging, predictive hedging focuses on hedging predicted cash flows from a firm's fundamental operations.

Analyzing hedging in the gold mining industry, Adam and Fernando (2006) found that the value contribution from selective hedging was minimal, only resulting in a larger cash flow volatility. Sanda et al. (2013) do, however, find evidence of increased profits among Norwegian power companies engaging in speculative hedging, subject to increased volatility. Since the focus of the first paper is on risk management and selective hedging is shown to increase risk, it only considers predictive hedging.

Dynamic hedging for a Norwegian hydropower producer: Electricity prices, inflow and currency risk

Joakim Dimoski[a,*], Sveinung Nersten[a,†]

[a]*Norwegian University of Science and Technology, Norway*

Abstract

In this paper, we present a global dynamic model for the risk management problem of a hydropower producer. The problem is formulated as a multistage stochastic linear program. We focus on market risk and include stochastic processes for spot and futures prices, reservoir inflow and foreign exchange rate. We present a sequential approach, obtaining optimal production and hedging decisions separately. To manage the risk, we allow for trading in currency forward and power futures contracts. The paper presents three main contributions to the field of hydropower hedging; the inclusion of currency risk and currency derivatives, modeling risk preferences using conditional risk mapping and introducing variables for accurate replication of the cash flow structure from a portfolio of financial contracts. We quantify the risk-reduction effect of currency hedging when there is currency risk. It is found to be moderate - currency hedging increases the 5% CVaR of the terminal discounted cash flows by 2.4%. We also find that including monthly power futures in the hedging strategy allows for precision hedging that can contribute to substantial reductions in risk.

Keywords: Risk management, hydropower production, currency risk, Markov processes, stochastic dynamic programming

3.1 Introduction

For a hydropower company, an important challenge lies in utilizing the water in their reservoirs optimally such that the value of their long-term production cash flows is maximized. Simultaneously, these cash flows are subject to large uncertainties associated with different types of risks. In this article, we focus on the market risk of a hydropower producer, and how it can be managed. Market risk includes risk factors such as price risk, exchange rate risk and interest rate risk. We also account for inflow risk. By utilizing financial derivatives like forward contracts and options, risk-averse companies can lower their market risk exposure in accordance with their risk preference.

The activity of reducing a firm's market risk is

*E-mail: joakim.dimoski@gmail.com
†E-mail: sveinung.nersten@gmail.com

referred to as hedging or risk management, and the literature presents multiple approaches to how this can be conducted. As shown in Stulz (1996), most older academic theory emphasizes that the target of risk management should be to minimize the cash flow variance. While acknowledging that variance minimization is a useful goal for companies seeking to manage their risk, he proposes that the primary target should rather be to provide protection against the possibility of costly lower-tail outcomes. This complies with the findings of Sanda et al. (2013). They find that multiple Norwegian hydropower companies aim for downside protection in their risk management practices. Other reasons for firms to manage their risk include reduced default risk, implying that the firm can benefit from a more flexible capital structure and larger debt capacity (Stulz, 1996), and tax advantages (Smith and Stulz, 1985).

Dupuis et al. (2016) classify hedging procedures into two categories, namely *static* and *dynamic* hedging. Actors using static hedging only trade hedging derivatives at one point in time with no subsequent rebalancing of their hedging portfolio. The objective in such models is to reduce the risk associated with the end-of-horizon cash flows. Static hedging for hydropower producers is proposed in multiple papers, e.g. Fleten et al. (2010). On the contrary, actors using dynamic hedging adjust their portfolio continuously as new market information becomes available. Dupuis et al. (2016) further present two sub-categories of dynamic hedging, *local* and *global*. While local hedging procedures focus on minimizing the short-term risk, that is, until the next rebalancing, global hedging procedures seek to minimize risk associated with all future cash flows. Examples of papers on local hedging in electricity markets include Zanotti et al. (2010) and Liu et al. (2010), while Mo et al. (2001a), Fleten et al. (2002) and Dupuis et al. (2016) propose global models. Fleten et al. (2002) show that a dynamic approach yields better results than a static approach for a hydropower producer. Dupuis et al. (2016) obtain similar results, taking the perspective of a risk-averse retailer.

Wang et al. (2015) analyze 18 different minimum-variance strategies for 24 commodity, currency and equity markets. They find that a naive hedging strategy (hedge ratio $= 1$) performs better or almost as well as the minimum-variance strategies in all of the tested markets. This shows that engaging in a sophisticated hedging process may have limited value added. They do, however, not include the electricity market, and use a static approach with a rolling window method instead of dynamic hedging.

In this paper, we present a global dynamic hedging model for a price-taking hydropower producer owning a single plant[1] that participates in the Nordic electricity market. We propose a sequential approach, first using a dynamic production planning model to obtain production policies that maximize the expected terminal cash flows. Then, the production decisions are stored and used as stochastic variables in a dynamic hedging model. Production policies are obtained using a modified version of the production model proposed by Dimoski et al. (2018), incorporating correlated inflow and spot price scenarios generated using the price of futures contracts. For the

[1]Note that hedging is typically undertaken on a company level, hedging the production of multiple hydropower plants. Generally, our hedging approach can be applied to a portfolio of hydropower plants as well.

hedging model, we allow the producer to trade in the monthly, quarterly and yearly power futures contracts available at the Nordic power derivatives market. As we take the perspective of a Norwegian producer and all prices in the Nordic market are given in euros (EUR), we also treat the currency spot exchange rate between EUR and the Norwegian krone (NOK) as stochastic. To hedge this uncertainty, we allow the producer to trade in currency forward contracts as well. To model risk preferences, we use the nested conditional value-at-risk (nested CVaR), proposed initially by Shapiro et al. (2013) and later used in Löhndorf and Wozabal (2017). We have a time horizon of two years and use semi-monthly granularity. Furthermore, we include transaction and tax costs.

We have three main contributions to the field of hydropower hedging. First, we include currency exchange rate as a stochastic variable, thereby incorporating this as a risk factor. To manage the foreign exchange risk, we also allow for trading in currency derivatives. Second, we model the risk preferences of the producer using conditional risk mapping and nested CVaR. While multiple papers use conditional risk mapping in other industries, there exists no research on this in the field of hydropower hedging to the authors' knowledge. Third, to replicate the cash flows from currency and power futures as realistically as possible, we propose using variables for confirmed future cash flows. It is necessary to include such variables to incorporate the uncertainty in the currency exchange spot rate. We also quantify the effect of removing the option to hedge currency risk and the effect of using a sequential approach instead of a simultaneous. Further, we test how our proposed hedging approach compares to an heuristical approach using hedge ratio requirements, based on the practice of a Norwegian hydropower producer.

Due to the large dimensionality of our problem, obtaining optimal decision policies is very computationally demanding using classic stochastic dynamic programming. Therefore, we use a modified version of stochastic dual dynamic programming (SDDP, Pereira and Pinto, 1991) known as approximate dual dynamic programming (ADDP, Löhndorf et al., 2013). SDDP and similar approaches are widely used in existing literature on hydropower scheduling (e.g. Mo et al., 2001b and Rebennack, 2015) and hedging (e.g. Fleten et al., 2002 and Iliadis et al., 2006). ADDP integrates SDDP with Markov processes, meaning that both decision problems must be formulated as Markov Decision Processes (MDP). 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, irrespective of its history. Similarly, in an MDP, all decisions are made based on the current state of the world. Treating hydropower production planning as an MDP is common in literature, as shown in Lamond and Boukhtouta (1996). Further, we discretize all state variables into a scenario lattice, as this is a requirement for using ADDP. To construct the lattice, we use the method proposed by Löhndorf and Wozabal (2017).

Some papers, e.g. Mo et al. (2001a), propose integrated models where production and hedging decisions are made simultaneously. However, Wallace and Fleten (2003) claim that it is favorable to treat production planning and risk management as sequential activities. Among others, they argue that it is not possible to increase the value of a power portfolio by trading hedging derivatives in an efficient market, as this can only be achieved

through a change in production. Therefore, a hydropower producer should first seek to obtain a production schedule that maximizes their expected cash flows, and then use hedging to reduce the risk of their portfolio to the desired level.

The proposed method for risk preference modeling, the nested CVaR, is based on conditional convex risk mapping (see Ruszczynski and Shapiro, 2006). Boda and Filar (2006) and Shapiro (2009) show that global hedging strategies aiming to reduce risk associated with terminal cash flows are time-inconsistent, meaning that decision policies in such models are affected by past gains and losses. However, this is not the case for the nested CVaR. The issue with time consistency of CVaR is also discussed in Godin (2016), who proposes using a similar version of the nested CVaR known as the conditional CVaR.

To generate scenarios for spot and power futures prices, we use a multivariate HJM model (Heath et al., 1992) explaining movements in a forward curve. In a liquid power market, the available future and forward contracts traded at a given time should represent the current time risk-adjusted market expectations for future spot prices. Furthermore, a high-resolution forward curve can be constructed using the price and delivery periods of all available futures contracts, as shown in e.g. Fleten and Lemming (2003), Benth et al. (2008) and Kiesel et al. (2018). Thus, a model explaining the evolution of a forward curve can both be used to find future spot prices and to calculate the price of futures contracts with different delivery periods.

We also incorporate stochastic processes for the risk factors inflow and currency. In hydropower-dominated systems, there is typically a negative correlation between inflow and electricity price, since the inflow largely determines the supply side of the market. This provides a *natural hedging* effect. We are also interested in investigating the magnitude of the currency risk, as well as how it relates to the other risk factors. A preliminary hypothesis is that the system price is negatively correlated with the EURNOK exchange rate as well. Our reasoning is that because all bid and ask orders in the Nordic market are placed by companies operating with their local currency, this should influence the system price, denoted in EUR/MWh.

The paper is organized as follows. In Section 3.2, we present relevant background information about hydropower risk management in the Nordic countries. Section 3.3 includes a description of the hedging problem, formulated as a multistage stochastic program. We also give an overview of the algorithms used to construct a scenario lattice and solve the decision problems. In Section 3.4, we show how the dynamics behind the risk factors can be modelled, and Section 3.5 is devoted to numerical results. Conclusions are made in Section 3.6.

3.2 Hydropower risk management

In this section, we give an overview of relevant aspects for risk management in the context of hydropower production in the Nordic countries. First, we underline important risk factors and derivatives that can be used to reduce the exposure of a hydropower company. Then, we present some relevant risk measures firms can use to quantify the extent of their risk exposure. The section is concluded with a brief overview of how risk management is practiced in the industry today.

3.2.1 Risk factors and mitigating derivatives

A hydropower company is exposed to a variety of financial risk factors which can negatively affect their production cash flows. The Basel Committee on Banking Supervision (2012) categorizes financial risk into credit risk, operational risk, liquidity risk and market risk. The credit risk, which is the potential that the counter-party will not live up to its contractual obligations, can be relevant for the hydropower producer if they enter long term OTC contracts. The operational risk is related to losses caused by inadequate or failed internal processes, people and systems, or from external events. This includes political risk, which Fleten et al. (2012) identify as an important risk factor for hydropower producers, as their cash flows are sensitive to changes in e.g. government regulations and tax rates. Liquidity risk is the risk that trading an asset adversely affects the market price. For hydropower producers, it is relevant for the EPAD futures, discussed below.

Neither the operational risk, credit risk nor the liquidity risk will be treated here - the focus of this article is the market risk. The market risk includes factors such as commodity price risk, currency risk and interest rate risk. Furthermore, there are two market risk factors specific to a Nordic hydropower producer: Inflow risk and area price difference risk. Fleten et al. (2012) identify inflow and price uncertainty as the major risks for a hydropower producer. To hedge the market risk, it is possible to use derivatives for the electricity price and the exchange rate which are traded in financial markets. Currently, there exists no liquid market for derivatives based on inflow uncertainty in the Nordic countries. However, as shown in Foster et al. (2015), there have been developed methods for pricing such contracts based on hydraulic and weather indices.

The primary trading place for Nordic hydropower producers is the Nord Pool day-ahead market, colloquially referred to as the spot market. This market is divided into several bidding areas, where the spot price of each area is calculated based on the supply-demand equilibrium of each area and constraints on transmission capacity between the areas. All prices are denoted in EUR/MWh. Nord Pool also calculates the system price, which is an unconstrained market clearing reference price for the entire Nord Pool market. Financial contracts for the Nordic market are traded at NASDAQ OMX Commodities Europe (hereby NASDAQ OMX). The market is purely financial in the sense that no physical energy is exchanged - only cash. This is in contrast to the day-ahead market, which is a physical market.

The system price risk can be hedged by trading futures and options at NASDAQ OMX. Unlike traditional forward and futures contracts, which are contracts for trading an asset at a specified point in time, the power futures traded at NASDAQ OMX have delivery periods spanning a day, week, month, quarter or year. Thus, their nature is more similar to financial swaps. As explained in Fleten et al. (2010), futures contracts are marked-to-market each day prior to the beginning of the delivery period. Contracts that are in delivery are settled daily based on the difference between the spot price and the last price of the futures contract before going into delivery.

Since the producer receives the area spot price for their production while the system spot price serves as the underlying of the traded power futures, there is basis risk between the area price and the power

futures. NASDAQ OMX also offers contracts for hedging the area price difference, whose reference is the difference between the system spot price and the price in a specific bidding area. These contracts are known as EPADs, and they are significantly less liquid than the power futures. Also, they are only available for certain bidding areas. As shown in Houmøller (2017), there is a high correlation between the hourly system price and the price of most Norwegian areas, averaging at approximately 0.89 for all areas combined in 2013-2016. This is above the limit (0.8) set by the IAS 39 accounting standard to qualify for hedge accounting, and therefore, Houmøller (2017) argues that it is sufficient for power producers in most Norwegian areas to disregard the area difference in their hedging strategy and hedge only with system price contracts.

The revenues of a Norwegian hydropower producer are generated in EUR, while the base currency of the producer is NOK. Because the exchange rate is fluctuating, there exists significant currency risk for the producer. The currency risk can be hedged using forward exchange contracts. The market for currency derivatives is highly liquid, with negligible bid-ask spreads for corporate customers. Currency forwards are over-the-counter (OTC) instruments, as they do not trade on a centralized exchange. Typically, an investor enters into a contract with a bank as the counterparty. Since all bid orders and ask orders in a price area are placed by companies operating with their local currency, the system price denoted in EUR/MWh should be influenced by the base currencies of the different areas. Therefore, we expect the system price to be negatively correlated with the EURNOK exchange rate, providing a natural hedging effect.

3.2.2 Risk measurement

Through risk measurement, a company can quantify its risk factors, and also assess the effect of their hedging strategy. For hydropower producers, risk measurement is typically based on the end-of-horizon cash flows, as these are simple to interpret (Fleten et al., 2010). Historically, standard deviation has been one of the most used metrics for risk management, as illustrated in Stulz (1996). While the standard deviation represents deviations from expected cash flows, both in positive and negative direction, Stulz (1996) argues that it is more interesting to consider risk associated with the downside of the cash flow distribution. A common metric for measuring this risk is the value-at-risk (VaR). For a given significance level α, the VaR of H discrete representations of the terminal cash flows h_i is defined as $VaR_\alpha = \min\{h_i \mid \sum_{j|h_j \leq h_i} 1/H \geq \alpha\}$. That is, the VaR for a significance level α is the cash flow that will be exceeded with a confidence level (probability) of $1-\alpha$. While VaR gives a risk manager information about their worst case scenario with a $1-\alpha$ confidence level, it gives no information about the cash flow distribution in the tail below the VaR. CVaR represents the cash flows a producer can expect if they fall below the level given by the VaR. A benefit of CVaR is that it better captures tail-effects, such as kurtosis and skewness. Mathematically, the CVaR of a discrete distribution with significance level α is given by $CVaR_\alpha = \sum_{j|h_j \leq VaR_\alpha} h_j/(H\alpha)$. As opposed to VaR, CVaR is also a coherent[2] risk measure, making it convenient to use in frameworks for risk management (Godin, 2016).

[2]Specifically, VaR does not qualify for the subadditivity axiom when the underlying loss distribution is non-normal. A risk measure $\Phi(X)$ is subadditive if $\Phi(X_1+X_2) \leq \Phi(X_1)+\Phi(X_2)$.

3.2.3 *Industry practice*

It is informative to see how risk management is performed in the hydropower industry, as this provides useful benchmarks. Therefore, we give a overview of some aspects of the industry practice, focusing on the Norwegian market specifically.

In their paper, Sanda et al. (2013) analyze the hedging strategies of 12 Norwegian hydropower companies, whose total production comprises approximately 34.4% of the total production in Norway. The authors find that none of these use an integrated model to obtain optimal production and hedging decisions. Five of the firms use a sequential approach, first obtaining the optimal production decisions in a separate model and then use these as the basis for their hedging decisions. The seven remaining firms use historical production scenarios to predict their future exposure.

Further, none of the firms analyzed by Sanda et al. (2013) use dynamic programming to obtain their hedging policy. They all use more static approaches. Eight of the firms use a hedge ratio approach in which their short positions in the financial market must be within a predefined range for a given time to maturity. Two of the firms have a minimum requirement for the VaR of their terminal cash flows, while the last two have no written hedging policy. Less than half of the firms include options and EPAD contracts in their hedging strategy, and in 11 out of 12 firms, futures contracts are the most used derivatives in terms of hedged volume. On average, approximately 90% of the traded derivative volume [MWh] among the firms was in quarterly and yearly contracts.

Stulz (1996) shows that many companies let their own market views influence their hedging practice, thereby taking speculative positions in their portfolios. He describes this as *selective hedging*. Sanda et al. (2013) show that selective hedging is present in many of the hydropower companies they study, even documenting negative hedge ratios for some of the companies. Adam and Fernando (2006) separate hedging into two components, namely *predictive* and *selective* hedging. As opposed to selective or speculative hedging, predictive hedging focuses on hedging predicted cash flows from a firm's fundamental operations. In this paper, we only focus on predictive hedging, and remove all possibilities for speculative trades.

3.3 Methods

In this section, we formulate the decision problems associated with production planning and hedging as multistage stochastic linear programs. We also introduce the algorithms used for discretizing all state variables into a scenario lattice and to obtain optimal decision policies.

3.3.1 *Assumptions*

We consider the problem faced by a price-taking hydropower producer participating in a deregulated market. The producer can compose a power portfolio consisting of long positions in physical production and short positions in financial futures contracts. Based on a broad set of endogenous and exogenous variables, like reservoir level, inflow, spot and forward prices, they must obtain both a production policy and a portfolio of derivatives satisfying their risk preferences. The producer owns a single hydropower plant, and to illustrate our approach, we consider a real plant located in the NO3 price area in Norway. The plant consists of two interconnected reservoirs and one turbine. The plant is mid-sized in terms of regulating capacity and production capacity, having a mean yearly

production of 191.3 GWh. Figure 3.1 illustrates some relevant properties of the plant.

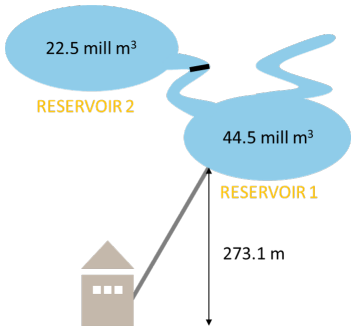

Figure 3.1: Properties of the hydropower plant

Following the findings of Wallace and Fleten (2003), we treat the problem of optimal production and hedging using a sequential approach. To obtain optimal production policies, we use a modified version of the reservoir management model proposed by Dimoski et al. (2018), which considered the same hydropower plant as this article. They treat the problem of obtaining optimal production decisions as a Markov decision process (MDP), which is common in most of the literature on medium-term reservoir management (e.g. Lamond and Boukhtouta, 1996). For the hedging problem, we propose using a dynamic approach, also modeling it as an MDP. While the findings of Sanda et al. (2013) show that none of the analyzed Norwegian companies use dynamic hedging, Fleten et al. (2010) argue that further research should be done within that field, motivating our approach.

Like Bjerksund et al. (2008), we assume that the decision maker participates in a complete market with no risk-free arbitrages. This means that all state transition probabilities can be represented by a unique martingale (risk-neutral) measure $\mathbb{Q}$. For both the scheduling model and the hedging model, we use a time horizon of approximately two years, which is in accordance with previous models on medium-term production scheduling (Wolfgang et al., 2009 and Abgottspon and Andersson, 2014) and hydropower risk management (Fleten et al., 2002 and Fleten et al., 2010). Further, we use semi-monthly granularity, meaning that each time stage has the length of half a month. Table A.1 in Appendix A.1 shows the time intervals that were used. A semi-monthly resolution was chosen such that the discrete stages could coincide conveniently with the delivery periods of monthly, quarterly and yearly futures contracts traded at NASDAQ OMX. Simultaneously, semi-monthly periods are short enough to be used in a medium-term scheduling model, which typically use weekly granularity.

For each time stage, multiple stochastic variables affect the decisions of the hydropower producer. In the production model, the only variables that are decision relevant are spot price $F_{t,t}$ and reservoir inflow $Y_{1,t}$ and $Y_{2,t}$. This is common in commercial software for production planning, as shown in Wolfgang et al. (2009) and Fleten et al. (2012). For the hedging model, we allow the producer to trade in monthly, quarterly and yearly power futures contracts whose delivery periods are within the chosen time horizon of $\hat{T} = 49$ semi-months. These contracts are chosen since they have high liquidity on NASDAQ OMX, and are the most common derivatives used by Norwegian hydropower producers (Sanda et al., 2013). Thus, additional state variables are the prices of six monthly, eight quarterly and one yearly futures contracts, denoted $F_{t,Mi}$ for $i = [1,...,6]$, $F_{t,Qj}$ for $j = [1,...,8]$ and $F_{t,Y1}$. The reference price of the futures contracts is the system spot price. As EPAD contracts are not available for the NO3 area, they are not included in the model. Due to the high empirical

correlation between the system and NO3 spot price, we also disregard the area price difference risk, and assume that the producer receives the system spot price and not the area spot price[3].

The production model disregards uncertainties in currency exchange rate, which is common in present-day production models. Nevertheless, we allow the producer to hedge the currency risk in the hedging problem. Therefore, the hedging model also includes the stochastic variables for spot exchange rate ($Q_{t,t}$) and for forward exchange rate at time t for maturity at time T ($Q_{t,T}$). Further, we denote by W_t the stochastic level of production used in the hedging model. The value of W_t is found using the production model, in which the production level w_t is a decision variable. Since we also include taxation effects, we let γ_c and γ_r denote the corporate and resource rent tax rate, respectively. Lastly, we include variable transaction costs c_F for trading in the power futures market. Transactions cost are negligible in the currency forwards market.

Further, we propose using the nested conditional value-at-risk (nested CVaR) to model the risk preferences of the producer. Unlike the *terminal* CVaR, which considers the CVaR of the total terminal cash flows, the *nested* CVaR considers the CVaR of the cash flows in all subsequent periods, irrespective of past cash flows. For a random variable X, significance level α and a weighting coefficient λ we define the function $\psi_{\lambda,\alpha}(X) = \lambda CVaR_\alpha(X) + (1-\lambda)\mathbb{E}(X)$, where $CVaR_\alpha(X)$ is defined as in Section 3.2.2. Using the formulation in Shapiro et al. (2013), the nested CVaR for a sequence of random variables $X_1, X_2, X_3, ...$ can be defined as

$$CVaR^{NEST}_{\alpha,\lambda}(X_1, X_2, X_3, ...) = X_1 + \psi_{\alpha,\lambda}(X_2 + \psi_{\alpha,\lambda}(X_3 + ...)) \quad (3.1)$$

In Section 3.3.2 and 3.3.3, we introduce decision variables for the linear programs. Unless otherwise stated, these are restricted to non-negative values.

