NTNU
Norwegian University of
Science and Technology

Structural Estimation Analysis in Hydropower Scheduling

Maren Boger
Einar Midttun Vestbøstad

Industrial Economics and Technology Management
Submission date: June 2016
Supervisor: Stein-Erik Fleten, IØT
Co-supervisor: Alois Pichler, IØT
Jussi Keppo, National University of Singapore

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

Problem Description

The goal is to understand how profitability factors and other factors affect production decisions for reservoir hydropower producers. How schedulers form expectations regarding future prices and inflow, is an interesting issue here.

Preface

This master thesis was conducted at the Norwegian University of Science and Technology (NTNU), Department of Industrial Economics and Technology Management. Our work falls within the Group of Financial Engineering.

We would specially like to thank our supervisor, Professor Stein-Erik Fleten for valuable discussions, his time and guidance. We would also like to thank Alois Pichler and Jussi Keppo for valuable discussions. Last but not least, we thank Kolsvik for making their production and reservoir data available for analysis in our thesis.

Abstract

When planning production, a hydro power reservoir managers need to form expectations for electricity prices in the future. When forming expectations, the Nordic electricity forward market is a useful tool for predicting how the underlying spot price will move. In this paper we develop a structural estimation model for a single agent hydropower producer in Norway. With this model, we analyse how primitives in the price process, related to the forward price, can be inferred from empirical data from actual production time series. By analyzing trends and patterns in observed time series we have approximately parametrized the state space transition. Central here is the connection we model between inflow and price, to capture dry- and wet year dynamics in the two. To demonstrate the model it has been applied to a specific hydro power plant in Norway. From the results we make a preliminary analysis of to what extent the producer uses forward information when planning production. The results indicate that this specific producer is inclined to take the forward price into account when planning, and that a forward price with 6 months to maturity is favored. An important byproduct of our model is the ability to calculate water values from the outputs.

Sammendrag

I produksjonsplanleggingsprosesser må vannkraftplanleggere gjøre seg opp forventninger for hvordan elektrisitetsprisen kommer til å utvikle seg i fremtiden. Disse forventningene baserer seg blant annet på informasjon fra forwardmarkedet i Norden. I denne studien utvikler vi en strukturell estimeringsmodell for en norsk vannkraftprodusent. Denne modellen brukes videre til å analysere hvordan skjulte parametere i prismodeleringsprosessene relaterer seg til empirisk data. Ved å analysere trender og mønstere i observerte tidserier finner vi en tilnærmet parameterisering for overgangen mellom tilstandsvariable. I disse overgangene er særlig sammenhengen mellom tilsig og pris viktig for å reflektere dynamikken fra tørr- og våtår. Vi har utprøvd modellen på et norskt vannkraftverk. Ut fra resultatene kan vi gjennomføre en innledende analyse av hvor mye vannkraftprodusenten vektlegger forwardinformasjon når de planlegger produksjonen. Resultatene indikerer at den analyserte vannkraftprodusenten er tilbøyelig til å bruke forwardinformasjon når de planlegger, og at forwardpriser med en 6 måneders tid til modning er foretrukket. Et nyttig biprodukt fra modellen er muligheten den gir for å gi en indikasjon på hva slags vannverdier kraftverket har opperert med.

Table of Contents

Chapter 1

Introduction

Managers of reservoir hydropower plants need to form expectations regarding future electricity prices when balancing immediate and future rewards from releasing water and generating electricity. Although surveys on forecasting methods employed by Fleten et al. (2008) shed light on how this happens, we analyse how primitives of this expectation process can be inferred from empirical data from actual production time series. Our main research goal is to apply structural estimation theory of Markov Decision Processes, developed by Rust (1987), to a hydropower planning problem. The model aims to work as a descriptive rather than a normative tool for the hydropower industry. In order to demonstrate the structural estimation model we apply it to a single-case hydropower plant. An example of insight gained from the demonstration, is to what extent a planning agent put emphasis on forward information when planning production. Another example is empirical insight in what expectations the hydro power producers use for future water values.

Structural estimation was first used on an optimal switching problem for buss engines by Rust (1987) and Rust (1994). The idea is that if we observe a set of states and actions taken by an agent, we can work backwards to infer the objective function of that agent, by maximizing the likelihood of matching the observed data. By maximizing the likelihood function, the analyst can obtain an understanding of parameters hidden in the economic model. In order to estimate the structural parameters in a stochastic dynamic programming problem, Rust used an algorithm he called the Nested Fixed Point (NFXP) algorithm. This algorithm has two parts, an outer loop that searches for the structural parameters with the maximum likelihood value, and an inner loop that solves the stochastic dynamic programming model given a value for the structural parameter. The NFXP algorithm is the one studied by Foss and Høst (2011). According to Su and Judd (2012) the NFXP algorithm is computationally demanding, because it iterates over all structural parameter values and then solves the underlying stochastic dynamic programming (SDP) model with high accuracy for each structural parameter value. To deal with this computational difficulty, they propose a new strategy, the Mathematical Program with Equilibrium Constraint (MPEC) approach. What this approach does differently from Rust's is that instead of having an inner loop that runs for each iteration of an outer loop, you have the SDP model as a constraint in the maximization of the likelihood. This way the only SDP that needs to be solved exactly is the one corresponding to the final estimate of the structural parameter, and not all of them. Because of the computational advantages the MPEC approach has over the NFXP algorithm, we utilize MPEC when applying structural estimation for a hydropower case.

Examples of topics where structural estimation has been utilized are optimal bus engine replacement (Rust, 1987), call center planning (Aksin et al., 2013), demand for durable goods (Rapson, 2014), effects from uncertainties on the willingness to invest in the oil drilling industry (Kellogg, 2014), investment timing for corn ethanol plants (Thome and Lin, 2015), evaluation of the likelihood of reactivation of temporarily closed oil and gas wells (Muehlenbachs, 2015) and maintenance and switching costs of peak power plants (Fleten et al., 2016) and (Johansen, 2015). For hydropower planning, Foss and Høst (2011) did an initial unsuccessful study where they resolved to focus on aspects of an algorithmic implementation of the structural estimation problem. Except from that study, structural estimation have yet to be applied for hydropower planning.

Hydropower planning problems are suitable to be treated as a stochastic dynamic problem, where a decision today change the reservoir levels and thereby affect future production opportunities. Here the dynamic properties of the planning problem can be accounted for together with multiple stochastic variables, such as inflow, price, forward price etc. A regulated market entails that power companies are obliged to supply electricity to the geographical area they are responsible for. Nandalal and Bogardi (2007), Tejada-Guibert et al. (1993), Maidment and Chow (1981), among others use stochastic dynamic programming for hydropower planning in regulated markets. In such a situation, a cost minimisation perspective is applicable and price can be treated as a deterministic variable. Norway was one of the first countries to deregulate their energy market, when the Energy act of June 1990 was introduced (Wolfgang et al., 2009). For unregulated markets, price can be treated as a stochastic variable in a stochastic dynamic programming problem. This has been explored by Fosso et al. (1999), Mo et al. (2001), Fleten et al. (2002) and Wolfgang et al. (2009), among others. For our case, we treat both inflow and price as stochastic input state variables in a stochastic dynamic programming problem.

We use a parametric approach suggested by Fleten et al. (2016) to deal with the transition probabilities in the structural estimation model. However they use a non-parametric approach in their analysis. The parametric approach involves time series modelling of state variables as Markovian processes and using them to calculate state space transitions. The goal of the time series modeling is to capture a sufficient amount of the state variable dynamics to make structural estimation a viable method for gaining insight in the hydropower planning industry. Time series forecasting approaches typically used for hydropower planning are: Forecasts based on trends and patterns in observed time series for price, simulation models who describe price formation and forecasts based on prices from the future market. These approaches are elaborated on in Section 3.2. Ultimately we have based our time series modeling on theory from Kolsrud and Prokosch (2010), combining several of these approaches. Here we model a connection between inflow and price that capture the dry year-/wet year dynamics of both processes and the dependency of price on inflow. The inflow-price connection has not been modelled like this before, and we believe it is a better explanation of reality than earlier models.

Our goal is to build a model that enables us to learn more about to what extent hydropower planners use forward prices when planning production. To what degree non financial firms actually use derivatives in practice is a topic for empirical analysis. Previous work on this includes Fleten et al. (2008), who, from a linear regression model, find that Norwegian hydropower producers seem to be using electricity forward (swap) prices when planning production. Näsäkkälä and Keppo (2008) and Fleten et al. (2002) arguments that electricity forward prices should be used in production planning for hydropower plants. One approach to inclusion of forwards is described by Mo et al. (2001) as a optimization model, tested for 38 reservoirs, where trades of forward contracts were incorporated in a utility function to hedge for very low price values. Fleten et al. (2002) on the other hand creates a hydropower planning model where the spot price scenarios are set so that the expected spot price equals the current market price for a forward with delivery at the corresponding period. According to Näsäkkälä and Keppo (2008), another commonly used approach is to compare the results of models, such as the EMPS model, with the forward curve. In the case of large differences, the value of the equilibrium model are adjusted to be closer to the forward curve values. Our model aims to provide an estimate of to what degree forward information are used to form expectations of future prices in the hydropower industry, as suggested done by Näsäkkälä and Keppo (2008), Fleten et al. (2002) and Fleten et al. (2008).

Based on the results from our structural estimation model we find that our model has the potential to let external analysts gain insight into water values used by producers when planning production. This is a valuable finding, since alternatives are time consuming and costly. From the results we also make a preliminary analysis, which indicates that the producer uses forward price information when planning, and that a forward price with 6 months to maturity is the most likely choice.

In Chapter 2 we present some major elements in structural estimation theory, including a basis for the parametric approach and the MPEC formulation. Chapter 3 lays out specifics on the context of a hydropower environment case. The aim of Chapter 4 is to put structural estimation in context of a general hydropower planning problem and to explain the background for specific model detail decisions. Chapter 5 covers model results and discusses future improvements for the model. Lastly, Chapter 6 has concluding remarks.

Chapter 2

Structural estimation

In this chapter we present the main elements of structural estimation fro Markov decision processes. Some assumptions for Markov decision processes are discussed in Section 2.1. Section 2.2 presents the Bellman equation for structural estimation. Section 2.3 presents the maximum likelihood estimation problem from the MPEC approach. Finally in Section 2.4 we use a parametric approach for calculating state space transitions and state space transition probabilities in the model.

2.1 Markov decision processes

Markov decision processes (MDP's) has been used in many theoretical studies for economic problems due to their applicability for most economic problems where decisions are made over time. MDP's have commonly been applied as normative models for how rational decision makers should behave. Later on, analysts have realized that MDP's also provide good empirical models for exploring how real-world decision-makers actually behave. Under a MDP assumption, where agents are considered rational, we can estimate the primitives, agent preferences and beliefs, applied in decision making. In order to do that, we intend to use structural estimation theory of Markov decision processes, developed by Rust (1987).

For a Markov decision process for structural estimation we define a state space, x_t, as a vector of state variables observed by both the agent and the analyst. By taking decision d_t at state x_t, we move to a new state x_{t+1}. We define this transition as the state transition function, $f(x_t, d_t)$, in

$$x_{t+1} = f(x_t, d_t) \tag{2.1}$$

The state transition probability is now denoted by $p(\mathrm{d}x_{t+1}|x_t, d_t)$, which is the probability of moving from state x_t to state x_{t+1}, given decision d_t (Gamba and Tesser, 2009).

Assumptions necessary for structural estimation for Markov decision processes are that agents really are rational expected-utility maximizers and that the problem can be regarded as stationary with a infinite horizon (Rust, 1994). In the case of stationarity, transition probabilities, discount factors and utility functions are the same for all t. In the case of infinite horizon for a stationary Markovian structure, the future is considered the same whether the agent is in state x_t or x_{t+k}, given that $x_t = x_{t+k}$. This implies that the only values who are affecting the hydropower release decision is the current state, x_t, combined with the associated probability to go to the state, x_{t+k}. By making an infinite horizon assumption, we avoid end of horizon effects affecting the decision rules. Thereby the assumption allow for a unique solution of the *Bellman equation* when the discount factor is set to a constant value between 0 and 1 (Rust, 1994).

2.2 Bellman equation

The valuation model of an individual agent can be described as in (2.2). It is formulated as a discounted discrete decision process, who has an unique solution when assuming the process is stationary and by assuming an infinite horizon. The *Bellman equation* is used to find the optimal policy from a *Markovian process*, meaning it only depends on the current state, X_t. Parts of the equation known as the *Bellman equation* are derived in (2.2)-(2.5).

$$V(x) := \max_{d \in D} \mathbb{E}_d \left(\sum_{t=0}^{\infty} \beta^t g(X_t, d_t) \middle| X_0 = x \right), \tag{2.2}$$

Here β denotes the discount factor and $g(X_t, d_t)$ is the profit function give state X_t and decision d_t. The expression $\mathbb{E}_d(\cdot|X_0 = x)$ denotes the expectation with respect to state transition probability $p(\cdot|x, d)$. Taking the immediate profit term outside the summation, and having the remaining future value conditional on the state in the next stage, gives

$$V(x) = \max_{d \in D} \mathbb{E}_d \left(g\left(X_0, d\right) + \beta \mathbb{E}_{d'} \left(\sum_{t=0}^{\infty} \beta^t g\left(X_{t+1}, d_{t+1}\right) \left| X_1 = x' \right. \right) \middle| X_0 = x \right) \tag{2.3}$$

As the expected future value is the definition of the value function (2.2) in the next stage, we can write (2.3) as

$$V(x) = \max_{d \in D} \mathbb{E}_d \left(g(X_0, d) + \beta V(X_1) \middle| X_0 = x \right) \tag{2.4}$$

Further on the immediate profit are moved outside the expectation, so that

$$V(x) = \max_{d \in D} \left(g(x, d) + \beta \mathbb{E}_d \left(V(X_{t+1}) \middle| X_t = x \right) \right) \tag{2.5}$$

In the Bellman equation, stated in (2.5), the maximum are calculated over all decisions, d, and depend s only on the present values of the state variables in x.

When so on assuming stationarity, the future looks exactly the same in time t as it does in time $t+k$, as long as $X_t = X_{t+k}$. Then X_{t+1} is distributed conditional on X_t in the same way as X_1 is distributed conditional on X_0, hence we express the conditional expectation based on X_0 from now on, the same way as Fleten et al. (2016). The Bellman function is then a fixed point for

$$V(x) = \max_{d \in D} \left(g(x; d) + \beta \mathbb{E}_d \Big(V(X_1) \Big| X_0 = x \Big) \right) \tag{2.6}$$

Error term
According to Rust (1994), when the SDP is stationary and has an infinite horizon, the optimal decision rule of the hydropower planner is a deterministic function of the observed states. This implies that if the analyst is able to observe all the states, the person would be able to perfectly predict the behavoir of each hydropower plant. As such a capability is not realistic, an error term is included in the form of an idiosyncratic shock, $\varepsilon(d)$. The $\varepsilon(d)$ can be interpreted as an unobservable, meaning a variable observed by the agent but not by the analyst. It is formulated as a vector with at least as many components as the number of possible decisions, D. The profit function is, in (2.7), written as an additive separable decomposition of the profit from observable decisions and the idiosyncratic shock, $\varepsilon(d)$, as done by Rust (1994) and Fleten et al. (2016).

$$g(x, \varepsilon; d) = g(x; d) + \varepsilon(d) \tag{2.7}$$

State transition with error term
When the error term is included, the state transition can be written as

$$(x_{t+1}, \varepsilon_{t+1}) = f(x_t, \varepsilon_t, d_t) \tag{2.8}$$

The state transition probability is then

$$\pi(\mathrm{d}x_{t+1}, \mathrm{d}\varepsilon_{t+1} | x_t, \varepsilon_t, d_t), \tag{2.9}$$

We assume the processes for x and ε to be conditionally independent, so

$$\pi(\mathrm{d}x_{t+1}, \mathrm{d}\varepsilon_{t+1} | x_t, \varepsilon_t, d_t) = \mathcal{E}(\mathrm{d}\varepsilon_{t+1} | x_{t+1}) p(\mathrm{d}x_{t+1} | x_t, d_t), \tag{2.10}$$

where $\mathcal{E}(\cdot)$ is the transition probability of ε (Rust, 1987).

