Aalto University
School of Science

Master's Programme in Mathematics and Operations Research

Forecasting hydro supply curves in the Nordic day-ahead electricity market

Jonas Edström

Master's Thesis
2025

Aalto University
School of Science

Author Jonas Edström

Title Forecasting hydro supply curves in the Nordic day-ahead electricity market

Degree programme Mathematics and Operations Research

Major Systems and Operations Research

Supervisor Prof. Ahti Salo

Advisor D.Sc. (Tech) Jari Järvi

Collaborative partner UPM Energy

Date 21/02/2025 **Number of pages** 67 **Language** English

Abstract

Accurate forecasting of electricity prices is essential for participants in the Nordic electricity market, enabling effective bidding strategies, better risk management, and efficient resource allocation. Electricity prices are often forecasted using fundamental models, which predict the supply from various electricity production technologies. Given that hydropower accounts for a large share of electricity production in the Nordics, modelling hydropower supply becomes a key element in forecasting electricity prices.

Hydropower supply is determined by the opportunity cost of production, commonly referred to as the water value, which reflects the trade-off between using stored water for immediate electricity generation or saving it for future use. Water values are computed by solving a large-scale stochastic dynamic optimization problem, accounting for uncertainties in future inflows, electricity prices and other market conditions. Water values vary between hydropower plants due to differences in reservoir capacities, inflow patterns and technical constraints. The water values of all hydro plants within a price area form the hydro supply curve that describes how much hydropower is supplied at different price levels. While the literature on computing water values for individual hydro plants is well developed, there is a lack of literature on how this knowledge can be applied to forecast the hydro supply curve.

This thesis aims to address gaps in the literature by providing a framework to forecast hydro supply curves over a short-term horizon up to six weeks. We present a forecasting model based on linear programming and the theoretical foundations from hydropower scheduling literature. Evaluating the model is challenging, as the true supply curve is unknown. Therefore, the model is primarily evaluated based on qualitative analysis of the predicted supply curves. The qualitative analysis indicates that the model provides a promising framework to model the hydro supply curve using price forecasts and hydro production as inputs. An attractive feature of the model is its relative simplicity, which enhances its interpretability and facilitates its application in real-world market scenarios.

Keywords Nordic electricity market, supply curve, hydropower, stochastic dynamic programming, water value, electricity price forecasting

Aalto-universitetet
Högskolan för
teknikvetenskaper

Författare Jonas Edström

Titel Prognostisering av vattenkraftens utbudskurvor på den nordiska dagen-före elmarknaden

Utbildningsprogram Mathematics and Operations Research

Huvudämne Systems and Operations Research

Övervakare Prof. Ahti Salo

Handledare TkD Jari Järvi

Samarbetspartner UPM Energy

Datum 21 februari 2025 **Sidantal** 67 **Språk** engelska

Sammandrag

Exakta elpriserprognoser är avgörande för aktörer på den nordiska elmarknaden, då det möjliggör effektiva budstrategier, bättre riskhantering och effektiv resursallokering. Elpriser prognostiseras ofta med hjälp av fundamentala modeller som förutsäger tillgången på el från olika produktionsteknologier. Eftersom vattenkraft står för en stor del av elproduktionen i Norden är modellering av vattenkraftens tillgång en viktig komponent vid prognostisering av elpriser.

Vattenkraftens tillgång bestäms av produktionskostnadens alternativkostnad, ofta benämnt vattenvärde, vilket reflekterar avvägningen mellan att använda lagrat vatten för omedelbar elproduktion, eller att spara det för framtida produktion. Vattenvärden beräknas genom att lösa ett storskaligt stokastiskt dynamiskt optimeringsproblem, som tar hänsyn till osäkerheter i framtida inflöden, elpriser och andra marknadsförhållanden. Vattenvärdena varierar mellan vattenkraftverk på grund av skillnader i reservoarers lagringskapacitet, inflödesmönster och tekniska begränsningar. Vattenvärdena för alla vattenkraftverk inom ett prisområde bildar vattenkraftens utbudskurva, som beskriver hur mycket vattenkraft som produceras vid olika prisnivåer. Även om litteraturen om vattenvärden för individuella kraftverk är väl utvecklad, saknas det litteratur om hur denna kunskap kan tillämpas för att prognostisera vattenkraftens utbudskurva.

Denna avhandling syftar till att fylla de luckor som finns i den befintliga litteraturen genom att erbjuda ett ramverk för att prognostisera vattenkraftens utbudskurva över en kortsiktig tidsperiod på upp till sex veckor. Vi presenterar en prognosmodell baserad på linjär programmering som bygger på de teoretiska grunderna från litteraturen om planering av vattenkraft. Att utvärdera modellen är utmanande, eftersom den verkliga utbudskurvan är okänd, och därför utvärderas modellen främst baserat på en kvalitativ analys av de förutspådda utbudskurvorna. Den kvalitativa analysen indikerar att modellen erbjuder ett lovande ramverk för att modellera vattenkraftens utbudskurva baserat på prisprognoser och vattenkraftsproduktion. En attraktiv egenskap hos modellen är dess relativa enkelhet som förbättrar tolkbarheten och underlättar tillämpningar i verkliga marknadsscenarier.

Nyckelord Nordiska elmarknaden, utbudskurva, vattenkraft, stokastisk dynamisk optimering, vattenvärde, elprisprognoser

Preface

Writing this thesis would not have been possible without the support from my colleagues at UPM Energy. I am especially grateful to Timo Javanainen and Matias Kinnunen for giving me this opportunity to deepen my knowledge within this topic. I also owe many thanks to Jari Järvi for numerous comments and improvements on the text, as well as my supervisor Ahti Salo for comments on the final version. While it is impossible to mention each of you by name, I am deeply grateful to everyone who has contributed with their time and insights.