3.3.2 Production planning problem

In this section, we formulate the production planning problem as a linear program. First, we define w_t as the time t spot production, denoted in [MWh]. When obtaining the optimal production plan, the objective of the production planner is to maximize their expected discounted terminal cash flows. Ignoring variable costs related to e.g. generator and turbine start-up, which is common in other papers concerning hydropower production planning (Wallace and Fleten, 2003), the cash flows earned by the producer can be set equal to the discounted revenues earned from physical sales. If the time t cash flows are given by $F_{t,t} \cdot w_t$ and β_t is a time-dependent discount factor, the value function V^P_t, which serves as the time t objective function of the MDP, is given by

$$V^P_t = F_{t,t} w_t (1 - \gamma_c - \gamma_r) + \beta_t \mathbb{E}[V^P_{t+1} \mid F_{t,t}, Y_{1,t}, Y_{2,t}, \pi_t] \quad (3.2)$$

Here, π_t denotes the decision policy at time t, and as the value functions suggest, cash flows from production are subject to both resource rent and corporate tax. Further, we let $v_{1,t}$ and $v_{2,t}$ denote the water level [m^3] in both reservoirs, $s_{c,t}$ the amount of water flowing from reservoir 2 into reservoir 1, and $s_{s,t}$ the amount of spilled water. Using this, the

[3] Note that we have estimated the parameters of a stochastic process for the area price difference in Appendix A.7.4.

volume balance in each reservoir is given by

$$v_{1,t} = v_{1,t-1} - w_t \cdot \kappa^{-1} + s_{c,t} + Y_{1,t} - s_{s,t} \quad (3.3)$$

$$v_{2,t} = v_{2,t-1} + Y_{2,t} - s_{c,t} \quad (3.4)$$

Here, κ is the energy coefficient $[MWh/m^3]$. While this is normally a function of e.g. turbine and generator efficiency, reservoir elevation and water density, it must be treated as a constant in order to keep the problem linear. This was also the case in Dimoski et al. (2018). Treating κ as a constant is in compliance with the EOPS model (SINTEF, 2017b), which is a common software used for medium-term production planning in the Nordic countries. Further, the upper and lower bounds for the water levels in reservoir 1 and 2 are summarized in (3.5). While there is no lower restriction in reservoir 1, reservoir 2 has a lower restriction in summertime set by local authorities. To prevent infeasibility scenarios, we add the slack variable $v^S_{2,t}$ to the last constraint. Violating this bound will result in a penalty cost given by $\upsilon \cdot v^S_{2,t}$, which is added to the objective function.

$$\begin{aligned} v_{1,t} \leq \overline{v_1}, \quad v_{2,t} \leq \overline{v_2} \\ v_{1,t} \geq 0, \quad v_{2,t} + v^S_{2,t} \geq \underline{v_{2,t}} \end{aligned} \quad (3.5)$$

At last, there is a production capacity constraint (3.6). Here, ξ is the maximum flow rate $[m^3/s]$ allowed through the turbine of the hydropower plant and ς_t is the number of seconds in a semi-month t.

$$w_t \leq \overline{w_t} = \xi \cdot \kappa \cdot \varsigma_t \quad (3.6)$$

By combining all expressions and constraints, the production planning problem can be summarized as solving the following subproblem at all time stages t

$$\begin{aligned} \max \quad & V^P_t(F_{t,t}, Y_{1,t}, Y_{2,t}, \pi_t) - \upsilon \cdot v^S_{2,t} \\ \text{subject to} \quad & (3.3), (3.4), (3.5), (3.6) \end{aligned}$$

3.3.3 Hedging problem

In this section, we formulate the hedging problem as a linear program. We must define variables and constraints for the trading of currency and power derivatives. In both cases, we include variables and balance constraints for tracking both *financial short positions* and *confirmed future cash flows*. Together, these are formulated such that the structure of the cash flows from the derivatives replicate their actual payoff structure. The model only allows for short positions in currency and power futures, excluding long positions. For hedging purposes, this should be sufficient, as the producer already has a long position in their physical production.

We assume that the exchange rate forward contract is settled at the contract maturity. Figure 3.2 illustrates an example of the cash flows from a currency forward contract with delivery in stage 8. No cash is exchanged until the maturity date, at which the difference in the forward and spot price is settled. We let $z_{t,T}$ denote the producer's total short position [EUR] at time t in currency forwards with maturity at stage T. New short positions entered at stage t for delivery in T are denoted $x_{t,T}$. This gives the following balance constraint for $t < T$, where $T \leq \hat{T}$

$$z_{t,T} = z_{t-1,T} + x_{t,T}, \quad T > t \quad (3.7)$$

Since a forward contract with instantaneous delivery is a spot trade, the balance constraint for $T = t$ is given by

$$z_{t,t} = z_{t-1,t}, \quad T = t \quad (3.8)$$

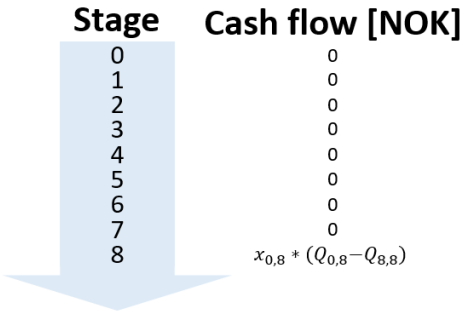

Figure 3.2: Cash flows from short position $x_{0,8}$ in currency forward contract with maturity $T = 8$ made at $t = 0$.

The currency forward rate is denoted $Q_{t,T}$ [NOK/EUR]. We let $y^C_{t,T}$ [NOK] denote the confirmed, positive cash flows that the producer is certain to receive at stage T given their trading activity in the currency forward market. For $T > t$, the balance for the confirmed part of the currency cash flows is given by (3.9). Note that cash flows from forward trading are subject to corporate tax rate γ_c only. This is in contrast to cash flows from physical production, which are subject to both resource rent tax rate γ_r and corporate tax rate γ_c.

$$y^C_{t,T} = y^C_{t-1,T} + x_{t,T} Q_{t,T} (1 - \gamma_c), \quad T > t \qquad (3.9)$$

In (3.9), the decision variable $y^C_{t,T}$ can be interpreted as the positive part of the cash flows that are to occur at maturity time T, as illustrated in Figure 3.2. When maturity is reached at $t = T$, the time t cash flows from currency hedging are given by (3.10). In (3.10), the negative cash flows from the forward positions are added to the positive ones to obtain the time t cash flows $y^C_{t,t}$. Note that $y^C_{t,t}$ can take both positive and negative values.

$$y^C_{t,t} = y^C_{t-1,t} - z_{t,t} Q_{t,t} (1 - \gamma_c), \quad T = t \qquad (3.10)$$

We move on to the variables and constraints for the trading of power futures. Power futures have a more complex cash flow structure than currency forwards, due to the following circumstances. First, they have a delivery period, which can be monthly, quarterly or yearly, instead of delivery at a specific point in time. Second, power futures have daily settlement. Before delivery, the settlement is based on the price change between two successive trading days. Within the delivery period, the settlement is based on the difference between the system spot price and the last price for which the contract was traded before entering into delivery. Contracts in delivery are not tradable. We replicate this structure as closely as possible assuming semi-monthly settlement. Figure 3.3 shows an example of the cash flows of a quarterly contract with a delivery period from stage 3 to 8. To receive the cash flows in NOK, they are multiplied by their respective spot exchange rate $Q_{t,t}$.

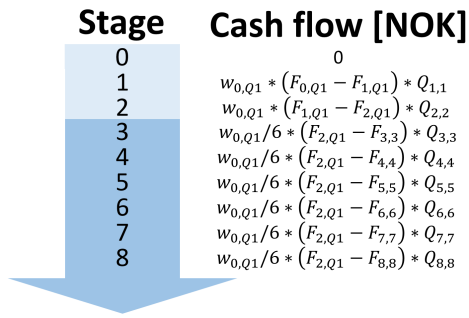

Figure 3.3: Cash flows from short position $w_{0,Q1}$ in a power futures contract with delivery in the upcoming quarter. The light blue part of the figures denotes time stages prior to the start of the delivery period, whereas the darker part denotes stages within the delivery period. During the delivery period, the quantity of the short position is divided by 6, as this is the number of stages covered by the contract delivery period.

We introduce decision variables for short positions. Let $u_{t,Mi}$, $u_{t,Qj}$ and $u_{t,Y1}$ [MWh] denote the total short position at stage t in futures contracts with delivery in i months, j quarters and one year, respectively. While the latter three variables denote the position in contracts that have not yet entered delivery, we must also store the short positions in the monthly, quarterly and yearly contracts that are

currently in delivery. These are denoted $u_{t,M}$, $u_{t,Q}$ and $u_{t,Y}$ [MWh], respectively. $w_{t,Mi}$, $w_{t,Qj}$ and $w_{t,Y1}$ [MWh] denote *new* short positions entered into at stage t. If the delivery period of a contract exceeds the model horizon $\hat{T} = 49$ semi-months, the corresponding decision variable ($w_{t,Mi}$, $w_{t,Qj}$ or $w_{t,Y1}$) is restricted to zero to ensure no trading.

We need to define balance constraints for short positions in power futures. The balances will be different depending on whether the time stage t represents the first or second part of a month, the beginning of a new quarter or the beginning of a new year. If t represents the second part of a month, the total short position will be given by the previous stage value plus new short positions for contracts not yet in delivery.

$$\begin{gathered} u_{t,M} = u_{t-1,M}, \quad u_{t,Q} = u_{t-1,Q}, \quad u_{t,Y} = u_{t-1,Y} \\ u_{t,Mi} = u_{t-1,Mi} + w_{t,Mi}, \quad i = [1,...,6] \\ u_{t,Qj} = u_{t-1,Qj} + w_{t,Qj}, \quad j = [1,...,8] \\ u_{t,Y1} = u_{t-1,Y1} + w_{t,Y1} \end{gathered} \tag{3.11}$$

When t represents the first part of a month, the contract that was the 1 month ahead (M1) in $t-1$ will enter into delivery. Further, M2 will become M1, and all other monthly contracts are shifted in the same manner. A new contract is introduced for delivery in six months (M6). If t represents the first part of the month, but not a new quarter, the constraints will be given by (3.12).

$$\begin{gathered} u_{t,M} = u_{t-1,M1}, \quad u_{t,Q} = u_{t-1,Q}, \quad u_{t,Y} = u_{t-1,Y} \\ u_{t,Mi} = u_{t-1,Mi+1} + w_{t,Mi}, \quad i = [1,...,5] \\ u_{t,M6} = w_{t,M6} \\ u_{t,Qj} = u_{t-1,Qj} + w_{t,Qj}, \quad j = [1,...,8] \\ u_{t,Y1} = u_{t-1,Y1} + w_{t,Y1} \end{gathered} \tag{3.12}$$

Using the same logic, the balance constraints for stages marking the beginning of a quarter, but not a new year, is given by (3.13).

$$\begin{gathered} u_{t,M} = u_{t-1,M1}, \quad u_{t,Q} = u_{t-1,Q1}, \quad u_{t,Y} = u_{t-1,Y} \\ u_{t,Mi} = u_{t-1,Mi+1} + w_{t,Mi}, \quad i = [1,...,5] \\ u_{t,Qj} = u_{t-1,Qj+1} + w_{t,Qj}, \quad j = [1,...,7] \\ u_{t,M6} = w_{t,M6}, \quad u_{t,Q8} = w_{t,Q8} \\ u_{t,Y1} = u_{t-1,Y1} + w_{t,Y1} \end{gathered} \tag{3.13}$$

If t represents the beginning of a year, the short position balance constraints are given by (3.14).

$$\begin{gathered} u_{t,M} = u_{t-1,M1}, \quad u_{t,Q} = u_{t-1,Q1}, \quad u_{t,Y} = u_{t-1,Y1} \\ u_{t,Mi} = u_{t-1,Mi+1} + w_{t,Mi}, \quad i = [1,...,5] \\ u_{t,Qj} = u_{t-1,Qj+1} + w_{t,Qj}, \quad j = [1,...,7] \\ u_{t,M6} = w_{t,M6}, \quad u_{t,Q8} = w_{t,Q8}, \quad u_{t,Y1} = 0 \end{gathered} \tag{3.14}$$

Having established the balance constraints for short positions in power futures, we can define variables and restrictions for the *cash flows* of the power futures portfolio. We first explain all variables, and define them mathematically afterwards. $y^F_{t,t}$ denotes the stage t cash flow from power futures trading. $y^F_{t,t}$ can take positive and negative values, and is part of the value function of the hedging problem (3.22). Since the currency spot rate at which the cash flows occur is not known in advance, $y^F_{t,t}$ is denoted in EUR, as opposed to $y^C_{t,t}$, which is denoted in NOK.

$y^F_{t,T}$ tracks the confirmed, positive part of the cash flows from power futures trading that will occur at $t = T$. $y^F_{t,T}$, where $T > t$, is not part of the value function and is only used to store the positive part of the cash flows that will occur in subsequent periods. In (3.21), the negative part of the cash flows is added the positive ones to obtain the time t cash flows $y^F_{t,t}$. $y^F_{t,t+1}$ stores two types of cash flows. The first type is related to changes in the value of the portfolio of contracts not yet in delivery. Since all contracts

are marked-to-market regularly, we need to store the forward prices in stage t to calculate the price changes in stage $t+1$. The second type of cash flows is associated with contracts in delivery. $y^F_{t,T}$ for $T > t+1$ stores the positive part of cash flows associated with contracts in delivery only. Because the longest delivery period spanned by any of the available contracts is 24 semi-months, it is only necessary to define $y^F_{t,T}$ for $T = [t+1, ..., t+24]$.

First, we consider the case where t represents the *first part of a month*. Then, the cash flow balances will be given by

$$\begin{aligned} y^F_{t,t+1} &= y^F_{t-1,t+1} + (1-\gamma_c)\Big(\sum_{i=1}^{6} u_{t,Mi}F_{t,Mi} \\ &\quad + \sum_{j=1}^{8} u_{t,Qj}F_{t,Qj} + u_{t,Y1}F_{t,Y1}\Big) \\ y^F_{t,T} &= y^F_{t-1,T}, \quad t+2 \le T \le t+23 \\ y^F_{t,T} &= 0, \qquad\qquad T = t+24 \end{aligned} \tag{3.15}$$

In this case, only $y^F_{t,t+1}$ is updated. This is because the next stage is within the same month as the current stage, meaning that no new contracts will enter into delivery. Therefore, the cash flows added to $y^F_{t,t+1}$ are only based on contracts not yet in delivery.

If t is the *second part of a month*, multiple contracts can potentially enter into delivery in the upcoming stage $(t+1)$. We must store the price and position of the futures contracts that enter delivery. The cash flows must be stored for the next two, six or 24 periods, depending on whether the contract which enters delivery is monthly, quarterly or yearly, respectively. To make the formulation more compact, we introduce the indicator functions $\mathbb{I}_Q$ and $\mathbb{I}_Y$. They will take the values 1 if the next stage $(t+1)$ marks the beginning of a new quarter and year, respectively, and 0 otherwise. Using this, the cash flow balances are given by

$$\begin{aligned} y^F_{t,t+1} = y^F_{t-1,t+1} + (1-\gamma_c)\Big(& \\ \frac{u_{t,M1}F_{t,M1}}{2} + \mathbb{I}_Q\frac{u_{t,Q1}F_{t,Q1}}{6} + \mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24} & \\ + \sum_{i=2}^{6} u_{t,Mi}F_{t,Mi} & \\ + (1-\mathbb{I}_Q)u_{t,Q1}F_{t,Q1} + \sum_{j=2}^{8} u_{t,Qj}F_{t,Qj} & \\ + (1-\mathbb{I}_Y)u_{t,Y1}F_{t,Y1}\Big) & \end{aligned} \tag{3.16}$$

$$\begin{aligned} y^F_{t,t+2} = y^F_{t-1,t+2} + (1-\gamma_c)\Big(\frac{u_{t,M1}F_{t,M1}}{2} & \\ + \mathbb{I}_Q\frac{u_{t,Q1}F_{t,Q1}}{6} + \mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}\Big) & \end{aligned} \tag{3.17}$$

$$\begin{aligned} y^F_{t,t+i} = y^F_{t-1,t+i} + (1-\gamma_c)\Big(\mathbb{I}_Q\frac{u_{t,Q1}F_{t,Q1}}{6} & \\ + \mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}\Big), \qquad i = [3, ..., 6] & \end{aligned} \tag{3.18}$$

$$\begin{aligned} y^F_{t,t+i} = y^F_{t-1,t+i} + (1-\gamma_c)\mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}, & \\ i = [7, ..., 23] & \end{aligned} \tag{3.19}$$

$$y^F_{t,t+24} = (1-\gamma_c)\mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24} \tag{3.20}$$

Based on this, we formulate the stage t cash flows from power trading ($y^F_{t,t}$) in (3.21). As was the case for currency forwards, the expression consists of the confirmed, positive cash flows saved in $y^F_{t-1,t}$ and all negative cash flows. Note that we must subtract $w_{t,Mi}, w_{t,Qj}$ and $w_{t,Y1}$ from the positions $u_{t,Mi}, u_{t,Qj}$ and $u_{t,Y1}$ to obtain the negative part of the cash flows associated with price changes in contracts prior to delivery, as these are based on the previous state positions. In addition, we also include variable transaction costs, represented by c_F [EUR/MWh].

$$
\begin{aligned}
y^F_{t,t} = y^F_{t-1,t} + (1-\gamma_c)\Big[& \\
& -(u_{t,M}/2 + u_{t,Q}/6 + u_{t,Y}/24)F_{t,t} \\
& \quad -\Big(\sum_{i=1}^{6}(u_{t,Mi} - w_{t,Mi})F_{t,Mi} \\
& \qquad + \sum_{j=1}^{8}(u_{t,Qj} - w_{t,Qj})F_{t,Qj} \\
& \qquad\quad + (u_{t,Y1} - w_{t,Y1})F_{t,Yi}\Big) \\
& - c_F\Big(\sum_{i=1}^{6} w_{t,Mi} + \sum_{j=1}^{8} w_{t,Qj} + w_{t,Y1}\Big)\Big]
\end{aligned}
\tag{3.21}
$$

For the hedging problem, we define the value function V^H_t as a linear combination of the stage t cash flows and $\psi_{\alpha,\lambda}(X)$. $\psi_{\alpha,\lambda}(X)$ is a risk measure with risk preference weighting λ and significance level α, defined in Section 3.3.1. The stage t cash flows is the sum of the cash flows from currency forward trading, power futures trading and spot production. W_t denotes the production in stage t [MWh], which is a stochastic variable in the sequential hedging approach.

$$
\begin{aligned}
V^H_t = {} & y^C_{t,t} + y^F_{t,t}Q_{t,t} + W_t F_{t,t} Q_{t,t}(1-\gamma_c-\gamma_r) \\
& + \beta_t \psi_{\alpha,\lambda}[V^H_{t+1} \mid F_{t,t}, Q_{t,t}, W_t, \pi_t]
\end{aligned}
\tag{3.22}
$$

We recall that λ adjusts the weighting between the CVaR and expected value of a random variable X. In this case X is the next stage value function V^H_{t+1}. Setting $\lambda = 0$ will thus turn (3.22) into the classic Bellman equation of dynamic programming associated with value maximization, while $\lambda = 1$ implies that (3.22) will maximize the stage t cash flows and the CVaR of V^H_{t+1}. For the sequential approach, setting $\lambda \neq 1$ makes little sense, as the expected profit from trading forward contracts is zero under the risk-neutral probability measure. Thus, one should set $\lambda = 1$ to solve the hedging problem in the sequential model. Adjusting λ does, however, make sense in a simultaneous model. Such a model can be formulated by combining the constraints and decision variables of the production and hedging problems and replacing the stochastic production level W_t with the decision variable w_t.

3.3.4 Building a scenario lattice

As in Löhndorf and Wozabal (2017), we build a scenario lattice by reducing the continuous Markov processes driving the state variables into discrete nodes. Each node contains one entry for each stochastic variable, which in our case are inflow, spot price, exchange rate and multiple forward prices. Generally, we keep the number of nodes per stage N_t in the lattice constant. In comparison, the number of nodes per stage in a non-recombining scenario tree grows exponentially with the number of time stages. Thus, the lattice approach allows for a higher number of time stages and scenarios while keeping the problem computationally feasible.

The nodes are found by a two-step algorithm proposed by Löhndorf and Wozabal (2017). The first step in constructing the lattice is finding the value of its nodes, which is done by first drawing a set of K Monte-Carlo simulations (S^k) of the underlying, correlated stochastic processes, where $k = [1, ..., K]$. We let $\overline{S_{t,n}}$ denote the nth node in time stage t. To find the value of the lattice nodes, we minimize the total squared distance between all N_t nodes and K simulations S^k_t for each time stage t.

The second step of the algorithm involves finding the transition probabilities. For two subsequent nodes $\overline{S_{t,n}}$ and $\overline{S_{t+1,m}}$, the transition probability p_{tnm} is initially found by counting the number of simulated paths whose stage t and $t+1$ states are located closest to the nodes $\overline{S_{t,n}}$ and $\overline{S_{t+1,m}}$ and dividing by the total number of paths that are

located closest to $\overline{S_{t,n}}$. As shown in Löhndorf and Wozabal (2017), a problem with this procedure is that the conditional expected next stage value of a state variable for $t < \hat{T}$ can slightly deviate from its true expected value, given by the continuous process. Such a bias can result in positive expected gains from trading in the futures market (or even arbitrage opportunities), although the continuous process used to generate the lattice indicates expected gains of zero. To avoid this and obtain an unbiased lattice, we use a procedure known as *backwards estimation*. For all stages $t < \hat{T}$, the entries in each node are adjusted such that their conditional expected successor node value is equal to the conditional mean of the true process. This is done iteratively, starting at $t = \hat{T}$ and working backward through the lattice. An extensive explanation can be found in Appendix A.2 or in Löhndorf and Wozabal (2017).

3.3.5 Solution method

For both the production and hedging problem, we use ADDP (Löhndorf et al., 2013) to obtain the decision policies $\pi_{t,n}$ in all lattice nodes. Like SDDP, ADDP is based on a stochastic extension of Benders decomposition and can be used to obtain approximate, near-optimal decision policies. When using ADDP, one of the main simplifications is that the value function is approximated to be a piece-wise linear, concave function of all resource variables (e.g., reservoir levels and financial short positions). In short, the value function is found by first drawing a given number of *forward passes* through the lattice, that is, a sequence of states. For each forward pass, the optimal decision policies are found by maximizing the approximate post-decision value functions. After each forward pass, a *backward pass* is performed, in which the approximated value functions are updated relative to the sampled sequence of states and all state decision policies. In practice, the approximate value function of each state is constructed by a set of supporting hyperplanes (linear constraints), where each pair of forward and backward passes results in the addition of a new hyperplane to the set.

3.4 Risk factor dynamics

In this section, we present the dynamics behind the stochastic variables driving the risk factors. We include stochastic variables for local inflow, currency spot and forward rates, and electricity spot and futures prices. All processes are estimated using average semi-monthly observations, with estimation windows as shown in Table 3.1. At last, we also show how the state variables are discretized using a scenario lattice, including the discretization of production decisions.

	Start	**End**
Electricity forward price	03.04.2011	31.12.2014
Inflow	01.01.1958	31.12.2014
EURNOK	04.01.1999	31.12.2014

Table 3.1: Data window for parameter estimation in the stochastic processes

3.4.1 Price process

We generate future scenarios of spot and futures prices using movements in a forward curve. As explained earlier, tradable forward and future contracts in the Nordic power market do not have delivery at a single point in time; instead, they have delivery periods stretching over a time period. A forward curve[4] aims to estimate the forward price for delivery at specific points in time, based on all contracts available in the market. In other words,

[4]The term forward curve might be confusing, as this paper considers futures contracts. However, as NASDAQ OMX only offered forward and not futures contracts for delivery periods exceeding one week before 2016, most literature uses the term forward curve.

the value of the curve at time T represents the price of a fictional forward contract with delivery exactly at time T. Multiple ways of constructing high-resolution forward curves are presented in the literature, e.g. by Fleten and Lemming (2003), Benth et al. (2007), Alexander (2008) and Kiesel et al. (2018).

We let $F_{t,T}$ denote the time t price of a forward contract with maturity at time T. For a forward contract with immediate delivery ($T = t$), the price of that contract is the current time spot price $F_{t,t}$. Thus, a stochastic process for the evolution of a forward curve can be used to generate future scenarios of its underlying, in this case, the spot price. The generated scenarios will also incorporate the seasonality of electricity prices, which Johnson and Barz (1999) found to be an essential characteristic of the market.