Value function with error term
With the help of (2.7) and (2.8) the value function with the idiosyncratic shock can be written as

$$V(x, \varepsilon) = \max_{d \in D} \{ g(x; d) + \varepsilon(d) + \beta \mathbb{E}_d \left(\int V(X_1, \varepsilon_1) \mathcal{E}(\mathrm{d}\varepsilon_1 | X_1) \middle| X_0 = x \right) \} \tag{2.11}$$

We define the expected value function

$$v(x,d) := \mathbb{E}_d\left(\int V(X_1,\varepsilon_1)\mathcal{E}(\mathrm{d}\varepsilon_1|X_1)\middle|X_0 = x\right), \qquad (2.12)$$

which is to be used as the constraint in the maximum likelihood estimator in the structural estimation part later, keeping in mind that $\mathbb{E}_d(\cdot|X_0 = x)$ denotes expectation with respect to $p(\cdot|x,d)$.
We can now put $V(X,\varepsilon)$ from (2.11) into the expected value function $v(x)$, i.e. take the expectation of (2.11)

$$v(x,d) = \mathbb{E}_d\bigg(\int \max_{d\in D}\{g(X_1;d) + \varepsilon_1(d) + \beta\cdot v(X_1)\}\mathcal{E}(\mathrm{d}\varepsilon_1|X_1)\bigg|X_0 = x\bigg) \qquad (2.13)$$

which is the fixed point equation for v.

Simplification using the Gumbel distribution
From extreme value theory we know that the normalized maximum converge to an extreme value distribution. Because of the maximization in the $v(x,d)$ function in (2.13) it is reasonable to specify ε as a mutually independent process of Gumbel variables, also independent from X. Fleten et al. (2016) remark that the extreme value distribution is closed under maximization. This allows for a simplification made by using the explicit formula for the conditional choice probability defined from (2.14) and (2.15).

$$\int \max_{d\in D}(\varepsilon(d) + c_d)\mathcal{E}(\mathrm{d}\varepsilon|x) = b\cdot\log\left(\sum_{d\in D}\exp\left(\frac{c_d}{b}\right)\right) \qquad (2.14)$$

Here the b can be interpreted as the degree of uncertainty in the Gumbel distribution, a proportional scaling parameter of the standard deviation of the distribution. Following Rust (1994) we use a value for the Gumbel scale parameter b that is normalized to 1. The term, c_d from (2.14) are defined to be

$$c_d := g(X_1;d) + \beta\cdot v(X_1) \qquad (2.15)$$

Using (2.14) and (2.15), we can write (2.13) as

$$v(x,d) = \mathbb{E}_d\left(b\cdot\log\left(\sum_{d'\in D}\exp\left(\frac{g(X_1;d') + \beta\cdot v(X_1)}{b}\right)\right)\middle|X_0 = x\right) \qquad (2.16)$$

Fixed point Bellman operator
For notational convenience we define the operator

$$t_\theta(v)(x,d) := \mathbb{E}_d\left(b\cdot\log\left(\sum_{d'\in D}\exp\left(\frac{g(X_1;d') + \beta\cdot v(X_1)}{b}\right)\right)\middle|X_0 = x\right) \qquad (2.17)$$

This way we can write the fixed point Bellman equation as.

$$v = t_\theta(v), \tag{2.18}$$

and use it as a constraint in the maximum likelihood estimation in the structural estimation formulation (Johansen, 2015)

2.3 Structural estimation

The structural parameters that are to be estimated are collected in a vector θ. Given a θ, the maximum likelihood estimation problem in the MPEC approach, described by Su and Judd (2012), is formulated as

$$\begin{aligned} \underset{\theta}{maximize} \quad & L(\theta, v_\theta, (X_n, d_n)_{n=1}^N) \\ s.t. \quad & v_\theta = t_\theta(v_\theta) \end{aligned} \tag{2.19}$$

Here the constraint v_θ is the fixed point Bellman equation (2.18). The objective function is the log-likelihood function, based on the likelihood function $l(\theta, v_\theta, (X_n, d_n)_{n=0}^N)$. The function l is the likelihood of observing date $(X_n, d_n)_{n=0}^N)$ given a certain θ. As a suitable maximum likelihood estimator we use the choice probability formula

$$P_v(d|x) = \frac{exp\left(\frac{g(x;d)+\beta\cdot v(x,d)}{b}\right)}{\sum_{d'\in D(x)} exp\left(\frac{g(x;d')+\beta\cdot v(x,d')}{b}\right)}, \tag{2.20}$$

which is given from the Gumbel distribution assumption for ε, as shown by Fleten et al. (2016). This is the probability of making decision d given state x. We can use (2.20) to write the likelihood function

$$l(\theta, v_\theta, (X_n, d_n)_{n=0}^N) = \prod_{n=1}^N P_v(d_n|X_n) \tag{2.21}$$

It follow from this that the log-likelihood function is

$$L(\theta, v_\theta, (X_n, d_n)_{n=0}^N) = \sum_{n=1}^N log(P_v(d_n|X_n)), \tag{2.22}$$

which is used in the maximum likelihood estimation in (2.19).

2.4 Estimation of conditional expectation

In the constraint of the maximum likelihood estimator (2.19), the operator $t_\theta(v_\theta)$ is an expectation conditional on $\{X_0 = x\}$. To be able to evaluate $t_\theta(v_\theta)$ we need to evaluate the conditional expectation. Fleten et al. (2016) suggest two different methods for computing the conditional expectation, one parametric approach using a time series model that follow an auto-regressive (AR) scheme and one non-parametric approach using kernel estimators. We are using the parametric approach in this paper, but expand upon their approach by having a set of AR processes that depend on each other instead of a single one.

State space transition
Given state space x_t, containing a set of state variables that follow an AR scheme, we write the state transition as

$$\boldsymbol{x_{t+1}} = f(\boldsymbol{x_t}, d_t) = \boldsymbol{A}\boldsymbol{x_t} + \boldsymbol{B}\boldsymbol{e} + \boldsymbol{c} \tag{2.23}$$

Here $\boldsymbol{x_{t+1}}$ and $\boldsymbol{x_t}$ are the state vectors at time $t+1$ and t respectively, $\boldsymbol{A}$ is a matrix with parameters for the state vector, $\boldsymbol{e}$ are the error terms of independent random variables, $\boldsymbol{B}$ is a matrix with parameters for the error terms and $\boldsymbol{c}$ is a vector with additional deterministic constants.

Some of the state variables might have correlated errors. In order to simulate multiple correlated errors, we first find their covariance matrix, Σ. Σ is then decomposed, using Cholesky decomposition, to obtain its lower triangular matrix L, such that $L * L^T = \Sigma$. By left-side multiplying L with uncorrelated standard Gaussian random errors, e', we can obtain the desired correlated random errors, e. Eq. (2.23) can then be written as

$$\boldsymbol{x_{t+1}} = f(\boldsymbol{x_t}, d_t) = \boldsymbol{A}\boldsymbol{x_t} + \boldsymbol{B}\boldsymbol{L}\boldsymbol{e'} + \boldsymbol{c} \tag{2.24}$$

Parametric Bellman equation
For notational purposes, we define a function

$$h(x) := b \cdot \log\left(\sum_{d \in D} \exp\left(\frac{g(x;d) + \beta \cdot v(x)}{b}\right)\right) \tag{2.25}$$

Since the state variables in x follow an AR scheme, we can write (2.17) as (2.26), reducing the conditional expectation to a simple expectation (Fleten et al., 2016).

$$t_\theta(v)(x,d) = \mathbb{E}_d(h(X_1)|X_0 = x) = \mathbb{E}h(Ax + c + BLe), \tag{2.26}$$

As we can see in (2.23), the state transition also includes the decision, which means the expectation, $\mathbb{E}$, in (2.26) is no longer conditional on the decision, d.
This new, simple, expectation can be further simplified by discretizing the error terms in e, and weighting the different discrete levels by an appropriate probability weight. Given N state variables, this discretization leads to

$$t_\theta(v)(x,d) \sim \sum_{i_1=1}^{M_1} \cdots \sum_{i_N=1}^{M_N} w_{i_1} \cdots w_{i_N} h(Ax + c + BL\tilde{e})), \tag{2.27}$$

which is our operator used as a constraint in the maximum likelihood estimation in (2.19). The discretized error terms, $\tilde{e}$, in (2.27) are described as

$$\tilde{e} = \begin{bmatrix} \tilde{e}_{i_1} \\ \cdot \\ \cdot \\ \cdot \\ \tilde{e}_{i_N} \end{bmatrix} \tag{2.28}$$

Here $\tilde{e}_{i_n}$ is a $M_n \times 1$ vector of center points for a M_n-point approximation of a normal distributed random variable and w_{i_n} is a $M_n \times 1$ vector with the corresponding probability weights, for $n = 1, \ldots, N$.

Until now we have described the background for a parametric approach to structural estimation. From now on we will attempt to apply (2.19), with (2.27) as a constraint, on a specific hydropower plant, using a set of parametric time series models.

Chapter 3

Hydropower

We are interested in demonstrating the structural estimation model described in Chapter 2 for a specific case of a single-agent hydropower producer. The structural estimation algorithm is estimating the log likelihood of the results from the decision model replicating observed data. In order to get statistically significant results we need to have access to sufficiently long time series. The observed time series also serve as a basis for the modelling of the different state variable processes in Section 4.1. To obtain an insight into the data that are relevant for our problem we first explain some of the characteristics of the Norwegian power market in 3.1. Then we go into more dept on how hydropower planning commonly is done in Norway in Section 3.2. Section 3.6 describe observed data for our hydropower planning case. Lastly Sections 3.3, 3.4, 3.5 and 3.7 elaborate on externally observed time series and how they interact.

3.1 The Norwegian power market

We study the Nordic power market. This market was the first to deregulate, with Norway as a first mover in 1991, making it the market with the most data to use for analysis of deregulated market behavior. Even if Nord Pool has expanded to many countries in Europe later, bidding mechanisms are still grouped in such a way that for Norway, only the Nordic countries are affecting the pricing. In a normal year, around half of the electricity production in the Nordic countries is from hydropower (NordPool, 2016) and (Kolsrud and Prokosch, 2010). When isolating Norway, about 95% of the total electricity is generated by hydropower (NVE, 2016). This implies that the market will be largely affected by hydrological changes and aggregated decision processes in the hydropower industry.

In this text the terms futures, forwards and swaps will be used interchangeably for a financial instrument where two parties agree today on the terms for selling or buying an asset some time in the future. Forward contracts are traded as a purely financial contract, while the spot market includes physical trades (Botterud et al., 2002) and (NordPool, 2016). Since its opening in 1995, the financial market for derivative products at Nord Pool has gone through considerable changes in terms of gradually changing the portfolio of products offered. Included in the changes are a transition from NOK/MWh to Euro/MWh and changes in the maturity times of the products offered. Futures are traded continuously during a trading day (Fleten et al., 2008). When the forward prices are set, the system price serves as a reference price, implying they are closely interlinked.

3.2 Hydropower planning in Norway

On the operations side of hydropower planning, Stochastic Dynamic Programming (SDP) are usually used for calculating water values, optimize power generation, manage reservoir levels and forecast prices. SDP has been used for this in Sweden and Norway for many years (Fleten et al., 2002). The decision to release water involves a trade off between reducing risk of spill by releasing water today and reduces risk of having to sell at low prices in the future if saving water. *Water values* are defined as the opportunity cost of using water immediately in stead of storing it (Botterud et al., 2010). The models named Vansimtap and EMPS are the most commonly used energy models by large hydropower producers in Norway. In the Vansimtap model, market price is exogenously given, usually from the EMPS model. In Vansimtap, both price and inflow are treated as stochastic variables in the short term planning problem (SINTEF, 2016b). For different hydropower producers, frequencies of running the model, adjustment for the amount of snow pack and short term inflow, and chosen data exchange vary when using the Vansimtap model (SKM-Energy, 2016).

Botterud et al. (2010) name information typically available to market participants as: weather forecasts (temperature, precipitation, wind), reservoir levels, inflow forecasts, current and expected future snow pack, hydro balance, fuel and emission prices, power plant and transmission line outages, market prices (history and future/forwards) and import/export. As for price forecasts, Botterud et al. (2010) claim them to commonly be bought from analytics companies, or created in-house. Therefore price information might typically vary for different market participants, as they might model it differently, prepare the available data differently, or interpret the results differently. The drivers utilized for price forecasting are therefore hidden from external analysts and only visible for the hydropower planner. We use structural estimation to gain insight in what drives the forecasting process for price in a case study of a hydropower planning process. Botterud et al. (2010) highlights 3 different approaches for forecasting future market price for hydropower production planning:

1. **Forecasts based on trends and patterns in observed time series for price**: This approach is forecasting in its simplest form and demands long time series in order to make reliable estimates of the stochastic parameters. An example of this is time series models using a discount rate and other market participants beliefs based on observed data and reduced-form regression models for how the price will develop.

2. **Use simulation models that describe price formations**: SINTEF (2016a) highlights this alternative, in the form of the EMPS model, to be the most commonly used model for price simulation in the hydropower industry. In this model, the

objective is to minimize the expected cost in the whole system of hydropower planners, given stochastic inflow scenarios. The entire dispatch system is simulated for different inflow scenarios, drawn from a randomized pool of observed time series (Wolfgang et al., 2009). The result, aggregated over all hydropower producers, with transmission constraints, is representative for water values. These water values are again used to simulate prices, given different inflow scenarios (SINTEF, 2016a), (Fosso et al., 1999) and (Botterud et al., 2010). As the EMPS model is a bottom-up model, it requires a substantial amount of planning resources. This typically involves preparing and analyzing data and running large SDP models.

3. **Forecasts based on prices from the future market**: Future prices are a representation of the general market consensus for future delivery of electricity. In an efficient market, the forward price reflect the value of future delivery of electricity. According to Näsäkkälä and Keppo (2008), production planning based on forward curve dynamics and value maximization can give expected earnings within 2.6% of the theoretical upper bound. To include forward price in a planning model who is approximating release decisions they create a regression model. Here they use the difference between the spot price today and forward prices for different maturity dates to help planners decide if they should produce today or wait for a later day. If the forward price is higher then the current spot price, the producers can expect higher prices in the future and therefore they have an incentive to save water.

Our model is meant to describe the main dynamics of the above mentioned price forecasting methods in order to enable analysis of the weighting on different strategies in observed production time series. In order to do that, we use a combination of a forecast based on trends and patterns in observed time series, with inclusion of dynamics from inflow and observed reservoir level, and forward information, mentioned in approach 2. As, mentioned, time series models are based on subjective discount rates and other not easily estimated market participant beliefs. Therefore, an inclusion of forward price data, representing the market consensus, should provide a valuable price model improvement. The connection between inflow and price simulations, used as a basis for our scheduling problem, is largely based on an empirical model by Kolsrud and Prokosch (2010). Based on their approach, we incorporate a connection between the combination of inflow and overall reservoir levels with modeled spot prices. The influence of inflow and reservoir state variables, in our price model, make our model more informative than approach 1, and makes the dynamics more similar to that from the EMPS model in approach 2. The model used in this study is meant to simplify scheduling procedures, as the most common approaches is based on complicated bottom-up models, such as the EMPS model (Kolsrud and Prokosch, 2010).

3.3 Electricity price data

As a basis for our time series models we have system spot price data from 01.01.1993 to 25.02.2016. When modelling the spot price later, we use the logarithm of the price. The reason for using log prices is twofold. First of all we want to scale the extremes in the spot price down, which for instance makes it easier to fit a seasonal function to the data. Using the logarithm of the price is also helps avoid negative values in price simulations. Both the spot price and the log of the spot price, is shown in Figure 3.1, in €$/MWh$ and log €$/MWh$ respectively. In order to calculate the system price, Nord Pool aggregate all purchase and sell orders from the Nordic region that are submitted before the daily deadline at noon. The intersection between the demand and supply curve is then used for calculating the resulting system price for each hour of the next day, making it a result of interactions between overall supply and demand.