Throughout the process of writing the thesis I have learned a lot about the complex structure that the electricity market is. This invaluable hidden knowledge that I have gained can be perfectly summarized by a quote from my former colleague.

"Nää hommat pitää vaan osata"

- Anonymous former colleague

Helsinki, 21/02/2025

Jonas Edström

Contents

Symbols and abbreviations

Symbols

c	Cost of thermal production
p	Day-ahead price
g	Thermal production
q	Hydro production
u	Spillage
v	Inflow
w	Reservoir level
α	Future profit
Φ	Future cost

Abbreviations

ATC	Available Transfer Capacity
CET	Central European Time
CHP	Combined Heat and Power
DoR	Degree of Regulation
EPAD	Electricity Price Area Difference
NTC	Net Transfer Capacity
RMSE	Root Mean Squared Error
SDDP	Stochastic Dual Dynamic Programming
SDP	Stochastic Dynamic Programming
SRMC	Short-Run Marginal Cost
TSO	Transmission System Operator
UMM	Urgent Market Message

1 Introduction

1.1 Background

Electricity is essential for modern society, industry, and consumers, driving economic growth, technological innovation, and daily convenience. Accurate electricity price forecasting is crucial for the efficient operation of energy markets, the stability of power grids, and the operations of electricity producers and consumers. Electricity producers seek to maximize profits by generating electricity when prices are high, while consumers and retail companies aim to purchase electricity when prices are low. It is not possible to store electricity in large quantities and therefore supply and demand must be continuously balanced. This special property of electricity distinguishes electricity markets from other commodity markets and increases the volatility of electricity prices. This makes electricity price forecasting a very challenging task.

There are several approaches to forecasting the day-ahead electricity price. Statistical models use past prices as well as external information, such as temperature and wind, to forecast the price. An example of a statistical model approach is time-series analysis. Computational intelligence models, such as artificial intelligence have also been used to forecast electricity prices. A third approach, which will be the focus of this thesis, is to use fundamental models. In fundamental models the supply of different generating technologies as well as demand are explicitly estimated. The price can then be predicted as the intersection point between the estimated supply and demand. Using fundamental market models to forecast electricity prices has a number of advantages. It provides an understanding of the underlying drivers of the price formation process. This helps in interpreting the model and understanding price movements. Fundamental models also allow studying different price scenarios based on different assumptions, which is vital for e.g. risk management and investment decisions.

This thesis focuses on the Nordic day-ahead electricity market, where electricity is traded the day before delivery and price is determined by supply and demand. Being one of the first deregulated electricity markets in the world, the Nordic market has been the subject of many studies on electricity price forecasting. Most of the electricity traded in the Nordic electricity market is sold in the day-ahead market. The day-ahead market is a blind auction, where market participants submit bids to sell or buy electricity at given price levels. The price of electricity is found at the equilibrium point between supply and demand. The structure of the day-ahead market is explained in more detail in section 2.2. As the day-ahead market is the most liquid, the price on the day-ahead market is usually taken as a reference price for traded electricity. The day-ahead price is therefore often the most important price for market participants.

To forecast prices using a fundamental model, the supply of different production technologies must be forecasted. A significant share of electricity production in the Nordic countries is produced by hydropower, which generates over 50% of the region's electricity. Because of the dominant position of hydropower, it is vital to understand the supply of hydropower to accurately forecast the electricity price. Hydropower is a flexible way to produce electricity, since water can be stored in reservoirs upstream from the power plant and used for production at a later time. The geographical features of the watercourse affect the possibilities to store water in reservoirs. In Norway, many reservoirs are capable of storing water for time periods up to multiple years. Hydro production can be adjusted quickly by increasing or decreasing the flow of water through the turbines. This flexibility provides hydro producers with possibilities to maximize revenue by optimally allocating their production to times with higher prices.

There is an opportunity cost associated with production of hydropower, as water used for immediate production cannot be used for future production. This opportunity cost will be referred to as the *water value*. The water value determines at what price a hydro producer is willing to use water for production of electricity. Water values are unique to every hydropower plant as they depend on the storage capacity of the reservoir connected to the plant as well as the filling level of the reservoir. Furthermore, the water value will depend on the hydro producer's expectation about future prices and therefore the water value is directly affected by all factors influencing the electricity price. When water values from multiple hydro plants are aggregated, a hydro supply curve is formed, describing how much hydropower is available for production at a given price. The hydro supply curve is increasing as more supply becomes available as the price increases. Predicting the hydro supply curve is challenging as numerous factors, such as weather, fuel prices, availability of transmission lines and maintenance influence the decisions of hydro producers. Understanding how these factors affect the supply curves is vital to understanding the market dynamics in areas with a large share of hydropower production.

1.2 Research objectives

This thesis has the following objectives:

1. Analyse the dynamics of water values in the Nordic electricity market and their impact on the hydro supply curve.

2. Propose a model to forecast short-term changes in the hydro supply curve based on information available to electricity market participants.

To analyse the main drivers of the hydro supply curves in the Nordic electricity market, the problem of optimal allocation of hydro resources will first be studied. The dynamics of the supply curve will be explained based on the scheduling problem. Based on the analysis, we aim to develop a modelling framework, which can be used as a reliable indicator of upcoming changes in the hydro supply curve. Due to the extreme difficulty in predicting the hydro supply curve, a model can be deemed satisfactory if it can correctly predict the direction of changes in the supply curve, without necessarily being able to predict the magnitude of the change. To analyse the hydro supply curves, historical data of hydro production, market prices and traded electricity volumes will be used.

2 The Nordic electricity market

This section explains the basics of electricity trading in the Nordic region. The history of the Nordic electricity market is briefly explained followed by the structure of the different sub-markets. In this thesis, the Nordic electricity market will refer to the interconnected markets in Norway, Sweden, Denmark and Finland.