To represent the evolution of the electricity forward curve we use the HJM framework, originally presented by Heath et al. (1992). As we use the risk-neutral measure and the initial investment of a forward contract is zero, the expected return of the forward contract must be zero. Like Koekebakker and Ollmar (2005), we let the volatility of a forward contract with maturity at T, $\sigma_{t,T}$, be a function of time to maturity $T - t = \tau$ only. The process explaining movements in the forward curve will then be given by[5]

$$\begin{aligned} \frac{dF_{t,T}}{F_{t,T}} &= \sigma_{t,T} dZ_{t,T} = \sigma_{\tau} dZ_{\tau,t} \\ \mathbb{E}(dZ_{\tau,t}, dZ_{\hat{\tau},t}) &= \rho_{\tau,\hat{\tau}} dt, \quad \tau, \hat{\tau} \in [\tau] \end{aligned} \tag{3.23}$$

Here, $dZ_{\tau,t}$ and $dZ_{\hat{\tau},t}$ are Wiener processes associated with forward contracts with time to maturity τ and $\hat{\tau}$, respectively. $dZ_{\tau,t}$ and $dZ_{\hat{\tau},t}$ are correlated by $\rho_{\tau,\hat{\tau}}$. $[\tau]$ denotes the set of all time to maturities where $\Delta t \leq \tau \leq \hat{T}$. Since our decision problem considers discrete time stages, (3.23) must be discretized. By using Ito's lemma and setting $dt = \Delta t$, the process can be written as

$$F_{t,T} = F_{t-\Delta t,T} \cdot \exp\Big(-\frac{1}{2}\sigma^2_{\tau+\Delta t}\Delta t + \sigma_{\tau+\Delta t}\sqrt{\Delta t}\varepsilon_{\tau,t} \Big) \tag{3.24}$$

Here, $\Delta Z_{\tau,t} = \sqrt{\Delta t}\varepsilon_{\tau,t}$, where $\varepsilon_{\tau,t} \sim N(0,1)$. We can modify (3.24) into an expression for the spot price as a function of $F_{t-\Delta t,t}$, given by

$$F_{t,t} = F_{t-\Delta t,t} \cdot \exp\Big(-\frac{1}{2}\sigma^2_{\Delta t}\Delta t + \sigma_{\Delta t}\sqrt{\Delta t}\varepsilon_{\Delta t,t} \Big) \tag{3.25}$$

In the final lattice, each node contains a forward curve covering the remaining periods of the model horizon. In other words, the forward curve at stage t has $\hat{T} - t$ discrete price points $F_{t,T}$, where $T = [t, ..., \hat{T}]$. The price with the largest time to maturity τ is removed in every transition. $F_{t,T}$ represents the price of a non-traded forward contract with delivery in time stage T, whose delivery period is a semi-month. The spot price $F_{t,t}$ is the first element of the curve. The price of the contracts traded in the market ($F_{t,Mi}$, $F_{t,Qj}$ and $F_{t,Y1}$) are obtained using the forward curve. They can be found by calculating the average price of the forward curve covering the delivery periods of the corresponding contracts. To construct the underlying curve used as input in the first time stage of the model, we use the method proposed by Fleten and Lemming (2003)[6].

To estimate the correlation matrix and the volatilities that describe the electricity forward curve dynamics, we must construct a set of semi-monthly log returns for all forward contracts with time to maturity $\tau \in [\tau]$. In our case, $[\tau] =$

[5] See Appendix A.7.1 for more extensive details and analyzes of the implemented price process.

[6] See Appendix A.3 for an explanation of the method for constructing forward curves.

$[1,...,48]$ semi-months. To calculate these returns series, Koekebakker and Ollmar (2005) propose constructing multiple high-resolution forward curves for a large set of historical trading days. For this, we adopt the method of Alexander (2008), which uses linear interpolation of forward prices in the market to estimate the forward curves[7]. Although the forward curves constructed using the method of Fleten and Lemming (2003) are smooth and continuous, we experience issues with unrealistic oscillations in the near end of some of the forward curves. The curves constructed using the method of Alexander (2008) are neither smooth nor continuous, but they resulted in a reasonable volatility curve, as shown in Figure 3.4. The log returns are calculated between the average semi-monthly values of the forward curves of two consecutive periods[8]. Using log returns allows us to aggregate returns over longer time periods by addition, an approach also used by Bjerksund et al. (2008).

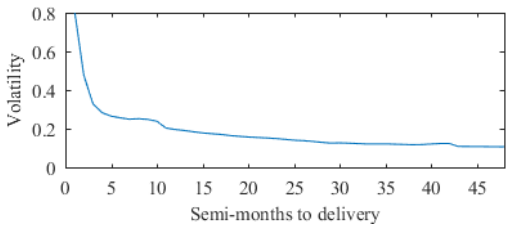

Figure 3.4: Annualized volatility curve σ_τ for electricity forward contracts.

Using the time series of returns, we can estimate the volatility curve for the term structure of forward prices. In Figure 3.4, the volatility curve can be understood as the volatility of returns of forward contracts with time to maturity τ. The volatility is monotonically decreasing with increasing time to maturity. This is called the Samuelson effect, originally proposed by Samuelson (1965). It means that forward prices tend to change more the closer they come to maturity. 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 forward prices.

3.4.2 Inflow process

Since the inflow time series provided from the case plant only contains data on aggregated inflow into both reservoirs, we must treat inflow as a single stochastic variable $Y_t = Y_{1,t} + Y_{2,t}$. In order to obtain $Y_{1,t}$ and $Y_{2,t}$, we have used the historical inflow split. We let $\zeta = 0.395$ denote the historical fraction of inflow flowing into reservoir 1, giving us $Y_{1,t} = \zeta Y_t$ and $Y_{2,t} = (1-\zeta)Y_t$.

We fit the geometric periodic autoregressive (GPAR) model suggested by Shapiro et al. (2013) to the inflow data for the hydropower plant[9]. Shapiro et al. (2013) found that a first-order periodic autoregressive process for the log-inflows provides a good description of their dataset. We find that it is suitable for our dataset as well. The inflow process is given by

$$Y_t = Y_{t-\Delta t}^{\phi_t} \exp\left(\hat{\mu}_t - \phi_t \hat{\mu}_{t-\Delta t} + \varepsilon_{Y,t}\right) \tag{4.19}$$

Here,

- Y_t is the inflow in period $t = 1,...,24$
- $\hat{\mu}_t$ is the mean log inflow in period $t = 1,...,24$
- ϕ_t is the time-dependent coefficient in the

[7] See Appendix A.4 for an explanation of the method for constructing forward curves.
[8] See Appendix A.5 for detailed procedure.
[9] See Appendix A.7.2 for an uncertainty diagram of the process.

autoregressive process in period $t = 1, ..., 24$

- $\varepsilon_{Y,t} \sim N(0, \sigma^2_{Y,t})$ is the error term representing the difference between the observed and predicted value in the autoregressive process
- $\sigma_{Y,t}$ is the time-dependent standard deviation of the error terms in semi-month $t = 1, ..., 24$

Since inflow Y_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, which is one of the prerequisites for representing a decision problem as a Markov decision process.

For a right-skewed inflow distribution such as the one that can be seen in Figure 3.5(a), a geometric process is better suited than an arithmetic process. It better captures the inflow dynamics, which can be extreme. Further, a geometric process does not allow for negative inflows. Shapiro et al. (2013) found the inflow distribution for Brazilian hydropower plants to be right-skewed as well, favoring a log transformation of the inflow observations.

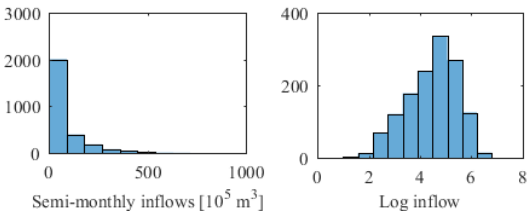

(a) Inflow (b) Log inflow

Figure 3.5: Historical distribution of semi-monthly inflow and log inflow for case plant.

The deviation of the log inflows from their mean, $\ln Y_t - \mu_t$, is represented as an AR(1) process. The suitability of a 1-lag process can be determined by investigating the partial autocorrelation of the historical data for $\ln Y_t - \mu_t$. 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 4.6 shows the partial autocorrelation of the $\ln Y_t - \mu_t$ time series. Similar to the findings of Shapiro et al. (2013), our dataset shows 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.

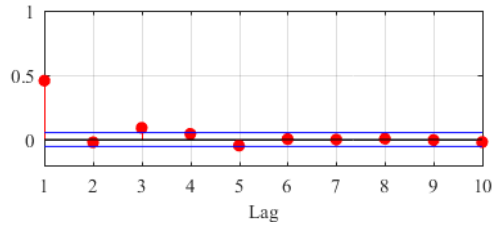

Figure 3.6: Partial autocorrelation of the $\log Y_t - \mu_t$ time series

3.4.3 *Exchange rate process*

Exchange forward rates can be priced using *covered* interest rate parity, as demonstrated by Aliber (1973). We adopt this method. Covered interest rate parity states that the interest rate differential between two currencies in the money markets should equal the differential between the forward and spot exchange rates. Otherwise, one could make arbitrage profits by simultaneously trading spot and forward[10]. Covered interest parity is given by

$$Q_{t,T} = Q_{t,t} \exp\left[(r - r_f)(T - t)\right] \tag{3.26}$$

Here, $Q_{t,T}$ is the forward exchange rate at time t with maturity at time T, $Q_{t,t}$ is the spot exchange rate, r is the domestic interest rate and r_f is the foreign interest rate.

In reality, covered interest rate parity holds closely, but not exactly (Levi, 2005). The failure to

[10] Such an arbitrage is demonstrated in Appendix A.6.

achieve exact covered interest parity could be due to transaction costs, political risks, and tax differences.

Further, we assume that *uncovered* interest rate parity holds. Uncovered interest rate parity states that the interest rate differential between two currencies in the money markets should equal the differential between the expected future exchange rate and spot exchange rates, as explained in Cumby and Obstfeld (1982). Uncovered interest rate parity is given by

$$\mathbb{E}(Q_{T,T}) = Q_{t,t} \exp\left[(r - r_f)(T-t)\right] \tag{3.27}$$

According to Bekaert et al. (2007), uncovered interest rate parity generally does not hold. The reason we include such a relationship is that under the risk-neutral probability measure, the forward exchange rate must equal the expected exchange rate. Since the upfront investment of a forward contract is zero, there can be no expected gains from entering into a contract under the risk-neutral probability measure.

To model the uncertainty in the exchange rate, we follow Huchzermeier and Cohen (1996) and use a geometric Brownian motion with drift equal to the interest rate differential, given by

$$\frac{dQ_{t,t}}{Q_{t,t}} = (r - r_f)dt + \sigma_Q dZ \tag{3.28}$$

Here, σ_Q is the annualized volatility of exchange rate returns. In discrete time, using $\varepsilon_{C,t} \sim N(0,1)$, we have that

$$Q_{t,t} = Q_{t-\Delta t,t-\Delta t} \exp\left((r - r_f - \frac{1}{2}\sigma_Q^2)\Delta t + \sigma_Q \sqrt{\Delta t}\varepsilon_{C,t}\right) \tag{3.29}$$

The EURNOK volatility is estimated using historical semi-monthly returns. The parameters of the currency process are shown in Table 3.2.

EURIBOR 3yr	**NIBOR 3yr**	**Initial EURNOK rate**	**Annualized volatility**
r_f	r	Q_0	σ_Q
0.13%	1.26 %	8.7035	5.7%

Table 3.2: Parameters of the currency process. Interest rates are logarithmic.

Figure A.6 in Appendix A.7.3 shows that the empirical forward curve dynamics of EURNOK exchange rate exhibit almost perfect correlation. Thus, for practical purposes, all currency risk is assumed to originate from the spot exchange rate, and a single factor model for the currency spot rate is sufficient. Once the spot rate is known, the forward rate will be deterministically given by the interest rate differential. This is in contrast to electricity markets, in which there is uncertainty both in the spot price and the forward curve. Due to the non-storable nature of electricity, there is no equivalent relationship between the electricity spot and forward prices.

3.4.4 Correlation between increments

We investigate the correlation between the increments of the stochastic processes. They are estimated using their historical Pearson correlation coefficients.

The electricity forward return series show a substantial degree of inter-correlation. This can be seen from the correlation matrix that is shown in Figure 3.7. Further, the correlation matrix shows that there is a clear decreasing trend in the correlation between contracts with larger maturity spreads. Using a 1% significance level, all intercorrelations of the electricity forward curve were found to be significant[11].

The increments of the inflow process and the

[11] The significance level of all correlations can be found in Appendix A.7.5.

price process show a weak negative correlation, ranging from -0.18 to -0.03 for different parts of the forward curve. Thus, there seems to have been a weak, natural hedging effect historically. However, all inflow-price correlations were found to be insignificant at both a 1% and a 5% significance level.

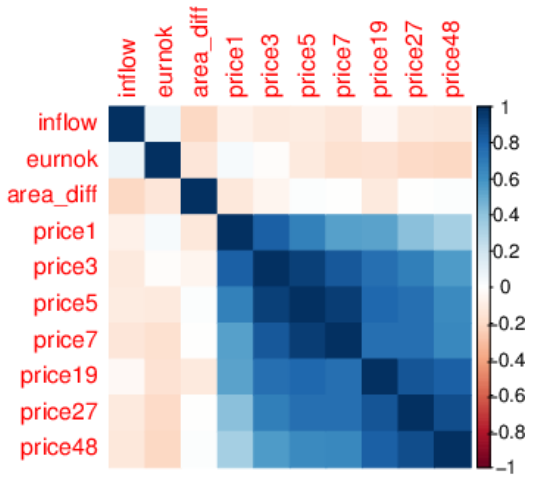

Figure 3.7: Correlation matrix for the increments of the stochastic processes. price1 denotes the first element of the forward curve (the spot price), and so on.

The increments of the currency process and the price process generally show a weak negative correlation, mainly between the currency increments and the long end of the electricity forward curve. For time to maturity $\tau \geq 6$ semi-months, the correlation coefficient ranges from -0.28 to -0.14 for different parts of the forward curve. Thus, there seems to have been a weak, natural hedging effect historically also here. The currency-price correlations were found to be insignificant at a 1% significance level. At a 5% significance level, they were significant for $\tau \geq 28$ semi-months.

3.4.5 Scenario lattice for state variables

To solve the MDP, we must discretize the Markov processes that describe movements in inflow, currency, and the electricity forward curve. We use a joint lattice for inflow and the forward curve with 100 nodes per stage. Each node at time t contains state variables for inflow, spot electricity price, and $48-t$ forward prices comprising the forward curve. The increments of these processes are correlated by their historical correlation coefficients, shown in Figure 3.7. The currency lattice is generated separately with ten nodes per stage, meaning that currency is treated as independent from inflow and price. We believe this is an acceptable approximation, as the price-currency relationship is not very strong in our dataset[12]. To create a combined lattice of currency, inflow, and prices, we calculate the cross-product of the price-inflow lattice and the currency lattice. This results in a combined lattice consisting of 49 stages and 1000 nodes per stage.

Due to the backwards estimation adjustment of the position of the nodes, there is some variation in the values assigned to a particular node when creating multiple lattices with the same processes. Thus, when we constructed the final lattice, we attempted to obtain initial node values that were consistent with the true input parameters. Further, we use the same lattice for all analyses to eliminate inaccuracies. Some of the true start values used in the processes, the start values in the final lattice and the standard deviation of the start values for multiple lattices are given in Table 3.3[13].

[12] Another reason for creating the currency lattice independently is that we were not able to create a high-quality joint lattice of spot price, forward prices, inflow, and currency, due to issues with the backwards estimation algorithm.

[13] A more comprehensive analysis of the lattice stability is found in Appendix A.8.

	True value	Std.dev	Lattice value
Inflow Y_0	80.60	3.83	75.94
Currency $Q_{0,0}$	8.704	0.0029	8.704
Spot $F_{0,0}$	25.53	0	25.53

Table 3.3: Comparison of true starting values, lattice values and the standard deviation of the ones constructed by the lattice for March 16, 2015

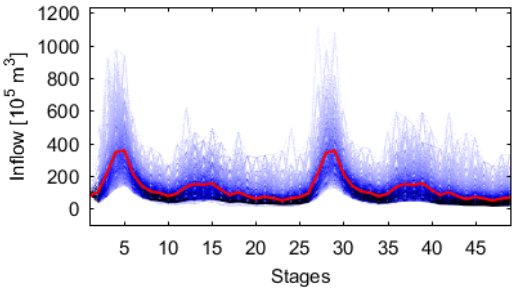

(a) Inflow lattice, 100 nodes

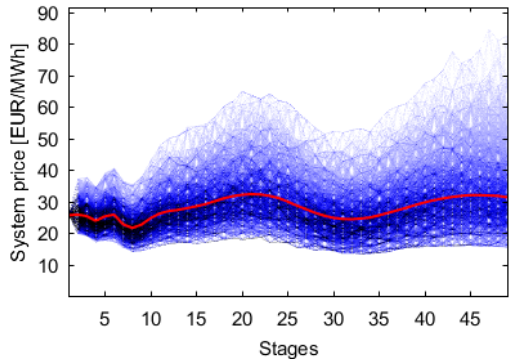

(b) Spot price lattice, 100 nodes

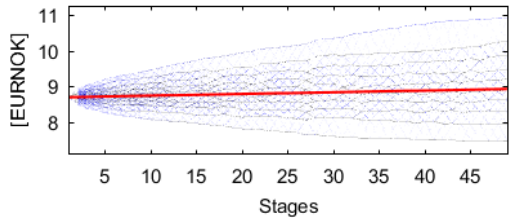

(c) Currency lattice, 10 nodes

Figure 3.8

For our case study, we use March 16, 2015 as the starting date, resulting in a model whose ultimate stage is in the end of March 2017. This is usually right before the start of the spring precipitation in Norway, and many reservoirs are close to empty. Therefore, we find it reasonable to disregard constraints on the terminal reservoir levels in the production model. As start values, we use the mean inflow recorded for the second half of March 2015 and the currency rate at this date. The initial state forward curve was constructed using the close prices of all monthly, quarterly and yearly contracts at the starting date. The scenario lattices for inflow, spot price and currency are plotted in Figure 3.8.

In the hydropower scheduling problem, the production in period t, w_t, is a decision variable. For the hedging problem, we treat production as an exogenous, stochastic variable W_t. To obtain the stochastic production levels, we first obtain the optimal production decisions in all nodes by running the production model. For this, we only use the price-inflow lattice, as we treat production decisions as independent from the currency exchange rate. Thus, we obtain a unique stochastic production level for all states of price and inflow. Next, we run 10^5 simulations of the production model, and calculate the production W_t in node $\overline{S_{t,n}}$ as the average of all simulated production decisions made in that particular node. During the 10^5 simulations, the least visited node had 127 visits. Figure 3.9(a) shows the simulated realizations of the production decisions obtained by the production model, and Figure 3.9(b) shows how the production decisions were discretized into a lattice with 100 nodes per stage. The red line, which represents the mean production at stage t, is almost identical in both figures. This ensures that the expected total production is changed only marginally (0.1% increase) by the discretization of the production decisions. The

mean yearly production values obtained by the production model, in the production lattice and historically are listed in Table 3.4.

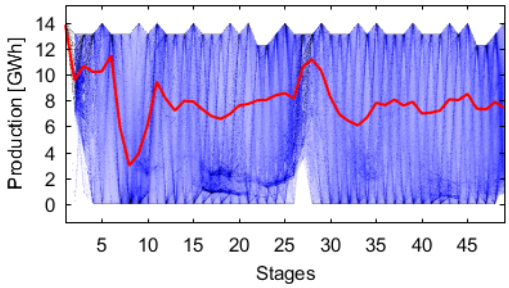

(a) Simulated production decisions

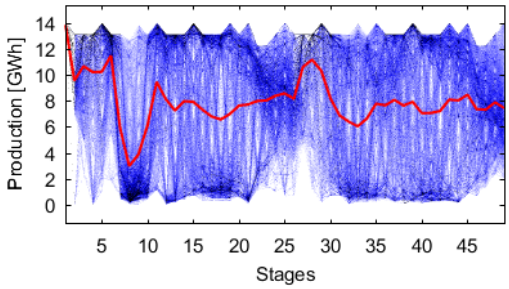

(b) Discretized production lattice

Figure 3.9

Production model	Production lattice	Historical
192.03	192.18	191.30

Table 3.4: Yearly mean production in production model, production lattice and historically in GWh. The former two values are based on $5 \cdot 10^5$ simulations

3.5 Numerical results

In this section, we summarize the numerical results of the hedging model. Unless otherwise stated, we use the sequential approach and treat production as a stochastic variable. For all analyzes, we have used the lattices presented in Section 3.4.5 and the coefficient values in Appendix A.9. Further, we have used maximum 500 pairs of forward-backward passes in the ADDP algorithm, and all simulated results are based on 10^5 simulations[14]. All code has been implemented in MATLAB and R, and we have used QUASAR, a general-purpose solver for stochastic optimization (Löhndorf, 2017), to construct the lattices and perform ADDP. Computations were conducted on a computer with 32 GB memory and 3.6 GHz speed.

3.5.1 Hedging performance

Here, we present the main results of the sequential hedging model. We are, in particular, interested in the risk reduction hedging can provide compared to the case of no hedging. Additionally, we also seek to explain the decision policies driving this reduction. To quantify the hedging effect, we use statistical measures of the discounted terminal cash flows over the two-year horizon. For different model specifications, we find the mean and standard deviation of the terminal discounted cash flows, in addition to the terminal VaR and CVaR at the significance levels 0.05 and 0.01. In Table 3.5, we have displayed the results for the cases with trading in all contracts, with trading in all contracts except for monthly power futures, and with no forward trading and only production. Recall that we use the nested CVaR to model risk preferences, which is different from the terminal CVaR. However, to the authors' knowledge, it is not possible to quantify the nested CVaR, favoring the terminal CVaR for risk measurement purposes.

As the terminal VaR and CVaR in Table 3.5 show, the dynamic hedging model clearly reduces the risk compared with the no hedging-case, utilizing currency forwards and power futures contracts. The mean terminal cash flow is slightly reduced, which can largely be explained by transaction costs. The risk reduction in terms of terminal risk measures

[14] Sensitivity analysis of results obtained using 500 passes and 10^5 simulations can be found in Appendix A.10.

is also evident when using only quarterly and yearly contracts, but it is higher when trading in monthly contracts is allowed. Note that the significance level α used in the *nested* CVaR does not necessarily maximize the *terminal* CVaR with the same significance level. This can be observed by comparing the results using $\alpha = 0.05$ and the results where $\alpha = 0.1$. The latter performs better both in terms of terminal CVaR(5%) and CVaR(1%).

Available contracts	None	All	All	All except Month	All except Month
α	-	0.05	0.1	0.05	0.1
Mean	40.50	40.26	40.23	40.35	40.35
Std	7.99	4.56	4.34	6.10	5.97
VaR(5%)	28.22	33.06	33.23	30.61	30.84
VaR(1%)	24.25	30.48	30.68	27.17	27.44
CVaR(5%)	25.82	31.45	31.66	28.53	28.75
CVaR(1%)	22.65	29.25	29.44	25.64	25.85
Mean prod. /yr.	192.18	192.17	192.08	192.22	192.15

Table 3.5: Mean value and statistical measures of the terminal cash flows in million NOK (M NOK) for a significance level α used in the nested CVaR expression. If all contracts are available, currency forwards and monthly, quarterly and yearly power futures can be traded. For reference, we also include the average yearly production [GWh] from each test.

The average optimal hedge ratios for power futures trading highly depends on the availability of monthly contracts. Table 3.6 shows the mean total hedge ratio of the four model variants, which is the mean ratio between the total short position in financial power contracts and the total amount of produced energy. It also includes the mean hedge ratio before the start of a new quarter and the start of the first year. The large difference between the quarterly and total hedge ratios for the cases including all hedging instruments indicates that a large quantity of monthly contracts is traded in the last three months ahead of delivery. Generally, the model prefers using monthly contracts, as the short position in such contracts constitutes 97% of the total trading volume [GWh] for the case where $\alpha = 0.10$. This might be because they allow for precision hedging of the expected production in a given month. They also allow for hedging of the most near-term cash flows.

