Dynamic Hedging for the Real Option Management of Hydropower Production with Exchange Rate Risks

Joakim Dimoski*, Stein-Erik Fleten*, Nils Löhndorf†[1], Sveinung Nersten*

*Department of Industrial Economics and Technology Management
Norwegian University of Science and Technology, NO-7491 Trondheim, Norway

†Luxembourg Centre for Logistics and Supply Chain Management
University of Luxembourg, L-1511 Luxembourg

Abstract

We study the risk management problem of a hydropower producer that participates in a wholesale electricity market and hedges risk by trading currency and power futures contracts. Our model considers three types of risks: operational risk due to future supply uncertainty, exchange rate risk for operations and trading in different currencies, and profit risks due to power price variability. We cast the problem as a Markov decision process and propose a sequential solution approach to handle the high complexity of the optimization problem. Our contribution is three-fold: first, we show how currency risk and currency derivatives can be included in real option models of hydropower generation; second, we accurately encode the cash flow structure from a portfolio of electricity and currency hedge contracts; and third, we compare optimization under a risk measure with often-used simple hedging strategies. For the case of a Norwegian hydropower producer, we quantify the reduction in risk through currency hedging where there is currency risk. We find that currency hedging leads to a moderate decrease in profit risk, and that including monthly power futures in the hedging strategy allows precision hedging that can contribute to a substantial risk reduction.

Keywords: OR in energy, stochastic programming, risk management, Markov processes, stochastic processes

[1]corresponding author, email: nils.loehndorf@uni.lu

1. Introduction

For hydropower producers with random natural inflows, a problem of practical relevance is to maximize cash flows from buying and selling power under operational as well as market risks. While managing price uncertainty as a type of market risk is well understood in the area of commodity storage (Murphy and Oliveira 2010, Nadarajah et al. 2015, Goel and Tanrisever 2017), other types of market risk are normally not taken into account. In particular, hydropower producers who participate in cross-border trades, like Canadian companies in the U.S. or Scandinavian companies in continental Europe, are exposed to additional exchange rate risks, as operations and trading take place in different currencies. In this article, we focus on hydropower operation under price and exchange rate risks as well as the risk of random natural inflows and how these risks can be mitigated using financial instruments.

Risk-averse companies can use financial instruments such as forward contracts and options to lower risk exposure to their risk preference. Reducing a company's market risk is referred to as hedging or risk management. Dupuis et al. (2016) distinguish between *static* and *dynamic* hedging and between *local* and *global* hedging. Companies use static hedging to hedge their portfolio at a point in time and without subsequent rebalancing. Dynamic hedging, in contrast, involves continuous adjustment of the portfolio as new market information becomes available. Local hedging focuses on minimizing short-term risk until rebalancing, whereas global hedging seeks to minimize risk associated with all future cash flows. Discussions of local hedging in electricity markets are in Zanotti et al. (2010) and Liu et al. (2010); global hedging models are discussed in Mo et al. (2001a), Fleten et al. (2002) and Dupuis et al. (2016). It is shown in Fleten et al. (2002) that dynamic hedging leads to better outcomes than static hedging for hydropower production.

In this article, we present a global dynamic hedging model for companies that operate hydropower assets. The proposed model builds on previous efforts of jointly modeling supply uncertainty and futures price dynamics to obtain more accurate values of water in a reservoir (Dimoski et al. 2018). We complement the previous model with a dynamic hedging model that accounts for trading in monthly, quarterly, and annual power futures contracts. Modeling futures trading along with operational decision is an attractive feature for risk management because all types of risks are then addressed by a single model as opposed to multiple models, which is often the case in practice.

We focus on the case of a price-taking hydropower producer who participates in a wholesale electricity market and who operates a single asset, but the model can be easily extended to handle a portfolio of assets. In addition to price risk, we consider uncertainty in currency spot exchange rates to address risk from cross-border trading. For example, prices in the Nordic electricity market are given in Euro (EUR). Norwegian power companies, however, operate in Norwegian Krone (NOK). We, therefore, allow the company to trade in currency forward contracts to hedge this uncertainty. We consider a planning horizon of two years in semi-monthly time increments, to achieve a better alignment of the time granularity of futures contracts with operational decision making. We also consider transaction and tax costs.

To explain forward curve movements and to generate scenarios for spot and power futures prices, we use a multivariate Heath-Jarrow-Morton (HJM) term-structure model (Heath et al. 1992). In an efficient market, available future and forward contracts traded at a given point in time represent the current risk-adjusted market expectations of future spot prices. A high-resolution forward curve can be constructed using the price and delivery periods of all available futures contracts. For example, see Fleten and Lemming (2003), Benth et al. (2008) and Kiesel et al. (2019). A model that explains the evolution of a forward curve can therefore be used to find future spot prices and to calculate the price of futures contracts for different delivery periods.

We also incorporate stochastic processes for natural inflows and exchange rates into the model. There is typically a negative correlation between natural inflows and electricity prices in the hydro-dominated systems of countries such as Norway, Canada, or Brazil. Electricity supply in these countries is largely determined by inflows. This provides a *natural hedging* effect. We are also interested in investigating the magnitude of exchange rate risk, and how it relates to the other risk factors. We presume that system price is negatively correlated with the EURNOK exchange rate, because all bid and ask orders in the Nordic market are placed by companies that operate using the local currency. This should influence the system price, which is denoted in EUR/MWh, and thus a negative correlation has a natural hedge effect.

Many power companies use simple strategies to hedge their risk. Wang et al. (2015) analyzed several minimum-variance strategies for a number of commodity, currency, and equity markets. They found that a naive hedging strategy (hedge ratio = 1) performs better or almost as well as the minimum-variance strategies in all the tested markets. However, the authors use a static approach and do not consider the electricity market. We propose a dynamic approach by formulating the decision problem as a multistage process and modeling risk preferences using the nested conditional value-at-risk (nested CVaR), see e.g., Shapiro et al. (2013). The nested CVaR is based on conditional convex risk mappings (see Ruszczynski and Shapiro (2006)).

Unlike other risk measures, nested CVaR ensures that risk preferences remain time-consistent under the optimal policy. Boda and Filar (2006) and Shapiro (2009) show that global hedging strategies such as the naive hedging strategy, which aim to reduce the risk of terminal cash flows, are not time-consistent. Static hedging for hydropower producers has also been studied in a number of articles, e.g. Fleten et al. (2010).

Tree-based stochastic programming and dynamic programming techniques are prohibitive in decision problems with many stages and a high-dimensional state space. We therefore resort to approximate dual dynamic programming (ADDP, Löhndorf et al. (2013)), which extends the acclaimed stochastic dual dynamic programming method of Pereira and Pinto (1991). Dual dynamic programming approaches are widely used in the literature on hydropower scheduling (e.g. Mo et al. (2001b) and Rebennack (2015)) and hedging (e.g. Iliadis et al. (2006)). ADDP allows objective coefficients to be stagewise-dependent random variables by discretizing the random data process to a scenario lattice (Löhndorf and Shapiro 2018). This property is important for modeling hedging strategies, as price processes are correlated in time. We use the method described in Löhndorf and Wozabal (2019) to construct lattices that are optimal in a certain sense but remain martingales.

Articles such as Mo et al. (2001a), Fleten et al. (2002) and Kouvelis et al. (2018) have proposed integrated models in which production and hedging decisions are made simultaneously. Wallace and Fleten (2003) claim that it is favorable to treat production planning and risk management as sequential activities. This is because it is not possible to increase the value of a power portfolio by trading hedging derivatives in an efficient market. Only a change in production can achieve a change in portfolio value. This implies that a storage operator should first seek to achieve a production schedule that maximizes expected cash flows and then use hedging to reduce the portfolio risk to the desired level.

This article makes three main contributions to the literature. First, we propose a model that integrates operational decisions for hydropower management with the replication of cash flows from currency and power futures. This involves, for example, dynamic listing and de-listing of futures contracts, dynamics of cash flows during delivery of power futures, transaction costs and a resource rent taxation issue that affects spot revenues from production but not from hedging. The model thereby integrates operational decisions and supply uncertainty with the financial hedging decisions that are made to mitigate market risks from price and exchange rate uncertainty. This allows us to quantify the benefit that currency hedging can provide to electricity traded in another currency. Second, as we cast the model as a Markov decision process, we can compare dynamic hedging (via the nested CVaR) with static hedging, which is popular among practitioners in risk management. Third, we compare the simultaneous approach of optimizing operational and hedging decisions with a sequential approach of separating operational and hedging decisions, which is often pursued in practice.

This article is organized as follows. In section 2, we present relevant background information on risk management in electricity markets, the emphasis being on hydropower production in Nordic countries. In section 3, we describe the hedging problem formulated as a Markov decision process. We also give an overview of the algorithms used to reduce the dimension of the stochastic processes and to solve the stochastic-dynamic decision problem. In section 5, we present how we modelled the different risk factors as stochastic processes. Section 6 presents the numerical results and a conclusion and an outlook on future work is given in section 7.

2. Risk management of hydropower assets

In this section, we provide an overview of relevant aspects of risk management for hydropower assets. Some aspects specifically apply to the Nordic power market, others are more general. First, we underline the important risk factors and the derivatives that can be used to reduce the risk exposure of the owner. Then we present relevant risk measures that firms can use to quantify their exposure to risk. The section concludes with a brief overview of current risk management practices in hydropower companies.

2.1. Risk factors and mitigating derivatives

Hydropower producers are exposed to a range of financial risk factors which can negatively affect their operational cash flows. The Basel Committee on Banking Supervision (2012) categorizes financial risk into credit risk, operational risk, liquidity risk, and market risk. Credit risk is the risk that the counter-party will not live up to their contractual obligations and can be relevant to owners who enter long term OTC contracts. Operational risk relates to losses that are due to inadequate or failed internal processes, people and systems, or external events. This includes political risk. Fleten et al. (2012) identify political risk as an important risk factor for hydropower producers, with hydropower cash flows being sensitive to changes of, for example, regulations on water usage and tax rates. Liquidity risk is the risk that trading an asset adversely affects market price. This is particularly relevant with respect to EPAD futures, as discussed below.

Operational risk, credit risk and liquidity risk will not be covered here. The focus of this article is on market risk. Market risk includes factors such as electricity price risk, currency risk, and interest rate risk. There are also two market risk factors that apply specifically to Nordic hydropower producers. These are risks in natural inflows that effect production yield as well as in area price differences. Fleten et al. (2012) identify inflow and price uncertainty as being major risk factors for hydropower producers. Power price and exchange rate derivatives which are traded in financial markets can be used to hedge market risk. However, there is currently no liquid market for inflow uncertainty derivatives in the Nordic countries. Foster et al. (2015) have developed methods for pricing these contracts based on hydrology and weather indices.

The Nord Pool day-ahead market – also referred to as spot market – is the primary market for Nordic power producers. The day-ahead market is split up into a number of price areas, the spot price of each area being calculated based on the supply-demand equilibrium of each area and transmission capacity constraints between areas. All prices are denoted in EUR/MWh. Nord Pool also calculates system price, which is an unconstrained market clearing reference price for the entire Nord Pool market. Financial contracts for the Nordic market are traded on NASDAQ OMX Commodities Europe. This market is a purely financial market where no physical energy is exchanged. The day-ahead market, by contrast, is a physical market.

System price risk can be hedged by trading futures and options on NASDAQ OMX. The power futures traded on NASDAQ OMX have delivery periods that span one day, one week, one month, one quarter, or one year. Traditional forward and futures contracts are, conversely, contracts for the trade of an asset at a specific point in time. Power futures are therefore more like financial swaps. As explained by 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.