Deregulation of the Nordic electricity market began with the Norwegian government's approval of the Energy Act in 1990, leading to the deregulation of the electricity market starting on January 1, 1991. Over the following years, the market expanded to include Sweden in 1996, Finland in 1998, and Denmark in 2000 [1]. The Nordic electricity market was the first deregulated international power market in the world. In a deregulated electricity market, power companies are not obligated to satisfy the demand of electricity in an area of responsibility. Instead, electricity is generated based on maximizing profit by trading electricity in the market [2]. Motivation for the electricity market reform was to improve resource utilisation and cost efficiency. A well functioning market would give better investment signals through Short-Run Marginal Costs (SRMC), increase equity between consumers and reduce geographical variations in prices [3].

The backbone of the electricity system is the transmission grid, enabling electric power to be transmitted from the producer to the consumer. Transmitting electricity over large distances is challenging as some of the electric power is lost as heat in the transmission cables. To minimize losses, high-voltage transmission lines are used to transfer electricity between local grids, which in turn use lower voltages. The national high voltage grids are owned by the Transmission System Operator (TSO) while regional and local grids are owned and operated by numerous smaller companies. The Nordic electricity market is divided into multiple price areas, as seen in figure 1. Norway is divided into five price areas, Denmark into two and Sweden into four with Finland consisting of only one price area.

In this thesis, we describe the Net Transfer Capacity (NTC) methodology for calculating the electricity transmission between areas. Flow-based market coupling, which was introduced in the Nordic area starting from October 30th 2024, is outside the scope of this thesis. The NTC transmission capacities in figure 1 represent the transfer capacity under ideal circumstances, i.e. assuming no maintenance or other exceptional situation affecting the grid capacity. The Available Transmission Capacity (ATC) for a given day is reported daily by the TSO and is used as the maximum quantity of electricity which can be traded on the given transmission line. The division into price areas is made based on bottlenecks in the transmission system, with transmission lines most likely to become bottlenecks forming the borders between price areas. To enhance market liquidity it is better to have fewer price areas [4]. Price areas have changed in the past. The current division of Norway into 5 price areas has been used since 2010. Sweden has been divided into the same price areas since 2011. Denmark has always been divided into the same two areas, while Finland has always been one price area [5].

The Nordic electricity market is a decentralized market, where the balancing of supply and demand ahead of delivery is not done by a system operator but rather handled by market mechanisms. The balancing of the grid during delivery is still done by a system operator, which in the Nordic countries is the TSO of the specific country. In a decentralized market, the price acts as a signal to market participants, ensuring balance between supply and demand ahead of delivery [4]. To ensure balance between electricity supply and demand in real time, electricity is traded in different markets based on the time between transaction and delivery of electricity, which ranges from years in the financial markets to seconds in the balance market. The timeline of the different markets are seen in Figure 2. Physical electricity can be traded earliest one day before delivery in the day-ahead market. After clearing of the day-ahead market, electricity can be traded in the intraday market up until the delivery hour, in which electricity is traded in the balancing market. The most liquid physical market is the day-ahead market, with 696 TWh of electricity being traded in the day-ahead market in the Nordic countries in 2023. This corresponds to 86% of all electricity produced and consumed. The high share of electricity traded in the day-ahead market makes the day-ahead price the most important reference price for physical electricity in the Nordic region. In comparison, the volumes traded in the intraday and balance markets were 21 TWh and 9 TWh [6]. The structure of the different markets are outlined in the following sections.

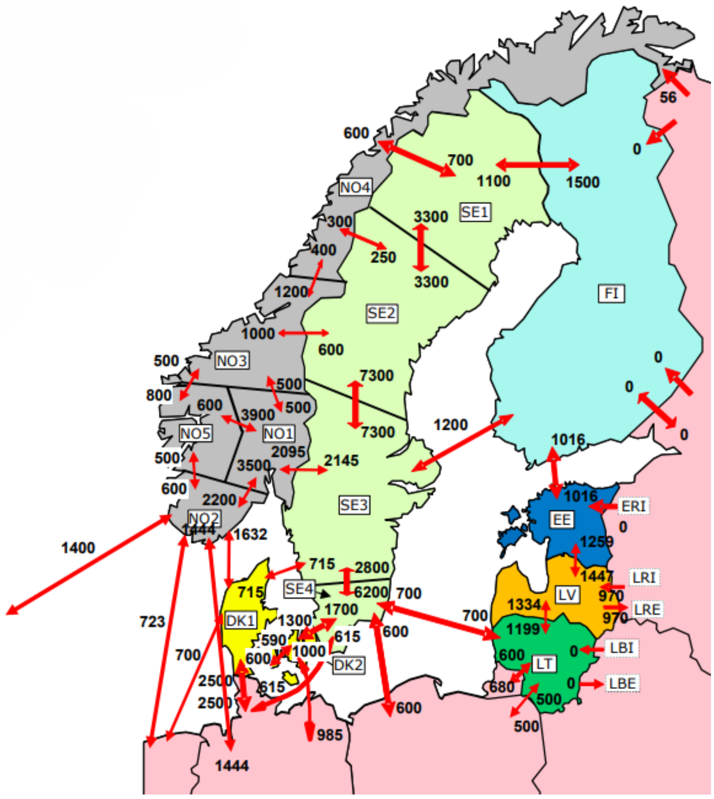

Figure 1: The price areas and the NTC on borders in the Nordic and Baltic countries as of January 1st 2024 reported by Nord Pool [5]. In addition to these connections, the 1400 MW Viking Link cable connects Denmark to Great Britain.