Available contracts	All	All	All except Month	All except Month
α	0.05	0.1	0.05	0.1
HR	1.543	1.582	0.479	0.514
HR Q	0.707	0.719	0.479	0.493
HR Y	0.256	0.274	0.172	0.184

Table 3.6: Mean hedge ratios, i.e. total short position in power futures divided by total expected long position in power production. The first row (HR) denotes the total hedge ratio, whereas the second denotes the mean ratios before the start of a new quarter. The third row denotes the mean hedge ratio before the start of the first year.

According to Sanda et al. (2013), a full hedge can be obtained with a hedge ratio of 0.53, due to taxation effects[15]. Compared to a hedge ratio of 0.53, the case allowing for trades in all contracts proposes a significant over-hedge. This might be an effect of modeling risk preferences using the nested CVaR. Recall that the cash flows from a short position in a power futures contract will be positive if the underlying expected spot price goes down and negative if it goes up. Assume that there is a positive correlation between the spot price and production in a semi-month t, i.e. that the production level is higher for high-price states than low-price states. Since the objective of the model is to maximize the nested CVaR, it will seek to increase the lower-tail cash flows as much as possible. Thus, a possible explanation for the over-hedge could be that the

[15]The derivation of this hedge ratio can be found in Appendix A.11.

model seeks to obtain a hedge ratio where the total cash flows from production and hedging are lifted as much as possible in the states where the price goes down, thereby increasing the CVaR, while simultaneously keeping the total cash flows in states where the price goes up above the lower-tail cash flows[16]. Also, as Table 3.6 shows, the largest amount of the over-hedged short position is entered less than three months before maturity with monthly contracts. Due to the short delivery period and low time to delivery of these contracts, the volatility of the cash flows from the over-hedged position will probably be moderate.

Available contracts	All	All	All except Month	All except Month
α	0.05	0.1	0.05	0.1
Hedge ratio	0.563	1.12	0.589	0.95
Maturity 1-5	0.397	0.560	0.684	0.806
Maturity 6-15	0.090	0.100	0.079	0.046
Maturity 16-30	0.192	0.186	0.053	0.049
Maturity 31-48	0.321	0.153	0.184	0.099

Table 3.7: Amount of trading in the currency forward market for different versions of the model. The first row denotes the currency hedge ratio, found by dividing the mean total short position in currency forward contracts minus taxes by the mean total cash flows from the entire portfolio of production and hedging, both in EUR. Thus, a full hedge will be obtained by a hedge ratio of 1. The next rows shows the percentage of trades performed in different intervals of time to maturities, denoted in semi-months.

The extent of trading in currency forwards is shown in Table 3.7. The table shows that the currency market is extensively used. Similar to the trading activity in power futures, currency forwards with shorter maturity times are preferred. However, the amount of trading in contracts with longer times to delivery is larger than for the power futures. We also observe that the extent of currency forward trading is larger when using $\alpha = 0.1$ in the nested CVaR expression, reaching volumes close to a full hedge position. In reality, we would expect the optimal level of currency hedging to be slightly lower, due to the negative correlation between price and currency that was not included in the final scenario lattices.

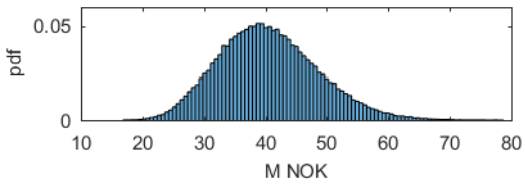

(a) Discounted terminal cash flows from production

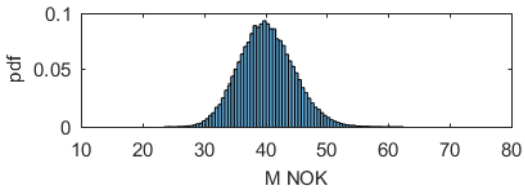

(b) Total discounted terminal cash flows (after hedging)

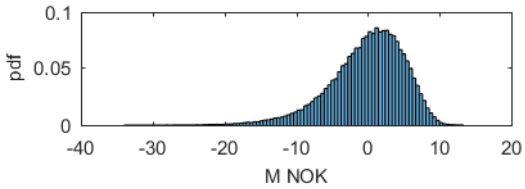

(c) Discounted terminal cash flows from power futures trading

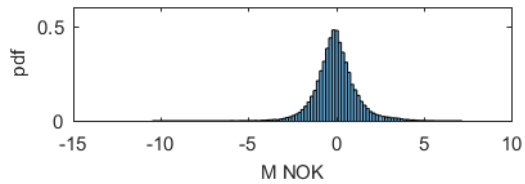

(d) Discounted terminal cash flows from currency forward trading

Figure 3.10: Cash flow distributions. From hedging model with all contracts and $\alpha = 0.10$

Further, we have plotted the distributions of the terminal cash flows from production, terminal cash flows after hedging, terminal cash flows from power

[16] A numerical example explaining this argument can be found in Appendix A.12.

futures and terminal cash flows from currency forwards in Figure 3.10. The cash flows are from the case in which we specify $\alpha = 0.1$. The hedging effect is evident when comparing the terminal cash flow distribution in Figure 3.10(a) and 3.10(b). The distribution of cash flows is visually compressed after hedging. Also note that the mean in Figure 3.10(c) and 3.10(d) is zero, which is correct assuming that there are no expected gains from trading in the forward markets.

3.5.2 Comparison of sequential and simultaneous approach

For comparison, we have also developed an integrated model that treats the problems of production and hedging simultaneously. We have tested this model for different values of the weighting parameter λ and significance level α, and analyzed the hedging performance of each version using the same risk measures as in Table 3.5. The results are depicted in Table 3.8. In compliance with the argumentation of Wallace and Fleten (2003), the results show that the simultaneous approach performs slightly worse compared to the sequential one. The main reason for this is that the model no longer treats production as a risk-neutral maximization problem, as the problem is now risk-averse. Thus, the production decisions in the simultaneous approach will seek to reduce the lower-tail distribution of the cash flows. The risk-aversion also affects the decision policies regarding the water level in reservoir 2, which has a large penalty cost associated with violating the summertime restriction. While the ratio of mean cash flows to production is quite similar in the simultaneous and the sequential approach, the risk-aversion results in a lower mean total production in the former case. This further leads to lower mean discounted cash flows, terminal VaR and CVaR. The simultaneous approach also has a higher standard deviation.

Model type	Sim.	Sim.	Sim.	Sim.	Seq.
Available contracts	All	All	All	All	All
α	0.05	0.1	0.05	0.1	0.1
λ	1	1	0.5	0.5	1
Mean	37.23	37.73	39.50	39.69	40.23
Std	4.60	4.77	4.57	4.64	4.34
VaR(5%)	30.30	30.46	32.45	32.48	33.23
VaR(1%)	28.13	28.06	29.95	29.99	30.68
CVaR(5%)	28.99	29.02	30.92	30.96	31.66
CVaR(1%)	27.13	27.07	28.75	28.80	29.44
Mean prod. /yr.	177.28	179.46	187.16	188.93	192.08
Mean CF / prod.	102.9	103.0	103.4	102.9	102.6

Table 3.8: Results of the simultaneous production and hedging model. For reference, we also include the statistical measures from the sequential model with $\alpha = 0.1$. The last metric denotes the ratio between the mean cash flows and mean production volume over the entire time horizon in [NOK/MWh].

3.5.3 Effect of currency hedging

One of the main contributions in this paper is the inclusion of currency risk and currency derivatives. To quantify the effect of currency hedging, we compare the hedging performance when currency forwards can be traded with the case in which currency derivatives cannot be traded. That is, we restrict hedging activity to power futures contracts only, and compare this to the case in which currency forward trading is allowed. Note that we include currency risk also when currency hedging is not allowed. However, we formulate the problem such that the optimal decision policies are independent of the uncertain currency rate. The results can be seen in Table 3.9. They show that when currency hedging is introduced, the standard deviation of the terminal cash flows decreases, and the terminal VaR and CVaR increase. Although currency hedging

reduces the market risk of the hydropower plant, the magnitude of the change is relatively small, less than 10% for all risk measures. This can be explained by the low volatility of the currency exchange process compared to that of the electricity price process.

	W/o currency hedging	**W/o currency hedging**	**W/ currency hedging**	**Change**
α	0.05	0.1	0.1	
Mean	40.43	40.39	40.23	
Std	4.99	4.82	4.34	-9.96%
VaR(5%)	32.45	32.71	33.23	1.59%
VaR(1%)	29.42	29.76	30.68	3.09%
CVaR(5%)	30.60	30.91	31.66	2.43%
CVaR(1%)	28.05	28.32	29.44	3.95%
Mean prod. /yr.	192.26	192.15	192.08	

Table 3.9: Effect of currency hedging on terminal cash flows. The last column denotes the percentual difference in the risk measures between the cases with and without currency hedging for $\alpha = 0.1$ in the nested CVaR expression.

3.5.4 Comparison with heuristics

As shown by Wang et al. (2015), simple hedging strategies might often yield better or equivalent hedging performance as more advanced procedures. Therefore, we have replicated the hedging strategy of a Norwegian hydropower producer, and compare its performance to the performance of our model. As shown in Sanda et al. (2013), most Norwegian companies use a heuristical approach with specific hedge ratio ranges for different times to maturity. Since the time horizon of our model is two years, we have used the hedging strategy of a firm whose hedging activity also begins two years prior to maturity. Figure 3.11 displays a slightly modified version of the lower and upper bounds of the firm's required hedge ratios for delivery in a given month, as shown in Sanda et al. (2013).

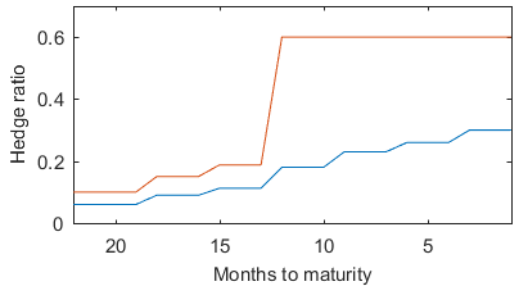

Figure 3.11: Required hedge ratio range of a real company. For example, in a given time stage, the expected hedge ratio (hedged volume divided by expected production) for the month that begins 12 months ahead must lie between 0.18 and 0.6.

To test the performance of the heuristical approach, we have included the hedge ratio ranges as constraints on financial short positions in the dynamic model. We have performed two separate simulations; one where the hedge ratio is set close to the lower restriction and one where it is closer to the upper restriction. As the currency hedging strategy of the firm is unknown, the option to trade currency derivatives has been removed. For the two aforementioned cases, the results are displayed in Table 3.10. From the perspective of maximizing terminal CVaR, the results show that a strategy close to the upper hedge ratio requirement is better than a strategy close to the lower. However, as stated in Sanda et al. (2013), the purpose of using a range instead of a specific requirement is to enable selective hedging. Therefore, the realized ratio will be adjusted based on the firm's market views. Further, the results show that the heuristical approach performs worse in terms of terminal risk measures than the proposed version allowing for all types of trading, but not by a large margin. Nevertheless, its performance is almost identical to the version omitting currency hedging for all risk measures except variance. These results might indicate that the over-hedging proposed by our approach, in which the largest short positions are

entered just before maturity, only has minor effects on the lower tail of the terminal discounted cash flows. This further illustrates the difference between the nested CVaR and the terminal CVaR, and that the former is not necessarily suitable for reducing risk measured by the latter. At last, the results show that the heuristical approach outperforms the one with no option to trade in monthly contracts, which underlines the benefit provided by precision hedging with monthly contracts.

Model type	Heuristic Lower	Heuristic Upper	Seq.	Seq.
Available contracts	All except currency	All except currency	All except currency	All
α	–	–	0.1	0.1
Mean	40.46	40.44	40.39	40.23
Std	5.88	5.10	4.84	4.34
VaR(5%)	31.40	32.50	32.71	33.23
VaR(1%)	28.52	29.63	29.76	30.68
CVaR(5%)	29.64	30.76	30.91	31.66
CVaR(1%)	27.21	28.33	28.32	29.44
Mean prod. /yr.	192.16	192.11	192.15	192.08
HR	0.302	0.584	1.505	1.582
HR Q	0.301	0.555	0.667	0.719
HR Y	0.259	0.421	0.255	0.274

Table 3.10: Results from hedging with an heuristic. For reference, we also include the statistical measures from the sequential model with $\alpha = 0.1$, both with and without the opportunity to trade in currency forwards. We also show the hedge ratios of each simulations, which are calculated in the same manner as in Table 3.6.

3.6 Conclusions

In this paper, we present a global dynamic model for hydropower risk management. We treat the problem sequentially. First, we obtain the optimal production decisions, and then we hedge the cash flows from production using currency forwards and electricity futures. Both the production planning problem and the hedging problem are modeled as Markov decision processes. We include correlated uncertainty in electricity spot and futures prices and reservoir inflow. The currency exchange rate is assumed to be independent. In order to reduce their risk exposure, the decision maker is allowed to trade in currency forward contracts and monthly, quarterly and yearly power futures. Risk preferences are modeled using the nested CVaR. The stochastic variables are discretized using a scenario lattice, and the resulting stochastic-dynamic decision problem is solved using ADDP. The results show that the dynamic hedging model substantially reduces the risk exposure compared to the case with no hedging, demonstrated by increasing the CVaR(5%) of the terminal cash flows by 23%.

We quantify the effect of including currency derivatives in the hedging strategy for a Norwegian hydropower producer. The effect is found to be moderate, as including currency derivatives results in a variance decrease of 9.96% and CVaR(5%) increase of 2.43% for the terminal cash flows. While not included in the scenario lattice, we find that there generally is a weak, negative correlation between the semi-monthly increments of the currency spot rate and the electricity forward curve. The magnitude of the correlation coefficient is larger for the long end of the forward curve, ranging from -0.28 to -0.14 for time to maturity $\tau \geq 6$ semi-months. Although these values are not very large and were found to be insignificant at a 5% significance level for $\tau \leq 28$ semi-months, they suggest that there might exist a weak, natural hedging effect.

Further, we have tested how a simultaneous approach in which we solve the problems of production planning and hedging simultaneously compares to the sequential one. In line with

the argumentation of Wallace and Fleten (2003), we find that the sequential outperforms the simultaneous both in terms of return and risk management, obtaining higher values for the mean, VaR and CVaR of the terminal discounted cash flows, in addition to a lower variance. This is mainly because the simultaneous model does not treat production as a risk-neutral maximization problem, as it considers risk-aversion.

We also investigate how the hedging model performs compared to a heuristical approach based on hedge ratio ranges, which is extensively used by Norwegian hydropower firms. This is done by implementing the hedge ratios of a company analyzed by Sanda et al. (2013) as constraints in the dynamic hedging model. The performance of the hedge ratio approach is slightly worse than the base model, (2.8% lower CVaR(5%)), but almost identical to the case with no currency trading. This suggests that a simple hedge ratio approach can be quite efficient, in line with the findings of Wang et al. (2015).

The results of the base model suggest that it is optimal to over-hedge expected production, primarily by using monthly futures contracts. This might be an effect of using the nested CVaR as the risk measure. The over-hedge might not have been experienced if we had modeled risk preferences based on the terminal cash flows instead. While the model considers over-hedging to be optimal. the performance of the hedge ratios approach suggests that the effect of over-hedging might be minor in terms of reducing the risk of the terminal cash flows. This further questions the suitability of the nested CVaR to represent risk preferences, and also suggests that more research should be conducted on risk measurement when using conditional risk mapping.

Both the base model, the heuristical version and the version without currency hedging perform better in terms of terminal risk measures than the approach with no option to trade in monthly power futures. Possible explanations for this include that monthly contracts allow for more precise hedging of the expected production in a given month than the contracts with longer maturity periods, as well as increased flexibility in the timing of trading due to their shorter delivery period. These findings suggest that hydropower producers can benefit from using monthly contracts to a greater extent. While the case company used in the heuristical analysis has historically made extensive use of monthly power futures contracts, their use among the other firms analyzed by Sanda et al. (2013) is marginal.

Nomenclature

Parameters

α	Significance level
β_t	Semi-monthly discount factor
γ_c	Corporate tax rate
γ_r	Resource rent tax rate
$\hat{\mu}_t$	Mean log inflow in period t
κ	Energy coefficient [MWh/m^3]
λ	Risk preference weighting coefficient
ϕ_t	Periodic auto-regressive coefficient in the GPAR inflow process
ρ	Correlation coefficient
σ_Q	Annualized volatility of exchange rate returns
σ_τ	Volatility of forward contract with τ time to delivery
$\sigma_{Y,t}$	Periodic coefficient in the GPAR inflow process
υ	Soft constraint violation penalty
ς_t	Number of seconds in semi-month t
ξ	Maximum discharge through the turbines of the hydropower plant [m^3/s]
ζ	Split factor for inflow
p_{tnm}	Probability for a transition between node n and m at time t
c_f	Transaction cost [EUR/MWh]
r	Domestic interest rate, NIBOR
r_f	Foreign interest rate, EURIBOR

Stochastic processes

$F_{t,Mi}, \quad i = [1,...,6]$	Price of monthly forward contract with delivery period i months ahead [EUR/MWh]
$F_{t,Qj}, \quad j = [1,...,8]$	Price of quarterly forward contract with delivery period j quarters ahead [EUR/MWh]
$F_{t,T}$	Forward price with maturity at time T [EUR/MWh]
$F_{t,t}$	Spot system price [EUR/MWh]
$F_{t,Y1}$	Price of yearly forward contract with delivery period in the upcoming year [EUR/MWh]
$Q_{t,T}$	Forward exchange rate with maturity at time T [EURNOK]
$Q_{t,t}$	Spot exchange rate [EURNOK]
W_t	Stochastic production [MWh]
$X, X_1, X_2, X_3, X_{t,l}$	Random variable
Y_t	Total inflow [m^3]
$Y_{1,t}, Y_{2,t}$	Inflow to reservoir 1 and 2, respectively [m^3]
Z_t	Wiener process

Decision variables

$s_{c,t}$	Amount of water flowing between reservoir 1 and 2 at time t [m^3]
$s_{s,t}$	Spilled water at time t [m^3]
$u_{t,Mi}$	Total short position [MWh] at time t in monthly contracts with i months to delivery
$u_{t,M}$	Short position [MWh] at time t in the monthly contract currently in delivery
$u_{t,Qj}$	Total short position [MWh] at time t in quarterly contracts with j quarters to delivery
$u_{t,Q}$	Short position [MWh] at time t in the quarterly contract currently in delivery
$u_{t,Y1}$	Total short position [MWh] at time t in year ahead contract
$u_{t,Y}$	Short position [MWh] at time t in the yearly contract currently in delivery
$v^S_{2,t}$	Slack variable for reservoir level 2 lower bound [m^3]
$v_{i,t}$	Reservoir volume [m^3]
w_t	Spot production [MWh]
$w_{t,Mi}$	New short position [MWh] entered at time t in monthly contracts with i months to delivery
$w_{t,Qj}$	New short position [MWh] entered at time t in quarterly contracts with j quarters to delivery
$w_{t,Y1}$	New short position [MWh] entered at time t in the year ahead contract
$x_{t,T}$	New short position [EUR] entered at time t in currency forwards with delivery time T
$y^C_{t,T}$	Confirmed future cash flow from currency forward contract [NOK]
$y^F_{t,T}$	Confirmed future cash flow from electricity forward contract [EUR]
$z_{t,T}$	Short position [EUR] at time t in currency forwards with delivery time T

Other variables

ε_t	N(0,1) distribution
ϕ	General risk measure given random variable X
π_t	Optimal decision policy at time t
$\psi_{\lambda,\alpha}(X)$	Risk measure with risk preference λ and significance level α given random variable X
H	Number of terminal cash flows h
K	Number of simulations to build lattice
L	Number of state variables in a lattice
N_t	Number of nodes at stage t
V_t^H	Value function in the hedging problem
V_t^P	Value function in the production planning problem
i	Number of months ahead of spot, index for terminal cash flow
j	Number of quarters ahead of spot
k	Simulation index, $k \in 1,...,K$
l	State variable index. $l \in 1,...,L$
n,m	Node index at time t, $n,m \in 1,...,N_t$
$\overline{h}$	Mean terminal cash flow
h_i	Terminal cash flow $i \in 1,...,H$
$\mathbb{E}$	Expectation
$\mathbb{I}_M, \mathbb{I}_Q, \mathbb{I}_Y$	Indicator function that is 1 if the next stage is the beginning of a month/quarter/year and 0 if not
$\mathbb{Q}$	Risk-neutral probability measure
$\overline{S}_{tn}$	The nth node at time t
S_t	Environmental state at time t
S_t^k	Simulation k at time t
$[\tau]$	Set of all time to maturities
$\hat{T}$	End of horizon time
$\tau, \hat{\tau}$	Time to maturity
T	Maturity time of a forward contract
t	Time
$CVaR_\alpha$	Conditional value at risk with confidence level $1-\alpha$
$CVaR_{\alpha,\lambda}^{NEST}$	Nested conditional value at risk with confidence level $1-\alpha$ and risk preference λ
VaR_α	Value at risk with confidence level $1-\alpha$

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

Joakim Dimoski[a], Sveinung Nersten[a,*]

[a]*Norwegian University of Science and Technology, Norway*

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. We use the price of forward contracts to forecast future spot prices, and use multiple factors to describe movements in price. Further, we include a short-term correlation between movements in electricity price and local inflow. Our main contribution is a comparison of 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 the case where price movements are in fact driven by multiple factors and correlated with local inflow. In both situations, we find the loss to be approximately 2-3 %. We have based our study on a Norwegian hydropower plant.

Keywords: Hydropower reservoir management, Markov decision process, multi-factor price process, price and local inflow correlation, stochastic dynamic programming

4.1 Introduction

The decision problem of hydropower producers, which seek to dispatch the water in their reservoir optimally, has existed for many years, and multiple approaches for formulating and modeling such problems have been proposed. Massé (1946) argues that deterministic models are not good enough, as they do not incorporate the flexibility a production planner has when it comes to the timing of production. Instead, one should use a flexible approach which can provide the hydropower producer with optimal decision policies for both the current and future states of the world, incorporating the uncertainty in future states. The flexible approach proposed by Massé (1946) is still relevant for how reservoir management is performed today.

*Lead author. Kolbjørn Hejes vei 1E, 7034 Trondheim, Norway. +47 458 23 304. sveinung.nersten@gmail.com

Inspired by Massé (1946) and multiple papers of more recent date, we aim to create a dynamic scheduling model for a price-taking hydropower production planner that participates in a deregulated power 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, we consider two stochastic variables (spot price and inflow), and we use a time horizon of two years and weekly granularity which is suitable for medium-term planning. 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).

Our contributions in this work include the use of a multi-factor price process, as opposed to existing models for reservoir management which often use single-factor processes to describe movements in price. We include a correlation coefficient between changes in price and local inflow, thereby treating them as dependent variables. Further, we quantify the loss in expected revenues if they assume price and local inflow to be independent when they are in fact correlated, and equivalently, the losses that occur if they use a single factor price process when price movements are in fact described by multiple factors.

To obtain optimal decision policies in each discrete state for the production planner, one can use stochastic dynamic programming as introduced by Bellman (1957). An issue with dynamic programming is the so-called curse of dimensionality, that is, the problem might become too complicated to solve when the state space and number of decision variables become too large. In order to avoid this, Pereira and Pinto (1991) introduce an algorithm for stochastic dynamic programming, a solution approach known as stochastic dual dynamic programming (SDDP). SDDP and similar approaches are widely used in existing literature on hydropower production scheduling, e.g. in Mo et al. (2001b) and Rebennack (2015). Löhndorf et al. (2013) introduce a framework that integrates SDDP with Markov processes, referred to as approximate dual dynamic programming (ADDP). 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 a Markov Decision Process (MDP), all decisions are made based on the current state of the world and its future expected states, irrelevant of all past states.

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. Therefore, we treat inflows and price movements as Markov processes and use a scenario lattice to discretize all future states and transition probabilities. To construct the lattice, we use the method proposed by Löhndorf and Wozabal (2017). We also use their method for solving stochastic dynamic programs, ADDP, to obtain all optimal decision policies.