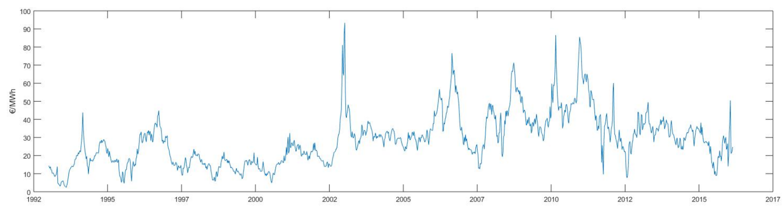

(a) Spot price

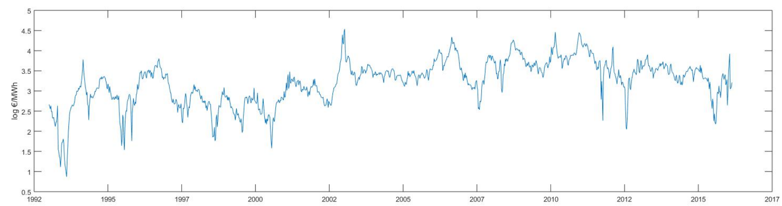

(b) Log price

Figure 3.1: (a) Spot price and (b) log price from 1993 to 2016

From the price plots we can see that there is seasonal dependencies in the price. The price is usually lower in the summer, and higher in the winter. The high electricity prices during winters in Norway are mainly driven by a substantial increase in demand because electricity is used for heating during cold winters in Norway. Also accumulated national inflow levels during winters are small, due to precipitation coming as snow, making electricity a scarce resource (Kolsrud and Prokosch, 2010). The low prices during the spring are influenced by the restriction imposed by the degree of regulation of Norwegian power plants. For power plants with a low degree of regulation, very high inflow levels will force them to produce, driving prices down. During this period, demand for heating is low, driving prices further down.

As most of the total electricity in Norway is generated from hydropower, production is easily changed on short notice. This makes the daily spot prices vary less than in purely thermal markets, as production shortages can be easily regulated. Seasonal fluctuations tend to be more prominant (Botterud et al., 2002). Nevertheless, spot prices commonly includes high prices or spikes. These can be explained by very low reservoir levels and they rarely last for more than a couple of days. As we increase our time resolution to a week, and prices is calculated as an average of the weekly price, the effect from these spikes are limited. Combined with the transformation to log prices the spikes are even more reduced.

We choose to use system prices instead of area prices. As the price forecast model we use compare information from the forward market with spot price forecasts and it is the system price which is used for pricing financial contracts, it is the most relevant price for comparison with forward information (Botterud et al., 2002) and (NordPool, 2016). What separates system price and forward price is that transmission capacities are not included in the system price calculations, while it is included for area prices. As long as no system congestion between bidding areas is present, the system price is the prevailing market price. The bidding area used for calculating the system price is constituted by Norway, Denmark, Sweden and Finland. On the other hand, by using area prices, we could have reflected the actual price expectations for the local producer more closely. Anyhow, our power plant area price is rarely affected by congestion restriction. Therefore, choosing to use the system price instead of the area price for production planning do not impose large differences.

3.4 Forward price data

Since we want to investigate how forward prices affect producers expectation of future prices, we need observed time series for the forward price. We have daily data for forwards with different time to maturity from 01.01.2006 to 31.12.2015. In order to obtain a weekly resolution we use the average weekly forward price. We can see the dynamics of a set of forward prices with different maturity dates in comparison to the spot price in Figure 3.2. As expected, the variance of forwards with a long time to maturity is lower than the ones with shorter times to maturity. Also, all of the forwards follow the realized spot price closely, with a slightly higher mean. This is expected as futures prices tend to be higher than spot prices due to risk premiums (Botterud et al., 2010).

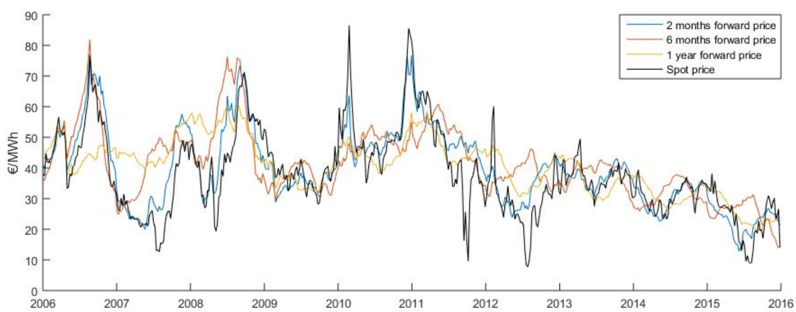

Figure 3.2: Comparison between forwards with different time to maturity, 2 months (blue), 6 months (red), 1 year (yellow) forward and spot price (black). From 2006 to 2016.

3.5 Aggregated reservoir data

We want to use the the deviation from aggregated reservoir levels over a large area as a factor in our price model. Weekly aggregated reservoir levels in Norway from 01.01.1993 to 25.02.2016, retrieved from NVE (2016), are shown in Figure 3.3. As we model the system price, which is based on bids from the entire Nordic region, reservoir levels from the entire region should be used. Anyhow, we only have access to data for the entire region from 01.01.2013 to 25.02.2016 (NordPool, 2016). When comparing the deviation from Norway alone and the entire Nordic region, we observe them to be highly correlated. We therefore conclude that Norwegian reservoir level deviation is a reasonable approximation for the entire region.

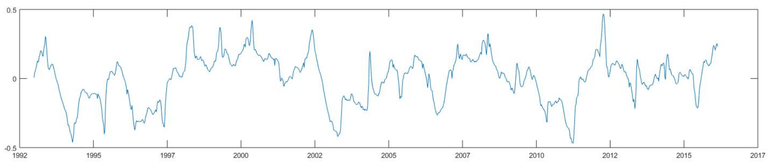

Figure 3.3: Deviation from normal overall reservoir level in Norway

3.6 Power producer data

Our study is based on observed time series of inflow, production and reservoir levels for a single hydropower station in Norway. Here we use data from Kolsvik, a reservoir owned by Helgelandskraft.

3.6.1 Production planning model

We want the planning model used in our analysis to be a reasonable approximation of what is done in real life for our case power plant. Also the model should be general enough to have the possibility of incorporating more hydropower planners later on. Assumptions we use in our model include: one-reservoir approximation, constant head assumption, sufficient reservoir flexibility assumption, sufficient production capacity assumption, price taker assumption, no marginal production cost and insignificant start-up and shutdown cost. In the appendix, Section 6.1, we elaborate on the mentioned assumptions and discuss if they are applicable for our case plant.

For production planning Helgelandskraft buys planning models from Markedskraft and SKM, suppliers of power market analyses and corresponding services for the energy markets. Both Markedskraft and SKM Energy use the detailed EMPS simulation model (Samkjøringsmodellen), developed by SINTEF, to forecast prices (Markeskraft, 2016). In Helgelandskraft, these forecasts are directly used in the VanSimTap model, provided by SKM Energy (SKM-Energy, 2016). As the EMPS model are the most commonly used, results from Kolsvik can give some initial insight into the industry standard. Also the Van-SimTap use an one reservoir approximation, the same as we use. Therefore Kolsvik is a suitable case for our analysis.

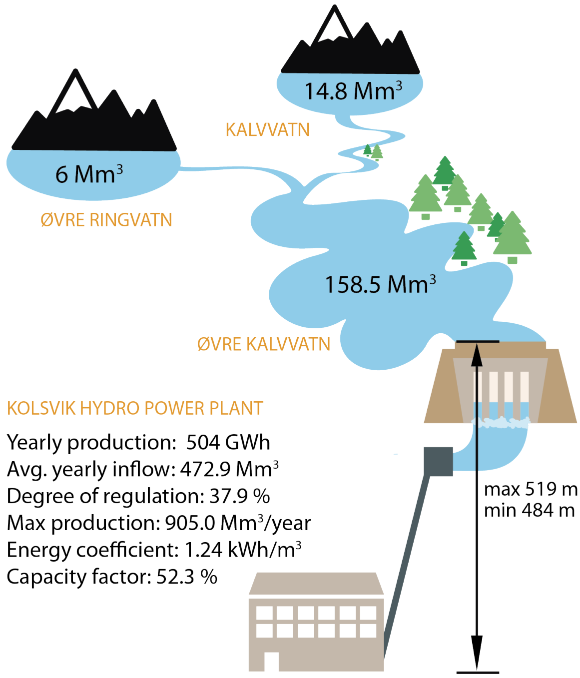

Figure 3.4: Overview of Helgelandskraft - Kolsvik

3.6.2 Inflow data

Daily and weekly inflow from 28.03.2000 to 31.12.2015, for the relevant catchment area, are shown in Figure 3.5. The data shows clear seasonal tendencies. The low periods can be explained by precipitation coming down as snow during winter, resulting in low inflow. The high periods can be explained by snow melting during spring and summer, resulting in high inflow and spring floods. There is also some small peaks in the fall, when there usually is a lot of rain in Norway. The occurrence of negative inflow levels, especially for daily inflow, indicates low measurement precision. For a weekly resolution, this will not be a prominent problem when we, as done by Kolsrud and Prokosch (2010), assume that producers correct the errors observed from negative values quickly after they occur. When that is the case, aggregating the data to a weekly resolution accounts for part of the measurement errors. The remaining negative values in the data set can be set to 0. When doing this, the difference between processed total seasonal inflow and original total seasonal inflow is relatively low, at only 0.37%. For a daily resolution, removing the negative values will result in a 1.16% difference in seasonal accumulated inflow for processed and unprocessed data. The reduction in difference when reducing the resolution indicates that Kolsrud and Prokosch (2010) are right when assuming the producers correct the measurement errors. Later on, when modeling the inflow time series, in Section 4.1, the same removal of negative values is done. When modelling inflow we also work with the absolute values. We could have used a transformation to log prices, similar to what we do for the price. Using the logarithm is however not an option, since the inflow is often zero. A different transformation would have been needed. Further work on this could possibly improve the model.

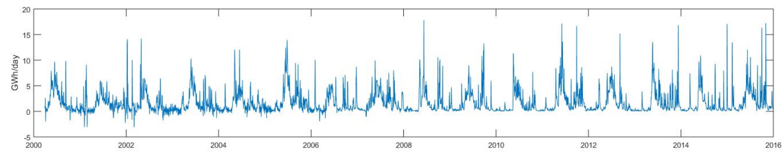

(a) Daily

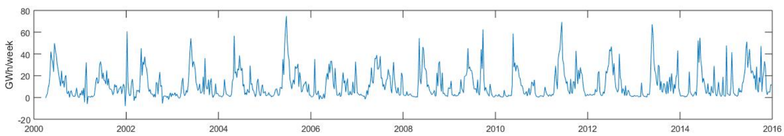

(b) Weekly

Figure 3.5: (a) Daily and (b) weekly inflow from 2000 to 2016

3.6.3 Reservoir data

We need information on the local reservoir levels as they are used as a factor in our price model. We have reservoir data measured each hour for 26.09.1996 to 31.12.2015 from Kolsvik. As we have a weekly resolution in our analysis, we convert the resolution of observed reservoir levels by using the average reservoir level each week and interpolating where reservoir data are missing. Weekly reservoir level data are plotted in Figure 3.6. The plot reveals a yearly seasonal trend, where reservoir levels are at their lowest during spring, for so to increase towards max during summer and fall. The seasonality is a result of the relationship between demand and water availability.

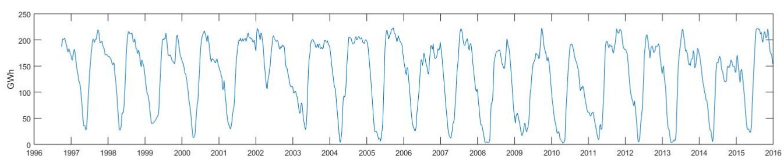

Figure 3.6: Local reservoir level, from 1996 to 2015

3.6.4 Production data

Production time series are used in the maximum likelihood estimation, where our model production results are compared to observed time series on production. For Kolsvik, the production data are given on an hourly basis from 01.01.1996 to 31.12.2015. When converting to weekly resolution, we simply sum all hourly production data for each week. From Figure 3.7, we observe that using three discretization levels for the production is a reasonable assumption for an hourly production. When aggregating the data for daily and weekly resolutions, the validity of the assumption decrease. Nevertheless, three production levels are still clearly visible, even at a weekly resolution. To cope with the large amount of memory needed for the structural estimation model, we find it necessary to use only two levels, which is the bare minimum.

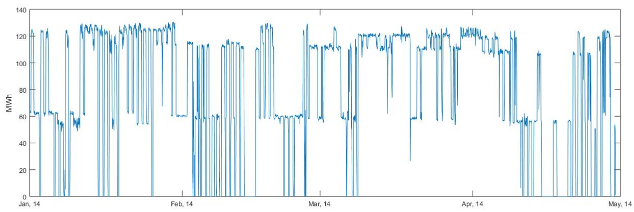

(a) Hourly

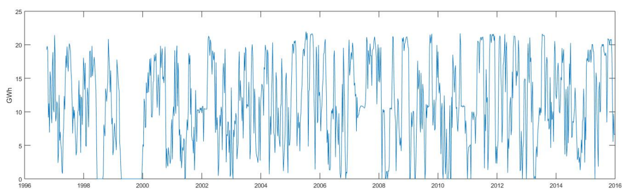

(b) Weekly

Figure 3.7: (a) Hourly production for the first quarter of 2014 and (b) weekly production from 1996 to 2016

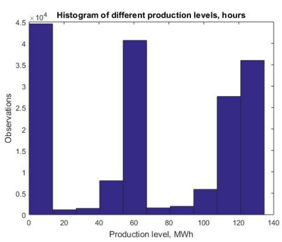

(a) Hours

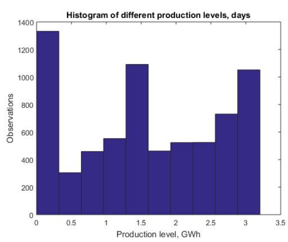

(b) Days

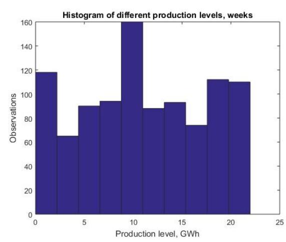

(c) Weeks

Figure 3.8: Histogram of different production levels, for hours, days and weeks

3.7 Correlation between data series

To form a basic understanding of the relationships between the different observed time series, later on accounted for in the time series model, we briefly presents how they correlate.

Correlation between inflow and reservoir deviation

High inflow in one region of Norway over a longer time will raise the total reservoir level in the country. Hence we can expect that there is a positive correlation between the accumulated inflow in one area of Norway and the reservoir level in all of Norway. More specifically, we look for a correlation between the cumulative local deviation from the normal historical inflow value, and the overall reservoir level deviation from the normal historical value in Norway as a whole. The amount of correlation will vary depending on where in Norway the power plant is, the size of the area and topography of the inflow catchment area. The correlation should anyhow always be positive. In Figure 3.9 we have plotted the deviation from normal cumulative local inflow for Kolsvik, and deviation from normal overall reservoir level in Norway. There is clearly a high correlation between the two time series, and the correlation coefficient, of 0.6578, is significantly different from zero.

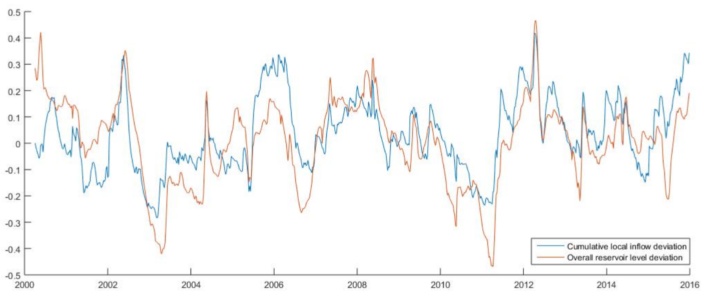

Figure 3.9: Deviation from normal cumulative local inflow (blue) and deviation from normal overall reservoir level in Norway (red), from 2000 to 2016

Correlation between spot price and reservoir deviation

As mentioned in Section 3.1, hydropower production accounts for about half of the total electricity generation in the Nordic region. And we showed in Section 3.5 that deviation from normal overall reservoir level in Norway is a good approximation for the deviation in all Nordic countries. We therefore expect the system spot price to be negatively correlated with the overall reservoir deviation in Norway. We show time series of the two in Figure 3.10.