A power producer receives the area spot price for their generation. As the system spot price serves as the underlying price for traded power futures, there exists a base risk between area spot price and power futures. NASDAQ OMX also offers contracts for hedging area price difference, the reference being the difference between system spot price and area spot prices. These contracts are known as EPADs. They are significantly less liquid than power futures and are also only available for certain areas. Houmøller (2017) has shown that there is a high correlation between the hourly system price and the price for most Norwegian areas. The average for all areas combined for 2013–2016 is approximately 0.89. This is above the limit of 0.8 set by the IAS 39 accounting standard for qualifying for hedge accounting. Houmøller (2017) therefore argues that it is sufficient that power producers in most Norwegian areas disregard area difference in their hedging strategy and hedge only using system price contracts.

Revenues of Norwegian power producers are generated in EUR. The producer's base currency however is NOK. Since exchange rates fluctuate, this means a significant currency risk for producers. Currency risk can be hedged using forward exchange contracts. The market for currency derivatives is highly liquid and bid-ask spreads for corporate customers are negligible. Currency forwards are over-the-counter (OTC) instruments and are not traded on a centralized exchange. An investor typically enters into a contract with a bank as counter-party. All bid orders and ask orders in a price area are placed by companies that operate using their local currency. The system price, which is denoted in EUR/MWh, should therefore be influenced by the base currencies of the different areas. We therefore expect system price to be negatively correlated with the EURNOK exchange rate, thereby providing a natural hedging effect.

2.2. Risk measurement

For hydropower producers, risk measurement is typically based on end-of-horizon cash flows, because these are simple to interpret (Fleten et al. 2010). Standard deviation has historically been one of the most widely used metrics (Stulz 1996), as it measures deviations from expected cash flows in the positive and negative direction. Nonetheless, Stulz (1996) argues that it is of greater interest to consider the downside risk of the cash flow distribution. A common metric for measuring downside risk is 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 $VaR_\alpha = \min\{h_i \mid \sum_{j|h_j \leq h_i} 1/H \geq \alpha\}$. Although VaR provides a risk manager with information on worst case scenarios, it provides no information about the distribution of cash flows below the VaR. The Conditional Value-at-Risk (CVaR), by contrast, measures expected cash flows below the VaR. Mathematically, the CVaR of a discrete distribution at significance level α is given by $CVaR_\alpha = \sum_{j|h_j \leq VaR_\alpha} h_j/(H\alpha)$. CVaR has a number of benefits over VaR:

1. CVaR better captures tail-effects such as kurtosis and skewness.
2. CVaR is a coherent[2] risk measure, which makes it particularly attractive in risk management (Godin 2016).
3. CVaR is a convex risk measure which makes it attractive for mathematical optimization.

[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 $\Phi(X_1 + X_2) \leq \Phi(X_1) + \Phi(X_2)$.

2.3. *Industry practice*

Examining how risk management is performed in the hydropower industry provides useful benchmarks. We therefore provide an overview of industry practices with focus on the Norwegian market.

Sanda et al. (2013) analyzed the hedging strategies of 12 Norwegian hydropower producers, their total production accounting for about 34.4% of total production in Norway. The authors find that none of the firms use an integrated model to obtain optimal production and hedging decisions. Five used a sequential approach by obtaining optimal production decisions first and then using these decisions for hedging. The seven remaining firms used historical production scenarios to predict future risk exposure.

All of the firms studied by Sanda et al. (2013) use static hedging approaches but none of them use dynamic hedging. Eight of the firms use a hedge ratio approach, their short positions in the financial market being required to 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. The remaining two have no written hedging policy. Less than half of the firms include options and EPAD contracts in their hedging strategy. Futures contracts are the most widely used derivatives in terms of hedged volume in 11 out of 12 firms. On average, approximately 90% of the traded derivative volume [MWh] of the firms was quarterly and annual contracts. None of the firms studied used a dynamic approach as hedging policy.

Stulz (1996) shows that hedging practices of many companies are influenced by their market views. These companies therefore take speculative positions in their portfolios, which Stulz (1996) refers to as *selective hedging*. Sanda et al. (2013) shows that many of the studied hydropower producers use selective hedging. The authors even observe negative hedge ratios in some companies. Adam and Fernando (2006) separate hedging into two components, *predictive* and *selective* hedging. Predictive hedging, unlike selective or speculative hedging, focuses on the hedging of predicted cash flows based on the firm's fundamental operations. In this article, we only focus on predictive hedging and ignore possible opportunities for speculative trades.

3. Model Formulation

In this section, we formulate the decision problems associated with production planning and dynamic hedging as Markov decision processes.

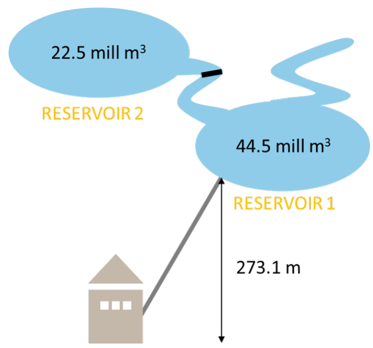

Figure 1: Topological setup, reservoir capacities and forebay elevation of the hydropower plant

3.1. Assumptions

We consider the problem faced by a price-taking hydropower producer who operates a single hydropower plant with substantial reservoir capacity. The producer participates in a deregulated electricity market and hedges long positions in physical production with short positions in financial futures contracts. The producer's objective is to find an asset-backed hedging strategy that matches risk preferences and accounts for risk in natural inflows, prices, exchange rates, as well as management decisions.

Although the model is generic, we consider an actual plant located in the NO3 price area in Norway (there are five price areas in Norway, NO3 being located in the center of the country). The plant consists of two interconnected reservoirs and one power station with one turbine. The plant is mid-sized in terms of reservoir size relative to annual inflow and in terms of generation capacity, with a mean annual production of 191 GWh. Figure 1 illustrates some relevant properties of the plant. Following Wallace and Fleten (2003), we model production planning and financial hedging as consecutive decision problems. We formulate each problem as a Markov decision processes (MDP), which is line with the literature on medium-term reservoir management (e.g. Lamond and Boukhtouta (1996)) as well as real option management of commodity storage (e.g. Nadarajah et al. (2015)).

In line with Bjerksund et al. (2011), we assume that the decision-maker participates in a complete market with no risk-free arbitrage. This means that all state transition probabilities can be represented by a unique martingale (risk-neutral) measure $\mathbb{Q}$. We use a time horizon of approximately two years for both, the production model and the hedging model, which is in line with previous models on medium-term hydropower planning (Wolfgang et al. 2009, Abgottspon and Andersson 2014) as well as models of hydropower risk management (Fleten et al. 2002, 2010)). We use semi-monthly time granularity so that each time stage is approximately two weeks (see Appendix A for how time is discretized). A semi-monthly resolution was chosen so that discrete stages coincide with delivery periods of monthly, quarterly and yearly futures contracts traded on NASDAQ OMX. Semi-monthly periods are also close enough to the time intervals that are typically used in medium-term hydropower planning.

Multiple random variables impact the decisions of the hydropower producer at each time stage. Random variables in the production model are spot price $F_{t,t}$ and reservoir inflows $Y_{1,t}$, $Y_{2,t}$.

In the hedging model, we allow the producer to trade in monthly, quarterly and yearly power futures contracts with delivery periods within the chosen time horizon of $\hat{T} = 49$ semi-months. These contracts have been chosen because of their high liquidity on NASDAQ OMX and because they are the most common derivatives used by Norwegian hydropower producers (Sanda et al. 2013). Random variables in the hedging model are therefore prices of six monthly, eight quarterly, and one yearly power futures, 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. EPAD contracts are not available for the NO3 area, and are therefore not included in the model. Due to high correlation between system and NO3 spot price, we disregard the area price difference, and assume that the producer receives the system spot price and not the area spot price. The hedging model also includes random variables for spot exchange rate ($Q_{t,t}$) and forward exchange rate at time t, for maturity at time T ($Q_{t,T}$).

Production decisions enter the hedging model as exogenous random variables, W_t. The value of W_t is found using the production model. We also include taxation effects, γ_c and γ_r denoting the corporate and resource rent tax rate. And finally, we include variable transaction costs c_F for trading in the power futures market. Transactions cost are negligible in the currency forwards market.

We propose to use the nested conditional value-at-risk (nested CVaR) to model the risk preferences of the producer. *Terminal* CVaR measures the CVaR of the terminal cash flow. As cash flows depend on realizations of randomness, risk preferences are state-dependent and hence not consistent across time. Time consistency can be restored by applying the *Nested* CVaR as risk measure which only considers the CVaR of possible future cash flows.

We define the function $\psi_{\lambda,\alpha}(X) = \lambda CVaR_\alpha(X) + (1-\lambda)\mathbb{E}(X)$ for a random variable X, significance level α and weight λ. The $CVaR_\alpha(X)$ is defined in Section 2.2. Following Shapiro et al. (2013), the nested CVaR for a sequence of random variables $X_1, X_2, X_3, ...$ is given by

$$CVaR_{\alpha,\lambda}^{NEST}(X_1, X_2, X_3, ...) = X_1 + \psi_{\alpha,\lambda}(X_2 + \psi_{\alpha,\lambda}(X_3 + ...)). \tag{1}$$

In Section 3.2 and 3.3, we formulate both, the production problem and the hedging problem, as Markov decision processes (MDPs).

3.2. Production planning problem

The objective of the production planner is to maximize the expected discounted terminal cash flows from power production and reservoir control. Denote w_t as power production decision [MWh] at time t. In line with the literature on hydropower planning, cash flows earned by the producer equal rewards earned from physical sales while variable start-up cost are ignored (Wallace and Fleten 2003). Denote $F_{t,t} \cdot w_t$ as revenue at time t and β_t is a time-dependent discount factor. Then, the value function of the MDP is given by

$$V_t^P = F_{t,t} w_t (1 - \gamma_c - \gamma_r) + \beta_t \mathbb{E}[V_{t+1}^P \mid F_{t,t}, Y_{1,t}, Y_{2,t}, \pi_t]. \tag{2}$$

Here, π_t denotes the decision policy at time t. Cash flows from production are subject to both resource rent and corporate tax. Denote $v_{1,t}$ and $v_{2,t}$ as reservoir volume $[m^3]$, $s_{c,t}$ as the amount of water flowing from reservoir 2 into reservoir 1, $s_{s,t}$ as the amount of spilled water, and κ as constant generation efficiency $[MWh/m^3]$. Using this notation, the volume balance in each reservoir is given by

$$v_{1,t} = v_{1,t-1} - w_t \cdot \kappa^{-1} + s_{c,t} + Y_{1,t} - s_{s,t} \tag{3}$$

$$v_{2,t} = v_{2,t-1} + Y_{2,t} - s_{c,t} \tag{4}$$

Treating κ as a constant corresponds to what is commonly assumed in long-term reservoir management, for example, in software such as EOPS that is widely used for medium-term production planning in the Nordic countries (SINTEF 2017). Both reservoirs have upper bounds. Reservoir 2 additionally has a time-dependent lower bound specified in the operating license issued by the regulator.

$$\begin{aligned} v_{1,t} \leq \overline{v_1}, \quad v_{2,t} \leq \overline{v_2} \\ v_{1,t} \geq 0, \quad v_{2,t} + v_{2,t}^S \geq \underline{v_{2,t}} \end{aligned} \tag{5}$$

To avoid running into infeasibilities, we introducing slack variable $v_{2,t}^S$. Violating this constraint will result in a penalty cost given by $\upsilon \cdot v_{2,t}^S$, which is added to the objective function.