We incorporate two stochastic state variables; spot price and inflow. EOPS (SINTEF, 2017b), which is one of the most common commercial programs for medium-term reservoir management for smaller systems in the Nordic countries, uses spot price scenarios generated using EMPS (SINTEF, 2017a). EMPS is a fundamental model which, among

others, can forecast spot prices in larger power systems by using historical scenarios of stochastic variables like area inflow and demand (Wolfgang et al., 2009). Mo et al. (2001b) show that there is a high correlation between the prices of successive weeks simulated 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), price scenarios in EOPS are still generated using EMPS, but the prices are now organized in a lattice using the scenarios directly instead of expressing them with an autoregressive process.

In contrast to how spot price scenarios are generated in EOPS, we generate them using movements in the price of forward contracts traded in the market. These movements are modeled using a multi-factor model, commonly referred to as an HJM model (Heath et al., 1992). 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 representing 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. Like Koekebakker and Ollmar (2005), we find the coefficients of the price process empirically by first constructing forward curves for many consecutive trading days, and then calculate daily deviations between the curves and use PCA to obtain multiple factors.

The other stochastic variable we consider 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 the use of autoregressive processes. In EOPS, the inflow for a local system is assumed to follow an ARIMA(1,1) process. Maceira and Damázio (2006) propose a periodic autoregressive process (PAR) for inflow in the Brazilian hydropower system. Since PAR allows for negative inflows and does not account for the skewness of the inflow distribution very well, Shapiro et al. (2013) propose to use geometric PAR models (GPAR). In GPAR, the deviations of the log inflows from their periodic mean are represented as an AR(1) process. We will adopt this approach in this paper.

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. (2001b). 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 found in Boger et al. (2017). Therefore, we include a correlation between movements in local inflow and the price of forward contracts in the stochastic processes.

The paper is organized as follows. In Section 4.2, we present the reservoir management decision problem as a mathematical program. We also give an overview of the stochastic processes used to describe the correlated movements of inflow and price and how these can be used to generate a scenario lattice. The section is concluded with a short description of ADDP, the framework used to solve the decision problem. In Section 4.3, we present a multi-reservoir hydropower plant in Norway on which we have tested our model. We also present the obtained process coefficients and correlation and show empirical results from running the model. The section is concluded with a calculation of the losses associated with using a single factor price process and from omitting the price-inflow correlation. Final conclusions are made in Section 4.4.

4.2 Methods

In the following sections, we will first formulate the decision problem associated with reservoir management as a mathematical program. Then, we will show how movements in the two relevant stochastic variables, spot price, and inflow, can be modeled. Further, we present how all future states of price and inflow can be discretized using a scenario lattice, and briefly present the solution method used to obtain optimal decision policies for each state.

4.2.1 Hydropower decision problem

In this part, we describe the problem faced by a price-taking hydropower production planner with multiple, interconnected reservoirs that participates in a deregulated market. 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.

The problem faced by the production planner is a stochastic dynamic decision problem, meaning that decisions must be made at different stages in time and 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 $Y_{b,t}$ into all reservoirs $b = [1,...,B]$. For

convenience, we denote $\hat{Y}_t = \{Y_{b,t} : b = [1,...,B]\}$ as a set of all inflows to all reservoirs at time t. Like Bjerksund et al. (2008), we assume that the decision maker participates in a complete market with no riskless arbitrage opportunities. 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. In a complete and arbitrage-free market, there would exist a unique risk-neutral, martingale measure $\mathbb{Q}$ that represents the risk-neutral probabilities of all future states for spot price and inflow.

Using the complete market and no-arbitrage assumption and denoting π_t as a decision policy at time t providing a cash flow of $CF_t = CF_t(P_t, \hat{Y}_t, \pi_t)$ and an appropriate discount factor $\beta < 1$, the expected discounted cash flows over a time horizon $\hat{T}$ are given by

$$\max_{\pi_t} \quad \mathbb{E}_{\mathbb{Q}}\left(\sum_{t=1}^{\hat{T}} \beta CF_t(P_t, \hat{Y}_t, \pi_t)\right) \tag{4.1}$$

Like Lamond and Boukhtouta (1996), we treat the reservoir management problem as a Markov decision process (MDP). The objective of MDPs is to obtain optimal decision policies (π_t) for all current and future states of the world. These policies maximize the value of all current and future cash flows, meaning that the policies do not only depend on their respective states, but also the space of potential future states. We denote by V_t the time t value of the current time cash flow and all future expected cash flows, and formulate it using the Bellman equation, first introduced by Bellman (1957)

$$V_t(P_t, \hat{Y}_t, \pi_t) = \max_{\pi_t} \quad CF_t(P_t, \pi_t) + \beta \mathbb{E}[V_{t+1}(P_{t+1}, \hat{Y}_{t+1}, \pi_{t+1} | P_t, \hat{Y}_t, \pi_t)] \tag{4.2}$$

Equation (4.2) 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 spot 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), cash flows can be set equal to revenues. Thus, (4.2) can be considered as the objective function of the decision problem. 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 4.3.2. Lastly, given a water density ρ and gravitational acceleration G, the

cash flow CF_t at time t can be written as

$$CF_t = P_t \cdot \rho \cdot G \cdot \varpi \cdot \sum_{b \in B} \left[H_{b,t} \cdot \sum_{i \in B, i \neq b} q_{bi,t} \cdot \eta_{bi,t} \right] \quad t = [1,...,\hat{T}] \quad (4.3)$$

We denote by $v_{b,t}$ the water level in $[m^3]$ in reservoir b at time t. Further, $s_{bi,t}$ 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

$$v_{b,t} = v_{b,t-1} - \sum_{i \in B, i \neq b} [w_{bi,t} + s_{bi,t}] + Y_{b,t} + \sum_{i \in B, i \neq b} s_{ib,t} \quad t = [0,...,\hat{T}], \quad b = [1,...,B] \quad (4.4)$$

Further, the problem faces multiple restrictions. All reservoirs are subject to a minimum and maximum level of water, denoted by $\underline{v_{b,t}}$ and $\overline{v_{b,t}}$. 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. Finally, 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, $\overline{s_{bi,t}}$. 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), $\overline{s_{bi,t}}$ will logically be 0. All these constraints can be summarized in (4.5)-(4.8).

$$v_{b,t} \leq \overline{v_{b,t}} \quad \text{for} \quad t = [1,...,\hat{T}], \quad b = [1,...,B] \quad (4.5)$$

$$v_{b,t} \geq \underline{v_{b,t}} \quad \text{for} \quad t = [1,...,\hat{T}], \quad b = [1,...,B] \quad (4.6)$$

$$q_{bi,t} \leq \overline{q_{bi}} \quad \text{for} \quad t = [1,...,\hat{T}], \quad b = [1,...,B], \quad i \in B, i \neq b \quad (4.7)$$

$$s_{bi,t} \leq \overline{s_{bi,t}} \quad \text{for} \quad t = [1,...,\hat{T}], \quad b = [1,...,B], \quad i \in B, i \neq b \quad (4.8)$$

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

$$\begin{aligned} &\max && V_t(P_t, \hat{Y}_t, \pi_t) \\ &\text{subject to} && (4.4), (4.5), (4.6), (4.7), (4.8) \end{aligned}$$

4.2.2 Electricity price process

We model spot price movements as a Markov process and use the price of forward contracts to forecast future spot prices. At time t, $F_{t,T}$ is the price of a forward contract traded in a market 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 process for the price development of forward contracts with different times to maturity can be used to represent 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 process will incorporate these expectations. An additional advantage of using all the prices of forward contracts traded in the market is that we include the seasonality of electricity prices, an important characteristic of electricity spot prices, as described by Johnson and Barz (1999). This implies

that the process does not need any deterministic function to account for seasonality.

As stated previously, we want to use a multi-factor process to describe movements in price. Including multiple factors should allow the process to better explain the real dynamics driving forward price movements, and thereby make better price predictions. Such a process can be formulated as a multi-factor extension of the HJM model, originally presented by Heath et al. (1992). An HJM model with I sources of uncertainty is given by

$$\frac{dF_{t,T}}{F_{t,T}} = \sum_{i=1}^{I} \sigma_{i,t,T} dZ_{i,t} \tag{4.9}$$

Here, $\sigma_{i,t,T}$ is the ith volatility function of a forward contract with maturity at time T, and $dZ_{i,t}$ is a source of uncertainty where $Z_{i,t}$ follows a Wiener process. Together, $\sigma_{i,t,T}$ for $i = [1,...,I]$ explain the dynamics driving the time t movement of a forward contract with maturity at time T. The ith volatility function is associated with the ith source of uncertainty, $dZ_{i,t}$. Since our decision problem considers discrete time stages, (4.9) must be discretized. By using Ito's lemma and setting $dt = \Delta t$, the process can be written as

$$F_{t,T} = F_{t-\Delta t,T} \cdot \exp\Big(- \frac{1}{2} \sum_{i=1}^{N} \sigma_{i,T-t+\Delta t}^2 \Delta t + \sum_{i=1}^{N} \sigma_{i,T-t+\Delta t} \sqrt{\Delta t} \varepsilon_{i,t} \Big) \tag{4.10}$$

Here, $\Delta Z_{i,t} = \sqrt{\Delta t}\varepsilon_{i,t}$ where $\varepsilon_{i,t} \sim N(0,1)$. Using $P_t = F_{t,t}$, we can modify (4.10) into an expression for the spot price P_t as a function of $F_{t-\Delta t,t}$, given by

$$P_t = F_{t,t} = F_{t-\Delta t,t} \cdot \exp\Big(- \frac{1}{2} \sum_{i=1}^{N} \sigma_{i,\Delta t}^2 \Delta t + \sum_{i=1}^{N} \sigma_{i,\Delta t} \sqrt{\Delta t} \varepsilon_{i,t} \Big) \tag{4.11}$$

4.2.3 *Estimating the volatility functions of the price process*

In the literature, multiple ways are proposed on how the volatility functions in (4.9) can be obtained. Koekebakker and Ollmar (2005) propose that they can be found empirically as a function of time to maturity, that is, on the form $\sigma_{i,t,T} = \Psi_i(T-t)$. In order to do so, we must construct a sufficiently large dataset of daily returns for multiple types of forward contracts $m = [1,...,M]$ with time to maturity $\tau_m = T - t$. Using τ as time to maturity, the volatility functions can be 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.

In the Nordic power market, tradable forward and future contracts 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. $F_{t_s,t}$ where $t > t_s$ denotes the value of a forward curve constructed at time t_s a for time t, and intends to represent the price of a fictional forward contract with delivery exactly at time t. Multiple ways of constructing forward curves are presented in the literature, e.g. by Fleten and Lemming (2003), Benth et al. (2007), Alexander (2008) and Kiesel et al. (2018). After constructing a set of curves for multiple consecutive days, we can use (4.12) to calculate daily returns at time t_j for contracts with time to maturity τ_a. This is 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).

$$x_{j,a} = \ln(F_{t_j,t_j+\tau_a}) - \ln(F_{t_{j-1},t_j+\tau_a}) \tag{4.12}$$

Here, $j = [2,...,J]$ and $a = [1,...,A]$, where J is the number of forward curves and A is the number of maturity dates for which we want to construct a dataset. 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

$$\mathbf{X_{J\times A}} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,A} \\ x_{2,1} & x_{2,2} & \dots & x_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ x_{J,1} & x_{J,2} & \dots & x_{J,A} \end{bmatrix} \tag{4.13}$$

Having found $\mathbf{X_{J\times A}}$, we use principal component analysis (PCA) to find the desired I volatility functions. PCA is an orthogonalization technique used to reduce the dimensionality of a dataset consisting of highly correlated variables. Mathematically, the principal components of $\mathbf{X_{J\times A}}$, whose correlation matrix is denoted $\mathbf{V}$, are given by $\mathbf{P} = \mathbf{X_{J\times A}W}$. 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. As shown in Clewlow and Strickland (2000), the volatility functions will then be given by $\sigma_{i,\tau_a} = \sqrt{\lambda_i} w_{ai}$, where $i = [1,...,A]$.

To reduce the dimensionality, we only include the volatility functions associated with the first I principal components. Typically, one would choose I such that the proportion of variance explained by the first I factors 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.

4.2.4 Inflow process

The inflow process is based on the geometric periodic autoregressive (GPAR) model presented by Shapiro et al. (2013). The authors found that a first-order periodic autoregressive process 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 Y_t is highly right-skewed. Therefore, they work with $\ln(Y_t)$ to obtain a distribution with less skew.

Let $\hat{\mu}_t$, $t = 1,...,52$ be the weekly averages of $\ln(Y_t)$ and $W_t = \ln(Y_t) - \hat{\mu}_t$ be the corresponding deviations. Shapiro et al. (2013) found that W_t could be described by an AR(1) process. (4.14) shows how the deviations of the log inflows from their mean can be described as an 1-lag autoregressive process.

$$W_t = \phi_0 + \phi_1 W_{t-1} + \varepsilon_{Y,t} \tag{4.14}$$

Here, ϕ_0 and ϕ_1 are parameters of the process, and $\varepsilon_{Y,t}$ 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 $\varepsilon_{Y,t} \sim N(0,\sigma_Y^2)$, where σ_Y is the standard deviation of the error terms. The parameters ϕ_0 and ϕ_1 are estimated by ordinary least squares regression. Because W_t observations are themselves deviations, ϕ_0 is highly insignificant. We set $\phi_0 = 0$ and use $\phi_1 = \phi$ from

this point on. Next, we find the log inflow, $W_t + \hat{\mu}_t$

$$\begin{aligned} W_t + \hat{\mu}_t &= \phi W_{t-1} + \varepsilon_{Y,t} + \hat{\mu}_t \\ &= \phi(W_{t-1} + \hat{\mu}_{t-1}) - \phi\hat{\mu}_{t-1} + \varepsilon_{Y,t} + \hat{\mu}_t \end{aligned} \quad (4.15)$$

The inflow Y_t can be expressed as a function of $W_t + \hat{\mu}_t$. We insert the obtained expression of $W_t + \hat{\mu}_t$ into $Y_t = \exp(W_t + \hat{\mu}_t)$, and get

$$\begin{aligned} Y_t = \exp(W_t + \hat{\mu}_t) = \\ \exp(\phi Y_{t-1}) \exp(-\phi\hat{\mu}_{t-1} + \varepsilon_{Y,t} + \hat{\mu}_t) \end{aligned} \quad (4.16)$$

By rewriting, we obtain the inflow process described by (4.18).

$$Y_t = \exp(\phi \ln Y_{t-1}) \exp(\hat{\mu}_t - \phi\hat{\mu}_{t-1} + \varepsilon_{Y,t}) \quad (4.17)$$

$$Y_t = \exp(\varepsilon_{Y,t}) \exp(\hat{\mu}_t - \phi\hat{\mu}_{t-1}) Y_{t-1}^{\phi} \quad (4.18)$$

We further allow the error term standard deviation σ_Y and the coefficient ϕ to be time-dependent. The final inflow process can then be expressed as

$$Y_t = \exp(\varepsilon_{Y,t} + \hat{\mu}_t - \phi_t\hat{\mu}_{t-1}) Y_{t-1}^{\phi_t} \quad (4.19)$$

where t is the week number and $\varepsilon_{Y,t}$ now follows the distribution $\varepsilon_{Y,t} \sim N(0, \sigma_{Y,t}^2)$. Since inflow Y_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 also follows a Markov process, which was one of the prerequisites for representing our decision problem as a Markov decision process.

4.2.5 Scenario lattice for spot price and inflow

To solve the MDP, we must discretize the exogenous Markov process that describes inflow and price movements. As in Löhndorf and Wozabal (2017), we do this by reducing the continuous Markov process to a discrete scenario lattice. In our case, each lattice node represents a state of both reservoir inflow and spot price. Generally, we keep the number of nodes per stage in the lattice constant. In comparison, the number of nodes per stage in a scenario tree grows exponentially with the number of time stages. Thus, the lattice approach allows for a higher number of time stages while still keeping the problem computationally feasible.

In order to construct a lattice, we use the method proposed by Löhndorf and Wozabal (2017). We denote N_t as the number of nodes at time t. Further, $\overline{S_{tn}} = \{P_{tn}, Y_{tn}\}$ denotes the nth state (or node) at time stage t, where P_{tn} is the state spot price and $Y_{tn} = \{I_{b,tn} : b = [1, ..., B]\}$ 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. In short, the lattice is constructed by first drawing a set of K Monte-Carlo simulations ($\hat{S}^k$) of spot price and inflow, where $\hat{S}_t^k = \{P_t^k, Y_t^k\}$ is the time t state of simulation k where $k \in [K]$. These simulations are drawn using (4.11) for the spot price and (4.19) for inflow, where the error terms $\varepsilon_{i,t}$ and $\varepsilon_{Y,t}$ are correlated with ρ_i. Afterwards, the location of all states $\overline{S_{tn}}$ is found by minimizing the Wasserstein distance between all N_t nodes and the K simulated draws $\hat{S}_t^k$ for each time stage t. Having located all nodes, the transition probabilities between two subsequent nodes $\overline{S_{tn}}$ and $\overline{S_{t+1,m}}$ can be found by looking at the number of simulated paths whose time t and $t+1$ states lie closest to the nodes $\overline{S_{tn}}$ and $\overline{S_{t+1,m}}$.

4.2.6 Solution method for optimization problem

We adopt the method known as approximate dual dynamic programming (ADDP) to find the near optimal decision policies π_{tn} in all nodes of the price and inflow lattice. ADDP was first introduced by Löhndorf et al. (2013). In principle, obtaining optimal decision policies for a Markov decision process should be possible

using traditional dynamic programming (DP) as introduced by Bellman (1957). Using the notation introduced in Section 4.2.5, the Bellman equation given in (4.2) can then be rewritten into

$$V_t(\overline{S_{tn}}, \pi_{tn}) = \max_{\pi_{tn}} \quad CF_t(\overline{S_{tn}}, \pi_{tn}) + \beta \mathbb{E}[V_{t+1}(S_{t+1}, \pi_{t+1} | \overline{S_{tn}}, \pi_{tn})] \quad (4.20)$$

A common problem with dynamic programming is the curse of dimensionality. It has been addressed by multiple authors, i.e. by Powell (2011). In this case, the main issue with a high-dimensional problem is that the decision space can become too large to find the optimal decisions 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 et al. (2013) 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). In short, the value function is found by first drawing a given number of *forward passes* through the lattice, that is, a sequences of states. For each forward pass, the optimal decision policies are found by maximizing the approximate post-decision value functions. After each forward pass, a *backward pass* is performed, where the approximated value functions are updated relative to the sampled sequence of states and all state decision policies. In practice, the approximate value function of each state is constructed by a set of supporting hyperplanes (linear constraints), where each pair of forward and backward passes results in the addition of a new hyperplane to the set. For a more detailed description of the ADDP algorithm, consult Löhndorf et al. (2013).

4.3 Results

In this section, we summarize the results of a case study conducted with data from a Norwegian hydropower producer. The Søa hydropower plant is presented in Section 4.3.1. In Section 4.3.2, we discuss the decision problem. Section 4.3.3 discusses the characteristics of the price process and its associated volatility functions. Then, we present the parameters of the inflow process in Section 4.3.4 and the price and inflow correlation in Section 4.3.5. Section 4.3.6 shows how we construct the scenario lattice using correlated Monte Carlo simulations. We present the expected revenues by applying the scheduling model in Section 4.3.7. In Section 4.3.8, we perform a backtest of our model compared to historical operations. Further, in Section 4.3.9 we analyze potential differences and losses in expected revenues with regards to different values of the price-inflow correlation. We perform a similar analysis in Section 4.3.10, considering the case were the number of factors I used in the forward price process is altered. All code and calculations are produced in MATLAB and R, except for the algorithms used to construct a lattice and the ADDP solver.

4.3.1 Case: Søa hydropower plant

We have received empirical data from the Søa 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 Norway. It consists of two reservoirs - Vasslivatn and Søvatn, and one Francis turbine. The discharge from the Søvatn reservoir to the Vasslivatn reservoir is controllable. In Table 4.2, we have listed the physical boundaries of both reservoirs. There is also a special summertime restriction that applies to Søvatn, which is set by local authorities. This restriction and its duration are also listed in Table 4.2. The outlet of the hydropower plant is in Hemnefjorden, which has an average head of -1 MASL (meters above sea level).

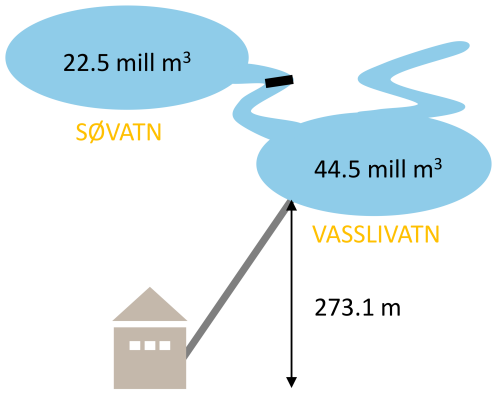

Figure 4.1: The Søa hydropower plant and the reservoir capacities. The elevation of 273.1m is the production-weighted average head difference between Vasslivatn and Hemnefjorden.

	Value	Unit
Maximum power capacity	36	MW
Mean yearly production	191.3	GWh
Avg. yearly inflow, total	311	mill m^3
Average inflow to Søvatn	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

Table 4.1: Characteristics of the Søa hydropower plant

The energy coefficient listed in Table 4.1 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.

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

Table 4.2: Water level constraints for Søa. All water levels are denoted in meters above sea level

4.3.2 Revised decision problem

In order to construct the lattice mentioned earlier and perform ADDP on Søa hydropower plant, we use QUASAR, a general-purpose solver for stochastic optimization (Löhndorf, 2017). To keep computation complexity at bay, we use a linear reformulation of the problem.

The number of reservoirs is $B = 2$, and we let $v_{1,t}$ and $v_{2,t}$ denote the water levels in Vasslivatn and Søvatn respectively. Also, $Y_{1,t}$ and $Y_{2,t}$ 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 Søvatn to Vasslivatn (denoted $s_{12,t}$ using notation from Section 4.2.1) is now denoted by $s_{c,t}$, and the amount of spilled water flowing from Vasslivatn to Hemnefjorden is denoted by $s_{s,t}$ (previously denoted $s_{1O,t}$). With these new notations, the general volume balance constraint from (4.4) can be written as (4.21) and (4.22) for Vasslivatn and Søvatn,

respectively.

$$v_{1,t} = v_{1,t-1} - w_t - s_{s,t} + Y_{1,t} + s_{c,t} \quad \text{for} \quad t = [1,...,\hat{T}] \quad (4.21)$$

$$v_{2,t} = v_{2,t-1} + Y_{2,t} - s_{c,t} \quad \text{for} \quad t = [1,...,\hat{T}] \quad (4.22)$$

Neither $s_{c,t}$ nor $s_{s,t}$ are restricted by an upper bound $\overline{s_{ij,t}}$, so this constraint is omitted from the revised problem formulation. Due to the linearity requirement, the cash flow expression CF_t defined by (4.3) can only be a function of one decision variable. While head elevation, turbine and generator efficiency are typically functions of one or more decision variables, we must model them as constants in order to keep the expression linear. Such simplifications are made in similar models for reservoir management, e.g., EOPS (SINTEF, 2017b). Madani and Lund (2009) also use a fixed head 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 Søa 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.

When we use constant values for head elevation and efficiency rate, the objective function of the optimization problem 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 $\rho GH\eta\varpi/\varsigma = \kappa$, where κ denotes the energy coefficient.