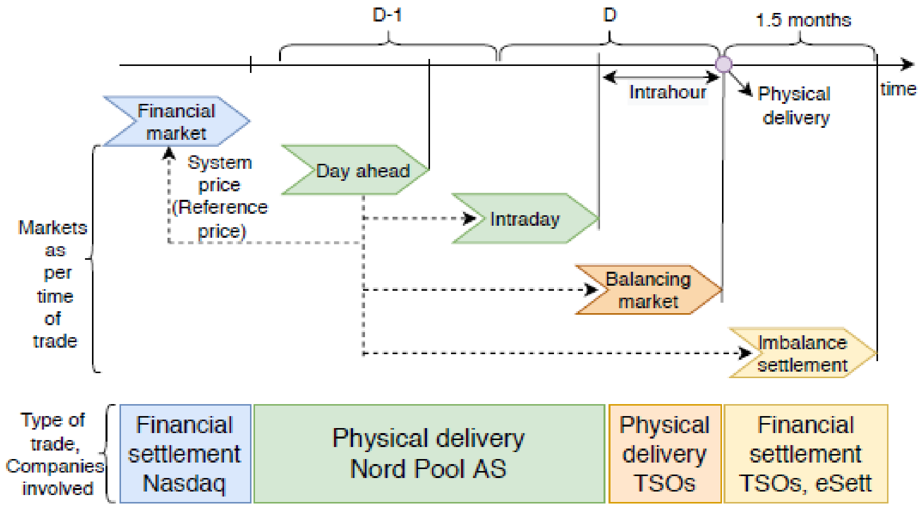

Figure 2: Timeline of Nordic electricity markets [7].

2.1 Financial market

The financial market for Nordic electricity is operated by Nasdaq. In the financial electricity market, futures and forward contracts are traded with maturities up to 10 years ahead. However, contracts with an expiry date more than 2-3 years into the future are generally not liquid. Contracts in the financial market are settled against the average price in the day-ahead market during the delivery period specified in the contract. The delivery periods can be a day, week, month, quarter or year. The most liquid futures products are the system products which settle against the system price over a given time period. The system price is a reference price for the whole Nordic region which is explained in section 2.2. Electricity Price Area Difference contracts (EPAD) are settled against the difference between the regional price and the system price. Contracts in the financial market are settled in cash against the average day-ahead price in the given period and do not include physical delivery of electricity.

Most participants in the financial markets are generation companies, electricity retailers and industrial companies, who aim to hedge their exposure to changing electricity prices. Electricity producers and consumers may also agree on bilateral future contracts, which have similar properties as the standardized contracts traded on Nasdaq. The difference is that the parties can agree on different collateral agreements, as the amount of collateral required for the market is significant.

Even if the financial market will not be the focus of this thesis, the prices in the financial market are relevant for hydropower production, as they represent the expected day-ahead price in the given delivery periods. Since the pricing of hydropower is also based on future price expectations, the financial market can be used to model the pricing of hydropower [8].

2.2 Day-ahead market

The focus of this thesis is the day-ahead market, which has the largest turnover volumes in the Nordic market. Since a large share of electricity is traded in the day-ahead market, the day-ahead price acts as a reference price for all power traded. It is the direct reference price in the financial markets and bilateral contracts are priced according to the expected day-ahead price. The day-ahead market in the Nordic countries is primarily operated by Nord Pool, with Epex Spot launching their platform for day-ahead trading in 2020.

The day-ahead market is organised as a daily blind auction, where market participants submit bids for electricity to be bought and sold each hour during the next day. Bids must be submitted before 12:00 Central European Time (CET) on the day before the delivery and the resulting market prices are usually available at 12:45 CET. The bids are in the form of price-quantity pairs for each hour. The price-quantity pairs describe how much electricity a producer is willing to sell at the given price, or how much electricity a consumer is willing to buy at the given price. The minimum price which can be bid is -500 €/MWh and the maximum price is 4000 €/MWh. All market participants bid their production or consumption in the price area where it is located, no special bids for cross border trading need to be submitted. The bids of each market participant represents the aggregation of all assets in the market participants portfolio. Thus, if a market participant has both production and consumption, they will only bid according to their net positions, no bids are sent for individual production or consumption units. Bids may also be given as block bids, in which a producer commits to producing a given quantity of electricity over multiple hours given that the average price in these hours is above a certain threshold price. The quantity of block bids submitted is typically much lower compared to the quantity of hourly bids.

Bids from all market participants within a price area are aggregated into a supply and a demand curve for the given price area. An example of aggregated supply and demand curves is given in Figure 3. To obtain the price for each area, the imported volume is added to the supply curve and the exported volume is added to the demand curve. The price is the intersection point of the supply and demand curves after the traded volumes are added. The day-ahead market employs marginal pricing, where all accepted bids are paid the market clearing price.

The imports and exports between areas are computed using the EUPHEMIA [9] algorithm, with the objective of maximizing social welfare given the constrained transmission capacities between areas. Social welfare is defined as the sum of consumer surplus, producer surplus and congestion rent. Consumer surplus is the difference between the price bid by consumers and the actual price they are paying. Similarly producer surplus is the difference in price received by the producer and the price they bid. Congestion rent is the product of the price difference between neighboring regions and the amount of electricity flowing on the transmission lines between the regions. Bids from all day-ahead exchanges across most European countries are given as input to the EUPHEMIA algorithm and the output is flows between areas and the area prices. If the transmission capacity is sufficient between two areas, the two market areas will be coupled and the price in both areas will be the same. In case the transmission capacity is insufficient, the areas will decouple with the price often being higher in the area with generation deficit. Transmission capacities available for the day-ahead market can change day-on-day based on maintenance works and forecasted production and consumption patterns.

Understanding the supply and demand curves is one of the most important factors when forecasting day-ahead prices using a fundamental market modelling approach. The target in the fundamental approach is to replicate the price formation by constructing supply curves for all production technologies as well as corresponding demand curves. The aggregated supply curve can then be constructed by combining the supply curves from all production technologies. The forecasted price is obtained by running a similar welfare optimization algorithm as EUPHEMIA and the predicted price is again found in the intersection of the forecasted supply and demand curves. This thesis focuses on one component of the aggregated supply curve, namely the aggregated supply curve of hydropower on a price area level.