Production capacity (6) is determined by the maximum flow rate $[m^3/s]$ of the plant's turbines, ξ, and the number of seconds in a semi-month, ς_t,

$$w_t \leq \overline{w_t} = \xi \cdot \kappa \cdot \varsigma_t. \tag{6}$$

The complete production planning problem at time t is given by

$$\begin{aligned} \max \quad & V_t^P(F_{t,t}, Y_{1,t}, Y_{2,t}, \pi_t) - \upsilon \cdot v_{2,t}^S \\ \text{subject to} \quad & (3), (4), (5), (6) \end{aligned}$$

3.3. Hedging problem

The objective of the dynamic hedging problem is to control risk exposure, which requires tracking trades in currency and power derivatives. We include variables and balance constraints for tracking financial short positions and committed future cash flows, which reflects their actual payoff structure. The model does not consider long positions in currency and power futures. Allowing only short positions should be sufficient for hedging purposes, as the producer has a long position in physical production.

We assume that the exchange rate forward contract is settled at contract maturity. Figure 2 illustrates an example of the cash flows from a currency forward contract with delivery in time 8. No cash is exchanged until the maturity date, when the difference in the forward and spot price is settled. We denote $z_{t,T}$ as the producer's total short position [EUR] at time t in currency forwards with maturity at time T. New short positions that enter at stage t for delivery in T are denoted by $x_{t,T}$. This gives the following balance constraint for $t < T$, where $T \leq \hat{T}$,

$$z_{t,T} = z_{t-1,T} + x_{t,T}, \quad T > t. \tag{7}$$

A forward contract with instantaneous delivery is a spot trade, and its balance constraint for $T = t$ is given by

$$z_{t,t} = z_{t-1,t}, \quad T = t. \tag{8}$$

The currency forward rate is denoted by $Q_{t,T}$ [NOK/EUR]. We let $y_{t,T}^C$ [NOK] denote the committed, 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 committed part

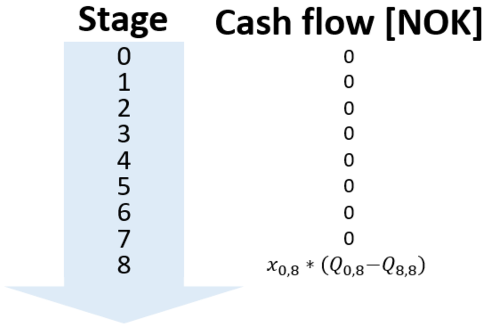

Figure 2: Cash flows from short position $x_{0,8}$ in currency forward contract with maturity $T = 8$ made at $t = 0$.

of the currency cash flows is given by (9). Note that cash flows from forward trading are only subject to corporate tax rate γ_c. This contrasts cash flows from physical production, which are subject to both resource rent tax rate γ_r and corporate tax rate γ_c,

$$y^C_{t,T} = y^C_{t-1,T} + x_{t,T} Q_{t,T}(1 - \gamma_c), \quad T > t. \tag{9}$$

Here, the decision variable $y^C_{t,T}$ can be interpreted as being the positive part of the cash flows that occur at maturity time T, as illustrated in Figure 2. At maturity in $t = T$, the time t cash flows from currency hedging are given by

$$y^C_{t,t} = y^C_{t-1,t} - z_{t,t} Q_{t,t}(1 - \gamma_c), \quad T = t \tag{10}$$

Here, negative cash flows from forward positions are added to the positive cash flows to obtain the time t cash flows $y^C_{t,t}$. Note that $y^C_{t,t}$ can take both positive and negative values.

Power futures have a more complex cash flow structure than currency forwards, since delivery periods are typically monthly, quarterly, or annual delivery bands for power, instead of delivery at a specific point in time. Also power futures use daily settlement. Before delivery, settlement is based on the price change between two successive trading days, and during delivery it 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, under the condition of semi-monthly settlement. Figure 3 shows an example of the cash flows of a quarterly contract with delivery period from stage 3 to 8. These are then multiplied by their respective spot exchange rate $Q_{t,t}$ to give cash flows in NOK.

The exposition of contract and cash flow dynamics for power futures is new to the literature,

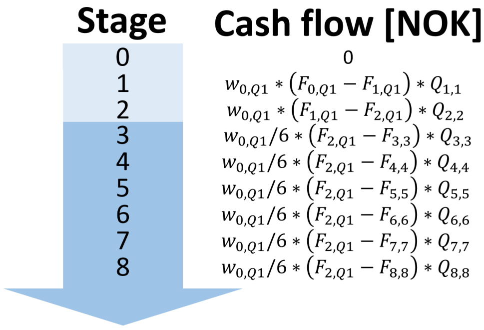

Figure 3: Cash flows from short position $w_{0,Q1}$ in a power futures contract with delivery in the upcoming quarter. The light blue part of the figure denotes time stages prior to the start of the delivery period, the darker part denoting stages within the delivery period. During the delivery period, the quantity of the short position is divided by 6, as this is the number of stages covered by the contract delivery period.

nevertheless, it is tedious, and we refer to Appendix B for details. We end up with expressions for the dynamics of futures positions as well as the stage t cash flows from power trading ($y^F_{t,t}$). It has a similar but more complex structure as the cash flows from currency forwards.

For the hedging problem, we define the value function, V^H_t, as a linear combination of the stage t cash flows and $\psi_{\alpha,\lambda}(X)$. Function $\psi_{\alpha,\lambda}(X)$ is a risk measure with risk preference weighting λ and quantile α, as defined in section 3.1. Stage t cash flows aggregate cash flows from currency forward trading, power futures trading, and spot production. Recall that, in the hedging problem, production W_t is an exogenous random variable.

$$V^H_t = y^C_{t,t} + y^F_{t,t} Q_{t,t} + W_t F_{t,t} Q_{t,t} (1 - \gamma_c - \gamma_r) + \beta_t \psi_{\alpha,\lambda}[V^H_{t+1} \mid F_{t,t}, Q_{t,t}, W_t, \pi_t] \tag{11}$$

Earlier, we explained how λ adjusts the weighting of CVaR and expected value of a random variable X. In our case, X is the next stage value function V^H_{t+1}. Setting $\lambda = 0$ is equivalent to maximizing the expected value, while setting $\lambda = 1$ involves maximization of the stage t cash flows and the CVaR of V^H_{t+1}. Setting $\lambda \neq 1$ makes little sense in the hedging problem with exogenous production W_t, as the expected profit from trading forward contracts is zero under the risk-neutral probability measure. Hence, we will use $\lambda = 1$ to solve the hedging problem when W_t is given. Adjusting λ does, however, make sense in a problem that models production planning and hedging simultaneously. Such a model can be formulated by combining constraints and decision variables of the production problem with those of the hedging problems, and replacing random variable W_t with decision variable w_t.

4. Solution Method

The two problems defined in the previous section are both discrete-time, continuous-state MDPs. As it is generally not possible to solve such problems exactly, we are going to use approximate dual dynamic programming (ADDP) to approximate the optimal policy of the continuous-state problem (Löhndorf et al. 2013, Löhndorf and Shapiro 2018, Löhndorf and Wozabal 2019). ADDP exploits the property that the state space of the MDP can be separated into a *resource state*, x_t, which is defined by the decision process (e.g., reservoir contents, cash balance, contract position), and an *environmental state*, ξ_t, which evolves independently of the decision process (e.g., futures prices, natural inflows, or production in the hedging problem).

Given that ξ_t is Markovian, $\mathcal{P}_t(F_t|F_2,\ldots,F_{t-1}) = \mathcal{P}_t(F_t|F_{t-1})$, dynamic programming equations can be written as

$$V_t(x_{t-1},\xi_t) = \max_{x_t\in\mathcal{X}_t(x_{t-1},\xi_t)} \left\{R_t(x_t,\xi_t) + \mathcal{V}_{t+1}(x_t,\xi_t)\right\},\ t=2,...,T, \tag{12}$$

where $R_t(\cdot,\cdot)$ is the (convex) immediate reward function and

$$\mathcal{V}_{t+1}(x_t,\xi_t) = \mathbb{E}_{\xi_{t+1}|\xi_t}\left[V_{t+1}(x_t,\xi_{t+1})\right],\ t=2,...,T, \tag{13}$$

with $\mathcal{V}_{T+1}\equiv 0$. We refer to functions $V_t(\cdot,\cdot)$ as value functions and $\mathcal{V}_t(\cdot,\cdot)$ as expected value functions.

The value functions, $V_t(x_{t-1},\xi_t)$, are concave in x_{t-1} if $\max\{R_t(x_t,\xi_t)+\mathcal{V}_{t+1}(x_t,\xi_t)|x_t\in\mathcal{X}_t(x_t,\xi_{t-1})\}$ are concave in x_{t-1}, which can be shown by induction in t going backwards in time. This holds, for example, for linear programming problems, but not for mixed-integer problems.

ADDP in a first step constructs a discrete scenario lattice of the environmental state by separating the outcome space of random variable ξ_t into a finite number of disjoint partitions. In the second step, ADDP approximates the value functions that are associated with the means of each partition from above using cutting planes. The resulting approximate value functions are piecewise constant in realizations of ξ_t and piecewise linear in the x_t and provide a policy for the continuous-state problem.

To construct the scenario lattice, at each stage $t=2,\ldots,T$, ADDP tries to find partition means, μ_{ti}, $i=1,...,K_t$, which are solutions of the following problem

$$\min_{\mu_{t1},\ldots,\mu_{tK_t}} \int \min_{i=1,\ldots,K_t} \|\xi_t-\mu_{ti}\|_p^p dP(\xi_t), \tag{14}$$

where $\|\cdot\|_p$ denotes the p-norm with $p \geq 1$. Given a solution $\mu_{t1}, ..., \mu_{tK_t}$ of (14), the space is partitioned into subsets Γ_{ti} consisting of points which are closest to μ_{ti}. The corresponding sets

$$\Gamma_{ti} := \{\xi_t : \|\xi_t - \mu_{ti}\|_p \leq \|\xi_t - \mu_{tj}\|_p,\ j = 1, ..., K_t,\ j \neq i\}, \tag{15}$$

represent the Voronoi partition associated with respective means μ_{ti}, $i = 1, ..., K_t$, which serve as the nodes of the lattice in t. Hence, we have $\mu_{ti} = \mathbb{E}[\xi_t | \xi_t \in \Gamma_{ti}]$.

As finding optimal means is an $\mathcal{NP}$-hard problem, Löhndorf and Wozabal (2019) propose to use stochastic approximation method which we will adapt here. The method draws S random sequences $(\hat{F}\xi_t^s)_{t=1}^T$, $s = 1, ..., S$, from the continuous process. With $(\beta_s)_{s=1}^S$ being a sequence of step sizes with $0 \leq \beta_s \leq 1$, $s = 1, ..., S$, and setting $p = 2$, means are updated recursively using

$$\mu_{ti}^s := \begin{cases} \mu_{ti}^{s-1} + \beta_k \left(\hat{\xi}_t^s - \mu_{ti}^{s-1}\right) & \text{if } i = \operatorname{argmin}\left\{||\hat{\xi}_t^s - \mu_{tk}^{s-1}||^2,\ k = 1, ..., K_t\right\}, \\ \mu_{ti}^{s-1} & \text{otherwise,} \end{cases} \tag{16}$$

with $\mu_{0i}^{s-1} \equiv 0$, for $i = 1, ..., K_t$, $t = 2, \ldots, T$, $s = 1, ..., S$. It can be shown that if the sequence $(\beta_s)_{s=1}^S$ fulfills $\sum_{s=1}^{\infty} \beta_s = \infty$ and $\sum_{s=1}^{\infty} \beta_s^2 < \infty$, then the means converge to local minimizers of (14).