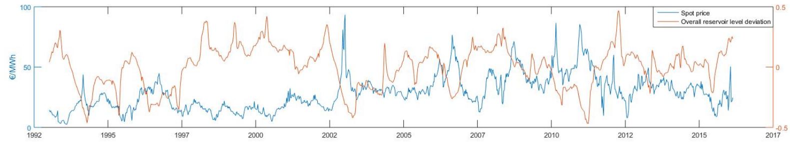

Figure 3.10: System spot price (blue) and deviation from normal overall reservoir level in Norway (red), from 1993 to 2016

From the figure we can see a trend that high overall reservoir deviation seem to lead to low spot price, and vice versa. We can also see that sudden drops or inclines in overall reservoir deviation may cause the reverse behaviour in spot prices. From this we conclude that there is a negative correlation between the two, something the correlation coefficient of -0.4241 confirms.

Correlation between inflow and spot price

Based on the two previous paragraphs we expect the cumulative local inflow deviation to also be negatively correlated with system spot price. Figure 3.11 shows the two time series, with inflow from Kolsvik. The figure doesn't show as strong dependency as Figure 3.10 did, in the previous section between price and overall reservoir deviation, which a correlation coefficient of -0.275 also shows. It is however significantly different from zero, and the figure also shows clear tendencies of negative correlation.

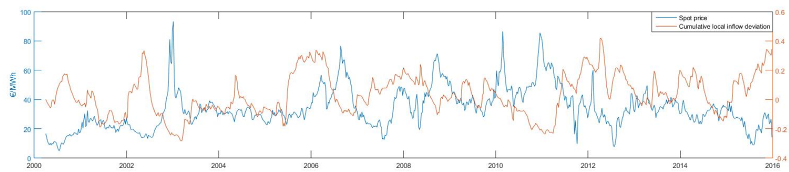

Figure 3.11: System spot price (blue) and deviation from normal cumulative local inflow (red), from 1993 to 2016

Correlation between spot price and forward price

Since the electricity forward price is written on the electricity spot price, we expect that they are highly correlated. This correlation vary with forward contracts with different maturity dates, the ones with shorter time to maturity being more correlated with the spot price than those with longer time to maturity. But all of them should still be positively correlated with the spot price. Figure 3.12 shows time series for the system spot price and a forward price with 2 months to maturity. The forward price closely follows the spot price, with the spot price having a slightly higher variance. The correlation coefficient is as high as 0.8692. (Fleten et al., 2008) points out that the seasonal component of the forward price is in total accordance with the behavior of the system price, as we also can see in Figure 3.12.

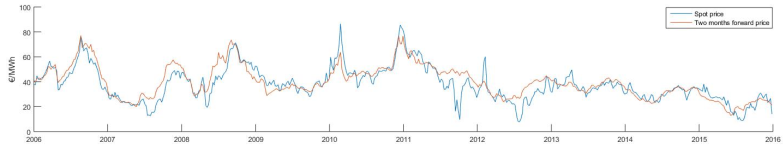

Figure 3.12: System spot price (blue) and 2 months to maturity electricity forward price (red), from 2006 to 2016

Correlation between forward price and reservoir deviation
Based on the previous paragraph, and the fact that spot price and overall reservoir deviation are negatively correlated, we expect the forward price and the overall reservoir deviation to also be negatively correlated. Figure 3.13 shows a forward with 2 months to maturity and the overall reservoir deviation, and a correlation coefficient of -0.456 confirms there is a significant correlation between the two.

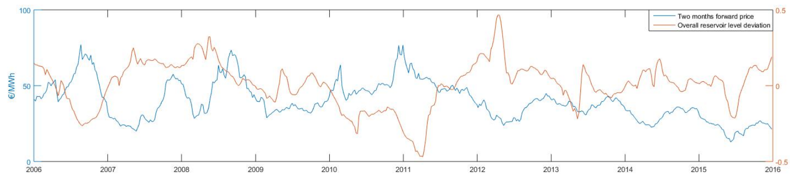

Figure 3.13: Electricity forward price with 2 months to maturity (blue) and deviation from normal overall reservoir level in Norway (red), from 2006 to 2016

We have now described the environment for hydropower planning and the details for a specific hydropower plant. This information will from now on be used together with the structural estimation model, introduced in Chapter 2.

Chapter 4

Structural estimation for hydropower

The aim of this chapter is to put structural estimation, defined in Chapter 2 in context of a general hydropower planning problem, described in Chapter 3, and to explain the background for specific model detail decisions. Figure 4.1 provides an overview of the model structure used for our analysis.

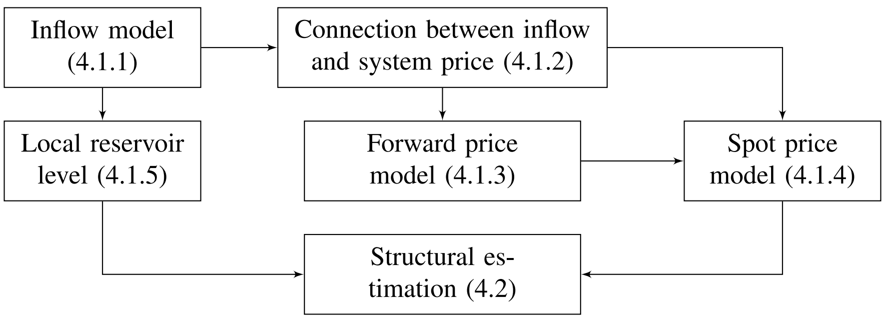

Figure 4.1: Overview of the model structure

The overall model use a combination of sub-models to analyze the decision-making done for a hydropower plant. In order to obtain the parametric approach for conditional expectations we model the state variables in Section 4.1.1, 4.1.2, 4.1.3 and 4.1.4 as *Markovian* time series. This enables us to use the time series in a parametric approach for the structural estimation problem in Section 4.2.

4.1 State variable processes

In this section we describe the different state variables in our hydropower case, and their state transition processes. The state variables used are the price, P_t, used in the profit function, and the reservoir level, S_t, and inflow, I_t, which are needed to know the reservoir level evolution. We also want a connection between inflow and price. Therefore we include the cumulative local inflow deviation, C_t, and the overall reservoir deviation, R_t, as state variables. Lastly, we have to model the forward price, F_t, which is closely linked to the structural parameter we want to estimate, and enters into the price process. Hence forward price is also a state variable. We assume inflow, overall reservoir deviation, price and forward price all to be auto-correlated with one lag, as elaborated in Appendix 6.1. This is in accordance with the *Markov property* in the *Markov decision process*, i.e. the state transition can only depend on the current state.

4.1.1 Inflow, I_t

To include both the serial correlation and the seasonal variations, we build a model consisting of two parts. As our aim is to duplicate the decision regime used by producers, such an inflow model is suitable in our case. In the same manner as done by Kolsrud and Prokosch (2010), we model inflow as a combination of an auto-regressive base process, X_t^I, and a seasonal function, $f^I(t)$.

$$I_t = f^I(t) + X_t^I \tag{4.1}$$

Inflow seasonality
The seasonal component is

$$f^I(t) = A_1^I cos(\frac{2\pi}{52}t + \phi_1^I) + A_2^I cos(\frac{4\pi}{52}t + \phi_2^I) + D^I, \tag{4.2}$$

where A is the amplitude, ϕ, the phase shift and a constant, D. This seasonal function is for a weekly resolution, hence the frequency is $\frac{2\pi}{52}$. The parameter values are estimated through a Fourier transformation, and can be found in Table 6.1, in Appendix 6.2. The resulting inflow seasonality can be seen in Figure 4.2. We can see that the amount of inflow in the spring, due to snow melting, vary from one year to the next. Hence we are unable to capture the total effect by only using a simple deterministic seasonal function. In order to predict the magnitude of the inflow coming in during the melting season information about the snow packs in the mountain could have been used. We do however not have data for snow packs, and thus we do not model such an effect.

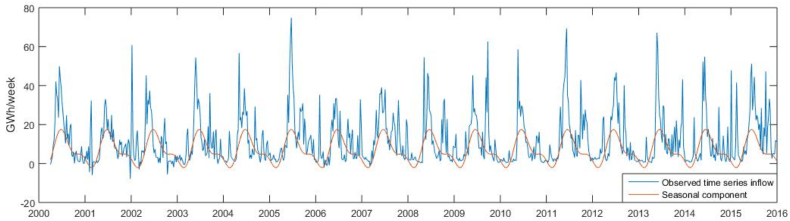

Figure 4.2: Observed inflow time series (blue) and seasonal component (red)

Base process

In order to obtain a Markovian inflow model, we use a simple auto-regressive model with one lag for the base inflow, X_t^I.

$$X_t^I = \gamma X_{t-1}^I + \varepsilon_t^I \tag{4.3}$$

In (4.3), γ is the auto-regressive coefficient, and ε_t^I is a normal random error with mean 0 and standard deviation σ^I. The parameters are estimated based on the observed inflow time series, by using the *estimate*-function in MATLAB. We find γ to be 0.5212 and σ^I to be 14.1. Figure 4.3 shows the resulting simulation based on the parameters, together with the observed base inflow time series.

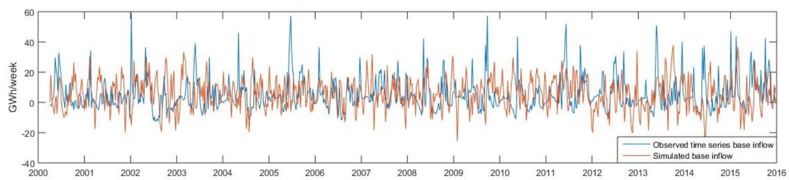

Figure 4.3: Observed base inflow time series (blue) and simulated base inflow (red)

Removal of negative inflows

In our observed inflow time series, described in Section 3.5, we see that the weekly inflow at times drops below zero. Negative inflow values can create problems in our SDP algorithm, since negative inflow would mean there is a possibility to reach a negative reservoir level. Hence we want to avoid this problem in our model. To do this we set all negative values to zero. Doing this will raise the mean slightly, something we fix by adjusting the seasonal component downwards the same amount.

Final process

Eq. (4.4) shows the inflow state variable transition expressed in terms of other state variables, in this case only itself.

$$I_t = \gamma I_{t-1} + \varepsilon_t^I + f^I(t) - \gamma F^I(t-1) \tag{4.4}$$

A simulation for the resulting inflow process, in comparison to the observed inflow time series, is plotted in Figure 4.4. The descriptive statistics in Table 6.4, in Appendix 6.2, shows the mean, standard deviation, minimum value, maximum value and the median for both the observed inflow time series and when simulating (4.4). Both the mean and standard deviation match quite well. The minimum value differ a little, but that is because we set the simulated values larger or equal to zero.

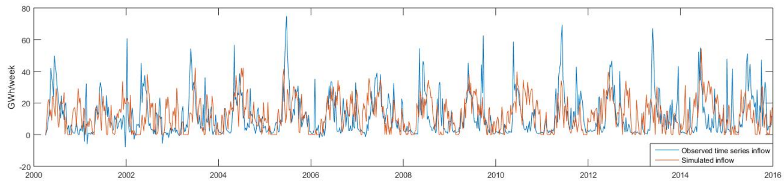

Figure 4.4: Observed inflow time series (blue) and simulated inflow (red)

4.1.2 Connection between inflow and system price

As mentioned in Section 3.1, there is a strong connection between local inflow, the overall reservoir level in Norway and the system price in the Nordic countries. This is because hydropower in these countries constitutes a large portion of total power production, and the fact that there is often similar weather across the region because it is a relatively small geographical area. Kolsrud and Prokosch (2010) present a model where this mentioned connection between local inflow, the aggregate reservoir level in Norway and the system price is accounted for. Instead of using the actual level, they use the deviation from the historical average as a scaling factor. In Section 3.7 we observed overall reservoir deviation to be positively correlated with local inflow and a negative correlation between overall reservoir level in Norway and the system price. As the local inflow cannot fully explain the overall reservoir deviation, we model the aggregate reservoir deviation as its own state variable, R_t, which is depend on the cumulative local inflow deviation, C_t, and has an additional error term. By doing this, we account for the difference between local inflow and aggregate reservoir level, resulting from differences between local inflow and inflow in other parts of the country (Kolsrud and Prokosch, 2010).

Deviation from normal cumulative inflow, C_t

Since overall reservoir deviation, R_t, depends on cumulative local inflow deviation, C_t, C_t needs to be a state variable. All state variables need to be Markovian. Hence we need to model C_t so that it only depends on the previous state. In order to achieve this, we model i_t as an exponentially weighted moving average. This way the cumulative inflow deviation in this state will only depend on local inflow in the previous state and the cumulative inflow deviation in the previous state.

The cumulative inflow is defined as

$$i_t := \sum_{j=1}^{\infty} \rho^{j-1} I_{t-j} \tag{4.5}$$

Here i_t is the cumulative local inflow, weighted heavily on recent inflow, and less on inflow information from long ago. I_{t-j} is the inflow at time $t-j$, and ρ is a weighting factor, deciding how much influence the cumulative inflow levels of last week should have on the levels in this weeks. Eq. (4.5) can be written as

$$i_t = I_{t-1} + \rho i_{t-1} \tag{4.6}$$

We would like i_t to reflect the cumulative inflow, approximately, the last half a year. First we calculate the cumulative inflow the last 26 weeks, i_t^*, from the inflow time series directly, i.e. i_t^* is always the sum of the inflow the last 26 weeks.

$$i_t^* = \sum_{k=1}^{26} I_{t-k} \tag{4.7}$$

We can then estimate ρ by fitting i_t to i_t^*. Doing this we find ρ to be 0.97.

Deviations from historical average cumulative inflow, C_t, is defined as

$$C_t := \frac{i_t - \bar{i}_t}{\bar{i}_t}, \tag{4.8}$$

where i_t is the cumulative inflow from (4.6), while $\bar{i}_t$ denotes the average of the observed cumulative local inflow over the corresponding sequence of days. Inserting (4.6) into (4.8) gives (4.9), which is the state variable transition for the cumulative local inflow deviation.

$$C_t = \frac{I_{t-1} + \rho\bar{i}_{t-1}C_{t-1} + \rho\bar{i}_{t-1}}{\bar{i}_t} - 1 \tag{4.9}$$

An example scenario for the resulting model from (4.9) is illustrated in Figure 4.5. In the descriptive statistics in Table 6.1, in Appendix 6.2, we can see that the values match quite well, but that the variance is a little higher for the simulation. This should however not be a problem later, since we discretize the cumulative local inflow deviation into only high and low values.

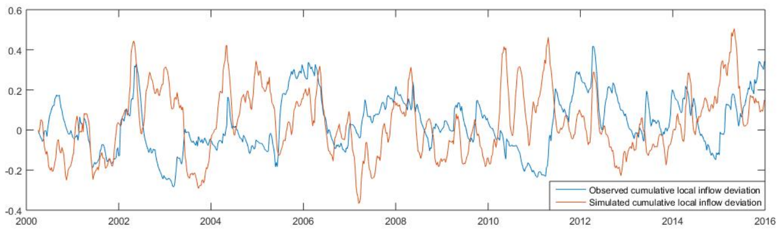

Figure 4.5: Deviation from normal cumulative local inflow, observed time series (blue) and simulation (red)

Deviation from aggregate reservoir level, R_t

Deviations from historical average aggregate reservoir level, R_{t-1}, is defined as

$$R_t := \frac{r_t - \overline{r}_t}{\overline{r}_t} \tag{4.10}$$

Here r_t is the aggregate reservoir level in Norway at day t, while $\overline{r}_t$ is the average of the observed aggregate reservoir levels in Norway at day, t, in a year. We use the same approach as Kolsrud and Prokosch (2010) to model the deviation from normal aggregate reservoir level in

$$R_t = \beta_1 R_{t-1} + \beta_2 C_{t-1} + \varepsilon_t^R \tag{4.11}$$

In (4.11), R_t is an ARX(1)-process depending on yesterdays aggregate reservoir level deviation, R_{t-1}, and yesterdays cumulative local inflow deviation, C_{t-1}, as an additional predictor. Here β_1 is the auto-regressive coefficient, β_2 is the coefficient for the predictor and ε_t^R is a normal random error with mean 0 and standard deviation σ^R. With the *estimate*-function in MATLAB we find β_1 to be 0.956, β_2 to be 0.0559 and σ^R to be 0.0247. Eq. (4.11) is the state transition for the aggregate reservoir level deviation. A resulting trajectory from eq. (4.11) together with the corresponding observed time series are shown in Figure 4.6. In the descriptive statistics in Table 6.1, in Appendix 6.2, we see that the values for observed and simulated process is quite similar. The simulation has a slightly higher mean, but we deem that acceptable.