In addition to the area prices, a so called system price is calculated for each hour in the day-ahead market. The system price is calculated by aggregating the supply and demand curves from all price areas in the Nordic region into a Nordic supply and demand curve, with exports and imports to areas outside the Nordic region added. The system price can be thought of as a Nordic price if no internal bottlenecks were present in the electricity grid. It acts as a reference price for electricity in the Nordic region. It is a theoretical price and no physical electricity is traded at the system price, unless a price area happens to have the same price as the system price. It is an important reference as many contracts in the financial market are settled according to the system price.

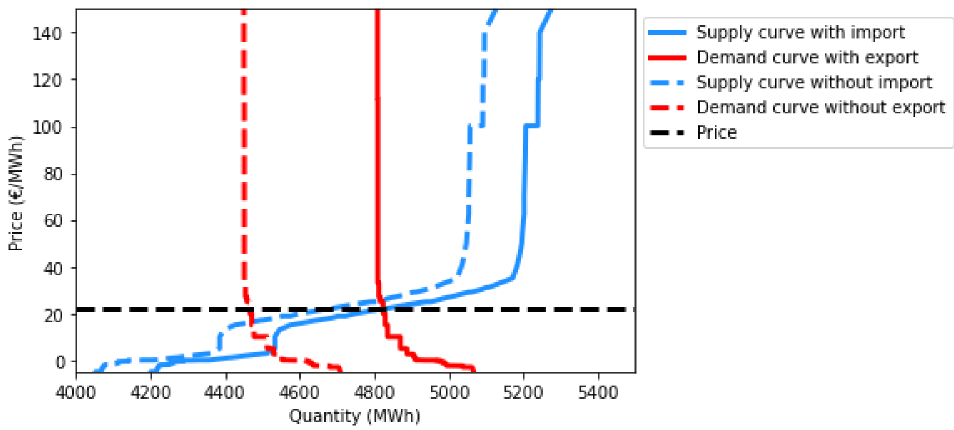

Figure 3: The supply and demand curves in Finland on the first hour of July 25th 2024. The dotted lines represent the supply curve obtained from hourly bids without accounting for export and import. The exported and imported volumes are added to the solid lines and price is formed at the intersection point of these curves.

2.3 Intraday market

After the clearing of the day-ahead market, power for the next day can still be traded in the intraday market. Trading in the intraday market is possible from the time the prices in the day-ahead market are published up until the gate closure time, which depending on the price area varies between 1 hour before start of delivery and the start of delivery [10]. The bids in the intraday market are hourly bids and block bids, which work similarly as in the day-ahead market. In contrast to the day-ahead market, the intraday market is cleared continuously. The bidding in the intraday market is similar to the financial market where the order book is updated in real time with bid-ask orders. The purpose of the intraday market is to balance the supply and demand closer to the time of delivery as producers and consumers gets more information about their available production capacity and electricity demand. Producers of renewable electricity from wind and solar usually cannot forecast their production perfectly one day before the delivery hour and they can update their production commitments in the intraday market. Because of the smaller traded volumes, the price volatility in the intraday market is higher than in the day-ahead market. Market participants are therefore incentivized to aim for an accurate production schedule from the day-ahead market to avoid having to balance their positions in the intraday market, which exposes the market participant to greater price risks. The higher volatility also provides opportunities for producers who can adjust their production closer to delivery. Producers of e.g. hydropower can sell more electricity in the intraday market if the price is higher than the price in the day-ahead market and buy back some of their already sold electricity if the price is lower in the intraday market.

2.4 Balancing market

The balancing and reserve markets in the Nordic countries are operated by the TSOs. Although supply and demand is balanced in the day-ahead and intraday markets, they do not guarantee operational security of the grid in real time. To ensure operational safety, the power system must be able to handle any single fault without resulting in load curtailment [7]. Sudden changes in electricity production or consumption can appear within the delivery hour because of e.g. plant outages or sudden changes in output from renewable energy sources. These deviations must be balanced to retain balance in the electricity grid. The balancing is done by the TSO buying or selling power in the balancing market. If there is a shortage of electricity, leading to a drop in the grid frequency, the TSO buys electricity from market participants. The market participant selling electricity can either be a producer increasing their production or a consumer decreasing their consumption. Conversely, if there is oversupply of electricity and the grid frequency is too high, the TSO can sell electricity to market participants. In this case a producer will decrease their production and a consumer will increase their consumption.

Some of the reserve capacity is procured by the TSO in capacity auctions, where producers and consumers promise the availability of regulation capacity to the TSO. Once a bid is accepted, the producer must make sure that they have production capacity available in case of up regulation and that they have sufficient level of production in case of down regulation. Most reserve markets in the Nordic countries use marginal pricing, meaning that producers are paid according to the highest and lowest bids in hours of up and down regulation, respectively. The reserve and balancing markets consist of several different products, defined by the speed at which the production, or consumption, must be adjusted. Furthermore, some reserve market products also require automatic activation. Because of the automation and fast response time required, hydropower is one of the best suited production technologies to participate in the balancing markets.

The reserve markets impact the prices in the day-ahead market by providing opportunity costs. An electricity producer can bid their production in the day-ahead market at a higher price, which can be justified by the price expectation in the reserve market. Capacity can also be sold in the capacity market before the gate closure in the day-ahead market. If a producer has already sold capacity, they cannot sell the same electricity as energy in the day-ahead market as this would remove the possibility to increase production on the request of the TSO. Since hydropower is the main technology used for balancing the grid in the Nordic countries, it is important to understand that the balancing market can affect the day-ahead bidding of hydropower.

3 Introduction to hydropower in the Nordic countries

Section 2 described how the day-ahead price depends on the supply and demand of electricity. Furthermore, the supply depends on the production technology. In this section, we focus on the different electricity production technologies in the Nordic countries and their impacts on the supply dynamics in the electricity market. We look into the role of hydropower in the market, with particular focus on the bidding behaviour of hydropower, which results in the hydro supply curve.