To estimate the transition probabilities between partitions at subsequent stages, we use the backwards estimation procedure proposed in Löhndorf and Wozabal (2019), which ensures that the discrete approximation of the stochastic process is a martingale.

With a given lattice, we can now compute an upper bound of the value functions using cutting planes by extending the SDDP method of Pereira and Pinto (1991) to scenario lattices, also referred to as Markov-Chain-SDDP (Löhndorf and Shapiro 2018).

Denote $\hat{V}_{t+1,n}(x_t, F_{t+1})$ as approximate value functions and $\hat{\mathcal{V}}_{tn}(x_t)$ as their expectations with $\hat{\mathcal{V}}_{t,0} \equiv 0$. At each iteration, $n = 1, ..., N$, we let ADDP draw one scenario from the lattice process, $(\hat{\xi}_t^n)_{t=1}^{T-1}$, to generate a sequence of sample decisions,

$$\hat{x}_t^n = \operatorname{argmax}_{x_t \in \mathcal{X}_t(\hat{x}_{t-1}^n, \hat{\xi}_t^n)} \left\{R(x_t, \hat{\xi}_t^n) + \hat{\mathcal{V}}_{t,n-1}(x_t, \hat{\xi}_t^n)\right\},\ t = 2, ..., T-1, \tag{17}$$

with $\hat{x}_1^n = \operatorname{argmax}\{R(x_1) + \hat{\mathcal{V}}_{t,n-1}(x_1) | x_1 \in \mathcal{X}_1\}$.

Note that a solution of the discretization only provides a policy for the set of scenarios given by the constructed lattice with nodes μ_{ti}, $i = 1, ..., K_t$, $t = 1, ..., T-1$, so that the respective upper bound is only valid for the *discretized* problem but not for the continuous one.

For a given set of sample decision, the algorithm can now construct cutting planes by recursively solving

$$\hat{V}_{tn}(\hat{x}_{t-1}^n, \mu_{ti}) = \max_{x_t \in \mathcal{X}_t(\hat{x}_{t-1}^n, \hat{\xi}_t^n)} \left\{ R(x_t, \hat{\xi}_t^n) + \hat{\mathcal{V}}_{tn}(x_t, \mu_{ti}) \right\}, \tag{18}$$

with $\hat{\mathcal{V}}_{Tn}(x_T, \mu_{Ti}) \equiv 0$ for $i = 1, ..., K_t$, $t = T, ..., 2$, $n = 1, ..., N$.

Since new cutting planes tighten the value function in regions that get explored under the current policy, it can be shown that the algorithm converges to the optimal policy in a finite number of steps if the problem is linear (Philpott and Guan 2008, Guigues 2016).

To approximate the value function under the Nested CVaR, we adopt the method described in Shapiro et al. (2013) who propose to change the probability weights of the cutting planes. We refer to Shapiro et al. (2013), Philpott et al. (2013), or Löhndorf and Wozabal (2019) for a detailed description.

5. Risk factor dynamics

In this section, we describe how we model the dynamics of the risk factors as stochastic processes. We include random variables for natural inflows, currency spot and forward rates, as well as for electricity spot and futures prices.

5.1. Price process

We use the Heath-Jarrow-Morton framework to model the evolution of the electricity price forward curve (Heath et al. 1992). Tradable forward and future contracts in the Nordic power market are not for delivery at a single point in time. Instead, they have delivery periods that stretch over months, quarters, and years. By contrast, a price forward curve estimates the forward price for delivery at specific points in time. The value of the curve at time T is therefore the price of a forward contract with delivery exactly at time T. To achieve this, we use the method of Fleten and Lemming (2003) to construct semi-monthly price forward curves. See Appendix C for more details.

In line with 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 $T - t = \tau$. The process that explains movements in the forward curve is then given by

$$\begin{aligned}\frac{dF_{t,T}}{F_{t,T}} &= \sigma_{t,T} dZ_{t,T} = \sigma_\tau dZ_{\tau,t} \\ \mathbb{E}(dZ_{\tau,t}, dZ_{\hat{\tau},t}) &= \rho_{\tau,\hat{\tau}} dt, \quad \tau, \hat{\tau} \in [\tau].\end{aligned} \tag{19}$$

See Appendix E for a more extensive analysis and details about the implemented price process.

5.2. Inflow process

We treat inflows into the reservoirs as dependent random variables $Y_t = \zeta Y_{1,t} + (1-\zeta) Y_{2,t}$, where ζ is constant, mostly because available data only contains aggregated inflows. We then fit the geometric periodic autoregressive (GPAR) model discussed in Shapiro et al. (2013) to the inflow data for the hydropower plant. Shapiro et al. (2013) showed that a first-order periodic autoregressive process provides a good fit to seasonal inflow data from Brazil, and we find that the model is also suitable for our case. See Appendix F for a detailed analysis as well as process parameters.

5.3. Exchange rate process

We follow Huchzermeier and Cohen (1996) and model exchange rate dynamics as geometric Brownian motion with drift equal to the interest rate difference,

$$\frac{dQ_{t,t}}{Q_{t,t}} = (r - r_f) dt + \sigma_Q dZ, \tag{20}$$

where σ_Q is the annualized volatility of exchange rate returns. All fitted model parameters can be found in Appendix G.

Since uncertainty in the EURNOK exchange rate forward curve is assumed to originate from the spot exchange rate, a single factor model for the currency spot rate is sufficient. Once the spot rate is known, the forward rate is given by covered interest rate parity (Aliber 1973). This is in contrast to electricity markets, where there is uncertainty both in the spot price and the forward curve.

5.4. Lattice Generation

Since inflows and electricity prices exhibit significant correlation, we generate a joint lattice of inflows and forward prices with 100 nodes per stage. Each node contains state variables for natural inflow as well as spot and 48 forward prices of electricity price. As correlation between these variables and shifts in exchange rate are independent, the currency lattice is generated

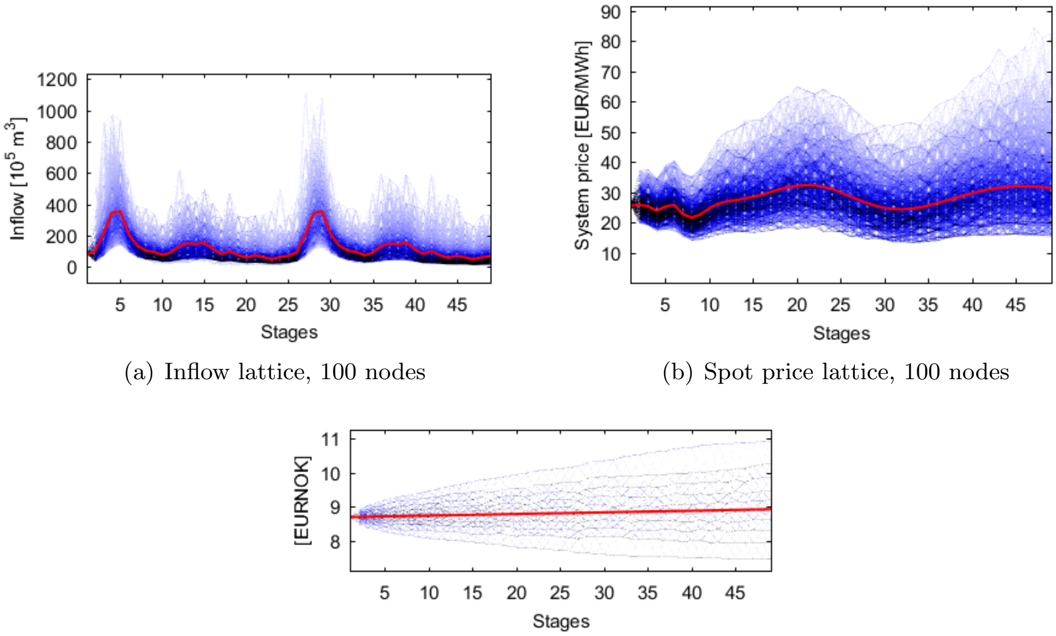

(a) Inflow lattice, 100 nodes

(b) Spot price lattice, 100 nodes

(c) Currency lattice, 10 nodes

Figure 4: Simulated time series (blue: individual scenarios, red: sample average

separately with 10 nodes per stage. To obtain a combined lattice of currency, inflow, and prices, we compute the cross-product of the price-inflow lattice and the currency lattice resulting in a final lattice that spans 49 stages and has 1000 nodes per stage. Figure 4 shows sample paths for inflow, spot price and currency taken from simulated state transitions of the scenario lattices.

We use March 16, 2015, as the starting date in our case study. This results in a model with an ultimate stage that is the end of March 2017, which is usually right before the start of the spring precipitation in Norway, and at a point in time when many reservoirs are close to empty. We therefore disregard constraints on terminal reservoir levels in the production model.

As starting values which define the root node of the lattice, we use the mean inflow recorded for the second half of March 2015 and the currency rate at this date, and the initial forward curve was constructed using the close prices of all monthly, quarterly and yearly contracts on this date.

Production in period t, w_t is a decision variable in the hydropower scheduling problem. For the hedging problem, we treat production as an exogenous, stochastic variable W_t. To obtain the stochastic production levels, we first compute the optimal production decisions by solving the production planning problem using the 100-node lattice. In a second step, we generate a sample of 10^5 production decisions by simulating the optimal policy and then insert production

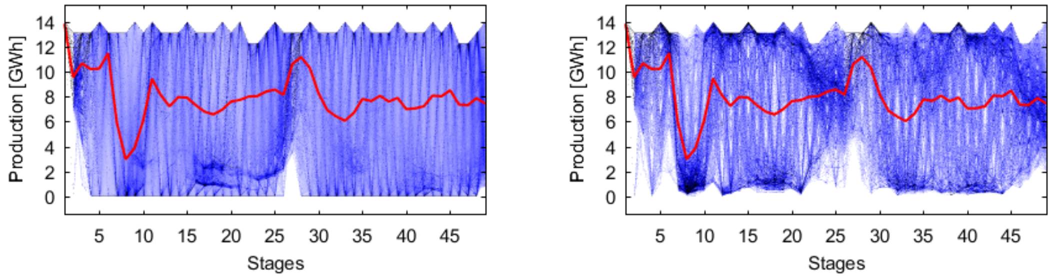

(a) Simulated production decisions

(b) Discretized production decisions

Figure 5: Simulated time series (blue: individual scenarios, red: sample average

as exogenous state variable W_t into the lattice by averaging all production decisions that are made at a particular node. In this way, we obtain a unique stochastic production level for all price-inflow states in the lattice.

Figures 5(a) and 5(b) show the simulated and discretized production decisions from the production model. As we can see, production means as well as the main structure of the process remain intact, which is demonstrates strong relationship between state variables and production decisions.

6. Numerical results

In this section, we summarize the numerical results of the hedging model, where production is treated as an exogenous random variable unless otherwise stated. The analysis uses the scenario lattices presented in Section 5.4 and the coefficient values in Appendix J. ADDP is run for 500 iterations which ensured convergence of all instances, and all simulated results are based on 10^5 simulations[3]. All code has been implemented in MATLAB and R. We used the QUASAR MATLAB package to model and solve the optimization problems (Löhndorf 2017). All computations were done on a computer with 32 GB memory and 3.6 GHz CPU speed.

6.1. Hedging performance

We begin our discussion with the hedging model. We are particularly interested in the risk reduction that hedging can provide, compared with no hedging. We also seek to explain the decision policies driving this reduction. Statistical measures of the discounted terminal cash

[3]Sensitivity analysis of results obtained using 500 iterations and 10^5 simulations can be found in Appendix K.