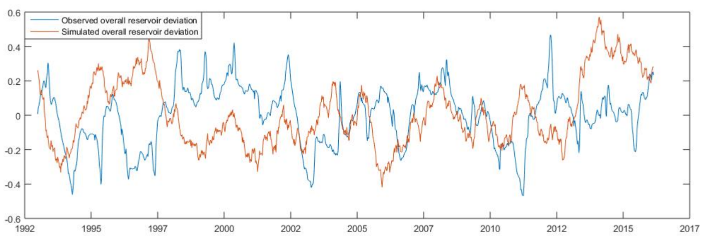

Figure 4.6: Deviation from normal overall reservoir level, observed time series (blue) and simulation (red)

4.1.3 Forward price, F_t

We want to investigate if forward prices affect producers expectation of future prices. To do so, we need to incorporate a process for the forward price into the price process. By letting the time to maturity for the forward be sometime into the future, T, we can test if the producer uses the forward price $F_{t,T}$, at time t, when planning for the future. When going from time t to time $t+1$, the time to maturity will change from T to $T+1$. This way the forward price process will always have a constant time to maturity, $T-t$. From here on we write the forward price $F_{t,T}$ as F_t, to simplify notation. The forward price process is modelled with a deterministic seasonal component, $f^F(t)$, and a base process (or error), X_t^F.

$$F_t = f^F(t) + X_t^F, \tag{4.12}$$

Seasonality, $f^F(t)$

The seasonal component for the forward price is

$$f^F(t) = A^F cos(\frac{2\pi}{52}t + \phi^F) + D^F, \tag{4.13}$$

where A^F is the amplitude, ϕ^F, the phase shift and D^F is a constant. As with the inflow, the seasonal function is for a weekly resolution. Hence the frequency is $\frac{2\pi}{52}$. The parameter values are estimated through a Fourier transformation, and are shown in Table 6.2, in Appendix 6.2. The resulting seasonal function for the forward price is plotted in Figure 4.7.

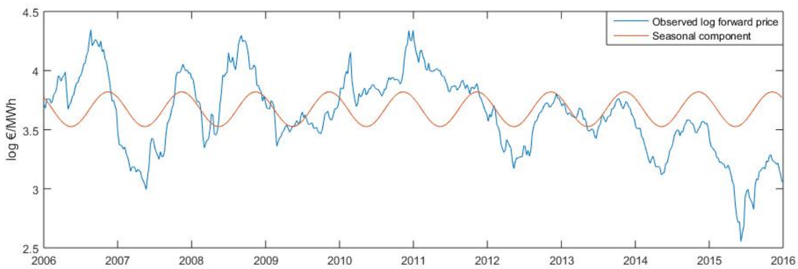

Figure 4.7: Observed 2 months to maturity forward price time series (blue) and seasonal component of the 2 months to maturity forward price (red)

Base process, X_t^F

The base process is modelled as an ARX(1)-process where X_t^F is dependent on the value of the previous day, X_{t-1}^F. In the Nordic market, the forward price dynamics depend on the observed hydrological situation at the time. This involves the current content of water and snow in the reservoirs, hydro inflow and prices (Audet et al., 2004) and (Botterud et al., 2010). Therefore the overall reservoir deviation, R_{t-1}, are included as an additional predictor in

$$X_t^F = \varphi_1 X_{t-1}^F + \varphi_2 R_{t-1} + \varepsilon_t^F \tag{4.14}$$

As we pointed out in the data description in Section 4.1.2, the aggregate reservoir deviation, R_t, is dependent on the cumulative local inflow deviation, C_t. Thereby, the inclusion of R_t, accounts for both the influence of inflow and reservoir levels. In (4.14) φ_1 is the auto-regressive parameter, φ_2 is the parameter for the predictor and ε_t^F is a normal random error term with mean 0 and standard deviation σ^F. The parameters are estimated using the *estimate*-function in MATLAB and subsequently manually fine tuned in order to better match the moments. The parameter values for the different forwards are shown in Table 4.1

Table 4.1: Parameter values for forward with 2 months-, 6 months- and 1 year to maturity

	φ_1	φ_2	σ^F
2 months	0.96	-0.02	0.08
6 months	0.97	-0.019	0.057
1 year	0.98	-0.0031	0.0451

Final process

The state variable transition for the forward price, expressed in terms of other state variables is then

$$F_t = \varphi_1 F_{t-1} + \varphi_2 R_{t-1} + \varepsilon_t^F + f^F(t) - \varphi_1 f^F(t-1) \tag{4.15}$$

Figure 4.8 shows a trajectory for the forward price with 2 months to maturity together with the corresponding observed time series. In Table 6.1, in Appendix 6.2, the descriptive statistics for the observed forward time series and simulation is shown, for time to maturity of two months, six months and one year. Both the means and the standard deviations are good matches for all the processes, while the min and max values are not perfect matches.

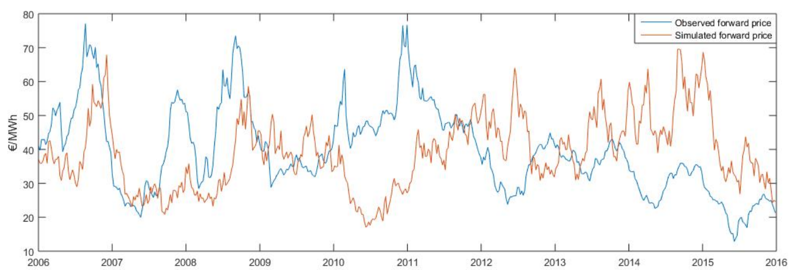

Figure 4.8: Observed 2 months to maturity log forward price time series (blue) and simulated 2 months to maturity log forward price (red)

4.1.4 Price, P_t

We define the entire price process as

$$P_t := f^P(t) + Y_t^P + \zeta(F_t - Y_t^P - f^P(t)) \tag{4.16}$$

It consists of a deterministic seasonal component, $f^P(t)$, and a stochastic base process, Y_t^P. The expression also has an adjustment factor, ζ, to include the forward price into the process. In our model, we use the forward price as the expected spot price. However, in reality the difference is not 0, as the forward often includes a risk premium. The forward price is by Näsäkkälä and Keppo (2008) described as the risk adjusted expected spot price. In our model we assume the risk premium to be so small it disappears in the noise of the base price process.

How much a producer decides to use the forward price in favor of the seasonal function and base process is denoted by the parameter ζ. If ζ is 1 the producer only uses the forward price to form expectations, if ζ is 0 a producer is indifferent towards the forward price. The emphasis placed on forward price information in production planning, ζ, is the main parameter we want to estimate.

To estimate the price process model, we first assume that the price process is explained without the forward price. Hence, when we estimate the parameters of the base process and seasonality function, we set ζ equal to zero. In that case the price process is simply reduced to only include the base process for the price Y_t^P plus the seasonality $f^P(t)$

$$P_t := f^P(t) + Y_t^P \tag{4.17}$$

Seasonality
The seasonal function for the price is

$$f^P(t) = A^P cos(\frac{2\pi}{52}t + \phi^P) + D^P, \tag{4.18}$$

where A^P is the amplitude, ϕ^P is the phase shift and D^P is a constant. The seasonal function for the price is with a weekly resolution, hence the frequency $\frac{2\pi}{52}$. The parameters is estimated with Fourier transformation, and are shown in Table 6.3, in Appendix 6.2. The resulting seasonal function, together with the observed price time series, can be seen in Figure 4.9.

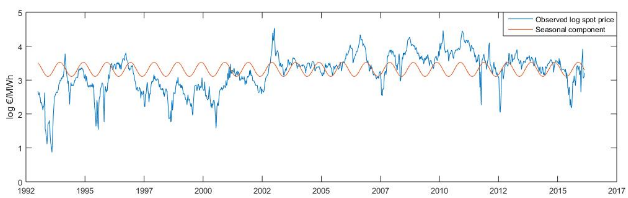

Figure 4.9: Observed log price time series (blue) and seasonal component of the log price (red)

State-space model

As remarked on in the data description Section 3.7, the spot price is negatively correlated with the aggregate reservoir deviation. Since the deviation from normal overall reservoir level is free of any seasonality, the correlation is incorporated in the base process of the price. Therefore we model the base process as a state-space model. A state-space model have an underlying AR(1)-process, X_t^P, which is unknown from an observers point of view. The underlying process X_t^P is

$$X_t^P = \delta X_{t-1}^P + \varepsilon_t^P, \tag{4.19}$$

where δ is the auto-regressive coefficient and ε_t^P a normal random error with mean 0 and standard deviation σ^P. The state space model also includes an observed process, Y_t^P, which is our base price process. The base price, (4.20), consists of the unknown underlying process X_t^P plus an additional predictor. In our case, the predictor is a term depending on the overall reservoir deviation, R_t.

$$Y_t^P = X_t^P + \eta R_t \tag{4.20}$$

The predictor term, ηR_t, will act as a mean reversion level for the price. The base process of the price, Y_t^P, will in the long run move towards the value of R_t, scaled by some parameter η. The scaling parameter, η, is effectively a measure of how much the price is affected by the aggregate reservoir deviation, which again is affected by local inflow, giving a connection between local inflow and price.

Using ηR_t as a mean reversion level for the base price was proposed by Kolsrud and Prokosch (2010). They estimated the η parameter by changing it to fit the correlation between Y_t^P and R_t to its historical level. We take a different approach, by formulating the base price as a state-space model. Based on the observed price time series, the parameters of the model is estimated via maximum likelihood, with the help of *Kalman filtering* in MATLAB (Durbin and Koopman, 2001). We find δ to be 0.96, η to be -0.3192 and σ^P to be 0.102. A simulation of the base price, the time series for the base price and the overall reservoir deviation time series are plotted in Figure 4.10.

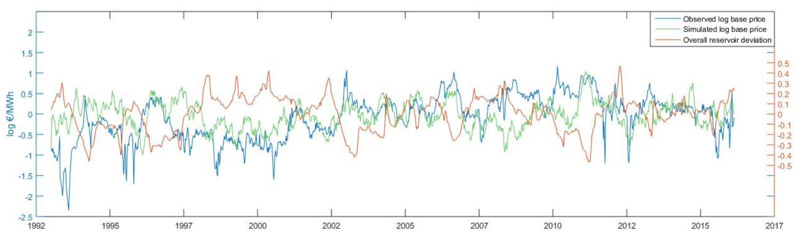

Figure 4.10: Observed log base price time series (blue), simulated log base price (green) and overall reservoir deviation (red)

Final process
The state variable transition for the price, expressed in terms of other state variables is

$$
\begin{aligned}
P_t = \delta P_{t-1} + \zeta(\varphi_1 - \delta)F_{t-1} + ((\beta_1 - \delta)\eta + \zeta(\varphi_2 - (\beta_1 - \delta)\eta))R_{t-1} + \\
(1-\zeta)\eta\beta_2 C_{t-1} + (1-\zeta)\varepsilon_t^P + \zeta\varepsilon_t^F + (1-\zeta)\eta\varepsilon_t^R + \\
f^P(t) - \delta f^P(t-1) + \zeta(f^F(t) - \varphi_1 f^F(t-1) - (f^P(t) - \delta f^P(t-1)))
\end{aligned} \quad (4.21)
$$

A trajectory of the resulting price process, in comparison to the observed price time series and the overall reservoir deviation time series is plotted in Figure 4.11.

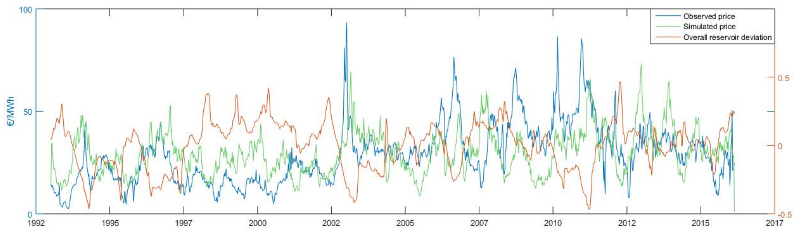

Figure 4.11: Observed price time series (blue), simulated price (green) and overall reservoir deviation (red)

We can see that the simulation reacts to the extremes in overall reservoir deviation in a similar manner as the observed price does. It is however hard to model the spikes in the price at the correct moments. This is because the observed price does not react uniformly to similar values of the overall reservoir deviation, making the extremes tough to predict. The unpredictable behavior can come from other factors, like higher demand due to lower temperature, but our model is not advanced enough to include such factors. Including simultaneity to react to influence from the demand in price is suggested in the assumptions section, 6.1, of the appendix as a possible improvement for the future.

Descriptive statistics for the price are in Table 6.1, in Appendix 6.2, for both the observed time series and the simulation. Here we can see that both the mean and the standard deviation is very close for observation and simulation.

4.1.5 Local reservoir level, S_t

Hydro reservoir balance
We also need to describe how the reservoir level is updated when going from one state to the next. The hydro reservoir balance in (4.22) describes this transition. In contrast to the other state variables, I_t, C_t, R_t, F_t and P_t, who are exogenous state variables in our model, S_t is a endogenously calculated state variable. This means it is dependent on the decision variable d.

$$S_t = min\{S_{t-1} + I_{t-1} - u(X_{t-1}, d_{t-1}), S^{max}\} \tag{4.22}$$

The reservoir level at time t, S_t, is the level at time $t - 1$, S_{t-1}, plus the inflow into the system from time $t - 1$ to t, I_{t-1}, minus the amount released from time $t - 1$ to t, $u(X_{t-1}, d_{t-1})$, as long as this sum is less than the maximum reservoir level allowed, S^{max}. Otherwise the reservoir level is equal to S^{max}. Any excess water will then spill over, and not be utilized. The release function is described in Section 4.2. The hydro reservoir balance could also reflect evaporation loss. In large reservoirs, evaporation affects are relatively small and they do not effect the nature of our problem. For that reason they are not included (Tejada-Guibert et al., 1993).

Reservoir storage constraint
Eq. (4.23) forces the reservoir level, S_t, to stay inside the lower and upper bounds, S^{min} and S^{max} respectively.

$$S^{min} \leq S_t \leq S^{max} \tag{4.23}$$

This reservoir constraint equation is not explicitly enforced in our model, but implicitly included through (4.22) for the upper bound, S^{max}, and (4.25) for the lower bound, S^{min}.

4.1.6 State discretization

Discretizing the state variables into a finite set of levels are necessary when implementing the model. The main reason for this is that we need to know the value function, t_θ in (2.27) at each possible state. Hence there needs to be a finite number of states in order to limit the memory usage to a feasible range for the solver. Increasing the number of states, i.e. the number of discrete levels for each state variable, has a huge impact on the memory usage. In order to capture as much of the dynamics of the problem as possible, the state space discretization should be as dense as possible, while staying below the memory limit of the solver.

With discretizing state variables we need to map values that fall outside the grid to the optimal grid point. This is done by simply rounding to the closest point. When discretizing the different state variables, it is important to remember that they interact with each other. This is especially important to take into consideration for the reservoir level transition, denoted in equation (4.22). If the sum of inflow and release at time t is not large enough to let the mapping of the reservoir level increase/decrease from time t to time $t + 1$, releasing water will not affect the reservoir level. That again means you can generate electricity without loosing any of the water in your storage. To avoid this effect we discretize the three variables incorporated in the state transition, inflow, reservoir level and release, on the same grid. That avoids the problem of producing without changing reservoir level, as well as not having to round to the closest grid point for the reservoir level, since the next state reservoir level always falls on the grid.

The fact that inflow, release and reservoir level need to lie on the same grid leads to some problems. The numeric difference between an empty and a full reservoir is large, compared to the size of weekly inflow levels. Therefore the reservoir level needs to be divided into a lot of levels for one unit of inflow and release in order to fall on the same grid. For the reservoir transition to work as it should, inflow, release and reservoir levels should also be discretized so that is captures the state transition dynamics.

As maintaining the dynamics of inflow, reservoir level and release are a major priority when discretizing the state space, the discretization of the other state variable might suffer. Variables like the price and the overall reservoir deviation are therefore severely limited in order to not exceed the memory limit. As weeks are also used as an additional state variable, memory usage is even more restrained. In our case the model is therefore forced to contain the minimum amount of discrete levels for all of the state variables to allow the program to run without further algorithm improvements.

4.2 Structural estimation

Based on the processes from Section 4.1, in combination with knowledge about the state space composition, we formulate a structural estimation model for a single agent hydropower producer. First we define what structural parameters we want to estimate, before we describe the specific profit function. After this we describe the state space transition in relation to the state variables in the previous section, and then we use this to finally get our fixed point Bellman operator for the hydropower producer case, which we use in the maximum likelihood estimation constraint from Section 2.19.