3.1 Electricity production

The main sources of electricity in the Nordics are hydro, nuclear, wind, solar and combined heat and power (CHP) (see Figure 4). We have made a distinction between two types of hydro production, run-of-river and reservoir. In addition there is also pumped hydro, where water can be pumped from the lower reservoir to the upper reservoir when electricity is cheap, and released through the turbine when electricity is expensive [11]. The share of pumped hydro is small, less than 2 TWh [12] per year, and therefore it has been included in the reservoir category. Dividing hydro production into these two categories gives better insights into the flexibility of hydro production in different price areas.

Most of electricity in the Nordics is generated by hydropower, as seen from Figure 4. The share of hydropower generation is especially large in all Norwegian price areas and price areas SE1 and SE2. The two Danish price areas, DK1 and DK2, as well as price area SE4 stand out as the only price areas with negligible hydro production volumes. In Norway, hydro production is the dominant production technology in all price areas. Most production is located in price areas NO2, NO4 and NO5. It can be observed that NO1, NO2 and NO3 have significant shares of run-of-river production, while almost all production in NO4 and NO5 is from reservoirs. On the other hand, the hydro production in Sweden is very much concentrated to the northern price areas SE1 and SE2. The share of hydro production in Finland is small compared to Norwegian and Swedish areas, but the total hydro production volume is still significant. Compared to other areas, Finnish hydro production is dominated by run-of-river, with almost no reservoirs with large storage capacity. With such large hydro resources in the Nordics, the supply of hydropower has a big impact on the day-ahead electricity price.

Many price areas are unbalanced when it comes to total production and consumption, with some areas being either large exporters or importers. The two price areas in northern Sweden, SE1 and SE2, have a large surplus of electricity from wind and hydro while the consumption is lower. The Norwegian areas NO2 and NO5 are also large net exporters of electricity. The areas with the largest net deficit are NO1 and SE4 located in the southern part of Norway and Sweden, respectively. These two areas are more populated and therefore have larger consumption, while production is limited. Electricity is thus often transferred from the northern parts of the Nordics to the south. Furthermore, the Nordic countries have strong interconnections with Central European countries, recall Figure 1. As the Nordic region as a whole often has a surplus of electricity production, electricity is often exported to Central Europe.

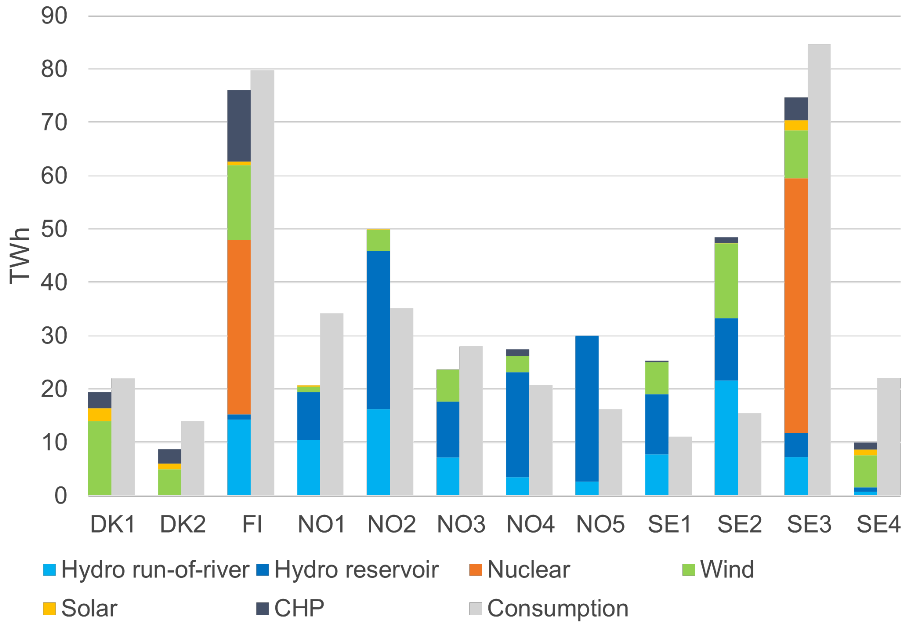

Figure 4: Electricity production by technology and consumption in each Nordic price area in 2023 reported by Volue [12].

There are, however, situations when the flow of electricity goes towards the Nordics, for example during times of high solar and wind production in Central Europe [13]. Also during cold winter days, the production capacity in the Nordics may not be enough to cover demand and imports are needed. This is especially the case in dry years when hydro production is limited.

The production system in Central Europe has a much larger share of fossil fuel-based electricity production compared to the Nordic countries. Central European prices are therefore often set by SRMC of thermal power plants. Due to the potential need for imports of thermally generated electricity from Central Europe, SRMC of thermal plants sets the opportunity cost for hydro producers in the Nordics [14]. In addition to Central Europe, the Nordic market is connected to the Baltic region and to the UK. Recently, many of the interconnectors between the Nordics and other countries have been built and therefore there is limited data available to study how they have influenced the prices in the Nordics.

The different production technologies can be divided into price dependent and price independent production. Price dependent production has short-run marginal costs and the production will depend on the market price. On the other hand, price independent production is assumed to be bid at zero or negative prices.

Electricity generated by nuclear, solar and wind power can be assumed to be price independent. Nuclear plants are often offered to the market at negative prices, since it is more expensive for the plant operators to reduce the power output than to sell the electricity at a negative price. Due to the long start-up times of nuclear power plants, they are referred to as base load power, often producing at maximum capacity independent of the current price. With the increasing share of renewable energy sources, the value of more flexibility of nuclear power plants have increased [15]. In particular, output from nuclear power plants is reduced during times of negative prices. Electricity production from wind and solar power is also price independent, as no fuels are used and running a wind turbine or solar plant comes with no extra cost for the operator, except negligible wear and tear of the plant. The output of wind and solar power is volatile as the wind speed and sunshine varies. Therefore, the volume of wind and solar power bid to the market exhibits large fluctuations. Production from run-of-river hydro plants can be assumed to be price independent as water can be either used for production or spilled, but not stored in a reservoir. The generation pattern of a run-of-river plant will depend on the inflow of water to the plant and not on the market price.