Available contracts	None	All	All	All except monthly	All except monthly
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 production per year	192.18	192.17	192.08	192.22	192.15

Table 1: 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.

flows over the two-year horizon were used to quantify the hedging effect. We found the mean and standard deviation of the terminal discounted cash flows and the terminal VaR and CVaR at the significance levels 0.05 and 0.01 for different model specifications. In Table 1, we show the results for the cases of trading in all contracts, of trading in all contracts except monthly power futures, and of no forward trading and only production. Recall that we use the nested CVaR to model risk preferences. This is different from the terminal CVaR. As it is not possible to quantify the nested CVaR using simulation (Philpott et al. 2013), we will resort to the terminal CVaR to measure risk.

The terminal VaR and CVaR in Table 1 show that dynamic hedging clearly reduces risk, when compared with the no hedging case, by utilizing currency forwards and power futures contracts. The mean terminal cash flow is slightly reduced, which can be largely explained by transaction costs. Risk reduction in terms of terminal risk measures is also evident when using only quarterly and annual contracts. Terminal cash flows are higher where trading in monthly contracts is allowed. Note that the α parameter that sets the quantile of the *nested* CVaR does not necessarily maximize the corresponding *terminal* CVaR. This can be observed by comparing the results for $\alpha = 0.05$ with the results for $\alpha = 0.1$. Letting $\alpha = 0.1$ performs better both in terms of CVaR(5%) and CVaR(1%).

Average optimal hedge ratios for power future trading depends highly on the availability of monthly contracts. Table 2 shows the mean total hedge ratio of the four model variants. This is the mean ratio of the total short position in financial power contracts for the total amount of produced energy. The table includes the mean hedge ratio before the start of a new quarter and before the start of the first year. The large difference between quarterly and total hedge ratios, for all hedging instruments, indicates that a large volume of monthly contracts is traded in the last three months prior to delivery. The model prefers monthly contracts, as short positions in monthly contracts account for 97% of total trading volume [GWh] where $\alpha = 0.10$. This could

Available contracts	All	All	All except monthly	All except monthly
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 2: 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.

be because these contracts permit precision hedging of expected production in a given month. They also permit the most near-term cash flows to be hedged.

According to Sanda et al. (2013), a full hedge be obtained with a hedge ratio of 0.53. This is due to taxation effects[4]. Compared to a hedge ratio of 0.53, the case allowing for trades in all contracts proposes a significant over-hedge. This over-hedging might be an effect of modeling risk preferences using the nested CVaR. Recall that 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. Hence, we assume that there is a positive correlation between spot price and production in a semi-month t, which means that production level will be higher for high-price states than for low-price states. The objective of the model is to maximize the nested CVaR. The model will therefore seek to increase the lower-tail cash flows as much as possible. A possible explanation for the over-hedge could therefore be that the model seeks to obtain a hedge ratio which raises the total cash flows from production and hedging as high as possible in the states where the price goes down. This increases the CVaR and simultaneously keeps total cash flows in states where the price rises above the lower-tail cash flows[5]. As shown in Table 2, the largest volume of over-hedged short positions enters less than three months before maturity. The volatility of the cash flows from the over-hedged position will therefore be moderate, due to the short delivery period and low time to maturity of these contracts.

The extent of trading in currency forwards is shown in Table 3. We can see that the currency market is used extensively. Currency forwards with shorter maturity times are preferred, which is similar to the trading activity in power futures. However, the amount of trading in contracts with longer times to maturity 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

[4]The derivation of this hedge ratio can be found in Appendix L.

[5]A numerical example explaining this argument can be found in Appendix M.

Available contracts	All	All	All except monthly	All except monthly
α	0.05	0.1	0.05	0.1
Hedge ratio	0.563	1.120	0.589	0.950
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: 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.

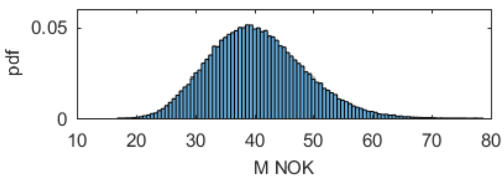

(a) Discounted terminal cash flows from production

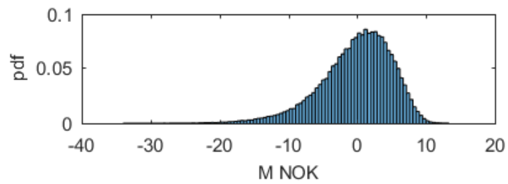

(b) Discounted terminal cash flows from power futures trading

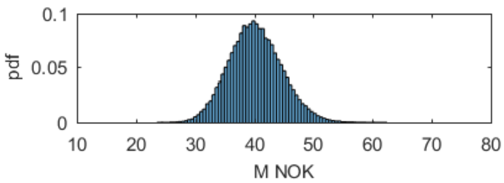

(c) Total discounted terminal cash flows (after hedging)

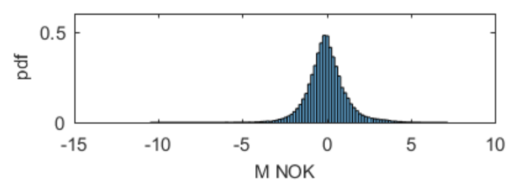

(d) Discounted terminal cash flows from currency forward trading

Figure 6: Cash flow distributions. From hedging model with all contracts and $\alpha = 0.10$

currency that was not included in the final scenario lattices.

Figure 6 plots the distributions of terminal cash flows from production, terminal cash flows after hedging, terminal cash flows from power futures and terminal cash flows from currency forwards, when $\alpha = 0.1$. The hedging effect is obvious when we compare the terminal cash flow distribution in Figure 6(a) with that in 6(c). Hedging makes the cash flow distribution narrower. Also note that means in Figure 6(b) and 6(d) are zero, which is confirms that there are no expected gains from trading in the forward markets.

6.2. Comparison of sequential and simultaneous approach

So far, we have considered production and hedging decisions sequentially. Let us now compare what happens when both decisions are made simultaneously in a single integrated model. We test the simultaneous approach for different values of λ and α and then studied the hedging performance of each variant based on the risk measures given in Table 1. The results in Table

Model type	Simultaneous	Simultaneous	Simultaneous	Simultaneous	Sequential
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 4: 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].

4 show that the simultaneous approach is at par with the sequential approach, which is in line with Wallace and Fleten (2003). The ratio of mean cash flows to production is quite similar in the simultaneous and the sequential approach. Risk-aversion, however, results in a lower mean total production in the simultaneous approach. This further leads to lower mean discounted cash flows, terminal VaR and CVaR. The standard deviation of the simultaneous approach is also higher.

6.3. Effect of currency hedging

One of the main contributions of this article is the inclusion of currency risk and currency derivatives in a real options model of hydropower production. We compare hedging performance where currency forwards can be traded with performance where currency derivatives cannot be traded, to quantify the effect of currency hedging. We therefore restrict hedging activity to power futures contracts and compare this to where currency forward trading is allowed. Note that we also include currency risk when currency hedging is not allowed. We formulate the problem such that the optimal decision policies are independent of the uncertain currency rate. The results of this analysis are summarized in Table 5. The table shows that the standard deviation of the terminal cash flows decreases when currency hedging is introduced, and terminal VaR and CVaR increase. Currency hedging does reduce the market risk of the hydropower plant. The magnitude of this change, however, is small: less than 10% for all risk measures. This can be explained by the low volatility of the currency exchange process compared with that of the electricity price process.

6.4. Comparison with heuristics

Wang et al. (2015) show that simple hedging strategies might yield a better or equivalent hedging performance than more advanced procedures. We therefore replicate the hedging strat-

α	**W/o currency hedging** 0.05	**W/o currency hedging** 0.1	**W/ currency hedging** 0.1	**Change**
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 5: Effect of currency hedging on terminal cash flows. The last column denotes the percent difference between the risk measures with and without currency hedging at $\alpha = 0.1$ in the nested CVaR expression.

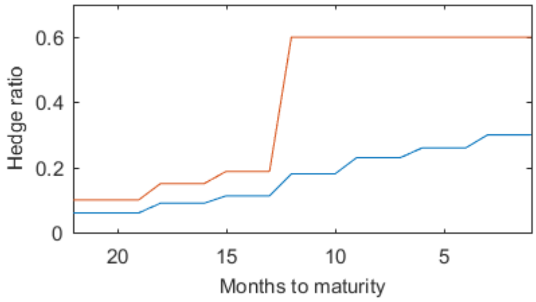

Figure 7: Required hedge ratio range of a real company. For example, in a given time stage, the expected hedge ratio (hedged volume divided by expected production) for the month that begins 12 months ahead must lie between 0.18 and 0.6.

egy of a Norwegian hydropower producer, and compare its performance to the performance of our model. Sanda et al. (2013) show that most Norwegian companies use a heuristic approach, with specific hedge ratio ranges for different times to maturity. The time horizon of our model is two years. We therefore use the hedging strategy of a firm whose hedging activity also begins two years prior to maturity. Figure 7 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).

To test the performance of the heuristic approach, we include the hedge ratio ranges as constraints on financial short positions in the dynamic model. We perform two separate simulations; one where the hedge ratio is set close to the lower bound and one where it is closer to the upper bound. The currency hedging strategy of the firm is unknown. The option to trade currency derivatives has therefore been removed. The results are displayed in Table 6. The results show that, from the perspective of maximizing terminal CVaR, a strategy close to the upper bound of the hedge ratio is better than a strategy close to the lower bound. The results also show that the heuristic approach performs worse in terms of terminal risk measures than the proposed model that allows for all types of trading, although the difference is small. The performance is nearly identical to the version omitting currency hedging for all risk measures

Model type	Heuristic lower	Heuristic upper	Sequential	Sequential
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 6: 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 2.

except variance. This indicates that the over-hedging proposed by our model, in which the largest short positions are entered just before maturity, has only minor effects on the lower tail of the terminal discounted cash flow distribution.

7. Conclusions

We present a global dynamic model for risk management of hydropower producers from the perspective of a Norwegian power company. We model the hedging problem as a sequential problem in which optimal production decisions are made first and hedging cash flows from production using currency forwards and electricity futures is done in a second step. Both, the production planning problem and the hedging problem, are modeled as Markov decision processes. We model electricity spot prices, futures prices, and natural inflows as correlated stochastic processes, whereas currency exchange rate is assumed to be independent. The decision maker is allowed to trade in currency forward contracts and monthly, quarterly, and annual power futures to reduce risk. Risk preferences are modeled using the nested CVaR. We use ADDP to solve the underlying stochastic-dynamic decision problem. The results show that the dynamic hedging model substantially reduces risk which is demonstrated by a 23% increase of the 5%-CVaR of terminal cash flows.

We find a moderate effect of including currency derivatives on the hedging strategy. Including currency derivatives results in decrease in variance of 9.96% and an increase in CVaR(5%) of 2.43% for the terminal cash flows. We find that there is a negative albeit weak 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 and ranges from -0.28 to -0.14 for time to maturity $\tau \geq 6$ semi-months. These values are not large and were found to be insignificant at a 5% significance level for $\tau \leq 28$ semi-months. However, they hint the existence of a weak, natural hedging effect.

We also compared the sequential approach with optimizing production and hedging simultaneously. We find that the sequential outperforms the simultaneous approach both in terms of returns and risk. The sequential approach achieves higher values of mean, VaR, and CVaR of the terminal discounted cash flows. The main reason for this is that the nested CVaR does not optimize what is being measured by our risk metrics. Alternative formulations if the model using the terminal CVaR, however, did not converge and the resulting suboptimal policies were worse than those with a nested CVaR. To arrive at a definitive answer, we must either find an alternative measure of risk that measures what is being optimized or find a way to speed up convergence when a terminal CVaR is used.