Structural parameters to be estimated
The structural parameters, or primitives, we want to estimate are collected in the set θ. As these parameters are changed we can estimate what value gives the closest fit to observed time series by maximizing the log likelihood value. In our hydropower specific case θ contains only ζ.

$$\theta = (\zeta) \tag{4.24}$$

Here ζ denotes to what degree the forward price information is taken into account when deciding on water release.

Release, $u(\cdot)$
The release, $u(X_t, d_t)$, is the amount of water, in GWh, produced when a decision, d_t, is taken. It is simply the decision, d_t, scaled with a production factor, Q. However, if there is not enough water in the reservoir to release as much as the producer wants to, the release is set to the current reservoir level, S_t, minus the lower limit for the reservoir level, S^{min}, plus the inflow, I_t. This way the release is never large enough to empty the reservoir below its lower limit. The lower restriction on the release is the reason $u(X_t, d_t)$ is dependent on the current state X_t.

$$u(X_t, d_t) = min\{S_t - S^{min} + I_t, d_t Q\} \tag{4.25}$$

Decision, d_t
At time t the single agent hydropower producer takes a decision, d_t, related to release. This decision is one level in the discrete decision space D. The reason the decision space needs to have discrete levels is that we also discretize the different state variables, including S_t and I_t, as mentioned in section 4.1.6. Since they are discretized, S_t and I_t stay on a grid, which means that for (4.22) to make sense, $u(X_t, d_t)$ needs to stay on the same grid, i.e. be discretized. For a similar reason, the amount of discrete levels, H, is dependent on the discretization of the other state variables. This is because a finer discretization on release, compared to for instance reservoir level, could mean a decision is made that does not affect the reservoir level transition, as mentioned in Section 4.1.6. The discretized decision space is then $d_t \in D = \{1, ..., H\}$.

Profit function, $g(\cdot)$
The profit function to be used in the constraint of (2.19) needs to be defined. In appendix 6.1 we assume no marginal production cost and no start-up cost. We also assume a constant head and that the single agent power producer is a price taker, hence the profits from production can be formulated as a simple linear relationship between price and produced electricity, without any cost part.

$$g(X_t; d_t) := P_t u(X_t, d_t) \tag{4.26}$$

State space transition
The different state variables in Section 4.1 are combined into the state space vector x_t.

$$x_t = (I_t, C_t, R_t, F_t, P_t, S_t) \tag{4.27}$$

Here I_t is the inflow, C_t the deviation from normal cumulative local inflow, R_t the deviation from normal overall reservoir level, F_t the forward price with maturity some specified time in the future, P_t is the price and S_t the local reservoir level.

The state transition function for going from state x_t to x_{t+1}, (2.1), is then for a single agent hydropower producer

$$x_{t+1} = f(x_t, d_t) = (I_{t+1}, C_{t+1}, R_{t+1}, F_{t+1}, P_{t+1}, S_{t+1}) \tag{4.28}$$

Using the state variable transitions described in Section 4.1, we can write (4.28) on vector form, like

$$\boldsymbol{x_{t+1}} = f(\boldsymbol{x_t}, d_t) = \boldsymbol{A_t x_t} + \boldsymbol{BLe'} + \boldsymbol{c_t} \tag{4.29}$$

Eq. (4.29) is the hydropower specific version of (2.24). Due to seasonality in some of the state variables, the matrices A and c in (2.24) need time-of-the-year indices. The matrices in (4.29) are described in detail in (4.30).

$$
\boldsymbol{x_{t+1}} = \begin{bmatrix} I_{t+1} \\ C_{t+1} \\ R_{t+1} \\ F_{t+1} \\ P_{t+1} \end{bmatrix}, \boldsymbol{x_t} = \begin{bmatrix} I_t \\ C_t \\ R_t \\ F_t \\ P_t \end{bmatrix},
$$

$$
\boldsymbol{A_t} = \begin{bmatrix} \gamma & 0 & 0 & 0 & 0 \\ \frac{1}{\bar{i}_{t+1}} & \frac{\rho \bar{i}_t}{\bar{i}_{t+1}} & 0 & 0 & 0 \\ 0 & \beta_2 & \beta_1 & 0 & 0 \\ 0 & 0 & \varphi_2 & \varphi 1 & 0 \\ 0 & (1-\zeta)\eta\beta_2 & ((\beta_1-\delta)\eta + \zeta(\varphi_2 - (\beta_1-\delta)\eta)) & \zeta(\varphi_1 - \delta) & \delta \end{bmatrix},
$$

$$
\boldsymbol{B} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & (1-\zeta)\eta & \zeta & (1-\zeta) \end{bmatrix},
$$

$$
\boldsymbol{e'} = \begin{bmatrix} \varepsilon_1 \\ 0 \\ \varepsilon_2 \\ \varepsilon_3 \\ \varepsilon_4 \end{bmatrix},
$$

$$
\boldsymbol{c_t} = \begin{bmatrix} f^I(t+1) - \gamma f^I(t) \\ \frac{\rho \bar{i}_t}{\bar{i}_{t+1}} - 1 \\ 0 \\ f^F(t+1) - \varphi_1 f^F(t) \\ f^P(t+1) - \delta f^P(t) + \zeta(f^F(t+1) - \varphi_1 f^F(t) - (f^P(t+1) - \delta f^P(t))) \end{bmatrix} \tag{4.30}
$$

$\boldsymbol{A_t}$ is the matrix containing all the parameters for the state variable transitions from Section 4.1. $\boldsymbol{c_t}$ contains the additional deterministic values, like the seasonal functions. $\boldsymbol{e'}$ is the error matrix, where ε_1, ε_2, ε_3 and ε_4 are independent standard Gaussian random errors, and correspond to ε_t^I from (4.3), ε_t^R from (4.11), ε_t^F from (4.14) and ε_t^P from (4.19) respectively. $\boldsymbol{B}$ is the matrix of parameters for the errors, while L is the lower triangular matrix of (4.31), which ensures that the errors are correlated correctly.

$$
\Sigma = \begin{bmatrix} cov(\varepsilon_t^I, \varepsilon_t^I) & 0 & cov(\varepsilon_t^R, \varepsilon_t^I) & cov(\varepsilon_t^F, \varepsilon_t^I) & cov(\varepsilon_t^P, \varepsilon_t^I) \\ 0 & 0 & 0 & 0 & 0 \\ cov(\varepsilon_t^I, \varepsilon_t^R) & 0 & cov(\varepsilon_t^R, \varepsilon_t^R) & cov(\varepsilon_t^F, \varepsilon_t^R) & cov(\varepsilon_t^P, \varepsilon_t^R) \\ cov(\varepsilon_t^I, \varepsilon_t^F) & 0 & cov(\varepsilon_t^R, \varepsilon_t^F) & cov(\varepsilon_t^F, \varepsilon_t^F) & cov(\varepsilon_t^P, \varepsilon_t^F) \\ cov(\varepsilon_t^I, \varepsilon_t^P) & 0 & cov(\varepsilon_t^R, \varepsilon_t^P) & cov(\varepsilon_t^F, \varepsilon_t^P) & cov(\varepsilon_t^P, \varepsilon_t^P) \end{bmatrix} \tag{4.31}
$$

The state variable transition for the reservoir level, (4.22), is not included in (4.30), but it is still part of the state transition on vector form in (4.29). The reason for not including it in the matrix form, is only because of the difficulty surrounding the min-operator. For convenience we therefore leave it out of (4.30), but we write it again here for consistency

$$S_t = min\{S_{t-1} + I_{t-1} - u(X_{t-1}, d_{t-1}), S^{max}\}$$

Achieving stationarity
In the general definition of structural estimation in Section 2.1 we have assumed stationarity for the Markov decision process, i.e. that the future looks exactly the same in state x_t at time t as it does in state x_{t+k} at time $t + k$, given that $x_t = x_{t+k}$. In the specific case of hydropower production, this does not hold true. The future does not look exactly the same as long as $x_t = x_{t+k}$ for all k. The seasonality in the different state variables causes the problem to be dependent on time. This is a problem, since time as a state variable will cause the total number of states to increase massively, especially for large time series. Foss and Høst (2011) propose a strategy for approximating the stationarity assumption for a hydropower problem. Since the seasonal part of the problem stays the same for different years, it will be sufficient to let the problem be conditional upon time of the year. Foss and Høst (2011) call this a weaker form of stationarity. They assume that the future looks the same from time t and time $t + k$, as long as $x_t = x_{t+k}$ for some k, where $k = \varphi i, i = 0, 1, 2, ...$ where φ is the amount of time units in a year (e.g. 52 for a weekly time index). If we set time of the year as a state variable, the problem is reduced from a non-stationary to an approximate stationary problem. When we for instance have a weekly resolution in our problem, we set $\varphi = 52$, implying that the future will look the same from time t and time $t+52i$, as long as $x_t = x_{t+52i}$. The Markov decision process is then stationary between years, for a given week. This will reduce the number of additional states due to seasonality from the total number of observation to only the amount of time units in a year. The state space will still increase a lot, but we find this way to handle the problem to be the best solution available.

Fixed point Bellman equation for hydropower
We can now use the newly stated approximate stationarity and the state space transition described earlier in this section to write a fixed point Bellman equation for a hydropower producer. In the hydropower producer case we have four random error terms, one for inflow, overall reservoir deviation, forward price and for price. Meaning that (2.27) will get a quadruple sum and corresponding weights, one for each random error term. Combining (2.27) with the state transition for our hydropower case (4.29), we get our hydropower-specific fixed point Bellman equation

$$t_\theta(v)(x, d, t) \sim \sum_{i=1}^{M_i} \sum_{r=1}^{M_r} \sum_{f=1}^{M_f} \sum_{p=1}^{M_p} w_i w_r w_f w_p h(A_t x + c_t + BL\tilde{e})), \tag{4.32}$$

where t_θ now is dependent upon time of the year. $\tilde{e}$ from (4.32) is described in (4.33). $\tilde{e}_i$, $\tilde{e}_r$, $\tilde{e}_f$ and $\tilde{e}_p$ are the center points for M_i-, M_r-, M_f- and M_p point approximations, respectively, of normal distributed random variables and w_i, w_r, w_f and w_p are the corresponding probability weights.

$$\tilde{e} = \begin{bmatrix} \tilde{e}_i \\ 0 \\ \tilde{e}_r \\ \tilde{e}_f \\ \tilde{e}_p \end{bmatrix} \tag{4.33}$$

Using the state space transition from (4.29), we can write the h-function in (2.27) as (4.34). $A_t^{(5)}$, $c_t^{(5)}$ and $(BL)^{(5)}$ represents the fifth row of each respective matrix in (4.30), and is the price part of the state transition process.

$$\begin{aligned} &h(A_t x + c_t + BL\tilde{e}) = \\ &\qquad b \cdot \log\left(\sum_{d \in D} \exp\left(\frac{(A_t^{(5)} x + c_t^{(5)} + (BL)^{(5)}\tilde{e}) \cdot u(A_t x + c_t + BL\tilde{e}, d)}{b}\right.\right. \\ &\qquad\qquad\qquad\qquad\qquad\qquad \left.\left. + \frac{\beta \cdot v(A_t x + c_t + BL\tilde{e}}{b}\right)\right) \end{aligned} \tag{4.34}$$

Eq. (4.34) can now be used to write the fixed point Bellman operator in (2.27) as (4.35), remembering that the decision, d, is part of the state transition (4.29).

$$\begin{aligned} &t_\theta(v)(x, d, t) \sim \sum_{i=1}^{5}\sum_{r=1}^{5}\sum_{p=1}^{5}\sum_{f=1}^{5} w_i w_r w_p w_f \\ &\qquad b \cdot \log\left(\sum_{d' \in D} \exp\left(\frac{(A_t^{(5)} x + c_t^{(5)} + (BL)^{(5)}\tilde{e}) \cdot u(A_t x + c_t + BL\tilde{e}, d')}{b}\right.\right. \\ &\qquad\qquad\qquad\qquad\qquad\qquad \left.\left. + \frac{\beta \cdot v(A_t x + c_t + BL\tilde{e})}{b}\right)\right), \end{aligned} \tag{4.35}$$

where the release function in (4.35) is explicitly

$$\begin{aligned} u(A_t x + c_t + BL\tilde{e}, d') = min\{&min\{S_t - S^{min} + I_t - u(x, d), S^{max}\} \\ &- S^{min} + \gamma I_t + (BL)^{(1)}\tilde{e} + f^I(t+1) - \gamma f^I(t), d'Q\} \end{aligned} \tag{4.36}$$

$(BL)^{(1)}$ represents the first row of BL, which is inflow.

Now we have established a base for structural estimation of hydropower. From this chapter, the next-state release function (4.36) describes how one decision depends on the previous one. Eq. (4.36), together with the rest of (4.35), denotes how the expectation, conditional on both state and decision, is handled in the parametric approach. Here (4.35) represent the hydropower-specific fixed point Bellman operator, used as the constraint in the maximum likelihood estimation. These equations are in the next chapter utilized in a demonstration of structural estimation for hydro power.

Chapter 5

Empirical Analysis

As mentioned in the introduction, the goal of this study is to examine the potential for applying structural estimation theory of Markov Decision Processes, developed by Rust (1987), to a single hydropower planning problem. In the case where such a model is viable, it can provide a descriptive base for analyzing unobservable parameters, or primitives, for hydropower policy generation. In our attempt we are using a parametric approach to deal with state transitions. The previous sections describe this model. In this section we want to go one step further by introducing some tentative results and use those to discuss the validity of our model. As stated by Rust (1994): "Put simply, since structural models can be falsified but never proven to be true, their predictions should always be treated as tentative and subject to continual verification".

In order to implement the main optimization problem in (2.19) we use AMPL (A Mathematical Programming Language), which is an algebraic modelling language used to solve complex mathematical optimization problems. In combination with AMPL we use the open source solver IPOPT (version 3.10.1), used for large-scale nonlinear optimization. This version of IPOPT is limited to a total memory usage of about $2Gb$, which affect the size of the total state space. This forces us to use the minimum number of discrete levels for all of the state variables. The total number of variables and constraints IPOPT uses is still very large, each of them being 16960.

Due to the large number of variables and constraint, the solve time is consequently long. The solver time increase because the ζ enters the model in the state transition equation in (4.1). In order to ensure linearity, we need to treat it as a parameter in stead of a variable in the implementation in AMPL. Therefore, the model has to be solved for each individual value of ζ, then to manually search for the ζ with the highest resulting log likelihood value. Solving the model for one value of ζ takes approximately 1 minute, which means solving for 26 values of ζ, which we have done, takes approximately 26 minutes. To find the parameters for the different state variable processes, and for general data processing and analysing, we use MATLAB.

We should point out that we do not have access to internal price forecasts for Kolsvik. This means that we, in what follows, assume that our model provides forecasts are close to what Kolsvik actually uses. That is not necessarily a correct assumption, and could affect the results.

5.1 Results

Water value comparison

If the structural estimation model is consistent with reality, we expect to observe some of the same dynamics for water values, for different weeks and reservoir levels, as given from the EMPS model. The water value is the expected marginal value of the energy stored in the reservoir, described in Section 3.2. In our model we can output the expected value function (4.35) for different reservoir levels throughout the year. By finding the derivative of the expected value with respect to the reservoir level, i.e. divide the change in expected value by the change in reservoir level, we can find the water values. This is a valuable finding, and an important by-product of our model. For an external agent, this is a cost effective and fast way for finding the marginal cost of production. In Figure 5.1 (a) we have plotted water values based on the value function from our model. They are plotted for each reservoir level and each week of the year. Figure 5.1 (b) displays a comparable plot of water values from the EMPS model.