Further, we only have available data on the aggregated inflow into both reservoirs, forcing us to treat inflow as a single stochastic variable $Y_t = Y_{1,t} + Y_{2,t}$. In order to obtain $Y_{1,t}$ and $Y_{2,t}$, we have used the historical inflow split given in Table 4.1. We let $\alpha = 0.395$ denote the historical fraction of inflow flowing into Vasslivatn, and thereby set $Y_{1,t} = \alpha Y_t$ and $Y_{2,t} = (1-\alpha)Y_t$. Also, since the water level in Søvatn is subject to a minimum restriction during the summer $\underline{v_{2,t}} > 0$, we must include a slack variable $v^S_{2,t}$ to account for cases in which 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 slack variable to the value function such that its value is kept at a minimum. By combining all the mentioned simplifications and adjustments, our decision problem at time t is reduced to

$$\begin{aligned}
\max \quad & V_t = P_t \cdot \kappa \cdot w_t - \upsilon \cdot v^S_{2,t} \\
& + \beta \mathbb{E}[V_{t+1} | P_t, Y_t, \pi_t] \\
\text{subject to} \quad & v_{1,t} = v_{1,t-1} - w_t - s_{s,t} + \alpha Y_t + s_{c,t} \\
& v_{2,t} = v_{2,t-1} + (1-\alpha) Y_t - s_{c,t} \\
& v_{1,t} \leq \overline{v_{1,t}} \\
& v_{2,t} \leq \overline{v_{2,t}} \\
& v_{1,t} \geq \underline{v_{1,t}} \\
& v_{2,t} + v^S_{2,t} \geq \underline{v_{2,t}} \\
& q_t \leq \overline{q}
\end{aligned}$$

where $\pi_t = \{w_t, v_{1,t}, v_{2,t}, s_{c,t}, s_{s,t}, v^S_{2,t}\}$. All coefficients and constant parameter values are given in Table 4.3. 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. Furthermore, we find the discount factor β using the risk-free rate r given by 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). Optimally, we would have used an estimate of the two-year maturity risk-free rate, but six months was the longest maturity available. The discount factor β is found using continuous compounding, given by $\beta = \exp(-r\Delta t)$.

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

Table 4.3: Model coefficients and constants

4.3.3 *Electricity spot price and forward curve dynamics*

The first step towards obtaining the volatility functions describing the forward curve dynamics is to construct historical forward curves. They are found by interpolating between forward prices as described by Alexander (2008). The dataset of this study includes forward prices for all trading days between April 28, 2011, to December 30, 2016, resulting in 1450 forward curves. The forward curves are constructed using closing prices of futures contracts traded at NASDAQ Commodities. These contracts are listed in Table 4.4.

Code	Length of delivery period	Trading period
ENOW	Week	1 - 6 weeks ahead
ENOM	Month	1 - 6 months ahead
ENOQ	Quarter	1 - 8 quarters ahead
ENOYR	Year	1 - 3 years ahead

Table 4.4: Electricity forward contracts traded on NASDAQ Commodities

Figure 4.2 shows the forward curve found for January 7, 2013. It is clear that the forward curves found using this method will be discontinuous in the points where we switch from one contract type to another, as illustrated in Figure 4.2. As can be seen in Figure 4.2, weekly contracts were 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 long end of the curve.

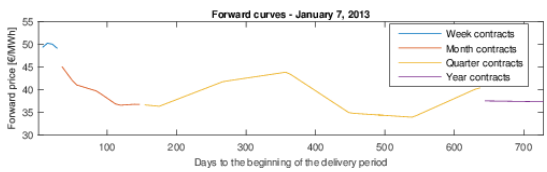

Figure 4.2: Forward curve created using linear interpolation

A time series of daily log returns is calculated for each relevant time to delivery, resulting in a 1449×104 matrix. To obtain a set of volatility functions that describe forward price movements, PCA is used on the returns time series as explained in Section 4.2.2. 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. 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 find the volatility functions $\sigma_{i,\tau}$ for all $i = [1,...,N]$ and time to maturity given by $\tau = [1,2,...,104]$ weeks.

Using the time series of returns, we can estimate an overall volatility curve for the term structure of forward prices, as well as the volatility functions associated with the principal components. In Figure 4.3, the overall volatility function can be understood as the volatility of returns of forward contracts with time to maturity τ. Since it represents the actual volatility of forward price returns, it will always be positive. The volatility functions associated with the principal components must, however, not be interpreted the same way as the overall volatility function, as they do not represent the volatility in terms of price movements of a single asset. They can also take negative values, as opposed to the overall volatility function.

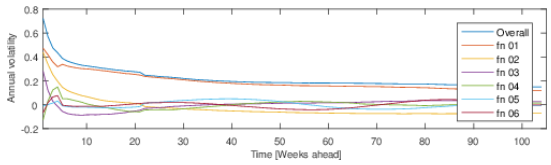

Figure 4.3: Volatility functions found by using method of linear interpolation. 'Overall' denotes the overall volatility curve, and fn i denote the volatility functions given by principal component i

As can be seen from the dark blue curve in Figure 4.3, the overall volatility is *monotonically decreasing*. 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.

The electricity forward return series show a substantial degree of inter-correlation. This can be seen from the correlation matrix that is shown in Figure 4.4. A high degree if 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.

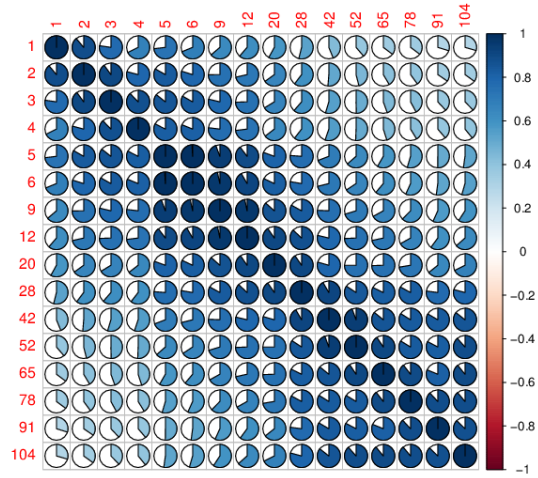

Figure 4.4: Correlation matrix associated with returns of forward contracts. The column and row names are both the number of weeks until the beginning of the delivery period

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. Only six principal components are needed to explain over 95% of the cumulative variance, as shown in Table 4.5.

Number of factors	1	2	3	4	5	6
Explained variance	0.73	0.88	0.92	0.93	0.95	0.96

Table 4.5: Proportion of explained variance for different numbers of explanatory factors

4.3.4 Inflow process parameters

We fit the geometric periodic autoregressive (GPAR) model suggested by Shapiro et al. (2013) to the inflow data for the Søa hydropower plant. The dataset consists of daily inflow observations for each day between January 1, 1958, and December 31, 2016. According to TrønderEnergi, 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.

Recall that the inflow is given by

$$Y_t = \exp\left(\varepsilon_{Y,t} + \hat{\mu}_t - \phi_t \hat{\mu}_{t-1}\right) Y_{t-1}^{\phi_t} \tag{4.19}$$

Here,

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

For a right-skewed distribution such as the one that can be seen in Figure 4.5, a geometric process is better suited than an arithmetic process. It better captures the inflow dynamics, which can be extreme. Further, a geometric process does not allow for negative inflows. Shapiro et al. (2013) found the inflow distribution for Brazilian hydropower plants to be right-skewed as well, favoring a log transformation of the inflow observations.

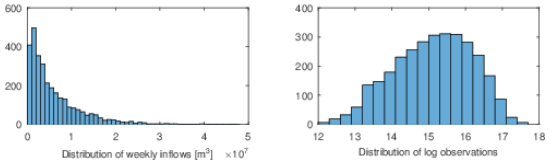

Figure 4.5: Distribution of inflow and log inflow observations

W_t, the deviation of the log inflows from their mean, is represented as an AR(1) process. The suitability of a 1-lag process can be determined by investigating the partial autocorrelation of the historical data for W_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 4.6 shows the partial autocorrelation of the W_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.

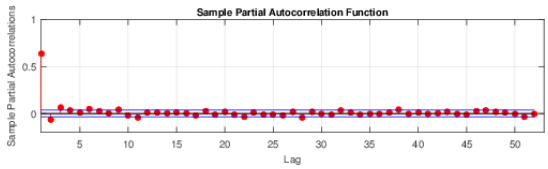

Figure 4.6: Partial autocorrelation of the W_t time series

The inflow process 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 ($\sigma_{Y,t}$). Figure 4.7 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.

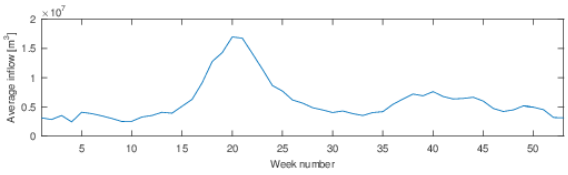

Figure 4.7: Average inflow for a certain week of the year

4.3.5 Price and inflow correlation

As previously mentioned, we consider the correlation between the electricity spot price and inflow to the hydropower plant. In Section 4.2.5, this was introduced as the correlation ρ_i between the error terms $\varepsilon_{i,t}$ and $\varepsilon_{Y,t}$ from Equation (4.11) and (4.19), where $i = [1,...,I]$ indicates a principal component. Since the first principal component explains 73.38% of the total variation, we choose only to calculate ρ_i for $i = 1$, and set $\rho_i = 0$ for $i \neq 1$. Therefore, ρ_1 is hereby denoted ρ. Mathematically, the correlation was calculated by estimating the historical correlation coefficient between the normalized error term of the inflow process, $\varepsilon_{Y,t}$, and the normalized first principal component ($\mathbf{p}_1$).

The error term in the inflow process 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 the first principal component, historically.

4.3.6 Monte Carlo simulations and lattice construction

In order to create a lattice, we had to run multiple parallel Monte Carlo simulations of spot price and inflow paths. The starting values of the price simulations included one current time spot price and the price of $\hat{T} - 1$ forward contracts. Since we use weekly granularity and a horizon of $\hat{T} = 105$ weeks, this requires 104 weekly forward contracts with time to delivery $\tau = [1,...,104]$. 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. Unlike the forward curves used to construct the volatility functions, which were discontinuous (see Figure 4.2), forward curves constructed using the method of Fleten and Lemming (2003) are both continuous and smooth. The method did, however, not provide us with plausible volatility functions, as we experienced issues with unrealistic oscillations in the near end for some of the forward curves.

Mathematically, the forward price of a contract

with delivery in a given week $W = [2,...,105]$ is calculated using the average value of the forward curve within the time interval of that particular week. For the weeks $W = [2,...,7]$, the weekly average value of the forward curve will be the price of the six weekly forward contracts sold in the market. 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 Søa 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 and the scope of this paper.

To construct the price and inflow scenario lattice, we have used 380.000 Monte Carlo simulations. The lattice consists of 100 nodes for all time stages except the starting one, giving a total count of 10401 nodes. Furthermore, each node has two entries, inflow and spot price. Figure 4.8(a) displays the spot price lattice with starting date January 7, 2013, while Figure 4.8(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.

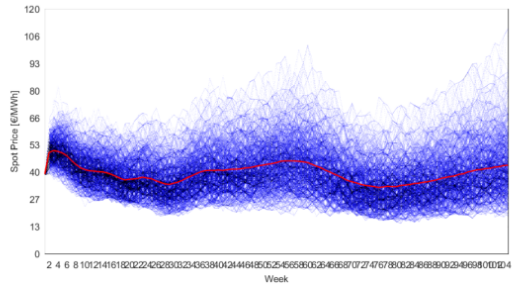

(a) 105-weeks spot price lattice

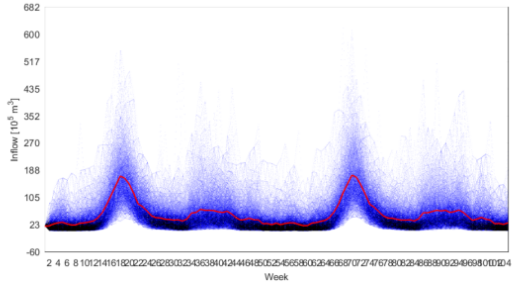

(b) 105-weeks inflow lattice

Figure 4.8: Spot price and inflow lattices constructed with data for January 7, 2013. The Y-axis for all plots denotes the time stages (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. In the case of the spot price lattice, the mean line is the initial forward curve.

4.3.7 Expected discounted revenues from hydropower plant

Having constructed the lattice, it is now possible to run the scheduling model. We run it for five different starting dates with different underlying forward curves and historical starting values for the reservoir levels. 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-neutral 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 production 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 4.6 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 Table 4.6, we also include one of the most important immediate results for the production planner; the value of w_1. We recall that w_t is the amount of water nominated for production at time stage 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 upcoming 105 weeks. We also include the average water dispatch $q_1 = w_1/\varsigma$, where $\varsigma = 604800s$ is the number of seconds per week.

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.

Parameter	**Unit**	**Jan 7 2013**	**Apr 8 2013**	**Jul 8 2013**
EDR	[M EUR]	15.93	15.95	14.36
EDR/Prod.	[EUR/MWh]	40.12	41.24	35.35
w_1	[M m^3]	0	5.21	10.28
q_1	$[m^3/s]$	0	8.62	17.00
Parameter	**Unit**	**Oct 7 2013**	**Jan 6 2014**	
EDR	[M EUR]	15.26	13.83	
EDR/Prod.	[EUR/MWh]	38.36	33.21	
w_1	[M m^3]	0	0	
q_1	$[m^3/s]$	0	0	

Table 4.6: 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.

An important question that arises is how one should handle the end level of the reservoir. We have not imposed any end level restrictions. Thus, there is no incentive to keep water in the reservoir at the end of the horizon. In some of the above simulations, e.g., the ones starting and ending in January, emptying the reservoir would probably be a poor decision in the reality since one would normally expect high prices in the upcoming periods. Emptying the reservoirs in the last time stage will result in expected discounted revenues that are slightly larger than what one would achieve in reality. Nevertheless, the end-of-horizon effects 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 two years. Note that for cases in which the time horizon ends during the spring when inflows are typically at their maximum, it is reasonable to allow emptying the reservoir as much as possible.

4.3.8 *Backtesting the production policy with realized price and inflow data*

A crucial analysis for assessing the performance of the scheduling model 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 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 Søa power plant generated discounted revenues of *10.89 million EUR* between January 7, 2013, and January 11, 2015, from 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 4.9 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 stage, 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 model performs well compared to reality also without emptying the reservoir in the last time stage. Our model did, however, utilize more of its water for production 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*.

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 relatively large amounts of 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 relatively 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 Søvatnet starts to apply. While the real-life production planner fills up Søvatnet 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.

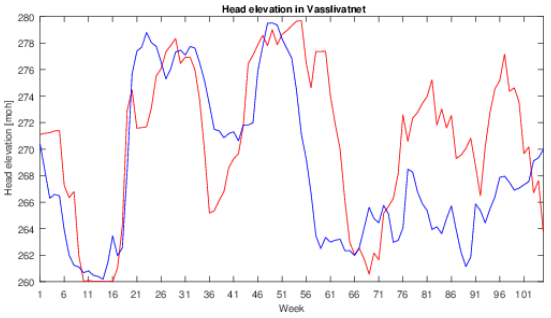

(a) Vasslivatnet

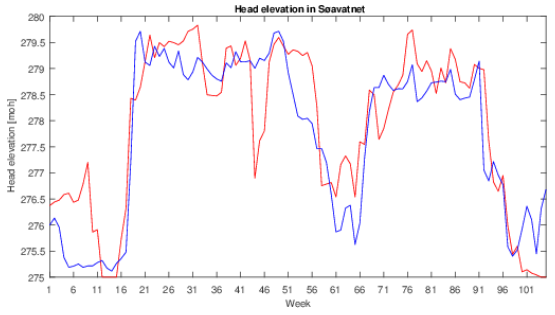

(b) Søvatnet

Figure 4.9: Modelled (red curve) and realized (blue curve) reservoir curves for Vasslivatnet and Søvatnet between January 7, 2013 and January 11, 2015.

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 that our model does not. Despite the circumstances discussed above, our model still manages to achieve similar results.

4.3.9 Loss calculations: Misspecified correlation coefficient

As previously stated, we incorporate a correlation between movements in price and local inflow. It is interesting to test the effect of introducing this feature, as it can tell us how models that assume no correlation perform compared to ours. Therefore, we have first calculated the expected revenues obtained when using different values of the correlation coefficient ρ. More importantly, we have also tested how decision policies obtained using $\rho = 0$ perform when inflow and price movements are in fact correlated, and what losses in expected revenues a plant can experience when this assumption is falsely made.

In order to test the effect of introducing the 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

discounted revenues of all three runs are displayed in Table 4.7 and mean optimal reservoir curves for Vasslivatn are displayed in Figure 4.10.

Correlation coefficient	[-]	0	-0.1765	-0.35
EDR	[M EUR]	16.10	15.93	15.83

Table 4.7: Expected discounted revenues (EDR) obtained using three different correlation coefficients ρ.

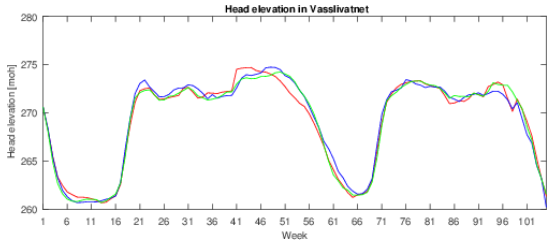

Figure 4.10: Average reservoir curves for Vasslivatnet found using three different values of the correlation coefficient ρ. The red curve denotes $\rho = 0$, the blue curve $\rho = -0.1765$ and the green curve $\rho = -0.353$.

By looking at the expected discounted revenues, we see that the higher we choose the correlation coefficient ρ, the larger are the expected discounted revenues. The correlation coefficient undoubtedly affects the results, implying that it must be estimated correctly.

Next, we 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 4.8, 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 Søa 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.

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

Table 4.8: 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 4.7.

4.3.10 Loss calculations: Number of factors in the price process

Further, up until now, we have used a price process 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 set of calculations similar to those in Section 4.3.9, but instead of testing different values of ρ, we use a different number of factors I in the price process. We investigate four different numbers of factors: 1 (that is, we use the overall volatility function [2]), 3 (91.53% explanation), 6 (96.04% explanation) and 10 (98.44% explanation). The obtained expected discounted revenues are shown in Table 4.9, and the mean optimal reservoir levels in Figure 4.11.

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

Table 4.9: Expected discounted revenues (EDR) obtained using different number of factors I to describe the underlying price process.

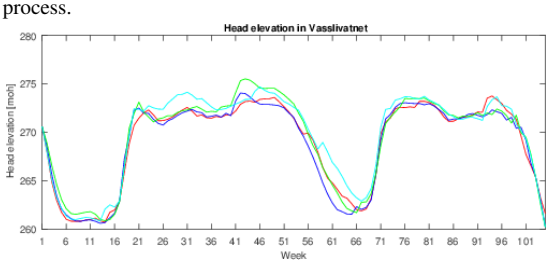

Figure 4.11: Average reservoir curves for Vasslivatnet obtained using different number of factors in the forward price process. The red curve is from a run with $I = 1$ factors, the dark blue one for a run with $I = 3$, the red curve for $I = 6$ and the light blue curve for $I = 10$.

The results in Table 4.9 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 stage. The optimal policies utilize the higher prices in the lattices with more factors, and the model thereby forecasts larger expected discounted revenues.

[2] This process disregards the forward curve intercorrelation.

As for the case with different values of ρ, it might be more interesting to test how the policies obtained using a simple one-factor price model perform when the price process can, in reality, be described using $I = [3,6,10]$ factors. We, therefore, redo the steps explained above for the case of different numbers of factors I instead of correlation coefficient ρ. The expected discounted revenues are displayed in Table 4.10. 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.

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

Table 4.10: Expected discounted revenues (EDR) calculated when using a policy where the number of factors is $I = 1$, but where the real stochastic price process is described by $I = [3,6,10]$. 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 I as in the stochastic process, as shown in Table 4.9.

4.4 Conclusions

In this paper, we present a medium-term model for reservoir management. We model the problem as a Markov decision process, and use a multi-factor model to describe changes in forward and spot prices. We incorporate a short-term correlation between the local inflow model and price model and solve the resulting stochastic-dynamic decision

problem using ADDP, which discretizes the Markov processes to a scenario lattice and then solves the problem by learning an outer approximation of the value function.

We find 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. Our analysis indicates that ignoring correlation can result in sub-optimal reservoir control decisions. In our case, we observe a decrease in expected revenues of 2.5% (that is, multiple 100.000 EUR yearly) if the correlation coefficient is, in fact, $\rho = -0.1765$.

Our analyzes also show that it is important to use multiple factors when describing price movements. The result of our case study is that solving the problem with a one-factor model when the true model has multiple factors decrease profits by about 2%. The number of factors will depend on the price data. We confirm the finding of Koekebakker and Ollmar (2005) that we need more factors than is typical of commodity price models. In our case, we use six factors which explain 96% variance.

Compared with historical production decisions, our model produced similar results, even though decisions were made in weekly granularity, whereas the planner makes planning decisions on a daily basis. Hence, using our model provides reliable production decisions if used for medium-term planning. 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.

Further work

Although both papers provide us with interesting results regarding the production and hedging problems, multiple aspects can be modified, improved or investigated. In this section, we propose some possible work that can be conducted by researchers or practitioners. As the hedging model was the primary focus of the master's thesis, it is emphasized more than the production planning model.

Test other methods for modeling risk preferences

It would be interesting to test the difference in hedging performance if we had modeled risk preferences using a terminal risk measure instead of the nested CVaR. Naturally, if we had maximized the terminal CVaR instead of the nested CVaR, the hedging performance measured in terms of the terminal CVaR should increase. Although the terminal risk measures will result in time inconsistent decision policies, using them can make sense from a budgetary point of view. Also, using the nested CVaR resulted in very large hedge ratios, which might not have been the case if we had used a terminal risk measure in the objective function of the hedging problem.

Include area price difference risk and mitigating derivatives

While EPAD contracts on price difference are not available for NO3, we did, in fact, formulate a stochastic process representing dynamics of the area price difference[1]. Due to computational limits and our already large state space, the area price difference risk was omitted from the final model. However, it would be interesting to include it as a risk factor in a hedging model to quantify the extent of the uncertainty. This would be especially interesting from the perspectives of companies located in bidding areas with lower correlations between the system and area spot price, e.g. in Denmark and Finland.

Treat price and currency as correlated state variables

In the hedging model, we treated the currency exchange rate as independent of changes in price, even though we found that their increments exhibit a weak, negative correlation. This was due to issues with the algorithm used to construct the multivariate lattices. By treating currency, prices, and inflow as correlated, the model could have included the natural hedge in the currency market, which could have resulted in different policies for optimal currency trading.

Combine hedge ratio heuristics with conditional risk mapping

When implementing the upper and lower hedge ratio requirements of the company analyzed by Sanda et al. (2013) as constraints in the hedging model, we tried to combine this version of the model with the nested CVaR. However, this resulted in issues with infeasibility. Nevertheless, we propose that such a combination of hedge ratio ranges and risk preference modeling should be further tested, as the policies obtained from such a model would be based on a time-consistent risk measure while simultaneously keeping the hedge ratios on the desired level. As illustrated, the performance of the hedge ratio approach was quite good compared to the

[1] See Appendix A.7.4 for the area price difference process.

nested CVaR approach in terms of reducing the risk of the terminal cash flows.

Combine hedging model with other scenario lattices

It should be trivial to use externally created scenario lattices instead of the ones used in this master's thesis in the sequential hedging model. Therefore, we propose testing the model with other scenario lattices for production, inflow, prices, and currency. An example could be to use price scenarios from EMPS (SINTEF, 2017a) and production and inflow scenarios from EOPS (SINTEF, 2017b). Another possibility is to test the hedging model with production scenarios from another hydropower plant, or even a wind turbine or a solar power installations.

Obtain more sophisticated forward curves and volatility functions

In both papers, we use forward curves constructed by the method of Alexander (2008) to obtain the volatility function explaining the evolution of the forward curve. Simultaneously, the initial forward curve used in the first stage of the production model was constructed by the method of Fleten and Lemming (2003). To avoid this combination, we recommend that one should try to use more sophisticated software for forward curve construction, e.g. Elvis Front Manager. Using a better forward curve construction method may result in both more correct process coefficients and better starting values.

Further investigate the parameter correlations

In both papers, we investigate different sorts of correlations between prices, inflow, and currency. For the production planner analyzed in the second paper, we see that the effect of omitting the correlation between increments of inflow and the first principle component driving movements in the forward curve could result in large cash flow reductions for the producer. We do, however, have a limited dataset, resulting in a 95% confidence interval of $[-0.28\ -0.06]$ for the correlation coefficient. Therefore, it is important to calculate the correlation more precisely, as a faulty specification might result in reduced cash flows.