We also investigate how the hedging model performs in relation to a heuristic approach that is based on hedge ratio ranges – a model that is used extensively by Norwegian hydropower producers. The performance of the hedge ratio approach is slightly worse than the base model, with a 2.8% lower CVaR(5%) but almost identical to the case where there is no currency trading. This implies that a simple hedge ratio approach can be quite efficient, which is 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. The performance of the alternative hedge ratio approach, however, suggests that the effect of over-hedging might be minor in terms of reducing the risk of terminal cash flows. This further questions the suitability of the nested CVaR to represent risk preferences, and also suggests that more research on the applicability of convex risk measures in optimization problems is needed.

Both the base model, the heuristic approach, as well as the version without currency hedging perform better in terms of terminal risk than the approach with no option to trade in monthly power futures. A possible explanation is that monthly contracts allow more precise hedging of the expected production in a given month than the contracts with longer maturity periods. More precise hedging also increases flexibility in the timing of trading due to their shorter delivery period.

References

References

Abgottspon, H., Andersson, G., 2014. Medium-term optimization of pumped hydro storage with stochastic intrastage subproblems, in: IEEE Power Systems Computation Conference.

Adam, T.R., Fernando, C.S., 2006. Hedging, speculation, and shareholder value. Journal of Financial Economics 81, 283–309.

Alexander, C., 2008. Pricing, Hedging and Trading Financial Instruments. volume III of *Market Risk Analysis*. John Wiley & Sons, Chichester, West Sussex, UK.

Aliber, R.Z., 1973. The interest rate parity theorem: A reinterpretation. Journal of Political Economy 81, 1451–1459.

Basel Committee on Banking Supervision, 2012. Core principles for effective banking supervision. URL: `https://www.bis.org/publ/bcbs230.htm`. bank for International Settlement.

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

Bjerksund, P., Stensland, G., Vagstad, F., 2011. Gas storage valuation: Price modelling v. optimization methods. The Energy Journal 32, 203–227.

Boda, K., Filar, J.A., 2006. Time consistent dynamic risk measures. Mathematical Methods of Operations Research 63, 169—186.

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

Dupuis, D., Gauthier, G., Godin, F., et al., 2016. Short-term hedging for an electricity retailer. The Energy Journal 37.

Fleten, S.E., Bråthen, E., Nissen-Meyer, S.E., 2010. Evaluation of static hedging strategies for hydropower producers in the Nordic market. The Journal of Energy Markets 3, 1–28.

Fleten, S.E., Keppo, J., Näsäkkälä, E., 2012. Risk management in electric utilities, in: Kouvelis, P., Dong, L., Boyabatli, O., Li, R. (Eds.), The Handbook of Integrated Risk Management in Global Supply Chains. 2 ed.. John Wiley & Sons, Hoboken, New Jersey.

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

Fleten, S.E., Wallace, S.W., Ziemba, W.T., 2002. Stochastic programming in energy, in: Greengard, C., Ruszczyński, A. (Eds.), Decision Making Under Uncertainty. The IMA Volumes in Mathematics and its Applications. Springer, New York. volume 128 of *IMA Volumes on Mathematics and Its Applications*, pp. 71–93.

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

Godin, F., 2016. Minimizing CVaR in global dynamic hedging with transaction costs. Quantitative Finance 16, 461–475.

Goel, A., Tanrisever, F., 2017. Financial hedging and optimal procurement policies under correlated price and demand. Production and Operations Management 26, 1924—1945.

Guigues, V., 2016. Convergence analysis of sampling-based decomposition methods for risk-averse multistage stochastic convex programs. SIAM Journal on Optimization 26, 2468–2494.

Heath, D., Jarrow, R., Morton, A., 1992. Bond pricing and the term structure of interest rates a new methodology for contingent claims valuation. Econometrica 60, 77–105.

Houmøller, A.P., 2017. Investigation of forward markets for hedging in the Danish electricity market. Technical Report. Houmoller Consulting. Middelfart, Denmark.

Huchzermeier, A., Cohen, M.A., 1996. Valuing operational flexibility under exchange rate risk. Operations Research 44, 100–113.

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

Johnson, B., Barz, G., 1999. Selecting stochastic processes for modelling electricity prices, in: Energy Modelling and the Management of Uncertainty. Risk Books.

Kiesel, R., Paraschiv, F., Sætherø, A., 2019. On the construction of hourly price forward curves for electricity prices. Computational Management Science 16, 345—369.

Koekebakker, S., Ollmar, F., 2005. Forward curve dynamics in the Nordic electricity market. Managerial Finance 31, 73–94.

Kouvelis, P., Pang, Z., Ding, Q., 2018. Integrated commodity inventory management and financial hedging: A dynamic mean-variance analysis. Production and Operations Management 27, 1052—-1073.

Lamond, B.F., Boukhtouta, A., 1996. Optimizing long-term hydro-power production using Markov decision processes. International Transactions in Operational Research 3, 223–241.

Liu, S.D., Jian, J.B., Wang, Y.Y., 2010. Optimal dynamic hedging of electricity futures based on Copula-GARCH models, in: IEEE International Conference, IEEE, Macau, China. pp. 2498–2502.

Löhndorf, N., Shapiro, A., 2018. Modeling time-dependent randomness in stochastic dual dynamic programming. European Journal of Operational Research 2, 650–661.

Löhndorf, N., Wozabal, D., Minner, S., 2013. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. Operations Research 61, 810–823.

Löhndorf, N., 2017. Quasar optimization software 2.3. URL: `http://quantego.com`.

Löhndorf, N., Wozabal, D., 2019. Gas storage valuation in incomplete markets. URL: http://www.optimization-online.org/DB_HTML/2017/02/5863.html. working paper.

Mo, B., Gjelsvik, A., Grundt, A., 2001a. Integrated risk management of hydro power scheduling and contract management. IEEE Transactions on Power Systems 16, 216–221.

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

Murphy, F., Oliveira, F.S., 2010. Developing a market-based approach to managing the US strategic petroleum reserve. European Journal of Operational Research 206, 488–495.

Nadarajah, S., Margot, F., Secomandi, N., 2015. Relaxations of approximate linear programs for the real option management of commodity storage. Management Science 61, 3054–3076.

Nasdaq Oslo ASA and Nasdaq Clearing AB, 2018. Fee list - commodity derivatives. URL: http://www.nasdaqomx.com/digitalAssets/107/107680_180301-joint--appendix-7--fee-list.pdf.

Pereira, M., Pinto, L., 1991. Multi-stage stochastic optimization applied to energy planning. Mathematical Programming 52, 359–375.

Philpott, A., Guan, Z., 2008. On the convergence of stochastic dual dynamic programming and related methods. Operations Research Letters 36, 450–455.

Philpott, A., de Matos, V., Finardi, E., 2013. On solving multistage stochastic programs with coherent risk measures. Operations Research 61, 957–970.

Rebennack, S., 2015. Combining sampling-based and scenario-based nested Benders decomposition methods: application to stochastic dual dynamic programming. Mathematical Programming 156, 343—-389.

Ruszczynski, A., Shapiro, A., 2006. Conditional risk mappings. Mathematics of Operations Research 31, 544—561.

Samuelson, P.A., 1965. Proof that properly anticipated prices fluctuate randomly. Industrial Management Review 6, 41–49.

Sanda, G.E., Olsen, E.T., Fleten, S.E., 2013. Selective hedging in hydro-based electricity companies. Energy Economics 40, 326—338.

Shapiro, A., 2009. On a time consistency concept in risk averse multistage stochastic programming. Operations Research Letters 37, 143–147.

Shapiro, A., Tekaya, W., da Costa, J.P., Soares, M.P., 2013. Risk neutral and risk averse stochastic dual dynamic programming method. European Journal of Operational Research 224, 375–391.

SINTEF, 2017. EOPS - one area power-market simulator. URL: http://www.sintef.no/en/software/eops-one-area-power-market-simulator/.

Stulz, R.M., 1996. Rethinking risk management. Journal of Applied Corporate Finance 9, 8–25.

Thorvaldsen, T., Tyssing, N., Samuelsen, J., 2018. Kraftverksbeskatning. . KPMG Norway. Oslo, Norway.

Wallace, S.W., Fleten, S.E., 2003. Stochastic programming models in energy, in: Ruszczyński, A.P., Shapiro, A. (Eds.), Handbooks in OR and MS. Elsevier Science B.V.. volume 10. chapter 10, pp. 637–677.

Wang, Y., Wu, C., Yang, L., 2015. Hedging with futures: Does anything beat the naïve hedging strategy? Management Science 61, 2870–2889.

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

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

Appendix A. Time discretization

All stochastic processes are estimated using average semi-monthly observations, with estimation windows given in Table A.7.

Table A.8 shows how the calendar year is discretized into 24 semi-months.

Appendix B. Cash flows and balancing of electricity contracts

The following details how positions in electricity futures are tracked, depending on trading, listing and de-listing, and varying contract delivery periods. We also explain the cash flows from this trading, including the effect of transaction costs, product listing, settlement before and during delivery, and taxes.

Let us introduce decision variables for short positions in power futures. Denote $u_{t,Mi}$, $u_{t,Qj}$, and $u_{t,Y1}$ [MWh] the total short position at stage t in futures contracts with delivery in i months, j quarters and 1 year. The last index denotes contracts that have not yet entered delivery. We use $u_{t,M}$, $u_{t,Q}$ and $u_{t,Y}$ [MWh] to denote short positions of contracts that are currently in delivery. Variables $w_{t,Mi}$, $w_{t,Qj}$, and $w_{t,Y1}$ [MWh] denote *new* short positions that enter at stage t. If the delivery period of a contract exceeds the model horizon $\hat{T} = 49$ semi-months, then the corresponding decision variable ($w_{t,Mi}$, $w_{t,Qj}$ or $w_{t,Y1}$) is set to zero to guarantee that no trading takes place.

	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 A.7: Data window for parameter estimation in the stochastic processes

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.8: Semi-monthly discretization. Semi-month 4 consists of 15 days in leap-years.

Balance constraints for short positions in power futures depend on whether 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, then the total short position is given by the previous stage value plus new short positions for contracts not yet in delivery.

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

When t represents the first part of a month, the contract that was 1 month ahead (M1) in $t-1$ goes into delivery, M2 becomes M1, M3 becomes M2, etc. 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, then balance constraints for month contracts are given by (B.2). The remaining relationships in (B.1) do not change, that is, they remain valid for all stages t.

$$\begin{gathered} u_{t,Mi} = u_{t-1,Mi+1} + w_{t,Mi}, \quad i = [1, ..., 5] \\ u_{t,M6} = w_{t,M6} \end{gathered} \tag{B.2}$$

Using the same logic, balance constraints for quarter contracts for stages marking the beginning of a quarter, but not a new year, are given by (B.3).

$$
\begin{aligned}
u_{t,Qj} = u_{t-1,Qj+1} + w_{t,Qj}, \quad j = [1,...,7] \\
u_{t,Q8} = w_{t,Q8}
\end{aligned}
\tag{B.3}
$$

If t represents the beginning of a year, then balance constraints for annual contracts are given by (B.4).

$$
u_{t,Y1} = 0 \tag{B.4}
$$

Having established balance constraints for short positions in power futures, let us define variables and restrictions for the portfolio of power futures. Denote $y^F_{t,t}$ as cash flow from power futures trading in t that can take positive and negative values, and is part of the value function of the hedging problem (11). The currency spot rate at which cash flows occur is not known in advance. $y^F_{t,t}$ is therefore denoted in EUR and $y^C_{t,t}$ is denoted in NOK.