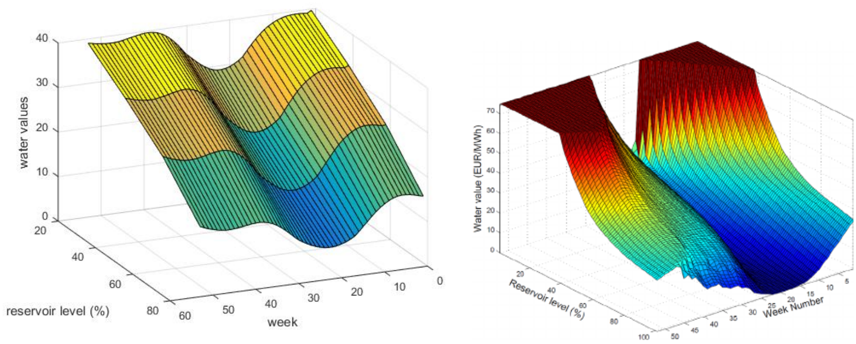

(a) Water values from our model for Kolsvik

(b) Water values from the EMPS model for a reservoir in south of Norway (Gebrekiros et al., 2013)

Figure 5.1: Water values plotted for different reservoir levels throughout a year

For a given reservoir level, the water values should be low in the spring, when there is a lot of inflow to the reservoir, since getting one more unit of water gives little extra value. However, in the winter, when the inflow levels are low, getting one more unit of water is very valuable. In Figure 5.1 we can clearly see similarities in the two plots, in the dynamics through the year, with lower water values in the spring and higher values in the winter. The minimum is around week 20 in both figures. After week 20 the values increase quickly to a maximum around week 40, before they decrease slightly through the winter until they drop down towards the minimum in the spring. This dynamic is consistent with our perception of reality.

The water values for our model should develop in a similar manner as it does for the EMPS-model at the extremes. When the reservoir is full, the water values should be very close to 0, i.e. an extra unit of water is almost worthless. At the other extreme, when the reservoir is close to empty and little inflow is coming in, an extra unit of water is valuable. This often occurs during winters. In Figure 5.1 (a) we see that the values have an inclination to be low when the reservoir is full and slightly higher when the reservoir is close to empty. We do however not see the same extreme instances as shown in the EMPS-model. Because of the limitations in memory we need to discretize the reservoir level quite coarsely, as mentioned in 4.1.6, which might leave us unable to explore what happens at the outer extremes. To explore this, our model might need a denser discretization grid.

Another trait we see in Figure 5.1 (b) is that the graph for the water values is convex, i.e. the water values increase more than linear when reservoir level decreases. This is not consistent with the water values from our model, in Figure 5.1 (a). Here the graph for the water values is close to linear. This could be an indication of problems in the model. Another reason could be that the results in Figure 5.1 (a) actually reflect how Kolsvik has operated, and that this deviate from the optimal strategy.

The water values from our model should be in the same range as the ones from the EMPS model. From Figure 5.1 we can see that the water values for our model are in a range of 10 to 40, while the EMPS water values are between 0 to over 75. This is somewhat consistent with our expectations. Apart from the extremes, the water values from our model are in the same range as the ones from EMPS. We can also plot the distribution of the price simulations from our model, which should be in the same range as the water values from our model. In Figure 5.2 we can see that the price simulation distribution has a mean at around 35 €$/MWh$. That is somewhat higher compared to our water values, which has a mean closer to 25 €$/MWh$. As simulated prices and water values in our model are in the same range, they indicate our model is behaving as expected. When also comparing the water values from our model with the EMPS model the observations provide are good indications that our model is working.

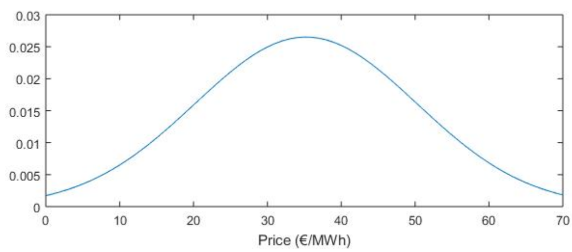

Figure 5.2: Distribution for the simulated price

Summed up, we see that the water values are low in the spring and high in the winter, as expected. At the extremes our results deviate from the comparable EMPS results, indicating that the level of discretization is too coarse. In addition the water values from our model misses the convexity from the ones from the EMPS model. On the other hand, the similarity between water values form our model and the EMPS model is a good indication that the model is working. Even though there are some inconsistencies in our results, we deem them as small enough to conclude that our model is quite sensible. We hereby treat the results as a satisfactory representation of our expectations of reality to be able to go on discussing the results from the maximum likelihood estimation.

Forward curve information results

One of the aims of our model exemplification was to estimate to what degree the price is adjusted towards the forward price, or if the forward price is used directly as expected price. The value of ζ which gives the best fit to observed data, the one with the highest log likelihood, can provide such an indication. As our model can be improved, the results could be influenced by such as the low state space resolution used. All conclusions should therefore be viewed as preliminary results and mostly as an example of how structural estimation results from a hydropower planning problem can be interpreted.

If our model is consistent with reality, we expect that ζ is in the range between 0 and 1. This is because we do not expect the producer to scale the price forecasts down if $\zeta < 0$ or up if $\zeta > 1$. Such values would mean the producer uses a price larger than the forward price or smaller than the forecasted spot price when planning. If our model is approximately in the desired value range, a shift between what forecast is emphasized would be expected. We can see from both Figure 5.3 and Figure 5.4 that a value of $0 < \zeta < 1$ approximately holds. A local max at a $\zeta > 1$ might indicate shortcomings in our price simulation model. A value over 1, might per example indicate that both the price and forward price forecast are underestimating the price value.

We would expect ζ to reflect the views of Näsäkkälä and Keppo (2008), Fleten et al. (2002) and Fleten et al. (2008). If such is the case we expect a value of ζ near 1 to give the highest likelihood. Such a value should imply that this specific hydropower producer uses information from the forward market when planning. In Figure 5.3 we plot the likelihood for different values of ζ, using a forward price with 2 months to maturity. Here we observe a global maximum likelihood value at about $\zeta = 0.5$. If the model is valid, such a value implies that the producer use a mixture of the two approaches. This is to some degree consistent with our expectations.

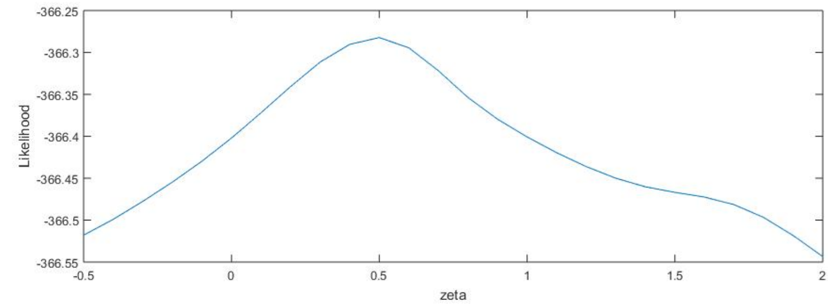

Figure 5.3: Likelihood for different ζ's, 2 months to maturity forward

To further investigate the use of forward price information by the producer, we can change the time horizon for the forward price. By changing the time to maturity, we expect there to be a change in optimal ζ due to the differences in the dynamics of the two forward prices. In Figure 5.4 we have plotted the likelihood for different ζ's for a forward contract with 6 months to maturity. Now there is a global maximum at $\zeta = 0.7$. In addition there is a local maximum around a ζ of 1.3. As we observe two maximums, the validity of the results are put to trial. Nevertheless, consistent maximum likelihood values close to $\zeta = 1$, instead of ζ close to 0 strengthens the indication towards a strong emphasis on using forward information in hydropower planning, given the validity of our model results.

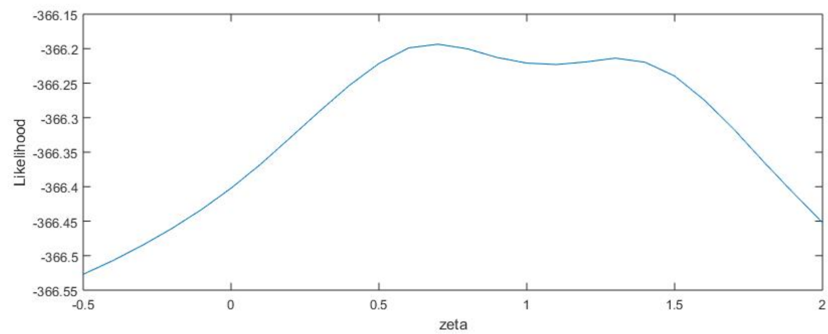

Figure 5.4: Likelihood for different ζ's for a forward contract with 6 months to maturity

Implication of changes in horizon of forward information

Another part of the model results we can explore is what forward horizon gives the best fit to the data. The forward with the highest maximum likelihood value should then be the type of forward price that is the most likely to have been used by the producer, and the corresponding ζ should tell us in to what degree the producer use the forward price information when planning. In Figure 5.5 we have plotted the likelihood for different ζ's, for forward contracts with time to maturity of 2 months, 6 months and 1 year respectively.

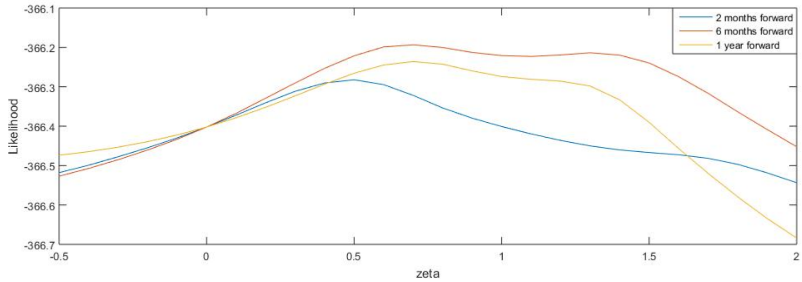

Figure 5.5: Likelihood for different ζ's, forward with 2 months- (blue), 6 months- (red) and 1 year (yellow) to maturity

We expect all the forward contracts to provide valuable data for the production planning, thereby give optimal ζ values close to 1. From Figure 5.5 we notice that this is the case. Even though the global maximum is a little smaller, the likelihood plot is weighted heavily towards a higher ζ. That is a further indication that the producers uses forward price information when planning for the future, since the likelihood for the different forward contracts is generally higher around $\zeta = 1$ than $\zeta = 0$.

We expect the highest likelihood value to be found for a forward price with time to maturity not too far into the future. Comparing the three plots for the likelihood, we find that highest maximum for the likelihood is for the 6 months to maturity forward price. It is slightly higher than for the 1 year to maturity forward, which again is higher than for the 2 months to maturity forward. The higher maximum likelihood value and the optimal ζ close to 1 indicates that the producer is inclined to look as far as 6 months into the future when planning. That is contrary to our expectations that the producer do not look to far into the future. The fact that the maximum likelihood is higher for the 6 months to maturity forward than the 1 year to maturity forward does however give an inclination towards a shorter horizon being used. A degree of regulation for Kolsvik of about 38% suggest that a forward with time to maturity of 6 months is not such a bad result. The degree of regulation tells us how much of the water, coming in through the year, the reservoir is able to store. A degree of regulation of 38% then means the reservoir of Kolsvik is able to hold water for almost 5 months without it spilling over, on average. Looking 6 months ahead when planning might consequently, not be that far off.

We can interpret these results as an indication that this producer is more inclined to use a price forecast similar to our 6 months forward price model in (4.12) than our price model without any influence from the forward price, as in (4.17). Consequently this indicates that our specific hydropower producer uses information from the forward market when planning for the future.

Implication of changes in interest rate

If we plot the likelihood for the same forward but use different interest rates, we can see which rate gives highest maximum likelihood and hence which it is more likely that the producer actually uses. Fleten et al. (2002) points out that modern financial theory dictates the appropriate discounting factor to be equal to the risk free rate for a risk neutral production planner. Bøckman et al. (2008) assumes this risk free rate for hydropower plants to be at 5.8%. We therefore expect the likelihood to be higher for an interest rate around this value. In Figure 5.6 we have plotted the likelihood for the forward contract with 6 months to maturity, with interest rates of 6%, 8% and 10% respectively.

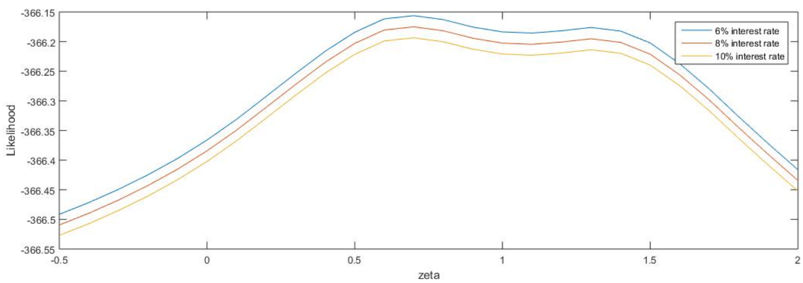

Figure 5.6: Likelihood for different ζ's, for 6% (blue), 8% (red) and 10% (yellow) interest rate, forward with 6 months to maturity

Here we can clearly see that the likelihood is higher for lower interest rates. This result is in accordance with our expectations that the industry uses generally low interest rates when planning. Our model is however not able to solve for lower than 6% interest rate, so we can not see what happens when we approach an interest rate of 0. The preliminary result in Figure 5.6 is nonetheless an indication that our model is working, since it predicts what we expect it to do.

5.2 Further Development

The model presented in this study is meant as a first attempt to utilize structural estimation for hydropower. To later on improve the validity of the result, the model is subject to further improvement. Due to the large number of variables and constraint, the solve time is consequently long and the state space severely limited. In further studies, memory usage should therefore be reduced, p.a. by using a non-parametric approach. The state variable processes can also be improved to better represent the underlying expectation processes. To further validate the model the reservoir dynamics should be simulated and to gain a broader industry insight the analysis should be extended to more hydropower producers.

Reduction of memory usage, by using a non-parametric approach
A very limiting shortcoming in our model is the restrictions on computational memory capacity, and the implications of this. In our approach we are forced to use as small a state space as possible, while still being able to describe the dynamics of the problem. A way to reduce memory usage could be to use a non-parametric approach to deal with

the conditional expectation in Section 2.4, instead of our parametric approach that uses the state variable processes in Section 4.1. Johansen (2015) uses an example of such an approach. He applies something called k-means clustering to reduce the number of state space transitions by grouping together observations into k clusters, and also compute the probability of moving between these clusters. This method could also be applied to our problem, and help reduce the state space. It could however mean loosing some of the dynamics in our model, especially the connection between inflow and price.

Improvements in state variable processes

In further studies, more details can be included and improvements done in the state variable processes in order to capture even more of the hydropower environment dynamics. An example of a possible change is to use a two-factor model where the two factors are correlated. This can be used instead of an one factor model for the forward price process, as proposed by Lucia and Schwartz (2002). In their process, the system price is an important factor in the model. While in our process, reservoir levels have a larger emphasize. However this raises the problem that the price is dependent on the forward price, while the forward price is dependent on the price. Lucia and Schwartz (2002) also mention that volatility is consistently different between cold and warm seasons. Therefore they propose a mean reverting diffusion process for further studies for volatility for the forward process. Also, including jumps in the spot price process can be an additional detail included in further studies. If so, such jumps would need to be Markovian, and could not, for instance, decay over time. Controlled environment experiments for what main drivers and attributes of time series processes are observed for hydropower plants can be done in future studies.

Simulate reservoir dynamics for model verification

We can verify our model further by making simulations of optimal decision policies, based on our processes in section 4.1. These simulated optimal decision policies can then be used as input into the structural estimation model. If the model is correct, the output should show results that are consistent with the input to the model. For instance if we were to assume a ζ of 1 when making policy simulations, the results from the maximum likelihood estimation should also show a ζ of 1. This however has to be a future endeavour. Because we model our structural estimation problem with the SDP embedded in the model implementation in AMPL, we can not easily extract the generated policies. To do so, we would have to built a separate stochastic dynamic programming model, which is beyond the scope of this paper. One could however build such a model based on our value function in Section 2.2 and the state variable processes in Section 4.1.

Extend the analysis to more hydropower producers

This preliminary study only includes a singular hydropower producer who is approximately utilizing the industry standard policy generation strategies. Therefore, we are at best getting a indication of practices in the general hydropower industry from studying this single producer. To get a more general perception of the wider industry behaviour, we suggest applying the model on as many producers as possible, choosing producers of different sizes and geographical locations to get an as general selection as possible. Such a model expansion should give the statistical significance to tell something about the overall practice in the hydropower industry. The model should be applicable to a wide range of producers, as long as the assumptions in Section (6.1) are met. When doing this, you would have to keep in mind that different producers have different inflow dynamics, which leads to different parameters for the processes related to specific hydropower plants. That could lead to a complicated model, where you would have to switch between a huge number of processes and parameters within the model when looking at a lot of different producers. It can however be done, if you were willing to take the time to do so.