Further, the first paper investigates the correlation between increments of inflow and futures contracts with different maturity dates. As the correlation matrix in Figure 3.7 shows, the correlation magnitude varied in a seemingly unstructured way for different times to maturity. The dynamics of price and inflow have been investigated in many papers, many of which propose more complex relationships than a constant Pearson's correlation coefficient. An alternative in our case could be to use different correlation matrices for different times of the year, e.g. for different seasons. For instance, inflow deviations in the spring might have a different impact on the forward curve than in other parts of the year. By using seasonal correlation matrices, these dynamics could have been captured better.

Extend the time horizon of the hedging problem

In Sanda et al. (2013), many of the analyzed firms start their trading activity multiple years prior to maturity. This suggests that the horizon of the hedging problem should maybe be extended. While the liquidity at NASDAQ OMX of yearly futures contracts with three or more years to delivery is low (Fleten et al., 2010), this is not the case for contracts with two years to maturity. Thus, the model horizon could be extended by

one additional year, incorporating the latter contract. This would, however, require a suitable process for describing the evolution of the forward curve in the long end. Using our approach, it is only possible to construct returns series explaining forward curve movements for the first two years of the curve. Thus, a new method for constructing forward curves with corresponding returns series is necessary to extend the problem horizon.

References

[1] Abgottspon, H. and Andersson, G. Medium-term optimization of pumped hydro storage with stochastic intrastage subproblems. *Power Systems Computation Conference (PSCC)*. IEEE, 2014, 1–7.

[2] Adam, T. R. and Fernando, C. S. Hedging, speculation, and shareholder value. *Journal of Financial Economics* 81(3) (2006), 283–309.

[3] Alexander, C. *Pricing, Hedging and Trading Financial Instruments*. Vol. III. Market Risk Analysis. Chichester, West Sussex, UK: John Wiley & Sons, Ltd, 2008.

[4] Aliber, R. Z. The interest rate parity theorem: A reinterpretation. *Journal of Political Economy* 81(6) (1973), 1451–1459.

[5] Anderson, R. W. and Danthine, J.-P. Hedging and joint production: Theory and illustrations. *The Journal of Finance* 35(2) (1980), 487–498.

[6] Bally, V. and Pages, G. A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli* 9(6) (2003), 1003–1049.

[7] Basel Committee on Banking Supervision. *Core Principles for Effective Banking Supervision*. 2012. URL: `https://www.bis.org/publ/bcbs230.htm`.

[8] Bekaert, G., Wei, M., and Xing, Y. Uncovered interest rate parity and the term structure. *Journal of International Money and Finance* 26(6) (2007), 1038–1069.

[9] Bellman, R. *Dynamic Programming*. Princeton Landmarks in Mathematics and Physics. Princeton, NJ, USA: Princeton University Press, 1957.

[10] Benth, F. E., Benth, J. S., and Koekebakker, S. *Stochastic modelling of electricity and related markets*. Singapore: World Scientific Publishing, 2008.

[11] Benth, F. E., Koekebakker, S., and Ollmar, F. Extracting and applying smooth forward curves from average-based commodity contracts with seasonal variation. *The Journal of Derivatives* 15(1) (2007), 52–66.

[12] Bjerksund, P., Stensland, G., and Vagstad, F. Gas Storage Valuation: Price Modelling v. Optimization Methods. 2008. URL: `https://brage.bibsys.no/xmlui/bitstream/handle/11250/163949/2008.pdf?sequence=1`.

[13] Boda, K. and Filar, J. A. Time Consistent Dynamic Risk Measures. *Mathematical Methods of Operations Research* 63(1) (2006), 169–186.

[14] Boger, M., Fleten, S.-E., Pichler, A., Keppo, J., and Vestbøstad, E. M. Backing out Expectations from Hydropower Release Time Series. *IAEE International Conference*. Singapore, 2017.

[15] Clewlow, L. and Strickland, C. *Energy Derivatives: Pricing and Risk Management*. London, UK: Lacima Publications, 2000.

[16] Cumby, R. E. and Obstfeld, M. *International interest-rate and price-level linkages under flexible exchange rates: A review of recent evidence*. 1982.

[17] Dimoski, J., Nersten, S., Fleten, S.-E., and Löhndorf, N. Hydropower reservoir management using multi-factor price model and correlation between price and local inflow. *IAEE International Conference*. Groningen, Netherlands, 2018.

[18] Dupuis, D., Gauthier, G., Godin, F., et al. Short-term Hedging for an Electricity Retailer. *The Energy Journal* 37(2) (2016).

[19] Flatabø, N., Haugstad, A., Mo, B., and Fosso, O. B. Short-term and Medium-term Generation Scheduling in the Norwegian Hydro System under a Competitive Power Market Structure. *EPSOM '98 Proceedings*. Switzerland: ETH Zürich, 1998.

[20] Fleten, S.-E., Bråthen, E., and Nissen-Meyer, S.-E. Evaluation of static hedging strategies for hydropower producers in the Nordic market. *The Journal of Energy Markets* 3(4) (2010), 1–28.

[21] Fleten, S.-E., Keppo, J., and Näsäkkälä, E. Risk Management in Electric Utilities. *The Handbook of Integrated Risk Management in Global Supply Chains*. Ed. by Kouvelis, P., Dong, L., Boyabatli, O., and Li, R. 2nd ed. Wiley, 2012.

[22] Fleten, S.-E. and Lemming, J. Constructing forward price curves in electricity markets. *Energy Economics* 25(5) (2003), 409–424.

[23] Fleten, S.-E., Wallace, S. W., and Ziemba, W. T. Hedging electricity portfolios via stochastic programming. *Decision Making Under Uncertainty. The IMA Volumes in Mathematics and its Applications*. Ed. by Greengard, C. and Ruszczynski, A. Vol. 128. New York, New York: Springer, 2002, 71–93.

[24] Fosso, O. B., Gjelsvik, A., Haugstad, A., Mo, B., and Wangensteen, I. Generation scheduling in a deregulated system. The Norwegian case. *IEEE Transactions on Power System* 14(1) (1999), 75–81.

[25] Foster, B. T., Kern, J. D., and Characklis, G. W. Mitigating hydrologic financial risk in hydropower generation using index-based financial instruments. *Water Resources and Economics* 10 (2015), 45–67.

[26] Godin, F. Minimizing CVaR in global dynamic hedging with transaction costs. *Quantitative Finance* 16(3) (2016), 461–475.

[27] Harrison, J. M. and Pliska, S. R. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and their Applications* 11(3) (1981), 215–260.

[28] Heath, D., Jarrow, R., and Morton, A. Bond Pricing and the Term Structure of Interest Rates A New Methodology for Contingent Claims Valuation. *Econometrica* 60(1) (1992), 77–105.

[29] Houmøller, A. P. *Investigation of forward markets for hedging in the Danish electricity market*. Tech. rep. Middelfart, Denmark: Houmoller Consulting, 2017.

[30] Huchzermeier, A. and Cohen, M. A. Valuing operational flexibility under exchange rate risk. *Operations Research* 44(1) (1996), 100–113.

[31] Iladis, N. A., Tilmant, A., Chabar, R., Granville, S., and Pereira, M. 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`.

[32] Iliadis, N., Pereira, V., Granville, S., Finger, M., Haldi, P.-A., and Barroso, L.-A. Benchmarking of hydroelectric stochastic risk management models using financial indicators. *Power Engineering Society General Meeting*. Montreal, Quebec, Canada: IEEE, 2006.

[33] Jaeck, E. and Lautier, D. Volatility in electricity derivative markets: The Samuelson effect revisited. *Energy Economics* 59 (Aug. 2016), 300–313.

[34] Johnson, B. and Barz, G. Selecting stochastic processes for modelling electricity prices. *Energy Modelling and the Management of Uncertainty*. Ed. by Kaminski, V. Risk Books, 1999. Chap. 1.

[35] Kiesel, R., Paraschiv, F., and Sætherø, A. On the construction of hourly price forward curves for electricity prices. *Computational Management Science* (2018), 1–25.

[36] Koekebakker, S. and Ollmar, F. Forward curve dynamics in the Nordic electricity market. *Managerial Finance* 31(6) (2005), 73–94.

[37] Kolsrud, C. W. and Prokosch, M. Reservoir Hydropower: The value of flexibility. Master thesis. Norwegian University of Science and Technology, June 2010.

[38] Lamond, B. F. and Boukhtouta, A. Optimizing long-term hydro-power production using Markov decision processes. *International Transactions in Operational Research* 3(3-4) (1996), 223–241.

[39] Levi, M. D. *International finance 4th edition*. Abingdon, Oxfordshire, UK: Routledge, 2005.

[40] Liu, S. D., Jian, J. B., and Wang, Y. Y. Optimal Dynamic Hedging of Electricity Futures Based on Copula-GARCH Models. *IEEE International Conference*. Macau, China: IEEE, 2010, 2498–2502.

[41] Löhndorf, N. *QUASAR Optimization Software 2.3*. 2017. URL: `http://quantego.com`.

[42] Löhndorf, N. and Wozabal, D. Indifference pricing of natural gas storage contracts. 2017. URL: `http://www.optimization-online.org/DB_HTML/2017/02/5863.html`.

[43] Löhndorf, N., Wozabal, D., and Minner, S. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. *Operations Research* 61(4) (2013), 810–823.

[44] Maceira, M. and Damázio, J. Use of the PAR (p) model in the stochastic dual dynamic programming optimization scheme used in the operation planning of the Brazilian hydropower system. *Probability in the Engineering and Informational Sciences* 20(1) (2006), 143–156.

[45] Madani, K. and Lund, J. R. Modeling California's high-elevation hydropower systems in energy units. *Water Resources Research* 45(9) (2009).

[46] Massé, P. *Les Reserves et la Regulation de I'Avenir dans la vie Economique*. Vol. I and II. Paris, France: Hermann, 1946.

[47] May, N. and Neuhoff, K. Financing Power: Impacts of Energy Policies in Changing Regulatory Environments. 2017. URL: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3046516`.

[48] McDonald, R. L. *Derivatives Markets*. 3rd ed. Pearson New International Edition. Harlow, Essex, UK: Pearson Education Limited, 2014.

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

[50] Mo, B., Gjelsvik, A., Grundt, A., and Kåresen, K. Optimisation of hydropower operation in a liberalised market with focus on price modelling. *Porto Power Tech Conference*. IEEE, 2001.

[51] Mo, B., Green, T., Warland, G., and Haugstad, A. *Dokumentasjon av ny prismodell til bruk i Vansimtap, Plansddp og Beta*. Tech. rep. Trondheim, Norway: Sintef Energiforskning AS, Sept. 2000.

[52] Modigliani, F. and Miller, M. H. The cost of capital, corporation finance and the theory of investment. *The American Economic Review* 48(3) (1958), 261–297.

[53] Nasdaq. *Fee list - Commodity derivatives*. 2018. URL: `http://www.nasdaqomx.com/digitalAssets/107/107680_180301-joint--appendix-7--fee-list.pdf`.

[54] Norwegian Ministry of Petroleum and Energy. *Ownership in the energy sector*. 2017. URL: `https://energifaktanorge.no/en/om-energisektoren/eierskap-i-kraftsektoren/`.

[55] Pereira, M. and Pinto, L. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52(1-3) (1991), 359–375.

[56] Powell, W. B. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Vol. 842. Princeton, New Jersey, USA: John Wiley & Sons, 2011.

[57] Rebennack, S. Combining sampling-based and scenario-based nested Benders decomposition methods: application to stochastic dual dynamic programming. *Mathematical Programming* 156 (2015), 343–389.

[58] Ruszczynski, A. and Shapiro, A. Conditional Risk Mappings. *Mathematics of Operations Research* 31(2) (2006), 544–561.

[59] Samuelson, P. A. Proof That Properly Anticipated Prices Fluctuate Randomly. *Industrial Management Review* 6(2) (1965), 41–49.

[60] Sanda, G. E., Olsen, E. T., and Fleten, S.-E. Selective hedging in hydro-based electricity companies. *Energy Economics* 40 (2013), 326–338.

[61] Shapiro, A. On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters* 37 (2009), 143–147.

[62] Shapiro, A., Tekaya, W., Costa, J. P. da, and Soares, M. P. Risk neutral and risk averse Stochastic Dual Dynamic Programming method. *European Journal of Operational Research* 224(2) (2013), 375–391.

[63] SINTEF. *EMPS - multi area power-market simulator - SINTEF*. 2017. URL: `https://www.sintef.no/en/software/emps-multi-area-power-market-simulator/`.

[64] SINTEF. *EOPS - one area power-market simulator*. 2017. URL: `http://www.sintef.no/en/software/eops-one-area-power-market-simulator/`.

[65] Smith, C. W. and Stulz, R. M. The determinants of firms' hedging policies. *Journal of Financial and Quantitative Analysis* 20(4) (1985), 391–405.

[66] Stulz, R. M. Rethinking Risk Management. *Journal of Applied Corporate Finance* 9(3) (1996), 8–25.

[67] Thorvaldsen, T., Tyssing, N., and Samuelsen, J. *Kraftverksbeskatning*. Tech. rep. Oslo, Norway: KPMG Norway, 2018.

[68] Wallace, S. W. and Fleten, S.-E. Stochastic Programming Models in Energy. *Handbooks in OR and MS*. Ed. by Ruszczyński, A. P. and Shapiro, A. Vol. 10. Elsevier Science B.V., 2003. Chap. 10, 637–677.

[69] Wang, Y., Wu, C., and Yang, L. Hedging with Futures: Does Anything Beat the Naïve Hedging Strategy? *Management Science* 61(12) (2015), 2870–2889.

[70] Wolfgang, O., Haugstad, A., Mo, B., Gjelsvik, A., Wangensteen, I., and Doorman, G. Hydro reservoir handling in Norway before and after deregulation. *Energy* 34(10) (2009), 1642–1651.

[71] Zanotti, G., Gabbi, G., and Geranio, M. Hedging with futures: Efficacy of GARCH correlation models to European electricity markets. *Journal of International Financial Markets, Institutions and Money* 20(2) (2010), 135–148.

Appendix

A.1 Discretization

Table A.1 shows how the calendar year is discretized into 24 semi-months.

Semi-month number	Start date	End date	Duration [days]
1	1.1	15.1	15
2	16.1	31.1	16
3	1.2	14.2	14
4	15.2	28.2 (29.2)	14 (15)
5	1.3	15.3	15
6	16.3	31.3	16
7	1.4	15.4	15
8	16.4	30.4	15
9	1.5	15.5	15
10	16.5	31.5	16
11	1.6	15.6	15
12	16.6	30.6	15
13	1.7	15.7	15
14	16.7	31.7	16
15	1.8	15.8	15
16	16.8	31.8	16
17	1.9	15.9	15
18	16.9	30.9	15
19	1.10	15.10	15
20	16.10	31.10	16
21	1.11	15.11	15
22	16.11	30.11	15
23	1.12	15.12	15
24	16.12	31.12	16

Table A.1: Semi-monthly discretization. Semi-month 4 consists of 15 days in leap-years.

A.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 squared distance between the nodes and a set of simulated state observations. Consider K Monte-Carlo simulations (S^k) where S_t^k is the time t state of simulation k where $k \in [1,...,K]$. Then, the squared distance between a simulated state S_t^k and its closest node $\overline{S_{t,n}}$ is given by the squared Euclidian distance $||S_t^k - \overline{S_{t,n}}||^2$. In order to avoid creating a unbalanced lattice, all state variables $l \in [1,...,L]$ are weighted with a weight $\theta_{t,l} = 1/\sigma_{X_{t,l}}$ when calculating the squared distance. Here $X_{t,l}$ is a random variable representing the possible time t realizations of the ith variable and $\sigma_{X_{t,l}}$ is its standard deviation. Without weighting, the lattice will be biased towards the state variables of larger magnitudes. In our case, the inflow magnitudes are very large at certain stages t.

It is difficult to find the optimal nodes $\overline{S_{t,n}}$ analytically. Therefore, we 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 $\nu = \{\nu_k : k \in [K]\}$. Further, we define $\overline{S_{t,n}^k}$ as

$$\overline{S_{tn}^k} = \begin{Bmatrix} \overline{S_{tn}^{k-1}} + \nu_k(S_t^k - \overline{S_{tn}^{k-1}}) & \text{if} \quad n = \arg\min_m\{||S_t^k - \overline{S_{t,m}^{k-1}}||^2, m \in [N_t]\} \\ \overline{S_{tn}^{k-1}} & \text{otherwise} \end{Bmatrix} \tag{A.1}$$

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

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 $\overline{S_{t,n}}$ to $\overline{S_{t+1,m}}$. Bally and Pages (2003) propose estimating the transition probabilities by

$$p_{tnm} = \frac{\sum_{k=1}^K \mathbb{I}_{\Gamma_{tn}}(S_t^k)\mathbb{I}_{\Gamma_{tm}}(S_{t+1}^k)}{\sum_{k=1}^K \mathbb{I}_{\Gamma_{tn}}(S_t^k)}, \quad n \in [N_t], m \in [N_{t+1}], t \in [\hat{T}-1] \tag{A.2}$$

Here, $\mathbb{I}_{\Gamma_{t,n}}(X)$ is an indicator function taking the value 1 if X is a part of the set $\Gamma_{t,n}$ and 0 otherwise. $\Gamma_{t,n}$ is the Voronoi decomposition of the nodes $\overline{S_{t,n}}$. In other words, $\Gamma_{t,n}$ is the set of simulated states S_t^k whose nearest node defined by the squared distance is $\overline{S_{t,n}}$. Mathematically, $\Gamma_{t,n}$ can be written as

$$\Gamma_{tn} = \left\{S_t^k : n = \arg\min_m\{||S_t^k - \overline{S_{t,m}}||^2, m \in [N_t]\}\right\} \tag{A.3}$$

As shown in Löhndorf and Wozabal (2017), a lattice constructed using (A.1) and (A.2) will be biased, as the expected successor node of a node at $t < \hat{T}$ will not equal the expected successor state given by the underlying processes. To obtain a lattice where this is the case, we use *backwards estimation*, as used in Löhndorf and Wozabal (2017). For a given node $\overline{S_{t,n}}$, its expected successor node is denoted $\mathbb{E}[\overline{S_{t+1}} \mid \overline{S_{t,n}}]$. In an unbiased lattice, the expected successor node of $\overline{S_{t,n}}$ must be given by the expression

$$\mathbb{E}[\overline{S_{t+1}} \mid \overline{S_{t,n}}] = \sum_{m=1}^{N_{t+1}} p_{tnm}\overline{S_{t+1,m}}, \quad m \in [N_{t+1}], t \in [\hat{T}-1] \tag{A.4}$$

The identity in (A.4) can be used to adjust the biased lattice so that it becomes unbiased. We let $S_{t,n,i}$ denote the value of the ith state variable included in the node $\overline{S_{t,n}}$. Then, the expected *predecessor* state of this node is denoted $\mathbb{E}[S_{t-1,i} \mid S_{t,n,i}]$. Using the conditional transition probabilities p_{tnm} found by (A.2), the adjusted value of all state variables in each node can be found using

$$S_{t,n,i} = \sum_{m=1}^{N_{t+1}} p_{tnm}\mathbb{E}[S_{t,i} \mid S_{t+1,m,i}], \quad i = [1,...,\dim(\overline{S_{t+1,m}})], m \in [N_{t+1}], t \in [\hat{T}-1] \tag{A.5}$$

When using (A.5), the nodes must be adjusted iteratively, starting at the end of the lattice and working backwards to its beginning. The nodes of the last time stage are fixed to their original values. Using this, one can adjust the value of the nodes in the penultimate stage by (A.5), and so on.

A.3 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. They model the curve discretely by finding one unique price for a set of constant time steps. We denote

$$\mathbb{S} = \left\{(t_{b,1}, t_{e,1}), ..., (t_{b,M}, t_{e,M})\right\} \tag{A.6}$$

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 $t_e = t_{e,M}$. We 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 (A.6).

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 $f(t_s) = [F_{t_s,t_b}, F_{t_s,t_{b+1}}, ..., F_{t_s,t_e}]'$ 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 $F^{ask}_{t,t_{b,j},t_{e,j}}$ and $F^{bid}_{t,t_{b,j},t_{e,j}}$ denote the ask and bid price of forward contract $j = [1, ..., M]$, where $(t_{b,j}, t_{b,j})$ denotes its delivery period. $D(t) = [D(t_b), D(t_b + 1), ..., D(t_e)]'$ denotes an underlying deterministic price function in vector form, and r is the model discount rate. Lastly, we denote $\omega \in \langle 0, \infty \rangle$ as the smoothness parameter. For high values of ω, 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 $f(t_s)$ is constructed by solving the minimization problem

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

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

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

Using the method of Lagrange multipliers, the problem can be reformulated to solving the matrix equation given by (A.7), where $\chi = [\chi_1, ..., \chi_M]'$ is the vector of Lagrange multipliers and $f(t_s)$ is the forward curve on vector form.

$$\begin{bmatrix} 2\mathbf{B} & \mathbf{A}' \\ \mathbf{A} & 0 \end{bmatrix} \cdot \begin{bmatrix} f(t_s) \\ \chi \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{F_{t_s}} \end{bmatrix} \tag{A.7}$$

In A.7, $\mathbf{A}$ 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 jth forward contract and $A_{j,t} = 0$ otherwise. $\mathbf{F_{t_s}} = [F_{t_s,t_{b,1},t_{e,1}}, ..., F_{t_s,t_{b,M},t_{e,M}}]'$ is a vector containing

the price of all M forward contracts traded on the market. The matrix $\mathbf{B}$ is $C \times C$ and given by (A.8).

$$\mathbf{B} = \begin{bmatrix} 1+\omega & -2\omega & \omega & 0 & 0 & 0 & 0 & \dots & 0 \\ -2\omega & 1+5\omega & -4\omega & \omega & 0 & 0 & 0 & \dots & 0 \\ \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 & 0 & \dots & 0 \\ 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 \\ 0 & \dots & 0 & 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega \\ 0 & \dots & 0 & 0 & 0 & \omega & -4\omega & 1+5\omega & -2\omega \\ 0 & \dots & 0 & 0 & 0 & 0 & \omega & -2\omega & 1+\omega \end{bmatrix} \tag{A.8}$$

A.4 Constructing forward curves by linear interpolation

Alexander (2008) 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 (2008) 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 method of Fleten and Lemming (2003), 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 $f(t_s)$.

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

$$t_{b,i} \leq t \leq t_{b,j} \tag{A.9}$$

By linear interpolation, the forward curve elements $F_{t_s,t}$ located in the interval $t \in (t_{b,i}, t_{b,j})$ are given by (A.10).

$$F_{t_s,t} = F_{t_s,t_{b,i},t_{e,i}} + \frac{t - t_{b,i}}{t_{b,j} - t_{b,i}} \cdot (F_{t_s,t_{b,j},t_{e,j}} - F_{t_s,t_{b,i},t_{e,i}}) \tag{A.10}$$

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

$$t_{e,i} - t_{b,i} = t_{e,j} - t_{b,j} \tag{A.11}$$

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 (A.9). 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.

A.5 Constructing returns series from forward curves

Having constructed forward curves for multiple consecutive days, one can create a returns dataset that can be used to find the volatility functions. 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 a volatility function on the form $\sigma_{t,T} = \sigma_\tau$, we must create returns series for a set of contracts with equal time to maturity $\tau = T - t$. $f(t_j)$ 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 (A.12) to calculate daily returns at time t_j for contracts with time to maturity τ_a.

$$x_{j,a} = \ln(F_{t_j,t_j+\tau_a}) - \ln(F_{t_{j-1},t_j+\tau_a}) \tag{A.12}$$

Here, $j = [2,...,J]$ and $a = [1,...,A]$, where J is the number of forward curves and A is the number of maturity dates for which we want to construct the returns dataset. 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

$$\mathbf{X}_{\mathbf{J}\times\mathbf{A}} = \begin{bmatrix} x_{1,1} & x_{1,2} & \dots & x_{1,A} \\ x_{2,1} & x_{2,2} & \dots & x_{2,A} \\ \vdots & \vdots & \ddots & \vdots \\ x_{J,1} & x_{J,2} & \dots & x_{J,A} \end{bmatrix} \tag{A.13}$$

A.6 Covered interest arbitrage

Covered interest rate parity can be derived by considering a reverse cash-and-carry arbitrage as presented in McDonald (2014). Suppose the $T-t$ maturity interest rate is r in the domestic currency and r_f in the foreign currency. Assume that $r > r_f$ and that an arbitrageur wants to borrow the foreign currency and invest in the domestic currency. She borrows at exchange rate $Q_{t,t}$. To hedge the position, she enters a long forward rate contract at forward rate $Q_{t,T}$. The cash flows from the reverse cash-and-carry strategy are shown in Table A.2.