Variable $y^F_{t,T}$ tracks the committed, positive part of the cash flows from power futures trading that will occur at time T. Variable $y^F_{t,T}$, where $T > t$ is not part of the value function. It is only used to store the positive part of the cash flows that will occur in subsequent periods. In (B.11), the negative part of the cash flows is added to the positive flows to obtain time t cash flows $y^F_{t,t}$. Variable $y^F_{t,t+1}$ stores two types of cash flows. The first is related to changes in the value of the portfolio of contracts not yet in delivery. All contracts are marked-to-market regularly. We therefore need to store the forward prices in stage t to calculate the price changes in stage $t+1$. The second type of cash flow is associated with contracts in delivery. Variable $y^F_{t,T}$ for $T > t+1$ only stores the positive part of cash flows associated with contracts in delivery. The longest delivery period spanned by any of the available contracts is 24 semi-months. It is therefore only necessary to define $y^F_{t,T}$ for $T = [t+1, ..., t+24]$.

We first consider the case where t represents the *first part of a month*. The cash flow balances will then be given by

$$
\begin{aligned}
y^F_{t,t+1} &= y^F_{t-1,t+1} + (1-\gamma_c)\Big(\sum_{i=1}^{6} u_{t,Mi}F_{t,Mi} + \sum_{j=1}^{8} u_{t,Qj}F_{t,Qj} + u_{t,Y1}F_{t,Y1}\Big) \\
y^F_{t,T} &= y^F_{t-1,T}, \quad t+2 \leq T \leq t+23 \\
y^F_{t,T} &= 0, \qquad\qquad T = t+24
\end{aligned}
\tag{B.5}
$$

Only $y^F_{t,t+1}$ is updated in this case. This is because the next stage is in the same month as the current stage, so that no new contracts go into delivery.

If t is the *second part of a month*, then multiple contracts can potentially enter into delivery in the upcoming stage $(t+1)$. The price and position of the futures contracts that go into delivery must thus be stored. The cash flows for the next 2, 6, or 24 periods must also be stored. How long they will be stored depends on whether the contract is monthly, quarterly, or yearly. We introduce the indicator functions $\mathbb{I}_Q$ and $\mathbb{I}_Y$ to make the formulation more compact. Function values are equal to 1 if the next stage $(t+1)$ marks the beginning of a new quarter and year, respectively, and 0 otherwise. With this definition, cash flow balances are given by

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

$$y^F_{t,t+2} = y^F_{t-1,t+2} + (1-\gamma_c)\Big(\frac{u_{t,M1}F_{t,M1}}{2} + \mathbb{I}_Q\frac{u_{t,Q1}F_{t,Q1}}{6} + \mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}\Big) \quad \text{(B.7)}$$

$$y^F_{t,t+i} = y^F_{t-1,t+i} + (1-\gamma_c)\Big(\mathbb{I}_Q\frac{u_{t,Q1}F_{t,Q1}}{6} + \mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}\Big), \quad i = [3, ..., 6] \quad \text{(B.8)}$$

$$y^F_{t,t+i} = y^F_{t-1,t+i} + (1-\gamma_c)\mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24}, i = [7, ..., 23] \quad \text{(B.9)}$$

$$y^F_{t,t+24} = (1-\gamma_c)\mathbb{I}_Y\frac{u_{t,Y1}F_{t,Y1}}{24} \quad \text{(B.10)}$$

In (B.11), we formulate the stage t cash flows from power trading $(y^F_{t,t})$. As with currency forwards, the expression consists of the committed, 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 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, because these are based on previous positions. We also include variable transaction costs, c_F [EUR/MWh].

$$\begin{aligned} y^F_{t,t} = y^F_{t-1,t} + (1-\gamma_c)\Big[&-(\frac{u_{t,M}}{2} + \frac{u_{t,Q}}{6} + \frac{u_{t,Y}}{24})F_{t,t} \\ &-\Big(\sum_{i=1}^{6}(u_{t,Mi} - w_{t,Mi})F_{t,Mi} + \sum_{j=1}^{8}(u_{t,Qj} - w_{t,Qj})F_{t,Qj} + (u_{t,Y1} - w_{t,Y1})F_{t,Yi}\Big) \\ &- c_F(\sum_{i=1}^{6} w_{t,Mi} + \sum_{j=1}^{8} w_{t,Qj} + w_{t,Y1})\Big] \end{aligned} \quad \text{(B.11)}$$

The cash flows from power trading (B.11) enters the value function for hedging.

Appendix C. Constructing Forward Curves Using Fleten and Lemming's Method

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

$$\mathbb{S} = \{(t_{b,1}, t_{e,1}), ..., (t_{b,M}, t_{e,M})\} \tag{C.1}$$

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 date t_s, and assume that the curve is constructed using M contracts with delivery intervals given in (C.1).

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 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 a seasonality function in vector form, and r is the model discount rate. Finally, 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

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

As an approximation, we set $r = 0$ and $D(t) = 0$. 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

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

Using the method of Lagrange multipliers, the problem can be reformulated to that of a system of equations given by (C.2), where $\boldsymbol{\chi} = [\chi_1, ..., \chi_M]'$ is the vector of Lagrange multipliers and $f(t_s)$ is the forward curve on vector form.

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

In (C.2), $\mathbf{A}$ 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 prices of all M forward contracts traded in the market. The matrix $\mathbf{B}$ is $C \times C$ and given by (C.3).

$$\mathbf{B} = \begin{bmatrix} 1+\omega & -2\omega & \omega & 0 & 0 & 0 & 0 & \dots & 0 \\ -2\omega & 1+5\omega & -4\omega & \omega & 0 & 0 & 0 & \dots & 0 \\ \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 & 0 & \dots & 0 \\ 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega & 0 \\ 0 & \dots & 0 & 0 & \omega & -4\omega & 1+6\omega & -4\omega & \omega \\ 0 & \dots & 0 & 0 & 0 & \omega & -4\omega & 1+5\omega & -2\omega \\ 0 & \dots & 0 & 0 & 0 & 0 & \omega & -2\omega & 1+\omega \end{bmatrix} \tag{C.3}$$

Appendix D. Constructing Forward Curves using Alexander's Linear Interpolation Method

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 electricity 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

$$t_{b,i} \leq t \leq t_{b,j} \tag{D.1}$$

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 (D.2).

$$F_{t_s,t} = F_{t_s,t_{b,i},t_{e,i}} + \frac{t - t_{b,i}}{t_{b,j} - t_{b,i}} \cdot (F_{t_s,t_{b,j},t_{e,j}} - F_{t_s,t_{b,i},t_{e,i}}) \tag{D.2}$$

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

$$t_{e,i} - t_{b,i} = t_{e,j} - t_{b,j} \tag{D.3}$$

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 (D.1). 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.

Appendix E. Electricity price process

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

We use the risk-neutral measure and an initial investment of a forward contract of zero. The expected return of the forward contract must therefore be zero. In line with 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 $T - t = \tau$. The process that explains movements in the forward curve is then given by[6]

$$
\begin{aligned}
\frac{dF_{t,T}}{F_{t,T}} &= \sigma_{t,T} dZ_{t,T} = \sigma_\tau dZ_{\tau,t} \\
\mathbb{E}(dZ_{\tau,t}, dZ_{\hat{\tau},t}) &= \rho_{\tau,\hat{\tau}} dt, \quad \tau, \hat{\tau} \in [\tau].
\end{aligned}
\tag{E.1}
$$

$dZ_{\tau,t}$ and $dZ_{\hat{\tau},t}$ are here Wiener processes associated with forward contracts with time to maturity τ and $\hat{\tau}$. $dZ_{\tau,t}$ and $dZ_{\hat{\tau},t}$ are correlated by $\rho_{\tau,\hat{\tau}}$. $[\tau]$ denotes the set of all time to maturities where $\Delta t \leq \tau \leq \hat{T}$. Our decision problem considers discrete time stages. Equation (E.1) must therefore be discretized. Using Ito's lemma and setting $dt = \Delta t$, the process can be rewritten

[6] See Appendix E for more extensive details and analysis of the implemented price process.

as

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

Here $\Delta Z_{\tau,t} = \sqrt{\Delta t}\epsilon_{\tau,t}$, where $\epsilon_{\tau,t} \sim N(0,1)$. We can modify (E.2) into an expression of the spot price which is a function of $F_{t-\Delta t,t}$, given by

$$F_{t,t} = F_{t-\Delta t,t} \cdot \exp\Big(-\frac{1}{2}\sigma^2_{\Delta t}\Delta t + \sigma_{\Delta t}\sqrt{\Delta t}\epsilon_{\Delta t,t}\Big) \tag{E.3}$$

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 dropped after every state transition. $F_{t,T}$ represents the price of a non-traded forward contract with delivery in period T, delivery period being a semi-month and the spot price $F_{t,t}$ being the first element of the curve. The price of the contracts $F_{t,Mi}$, $F_{t,Qj}$ and $F_{t,Y1}$ are obtained by calculating the average price of the relevant part of the forward curve. We used the method proposed by Fleten and Lemming (2003) to construct the underlying curve, which is used as input in the first time stage of the model[7].

A set of semi-monthly log returns for all forward contracts with time to maturity $\tau \in [\tau]$ must be constructed to estimate the correlation matrix and the volatilities that describe the electricity forward curve dynamics. In our case, $[\tau] = [1, ..., 48]$ semi-months. Koekebakker and Ollmar (2005) propose constructing multiple high-resolution forward curves for a large set of historical trading days, to calculate these returns series. We adopt the method of Alexander (2008), who uses linear interpolation of forward prices in the market to estimate the forward curves[8]. The forward curves constructed using the method of Fleten and Lemming (2003) are smooth and continuous. However, we experience issues with unrealistic oscillations at the near end of some of the forward curves. The curves constructed using the method of Alexander (2008) are neither smooth nor continuous, but resulted in a reasonable volatility curve, as shown in Figure E.8. The log returns are calculated between the average semi-monthly values of the forward curves of two consecutive periods. Using the time series of returns, we estimate the volatility curve for the term structure of forward prices. The volatility curve in Figure E.8 can be understood as being the volatility of returns of forward contracts with time to maturity τ. The volatility monotonically decreases with increasing time to maturity. This is called the Samuelson effect (Samuelson 1965). The effect means that forward prices tend to change more

[7]See Appendix C for an explanation of the method for constructing forward curves.
[8]See Appendix D for an explanation of the method for constructing forward curves.

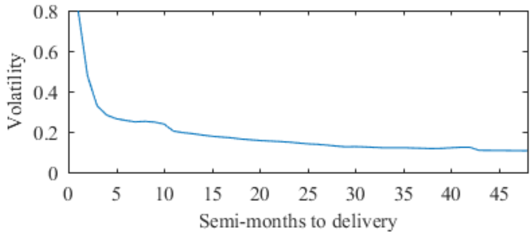

Figure E.8: Annualized volatility curve σ_τ for electricity forward contracts.

	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 E.9: Spot price volatility validation, using semi-monthly returns

the closer they are to maturity. The mechanism behind this phenomenon is that an information shock that affects short-term price has an effect on the succeeding prices which decreases as time to maturity increases. Weather forecasts are an example of information that one would expect to only have a short-term effect on electricity forward prices.

To validate 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 E.8. In other words, it is the volatility of the fictional one semi-month ahead forward contract. As explained in Section 5.1, 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 E.9 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.

Figure E.9 and E.10 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.