Electricity generated by thermal power and hydro reservoirs can be assumed to be price dependent. In this thesis, thermal plants refer to power plants fueled by oil, gas, coal or some other fuel. The SRMC of thermal generation is determined by the fuel prices, CO2 emissions allowance prices and plant efficiency. Most of the thermal based electricity production in the Nordics comes from CHP. Since CHP plants also produce heat, the decision to start up a CHP plant is dependent on the need for heating, which in turn depends on outdoor temperature. This makes it difficult to predict the price at which CHP is bid to the day-ahead market, as production planning is not done only based on price, but also based on heat load.

Pricing of reservoir hydro production is more difficult as water is a free but limited resource. Since water can be stored in the reservoir, production can be delayed to a future time with higher prices. We will see that the production volumes from hydro reservoir plants is strongly dependent on price and mostly independent of inflow to the plant.

3.2 Hydropower production

Hydropower relies on the potential energy of water at higher altitude being converted into electrical energy. A schematic picture of a hydropower plant is presented in Figure 5. The height difference between the upper and lower reservoir is called the plant head. As the potential energy of water is linearly dependent on the height difference, a hydro plant with larger head can produce more electricity from the same amount of water compared to a plant with lower head. The efficiency of the conversion from potential energy to kinetic energy depends on the design of the turbine and how the tailwater is transported away from the turbine. Hydro plants often have multiple turbines, which may have different characteristics, such as different efficiency curves.

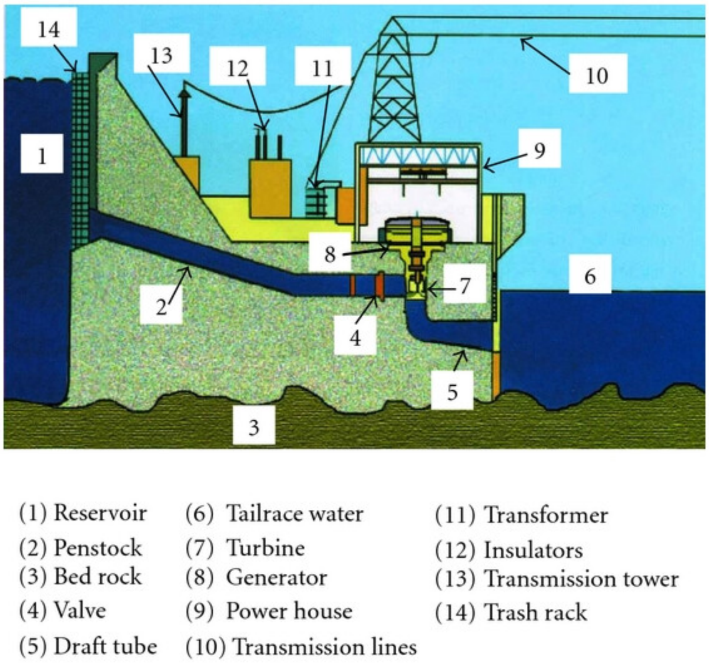

Figure 5: Schematic picture of a hydropower plant [16].

Hydropower is often split into three categories, run-of-river, reservoir and pumped hydro. In this thesis we do not explore pumped hydro in more detail and restrict ourselves to run-of-river and reservoir hydro. Run-of-river plants are not connected to an upstream reservoir and water can therefore not be stored for longer time periods. The distinction between run-of-river and reservoir production can be ambiguous and production volumes can vary depending on the source. One definition provided by European Network of Transmission System Operators for Electricity (ENTSO-E) is that run-of-river plants can have a maximum of 24 hours of storage [17]. Therefore, production at run-of-river plants is very much restricted to the flow of the river as too much or too little production would lead to flooding upstream or downstream from the plant. Typically, production from run-of-river plants is the highest during the spring flooding and is largely dependent on water inflow in the summer. The head of run-of-river plants is usually lower than plants with reservoirs and therefore a run-of-river plant usually requires more water to produce the same amount of electricity as a reservoir plant.

Hydro reservoirs can store water to be used at a later time. In mountainous areas larger dams can be constructed using the topography of the environment and thus large quantities of water can be stored. The distance between the reservoir and the power plant can be very large, even several kilometers, with water being transported through tunnels to the plant. Hydro reservoirs have a highest and lowest allowed water level, which may change during the time of the year. For example the lowest allowed water may be higher during summer for recreational purposes. The allowed water levels set a limit to how much water the operators of the hydro plant can store. If the water level gets close to the upper reservoir limit, the producer will have to start producing in order to lower the reservoir. This may be unprofitable as the producer does not necessarily receive the maximum price for the produced electricity. If the reservoir gets too low, the producer will be forced to stop production until the water level rises. This may also affect the profits of the producer negatively as production may be curtailed at times of high prices.

The size of reservoirs can vary from being able to store water from weeks to multiple years. To better understand the hydro system, it is interesting to look at the distribution of storage capacity between different plants. The storage capacity can be characterized by the degree of regulation (DoR), which is the ratio between storage capacity and average yearly inflow. If the DoR is above one, it means that the reservoir is able to store more than the average yearly inflow. The DoR is computed by summing all reservoirs upstream from the power plant and summing all the inflows to these reservoirs [18]. Thus, the DoR can be large even if the reservoir immediately above the plant is small, as there may be larger reservoirs further upstream. The DoR of run-of-river plants can also be large if the inflow to the run-of-river plant is controlled by a hydro plant upstream.