Time	t	T
Borrow $Q_{t,t}\exp\Big(-r_f(T-t)\Big)$ in foreign currency	$+Q_{t,t}\exp\Big(-r_f(T-t)\Big)$	$-Q_{T,T}$
Lend $Q_{t,t}\exp\Big(-r_f(T-t)\Big)$ in domestic currency	$-Q_{t,t}\exp\Big(-r_f(T-t)\Big)$	$Q_{t,t}\exp\Big((r-r_f)(T-t)\Big)$
Enter long forward rate contract	0	$+Q_{T,T}-Q_{t,T}$
Sum of cash flows	0	$+Q_{t,t}\exp\Big((r-r_f)(T-t)\Big)-Q_{t,T}$

Table A.2: Cash flows at time t and $T > t$ from a reverse cash and carry arbitrage in the money markets

The portfolio of positions in Table A.2 is risk-less and requires zero investment. If we assume no arbitrage, the forward rate $Q_{t,T}$ must be priced such that the total cash flow at time T is zero.

$$\begin{aligned} Q_{t,t}\exp\Big((r-r_f)(T-t)\Big) - Q_{t,T} &= 0 \\ Q_{t,T} &= Q_{t,t}\exp\Big((r-r_f)(T-t)\Big) \end{aligned} \tag{A.14}$$

Covered interest rate parity is given by (A.14). Note that $Q_{t,T}$ is independent of the market expectation of the future exchange rate.

A.7 Process parameter validation

In this section, we validate the parameters that were used in the underlying stochastic processes. We investigate certain properties of the price, inflow and currency process. Additionally, we propose a mean-reverting process to represent the area system price difference. Table A.3 shows the estimation windows that were used.

	Start	**End**
Electricity forward price	03.04.2011	31.12.2014
Electricity spot area and system price	01.01.2009	31.12.2014
Inflow	01.01.1958	31.12.2014
EURNOK	04.01.1999	31.12.2014

Table A.3: Data window for parameter estimation in the stochastic processes

A.7.1 Price process

To verify the price process, we compare the historical and modeled spot price volatility. Since the spot price volatility in the price process is estimated using forward price data only, there might be a discrepancy with the volatility estimated using spot price data. The modeled spot price volatility is given by the most near-term part of the forward curve shown in Figure 3.4. In other words, it is the volatility of the fictional one semi-month ahead forward contract. As explained in Section 3.4.1, the HJM framework considers seasonality, while still incorporating stationary returns distributions. We need to exclude seasonality effects when we estimate the volatility using historical spot price data. Therefore, the historical spot price volatility is calculated using seasonality-normalized log returns between semi-monthly average prices.

Table A.4 shows that the deviation between the historical and modeled volatility of the electricity spot price is relatively small. This strengthens the soundness of the price model. We have chosen not to adjust the volatility of the price process because this would not only affect the spot price volatility but the volatility of the entire forward curve as well.

	Historical spot price volatility	**Model spot price volatility**
Standard deviation of semi-monthly returns	0.17	0.16
Annualized volatility	0.84	0.81

Table A.4: Spot price volatility validation, using semi-monthly returns

Figure A.1 and A.2 show the scenario lattices for the spot price and forward prices with 3, 6 and 12 months to delivery. Notice the Samuelson (1965) effect - the volatility is decreasing with increasing time to maturity.

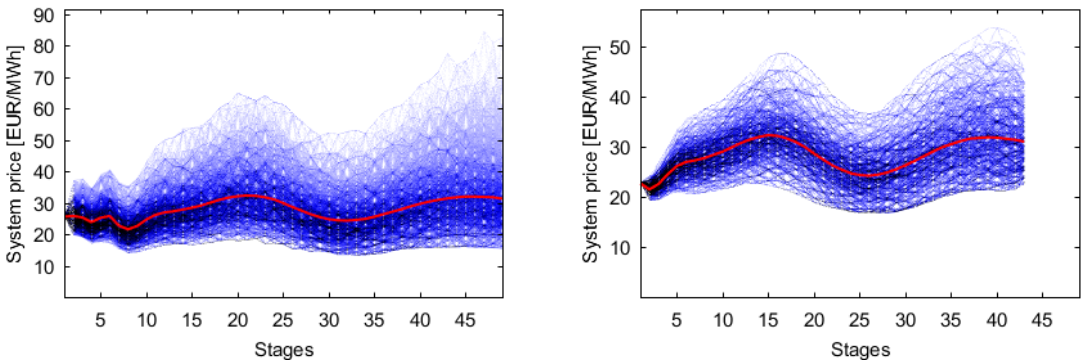

(a) Spot price lattice

(b) 3 months to delivery lattice

Figure A.1

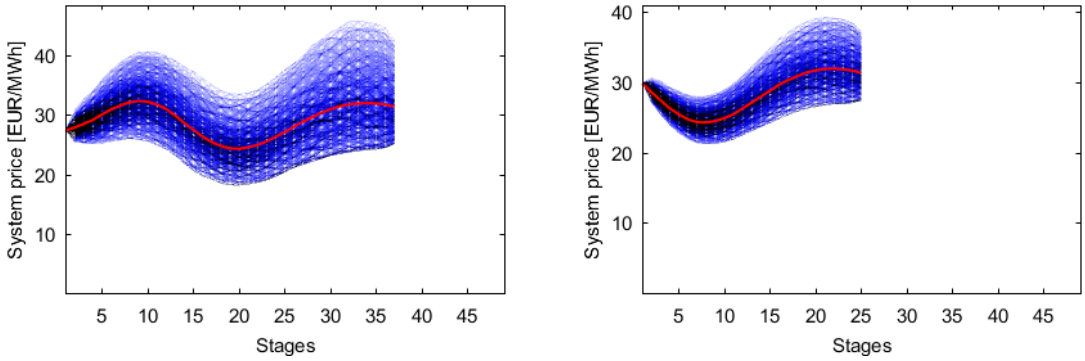

(a) 6 months to delivery lattice

(b) 12 months to delivery lattice

Figure A.2

Figure A.3 shows some of the stage t quantiles along with the mean for the spot price. Uncertainty in future states increases with time.

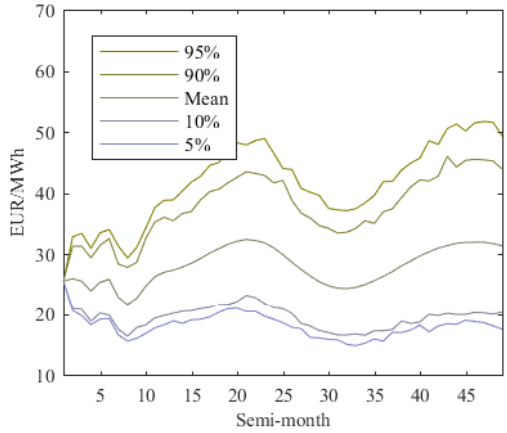

Figure A.3: Spot price uncertainty diagram

A.7.2 Inflow process

Figure A.4 shows some of the stage t quantiles along with the mean for the inflow. The inflow process is estimated such that the average standard deviation of the inflows equal the historical average standard deviation of inflows.

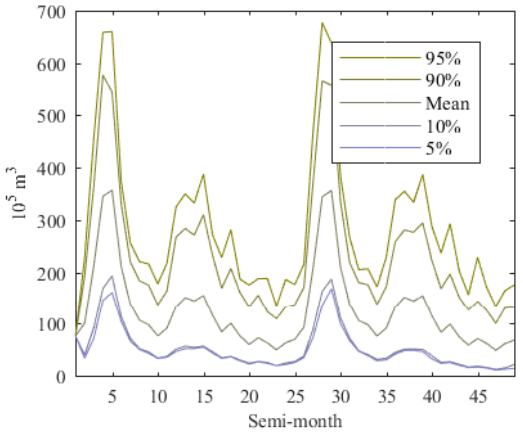

Figure A.4: Inflow uncertainty diagram

A.7.3 Currency process

Figure A.4 shows some of the stage t quantiles along with the mean for the EURNOK rate.

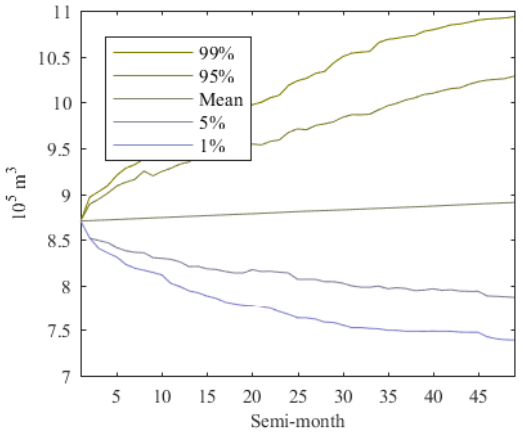

Figure A.5: Currency uncertainty diagram

To validate that covered interest parity holds, we analyze the dynamics of the EURNOK forward curve empirically. Figure A.6 shows that the forward curve dynamics of EURNOK exchange rate exhibit almost perfect correlation. For practical purposes, all currency risk can be considered to originate from the spot exchange rate, and a single factor model for the currency forward curve is sufficient for our purposes.

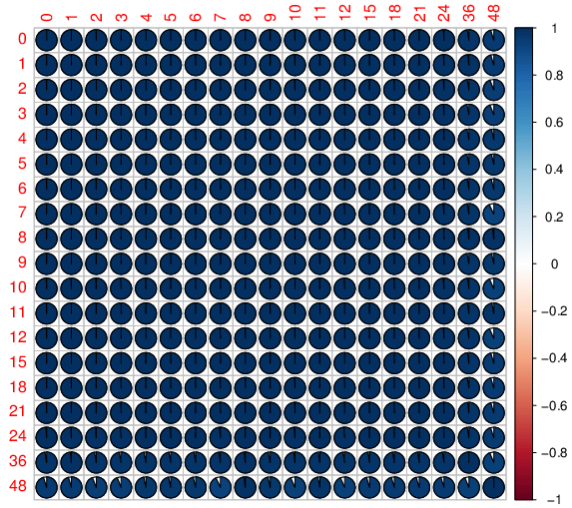

Figure A.6: EURNOK forward rate correlation matrix, showing the correlation of returns for currency forward contracts with number of months to maturity as indicated by the red numbers in the figure. The forward rate data is collected from the SIX Group. The dataset consists of daily observations from 2015-03-23 to 2018-03-12.

A.7.4 Area difference process

Since the producer receives the area spot price for their production while the system spot price serves as the underlying of the traded power futures, there is basis risk between the area price and the power futures. We have estimated a mean-reverting process to represent the area price difference risk. The area price $F_{a,t}$ can be expressed as a function of the system price $F_{t,t}$ and the log of the area price system price ratio, U_t in the following way:

$$U_t = \ln\left(\frac{F_{a,t}}{F_{t,t}}\right) \implies F_{a,t} = F_{t,t}\exp(U_t) \tag{A.15}$$

We assume that the logarithm of the area price system price ratio, U_t, follows a mean-reverting process. In particular, it follows the Ornstein-Uhlenbeck process shown in (A.16)

$$dU_t = [\eta(\bar{U} - U_{t-dt})]dt + \sigma_U dZ \tag{A.16}$$

Here, η is the rate of mean-reversion and $\bar{U}$ is the long term mean value of the log ratio. σ_U denotes the volatility of U. We discretize (A.16) and get

$$\Delta U_t = [\eta(\bar{U} - U_{t-\Delta t})]\Delta t + \sigma_U\sqrt{\Delta t}\varepsilon_t \tag{A.17}$$

Here, $\varepsilon_t \sim N(0,1)$. (A.17) is rewritten into a form suitable for least squares regression analysis.

$$\Delta U_t = \eta_0 + \eta_1 U_{t-\Delta t} + \eta_2\varepsilon_t \tag{A.18}$$

η_0 and η_1 and η_2 are coefficients to be estimated by regression, following the procedure in Clewlow and Strickland (2000). By substitution, it can be shown that

$$\eta = -\frac{\eta_1}{\Delta t} \qquad \bar{U} = \frac{\eta_0}{\eta\Delta t} \qquad \sigma_U = \frac{\eta_2}{\sqrt{\Delta t}} \tag{A.19}$$

In our case, we found the parameters shown in Table A.5. Further, all correlations with price and currency increments were insignificant at a 5 % significance level. The correlation between the increments of the area price process and the inflow process was, however, significant at $\alpha = 0.05$. It was found to be -0.21.

η	10.06971
$\overline{U}$	0.030995
σ_U	0.327369

Table A.5: Area difference process parameters for semi-monthly granularity

A.7.5 Significance of correlation matrix

Table A.6 shows the significance levels of the correlations shown in Table 3.7. For illustration, we also include the area price difference.

	inflow	eurnok	area_diff	price1	price3	price5	price7	price19	price28
inflow									
eurnok	0.08								
area_diff	-0.21*	-0.14							
price1	-0.08	0.03	-0.13						
price3	-0.12	-0.01	-0.06	0.81****					
price5	-0.10	-0.11	0.01	0.67****	0.93****				
price7	-0.14	-0.16	0.00	0.55****	0.84****	0.94****			
price19	-0.03	-0.16	-0.11	0.53****	0.76****	0.78****	0.75****		
price28	-0.15	-0.22*	0.02	0.40****	0.66****	0.73****	0.75****	0.82****	
price48	-0.12	-0.21*	0.01	0.34**	0.56****	0.63****	0.65****	0.82****	0.88****

Table A.6: $p < .0001$, ****; $p < .001$, ***; $p < .01$, **; $p < .05$, *

A.7.6 Variability of state variables

Having constructed uncertainty lattices for spot price, production and currency, we can quantify the variability of each risk factor. Table A.7 shows the quantiles of the distributions for mean spot price, mean EURNOK rate and yearly production over the time horizon of 49 semi-months. Clearly, price and production volume have the largest variability, suggesting that these are the major risk factors for the hydropower producer. However, this might be a hasty generalization as the present value of production revenues is not necessarily proportional to the mean spot price or mean EURNOK rate. This is because of discounting effects and since production is not evenly distributed throughout the year.

	Mean spot price [EUR/MWh]	Production [GWh/year]	Mean EURNOK
Mean	27.99	192.18	8.80
Std	4.55	20.90	0.44
Q1%	-30.5 %	-25.6 %	-11.0 %
Q5%	-24.3 %	-18.2 %	- 7.8 %
Q95%	+29.6 %	+17.6 %	+ 8.4 %
Q99%	+43.8 %	+24.4 %	+12.7 %

Table A.7: We investigate the mean spot price, production per year and mean EURNOK over the time horizon of 49 semi-months. The table shows the mean, standard deviation, the deviation in percent from the mean value for the 1%, 5%, 95% and 99% quantile. E.g., the number in the first column and third row show how much the 1% quantile of the mean spot price distribution deviates from the mean of the mean spot price over 49 semi-months.

A.8 Lattice stability

Due to the backwards estimation adjustment of the position of the nodes, there is some variation in the values assigned to a particular node when creating multiple lattices with the same processes. Table A.8 show how much the starting values for some state variables vary between each time a lattice is created.

	True value	Std.dev	Std.dev as a fraction of the true value
Inflow	80.604	3.828	4.75 %
Currency	8.704	0.0029	0.03 %
Spot	25.529	0	0.00 %
1 semi-month to delivery	25.936	0.124	0.48 %
2-48 semi-months to delivery	-	-	0.13 %*

Table A.8: Lattice stability. Values are the starting values for the state variables. * Average

A.9 Coefficient and parameter values

For all simulations in the hedging article, we have used the coefficient and parameter values given in Table A.9. The tax rates γ_r and γ_c (Thorvaldsen et al., 2018) and transaction costs c_F (Nasdaq, 2018) are correct as of June 6, 2018. As in Dupuis et al. (2016), the variable transaction costs for trading at NASDAQ OMX are given by the sum of the market trading fee (0.0045 EUR/MWh) and clearing fee (0.0099 EUR/MWh). This clearing fee is applicable if the total quarterly volume cleared by the firm is below 3 TWh. The time-dependent discount factor β_t is given by $\beta_t = \exp(-r\Delta t)$, where Δt denotes the length of the semi-month t. The energy coefficient κ is based on the empirical relationship between production and water dispatch, and calculations considering the mean empirical reservoir level and turbine/generator efficiency rate. It has been found to be slightly lower than the one currently used by the plant.

Coefficient/ Parameter	Explaination	Unit	Value
$\overline{v_{1,t}}$	Upper bound for reservoir volume in reservoir 1	Mm^3	44.5
$\overline{v_{2,t}}$	Upper bound for reservoir volume in reservoir 2	Mm^3	22.5
$\underline{v_{2,t}}$	Lower bound for reservoir volume in reservoir 2 between October 16 and May 24	Mm^3	0
$\underline{v_{2,t}}$	Lower bound for reservoir volume in reservoir 2 between May 25 and October 15	Mm^3	15.05
κ	Energy coefficient	kWh/m^3	0.63
ξ	Maximum allowed water flow in turbine	m^3/s	17
r	Continuously compounded annual risk free interest rate used in discount factor, given by 3-year NIBOR	–	0.0126
c_F	Transaction costs for trading power futures at NASDAQ OMX	EUR/MWh	0.0144
γ_r	Resource rent tax rate	–	0.357
γ_c	Corporate tax rate	–	0.23
ζ	Inflow split coefficient	–	0.395

Table A.9: Coefficients and constants used in hedging paper

A.10 Sensitivity analysis of hedging results

In this section, we present some results illustrating the sensitivity of using 500 forward-backward passes and 10^5 iterations. This has been done by performing six separate runs of the hedging model with $\alpha = 0.1$ and trading in all contracts, and comparing the mean and standard deviation of the main statistical measures included in e.g. Table 3.5. These results are summarized in Table A.10. The computational time of each run is approximately 4.5 hours, and its memory usage is close to the maximum capacity of the computer that is used.

	Mean CF	**Std CF**	**VaR**(5%)	**VaR**(1%)	**CVaR**(5%)	**CVaR**(1%)
Mean	40.20	4.33	33.35	30.81	31.79	29.60
Std	0.106	0.050	0.145	0.179	0.165	0.204

Table A.10: Sensitivity of statistical measures based on six separate runs.

As Table A.10 shows, the obtained results are subject to a certain degree of uncertainty. The standard deviations indicate that all statistical measures are stable in the first two digits, while there is some uncertainty in the third digit. This indicates that the results are sufficiently precise to assess and compare the general risk performance of the model variants, but increasing the number of passes and simulations would result in more precise results.

A.11 Optimal tax-neutral hedge ratio

In this section, we deduce the calculation of the optimal tax-neutral hedge ratio found in Sanda et al. (2013). We want to hedge the cash flows from production. The deviation from the expected cash flows from production in a stage T is given by

$$(F_{T,T} - \mathbb{E}(F_{T,T})) \cdot W_T(1 - \gamma_c - \gamma_r) \tag{A.20}$$

Here, $F_{T,T}$ is the spot price at time T, W_T is the production volume, γ_c is the corporate tax rate and γ_r is the resource rent tax rate. Further, the cash flow from a forward contract is given by

$$(F_{t,T} - F_{T,T}) \cdot u_{T,T}(1-\gamma_c) \tag{A.21}$$

where $F_{t,T}$ is the forward price at time t for delivery at time T and $u_{T,T}$ denotes the short position in the forward contract. Further, we assume that the expected spot price equals the forward price and set $\mathbb{E}(F_{T,T}) = F_{t,T}$. We also disregard the uncertainty in production W_T. We want the cash flow deviations from the long position in production and short position in the financial market to perfectly offset each other. Then, we get that the optimal tax-neutral hedge ratio $\frac{u_{T,T}}{W_T}$ is given by (A.22).

$$\begin{aligned}
(F_{T,T} - F_{t,T}) \cdot W_T(1-\gamma_c-\gamma_r) + (F_{t,T} - F_{T,T}) \cdot u_{T,T}(1-\gamma_c) &= 0 \\
(F_{T,T} - F_{t,T}) \cdot W_T(1-\gamma_c-\gamma_r) &= -(F_{t,T} - F_{T,T}) \cdot u_{T,T}(1-\gamma_c) \\
\frac{u_{T,T}}{W_T} &= \frac{1-\gamma_c}{1-\gamma_c-\gamma_r}
\end{aligned} \tag{A.22}$$

Using $\gamma_c = 0.23$ and $\gamma_r = 0.357$, we get that $\frac{u_{T,T}}{W_T} = 53.6\%$. This is slightly lower than the hedge ratio of 58.3 % calculated by Sanda et al. (2013), which was based on the Norwegian tax regime of 2010, when the corporate tax rate was higher and the resource rent was lower.

A.12 Numerical example of cash flows from overhedged position

In this section, we present a numerical example illustrating our argumentation for the large short positions suggested by the hedging model. We consider a simple two-stage problem with two possible scenarios for price and production in the second stage $t = 1$. As assumed, we consider a positive correlation between production level and spot price. The problem is illustrated in Figure A.7.

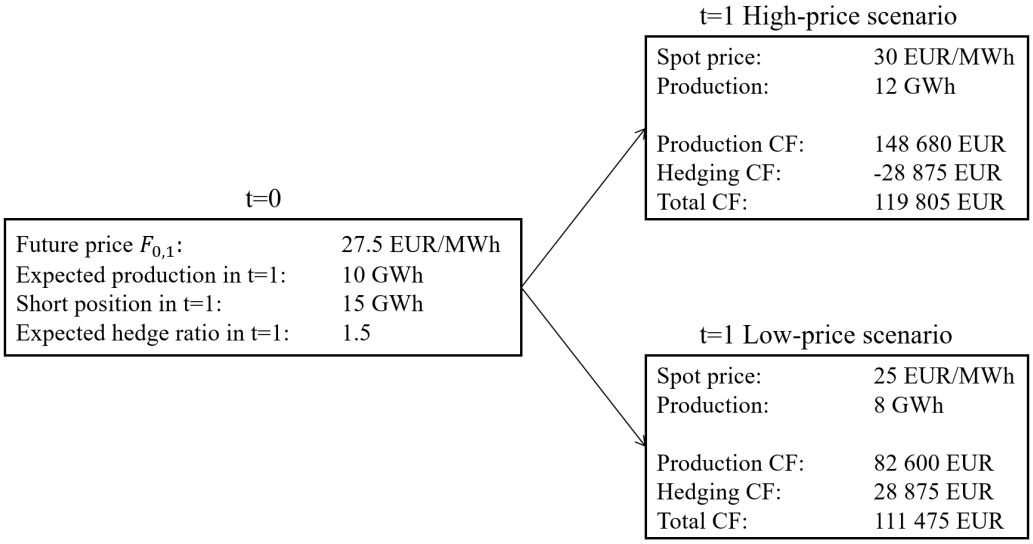

Figure A.7: Two-stage problem with stochastic price and production levels. The cash flows from production are subject to resource rent γ_r and corporate γ_c tax, whereas the cash flows from hedging are subject to corporate tax only. Transaction costs are disregarded.

As the figure shows, the expected production in stage $t = 1$ is 10 GWh, meaning that a hedge ratio of 1.5 can be achieved with a short position of 15 GWh in the futures market. For simplicity, we assume that the position is entered in stage $t = 0$ and only settled in stage $t = 1$, i.e., its delivery period only covers the time stage $t = 1$. If hedging is disregarded, the potential low-price scenario cash flows earned by the producer will be 82 600 EUR. However, when the large short position is included, the total cash flows in the low-price scenario are increased by 28 875 EUR. This increase is due to the decrease in spot price from its expected value, i.e., the futures price $F_{0,1}$. Thus, the over-hedged position results in low-price scenario cash flows of 111 475 EUR, meaning that the CVaR is lifted. The cash flows in the high-price scenario are decreased by the same magnitude, but the total cash flows still remain larger than in the low-price scenario.