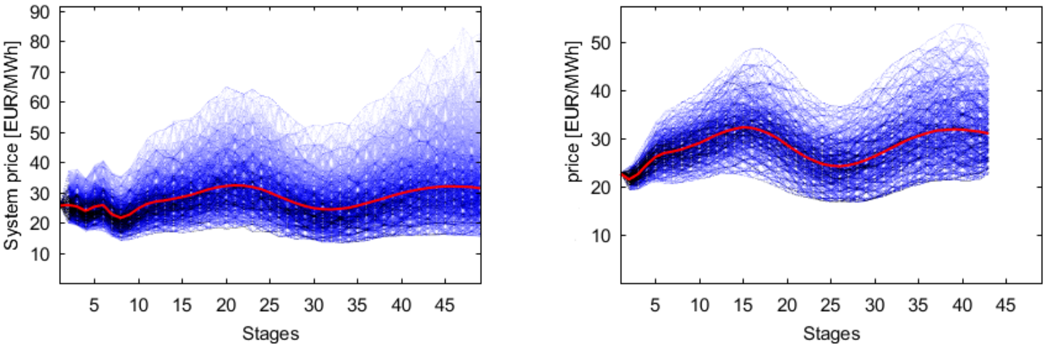

(a) Spot price lattice

(b) 3 months to delivery lattice

Figure E.9: Lattice for spot price and 3 month to delivery forward prices.

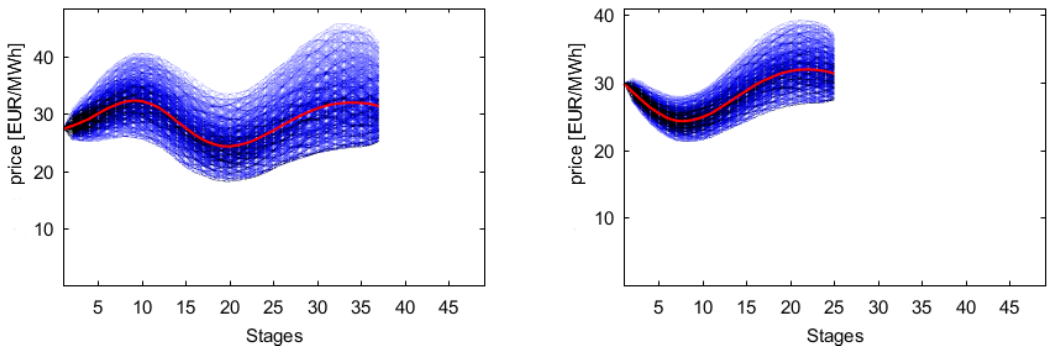

(a) 6 months to delivery lattice

(b) 12 months to delivery lattice

Figure E.10: Lattice for 6 and 12 months to delivery forward prices.

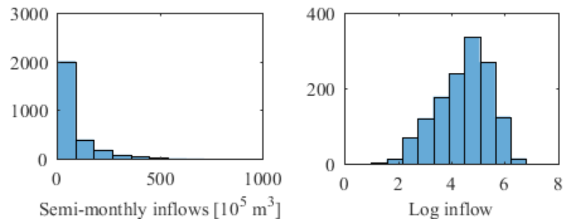

(a) Inflow (b) Log inflow

Figure F.11: Historical distribution of semi-monthly inflow and log inflow for the case plant.

Appendix F. Inflow process

We follow Shapiro et al. (2013) and model natural inflows as a geometric periodic autoregressive process which is given by

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

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

Inflow Y_t is a function of its first lag only. Future values of inflow are, therefore, only dependent on their current value and not the entire history. Inflow therefore follows a Markov process, one of the prerequisites for representing a decision problem as a Markov decision process.

A geometric process is better suited than an arithmetic process to a right-skewed inflow distribution, such as the one in Figure 11(a). A geometric process furthermore does not allow for negative inflows. Shapiro et al. (2013) found the inflow distribution for Brazilian hydropower plants to also be right-skewed, favoring a log transformation of the inflow observations. 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 lag-1 process can be determined by investigating the partial autocorrelation of the historical data for $\ln Y_t - \mu_t$. Partial autocorrelation is the correlation of a time series with its own lagged variables, but with the correlation effects of the values of the time series at all

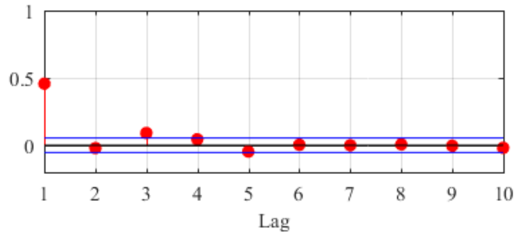

Figure F.12: Partial autocorrelation of the $\log Y_t - \mu_t$ time series

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 G.10: Parameters of the currency process. Interest rates are logarithmic.

shorter lags being removed. Figure F.12 shows the partial autocorrelation of the $\ln Y_t - \mu_t$ time series. Our dataset shows a high value at lag 1 and insignificant values for larger lags, which is similar to the findings of Shapiro et al. (2013). This indicates that it is sufficient to include only one lag in the autoregressive model.

Appendix G. Exchange rate process

We model exchange rate dynamics as geometric Brownian motion with drift equal to the interest rate difference,

$$\frac{dQ_{t,t}}{Q_{t,t}} = (r - r_f)dt + \sigma_Q dZ, \tag{G.1}$$

where σ_Q is the annualized volatility of exchange rate returns.

In discrete time, using $\epsilon_{C,t} \sim N(0,1)$, we have

$$Q_{t,t} = Q_{t-\Delta t, t-\Delta t} \exp\Big((r - r_f - \frac{1}{2}\sigma_Q^2)\Delta t + \sigma_Q \sqrt{\Delta t}\epsilon_{C,t}\Big). \tag{G.2}$$

EURNOK volatility is estimated using historical semi-monthly returns. The parameters of the currency process are shown in Table G.10.

Appendix H. Correlation between random variables

Considering the correlation matrix shown in Table H.11, we see that electricity forward returns exhibit a substantial degree of pairwise correlation. There is also a clear trend of decreasing correlation between contracts with larger time spreads between maturities. The

	inflow	eurnok	F1	F3	F5	F7	F19	F28
inflow								
eurnok	0.08							
F1	-0.08	0.03						
F3	-0.12	-0.01	0.81****					
F5	-0.10	-0.11	0.67****	0.93****				
F7	-0.14	-0.16	0.55****	0.84****	0.94****			
F19	-0.03	-0.16	0.53****	0.76****	0.78****	0.75****		
F28	-0.15	-0.22*	0.40****	0.66****	0.73****	0.75****	0.82****	
F48	-0.12	-0.21*	0.34**	0.56****	0.63****	0.65****	0.82****	0.88****

Table H.11: Correlation matrix of random increments. F1, F3,..., F48 denotes a forward contract with 1, 3,..., 48 semi-months to maturity. $p < .0001$, ****; $p < .001$, ***; $p < .01$, **; $p < .05$, *

analysis reported in Table H.11 shows that all inter-correlations of the electricity forward curve significant are significant at the 1% confidence level.

Increments of inflows and prices exhibit a weak negative correlation which ranges from -0.18 to -0.03. This reflects the hydro domination as power source in this market; a high inflow means high resource availability and thus possibly lower prices. However, all inflow-price correlations were found to be insignificant at both the 1% and 5% significance level.

Increments of the currency process and the price process exhibit only a weak negative correlation with correlation ranging from -0.28 to -0.14 for time to maturity $\tau \geq 6$ for different parts of the forward curve. The natural hedging effect therefore seems to be weak as well.

Currency-price correlations are only significant at the 5% significance level for $\tau \geq 28$ semi-months.

Appendix I. Variability of State Variables

Having constructed uncertainty lattices for spot price, production and currency, we can quantify the variability of each risk factor. Table I.12 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 premature conclusion as the present value of production revenues is not necessarily proportional to the mean spot price or mean EURNOK rate; production is not evenly distributed throughout the year.

Appendix J. Coefficients and Parameter Values

For all simulations, we used coefficient and parameter values given in Table J.13. The tax rates γ_r and γ_c (Thorvaldsen et al. 2018) and transaction costs c_F (Nasdaq Oslo ASA and

	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 I.12: 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.

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 J.13: Coefficients and constants used for numerical studies

Nasdaq Clearing AB 2018) are correct as of 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.

Appendix K. 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 1. These results are summarized in Table K.14. The computational time of each run is approximately 4.5 hours,

	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 K.14: Sensitivity of statistical measures based on six separate runs.

and its memory usage is close to the maximum capacity.

As Table K.14 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.

Appendix L. 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 cash flows from reservoir operations, as a deviation from the expectation, in stage T is given by

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

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

$$(F_{t,T} - F_{T,T}) \cdot u_{T,T}(1 - \gamma_c) \tag{L.2}$$

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 offset each other. Then, we get that the optimal tax-neutral hedge ratio $\frac{u_{T,T}}{W_T}$ is given by (L.3).

$$\begin{gathered}(F_{T,T} - F_{t,T}) \cdot W_T(1 - \gamma_c - \gamma_r) + (F_{t,T} - F_{T,T}) \cdot u_{T,T}(1 - \gamma_c) = 0 \\ (F_{T,T} - F_{t,T}) \cdot W_T(1 - \gamma_c - \gamma_r) = -(F_{t,T} - F_{T,T}) \cdot u_{T,T}(1 - \gamma_c) \\ \frac{u_{T,T}}{W_T} = \frac{1 - \gamma_c}{1 - \gamma_c - \gamma_r}\end{gathered} \tag{L.3}$$

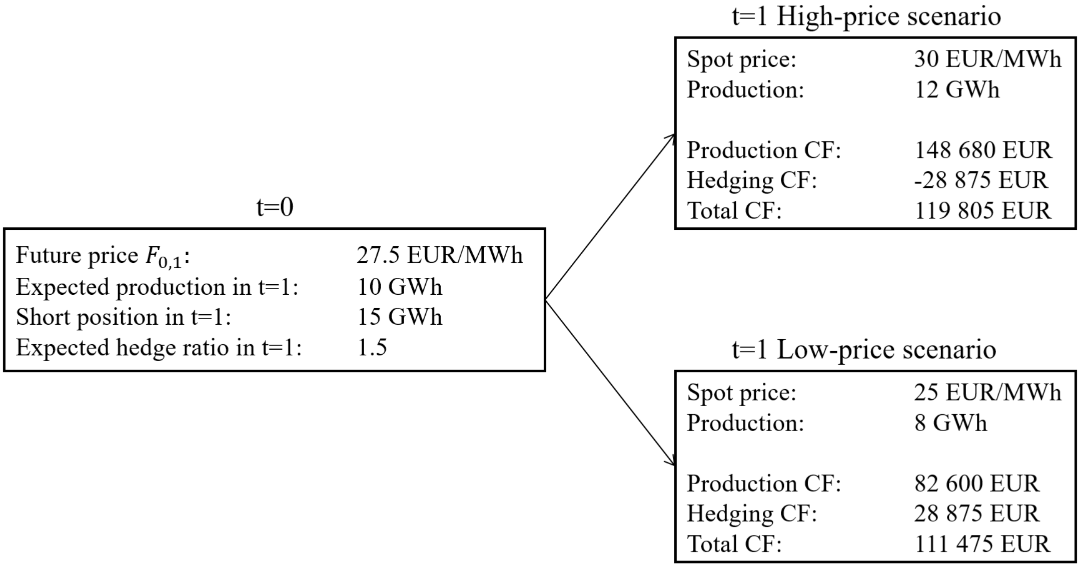

Figure M.13: Two-stage problem with stochastic price and production levels. The cash flows from production are subject to resource rent γ_r and corporate γ_c tax, whereas the cash flows from hedging are subject to corporate tax only. Transaction costs are disregarded.

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.

Appendix M. Numerical Example of Cash Flows from Overhedged Positions

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 M.13. 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 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.