Chapter 6

Conclusion

In this paper, we develop a structural estimation model for the hydropower planning industry and demonstrate it on a single hydropower reservoir. By using our structural estimation model, we analyse how primitives in the decision process can be inferred from empirical data from actual production time series. To exemplify the analysis of primitives, we analyse to what extent forward information is used when planning production.

Our model has the potential to let external analysts gain insight in water values used by producers when planning production. This is a valuable finding for potential external agents looking at the company from the outside. This could p.a. be relevant for regulatory bodies, as they could calculate the marginal cost of production and thereby improve their tools for analyzing and controlling if the producer exploit their market power. The preliminary results from the water value calculation shows some inconsistencies with our expectations, but we deem them small enough to conclude that our model is sensible. Especially the shape of the water value plot in Figure 5.1 (a) for different reservoir levels throughout the year supports this conclusion.

The ζ's we find for different forwards have a maximum likelihood in the range between 0 and 1, which is consistent with what we expect and what they should be. Our results show a clear indication that the ζ which gives the highest likelihood is located closer to 1 than to 0. Based on this we conclude that the producer uses forward price information to form expectation for the price. By comparing forward contracts with different time to maturity, we find that a forward with time to maturity of 6 months has the highest likelihood value. The producers reservoir is able to store water for up to 5 months before spilling over, hence using a forward price with time to maturity of 6 months, when planning, is quite consistent with expectations. We also find the likelihood for different interest rates. The results show that an interest rate of 6% has higher likelihood than a higher interest rate of, for instance, 10%, which is as expected and in line with what producers actually do. These results are indications that our model is working, and able to estimate underlying parameters in how hydropower producers form expectations of the future.

From structural estimation of hydropower plants, the scope of the analysis has the potential to be broadened to also include marginal cost, start-up costs, increasingly detailed price simulations, varying effects from overall reservoir level, etc. Further studies should also be done to validate the model by doing a simulation of optimal decisions. These simulations can then be used as inputs to the model, in order to better analyse the validation results. The model should be applied to a general selection of producers, in order to obtain results for the overall hydro power industry. To deal with the large state space and high memory use, the model could be altered by using a non-parametric approach to handle transition probabilities, instead of the parametric approach we used.

Bibliography

Aksin, Z., Ata, B., Emadi, S. M., Su, C.-L., 2013. Structural estimation of callers' delay sensitivity in call centers. Management Science 59 (12), 2727–2746.

Audet, N., Heiskanen, P., Keppo, J., Vehvilainen, I., 2004. Modeling electricity forward curve dynamics in the nordic market. In: Bunn, D. (Ed.), Modelling Prices in Competitive Electricity Markets. John Wiley & Sons, Inc., pp. 251–265.

Botterud, A., Bhattacharyya, A., Llic, M., 2002. M.futures and spot prices - an analysis of the scandinavian electricity market. Proceedings of the 34th Annual North American power Symposium.

Botterud, A., Kristiansen, T., Ilic, M., 2010. The relationship between spot and futures prices in the nord pool electricity market. Energy Economics 32, 967–978.

Bøckman, T., Fleten, S.-E., Juliussen, E., Langhammer, H. J., Revdal, I., 2008. Investment timing and optimal capacity choice for small hydropower projects. European Journal of Operational Research 190, 255–267.

Durbin, J., Koopman, S., 2001. Time Series Analysis by State Space Methods. Oxford University Press.

Fleten, S.-E., Haugom, E., Pichler, A., Ullrich, C. J., 2016. Structural estimation of switching options. Operations Research.

Fleten, S.-E., Keppo, J., Lumb, H., Weiss, V., 2008. Derivative price information use in hydroelectric scheduling. Working paper.

Fleten, S.-E., Wallace, S. W., Ziemba, W. T., 2002. Hedging electricity portfolios via stochastic programming. In: et al., C. G. (Ed.), Decision Making Under Uncertainty. Springer-Verlag, New York, pp. 71–91.

Foss, M., Høst, A., 2011. Hydroelectric real options: A structural estimation approach. Master thesis, Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management.

Fosso, O. B., Gjeldsvik, A., Haugstad, A., Mo, B., Wangsteen, I., 1999. Generating scheduling in deregulated system. the norwegian case. IEEE Transactions on Power Systems 14 (1), 75–81.

Gamba, A., Tesser, M., 2009. Structural estimation of real options models. Journal of Economic Dynamics Control 33, 798–816.

Gebrekiros, Y., Doorman, G., Jaehnert, S., Farahmand, H., 2013. Bidding in the frequency restoration reserves (frr) market for a hydropower unit. Conference: Innovative Smart Grid Technologies Europe.

Johansen, M., 2015. Structural estimation analysis of the strategic switching of peak generators. Master thesis, Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management.

Johnsen, T., 2001. Demand, generation and price in the norwegian market for electric power. Energy Economics 23, 227–251.

Kellogg, R., 2014. The effect of uncertainty on investment: Evidence from texas oil drilling. American Economic Review 104 (6), 1698–1734.

Kolsrud, C. W., Prokosch, M., 2010. Reservoir hydropower: The value of flexibility. Master thesis, Norwegian University of Science and Technology, Department of Industrial Economics and Technology Management.

Lucia, J. J., Schwartz, E. S., 2002. Electricity prices and power derivatives: Evidence from the nordic power exchange. Review of Derivatives Research 5, 5–50.

Maidment, D. R., Chow, V. T., 1981. Stochastic state variable dynamic programing for reservoir systems analysis. Water Resources Research 17 (6), 1578–1584.

Markeskraft, 2016. Markedskraft.
URL `https://www.markedskraft.com/Home/Customer-segments/Power-producers/`

Mirza, F., Bergland, O., 2012. Transmission congestion and market power: The case of the norwegian electricity market. The Journal of Energy Markets 5 (2), 59–88.

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

Muehlenbachs, L., 2015. A dynamic model of cleanup: Estimation sunk costs in oil and gas production. International Economic Review 56 (1), 155–185.

Nandalal, K. D. W., Bogardi, J. J., 2007. Dynamic Programming Based Operation of Reservoirs: Applicability and Limits. Cambridge University Press.

NordPool, 2016. Norwegian water resources and energy directorate.
URL `http://www.nordpoolspot.com/`

NVE, 2016. Norwegian water resources and energy directorate.
URL http://vannmagasinfylling.nve.no/Default.aspx?ViewType=AllYearsTable&Omr=NO

Näsäkkälä, E., Keppo, J., 2008. Hydropower with financial information. Applied Mathematical Finance 15 (5-6), 503–529.

Rapson, D., 2014. Durable goods and long-run electricity demand: Evidence from air conditioner purchase behavior. Environmental Economics and Management 68, 141–160.

Rust, J., 1987. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica 55 (5), 999–1033.

Rust, J., 1994. Structural Estimation of Markov Decision Processes. Vol. 4.

SINTEF, 2016a. Emps multi area power market simulator.
URL https://www.sintef.no/en/software/emps-multi-area-power-market-simulator/

SINTEF, 2016b. Vansimtap software.
URL http://www.sintef.no/programvare/vansimtap/

SKM-Energy, 2016. Vansimtap production planning.
URL http://www.skmenergy.com/production-planning/vannsimtap

Su, C., Judd, K., 2012. Constrained optimization approaches to estimation of structural models. Econometrica 80 (5), 2213–2230.

Tejada-Guibert, J. A., Johnson, S. A., Stedinger, J. R., 1993. Comparison of two approaches for implementing multireservoir operating policies derived using stochastic dynamic programming. Water Resources Research 29 (12), 3969–3980.

Thome, K., Lin, C.-Y. C., 2015. Investment in corn-ethanol plants in the midwestern united states.

Wolfgang, O., Haugstad, A., Mo, B., Gjeldsvik, A., Wangsteen, I., Doorman, G., 2009. Hydro reservoir handling in norway before and after deregulation. Energy 34 (10), 1642–1651.

Appendix

6.1 Hydropower planning assumptions

- **One-reservoir approximation**: We assume an one-reservoir approximation when planning production. This will save us computational complexity, as solving the Stochastic Dynamic Programming problem, with the inclusion of more than one reservoir, results in a large amount of states, usually referred to in literature as the curse of dimensionality (Nandalal and Bogardi, 2007). Hence it is advantageous to choose a plant that has one single reservoir, or where they use such an approximation when planning production. As shown in figure 3.4, Kolsvik has three main reservoirs, in addition to several smaller lakes and rivers. The combined maximum capacity of the two smaller reservoirs is only 13.1% relative to the largest reservoir. Since one reservoir is much larger than the others, we assume an one-reservoir approximation is suitable. Also, Bente Sund, production planner for the reservoir, has confirmed that they use a one-reservoir approximation, when using the already mentioned Vansimtap model. Therefore we deem an one-reservoir approximation suitable for our case company in particular. Such an assumption allows for inclusion of more hydropower plants in future studies.

- **Constant head assumption**: Possible head changes affects the value of the energy equivalent, in (6.1), denoting how much energy is stored in each m^3 of water in kWh/m^3. We are assuming a constant energy equivalent and thereby also a constant head in our model. The parameter η is the power plant efficiency in $\%$ and γ denotes the water density in kg/m^3. A change in head would also affect the production capacity, as seen in (6.2), denoting how much energy the turbines are capable of producing at max release. The parameter Q here denotes the discharge capacity of the hydropower plant in m^3/s. Head is defined as the difference in height between the reservoir and the turbine in m, thereby it changes as reservoir level changes. In order to be able to assume a constant head, our data should have negligible change in head from max to min levels in reservoir filling. The assumption is more reasonable the higher the head is compared to reservoir level difference. In our case company, the change in head is less than 10% of the average head, which we deem low enough to ignore. For each new hydropower plant included in the analysis, the change in head should be checked to assure the applicability of a constant head assumption.

$$e = \frac{1}{3.6 \cdot 10^6} \cdot \gamma \cdot g \cdot H \cdot \eta \tag{6.1}$$

$$\overline{U} = e \cdot Q \cdot \frac{3600}{1000} \tag{6.2}$$

- **Sufficient reservoir flexibility**: In order for the hydropower plant to be relevant for analysis, it should have sufficient flexibility in terms of storage and production capacity. From Figure 3.4, we see that Kolsvik have a degree of regulation of $37,9\%$, which means that a third of the yearly inflow can be stored in the reservoir. The degree of regulation of a reservoir provides an indication on how much water the plant is able to store. If the aggregated inflow minus production exceeds the reservoir volume, the inflow will be non-storable and have to be used continuously. A too low degree of regulation will lead to a reduced long term flexibility, thereby limiting the relevance of a SDP model. We deem a degree of regulation of $37,9\%$ to be high enough for Kolsvik to make a relevant case.

- **Sufficient production capacity**: The yearly capacity factor tells us how much of the incoming inflow the turbines are able to produce electricity from. It is measured in average yearly inflow divided by production capacity. The capacity factor for our case reservoir is $52,3\%$ and just over 200% during flooding season, lasting for about two months. If the accumulated inflow over a period raise the capacity factor above 100% at the same time as the reservoir levels are high, the reservoir will face a risk of spillage and thereby wasted water. Our case power plant do not face any major limitations concerning production capacity, as long as they make sure they have sufficient room in reservoirs during the spring flood.

- **Price taker assumption**: We also assume our case company to be a price taker. A price taking assumption is common when modeling Norwegian power plants. The price taking assumption is a necessary condition for a free market to be economically efficient. Such an assumption also avoid considerable modeling complexity as a result of oligopoly (Fosso et al., 1999). Total electricity production from hydropower in Norway is $132,3\ TWh$ (NVE, 2016). As the maximum production at Kolsvik is $504\ GWh$, Kolsvik will have a market share of only $0,38\%$. This indicates that they will not have any market power. Also, as the Norwegian electricity market has about 70 generating companies, a single producer will likely not have large enough share of the market to have a significant impact on prices. Mirza and Bergland (2012) find that market power for Norwegian hydropower planners at the most extreme never exceeds 1%. With such a low possible market power, we find a price taker assumption reasonable enough as a general assumption in the Norwegian hydropower industry. On the other hand Johnsen (2001) states that for each individual, price is given, while at the market level the price is endogenous. This is because all market players in the hydropower industry to some degree respond to the same rational market mechanisms. Therefore Johnsen (2001) developed a model for simultaneous determination of supply, demand and price in the competitive Norwegian electricity market. For later studies a simultaneous modeling approach can be a more reasonable assumption for price modeling.

- **No marginal production cost**: We assume the producer is strictly profit maximizing. Fosso et al. (1999) state that, in a deregulated market, a generation company has in principle no other objective then to produce and sell electricity with a maximum profit. This implies that we can assume no production requirements or start-up costs for our scheduling problem, and for most Norwegian hydropower plants. Formulation of the profit function as a strictly profit maximizing function is therefore reasonable.

- **Insignificant start-up and shutdown cost**: As the granularity of the model is one week at its finest, start-up and shutdown costs become insignificant for our hydropower planning problem.

- **AR-1 assumption**: We assume inflow, overall reservoir deviation, price and forward price all to be auto-correlated with one lag. As we use the Markovian property, this assumption is a model necessity. All the different time series exhibit varying degree of auto-correlation. For inflow time series monthly the assumption holds as weekly and especially daily and hourly inflows generally exhibit high serial correlations. Because we are working with a daily (and weekly) resolution we find it reasonable to incorporate the serial correlation attribute in our inflow model. For price modeling, Botterud et al. (2010) state that analysis they have done of price scenarios shows that prices from one week are strongly dependent on observed prices from the foregoing week, therefore an AR(1)-assumption is reasonable for our price process. For inflow modeling, Nandalal and Bogardi (2007) claim that the Markov-I assumption is commonly used, as this sufficiently reflect reality, without being overly complex. Also Botterud et al. (2010) mention inflow time series to commonly be modeled with a multivariate first order auto-regressive model. By using the auto-correlation function in MATLAB, we find all mentioned time series to be highly auto correlated. Therefore an AR-2 assumption could increasingly reflect reality, but entails increasingly computational complexity (Nandalal and Bogardi, 2007). Also an AR-2 assumption does not suit our model structure.

6.2 Descriptive statistics and parameter values

Below are the parameters estimated for the different seasonal function, for inflow, $f^I(t)$, in Table 6.1, forward price, $f^F(t)$, in Table 6.2 and the price, $f^P(t)$, in Table 6.3.

Table 6.1: Parameters for inflow seasonal function

D^I	A_1^I	A_2^I	ϕ_1^I	ϕ_2^I
11.9786	7.6571	3.8396	8.6783	3.7237

Table 6.2: Parameters for forward price seasonal function, for 2 months, 6 months and 1 year forward

	D^F	A^F	ϕ^F
2 months	3.6727	0.1467	5.3249
6 months	3.7118	0.1610	5.3249
1 year	3.6836	0.1125	5.3249

Table 6.3: Parameters for price seasonal function

D^P	A^P	ϕ^P
3.5836	0.1739	12.0945

In Table 6.4 are the discriptive statistics for all the state variables, both for the observed time series and for the process simulations.

Table 6.4: Descriptive statistics for the state variable processes, for both observed time series and simulations

		Mean	Stdev	Min	Max	Median
Inflow	Obs	11.9776	13.1096	-7.8992	74.8820	6.6924
(GWh)	Sim	12.0983	13.0358	0	65.3079	8.2314
Cumulativ inflow deviation	Obs	0.0320	0.1373	-0.2839	0.4179	0.0150
(%)	Sim	0.0467	0.2200	-0.5244	0.9628	0.0147
Overall reservoir deviation	Obs	0.0009	0.1721	-0.4690	0.4668	0.0235
(%)	Sim	0.0545	0.2120	-0.4158	0.5064	0.0563
2 months forward	Obs	3.6316	0.3279	2.5537	4.3453	3.6433
(log €/MWh)	Sim	3.6307	0.3165	2.7223	4.6126	3.5996
6 months forward	Obs	3.6798	0.2834	2.6358	4.4060	3.6941
(log €/MWh)	Sim	3.6525	0.2786	2.9539	4.6788	3.7428
1 year forward	Obs	3.6721	0.2426	2.9774	4.1016	3.7257
(log €/MWh)	Sim	3.6710	0.2404	3.1802	4.1814	3.6436
Price	Obs	3.5629	0.3916	2.0483	4.4609	3.5918
(log €/MWh)	Sim	3.5534	0.3883	2.3800	4.2385	3.2614