The DoR for all hydropower plants in Norway is available from the Norwegian Water Resources and Energy Directorate (NVE) [18]. Histograms of the DoR in Norwegian areas are shown in Figure 6, where we have categorized each hydropower plant into five groups based on DoR and summed up the total production capacity within the group. No similar public data is available for Sweden and Finland to the knowledge of the author, which makes it harder to draw conclusions about their hydropower storage capacity. The DoR reported by NVE correspond well to the shares of run-of-river production in figure 4. The largest shares of run-of-river production with low DoR is found in price areas NO1 and NO3. On the other hand, most hydro plants with a high DoR are located in price areas NO2, NO4 and NO5. In price areas NO2, NO4 and NO5 most hydro plants have somewhat high DoR at above 0.25. Most of the very flexible production with DoR above 1 is located in price areas NO2 and NO4. In conclusion, the storage capacity can be very different in different price areas, which is an important consideration in the pricing of hydropower.

Another important characteristic of hydropower production is that there are hydrological connections between multiple hydropower plants. Many hydropower plants belong to the same water system, which adds more complexity to the operational decisions of hydro producers. To illustrate the complex topology of a cascaded river system, a schematic picture of Lule River in Northern Sweden is shown in Figure 7. The river system consists of two river branches, some smaller rivers and several reservoirs. Water released from upstream plants becomes inflow to downstream plants. This also allows for run-of-river plants to be operated according to the price. Even if the storage capacity of a run-of-river plant is practically zero, water can be stored in a reservoir upstream and then released as inflow to the run-of-river plant when the price is higher. This partially explains why hydro production correlates very well with the day-ahead price also in price areas with more run-of-river plants than reservoir plants. On the other hand, there can be a significant delay between the release of water from the upstream plant and the water arriving at the downstream plant. Depending on the geographical distance and topology, the time span can be several hours. This is

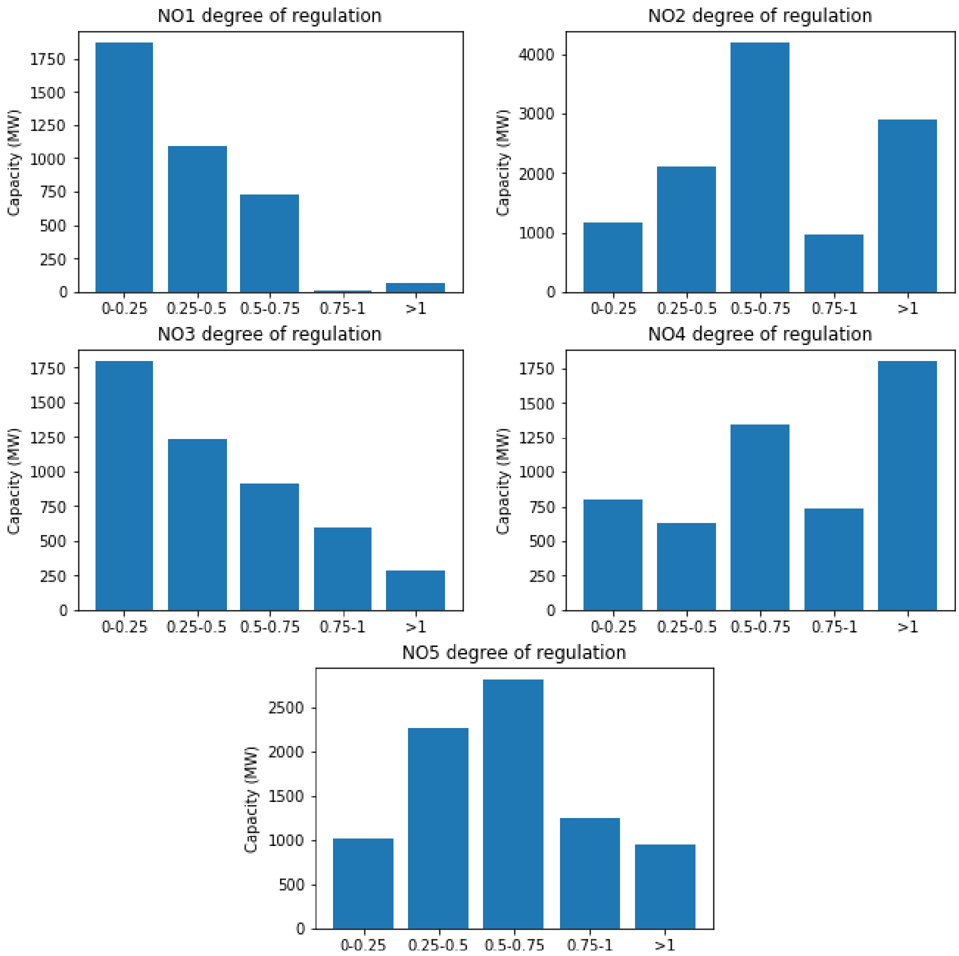

Figure 6: Production capacity by DoR for all five price areas in Norway.

exemplified by the Lule River (Fig. 7) where the availability of water at the downstream run-of-river stations can be regulated by the bigger reservoirs Suorva and Tjaktjajaure. It can therefore be argued that the division into run-of-river and reservoir plants cannot be taken as a perfect measurement of the flexibility in a hydro system.

3.3 Hydrology

In the analysis of hydroelectric power systems, it is standard procedure to convert all units related to hydropower production into energy equivalents. Data on reservoir storage, precipitation, inflow and snow are given as units of energy. Energy content of stored water is calculated using the efficiency of all downstream hydro plants. Thus, the energy equivalent of one unit of water depends on the location of the water in the system. A unit of water at higher altitude, which will pass through many downstream plants, will have a higher energy equivalent compared to a unit of water at lower altitude with a smaller number of hydropower plants downstream. The characteristics of hydropower plants also affect the energy content of water as e.g. a plant with high

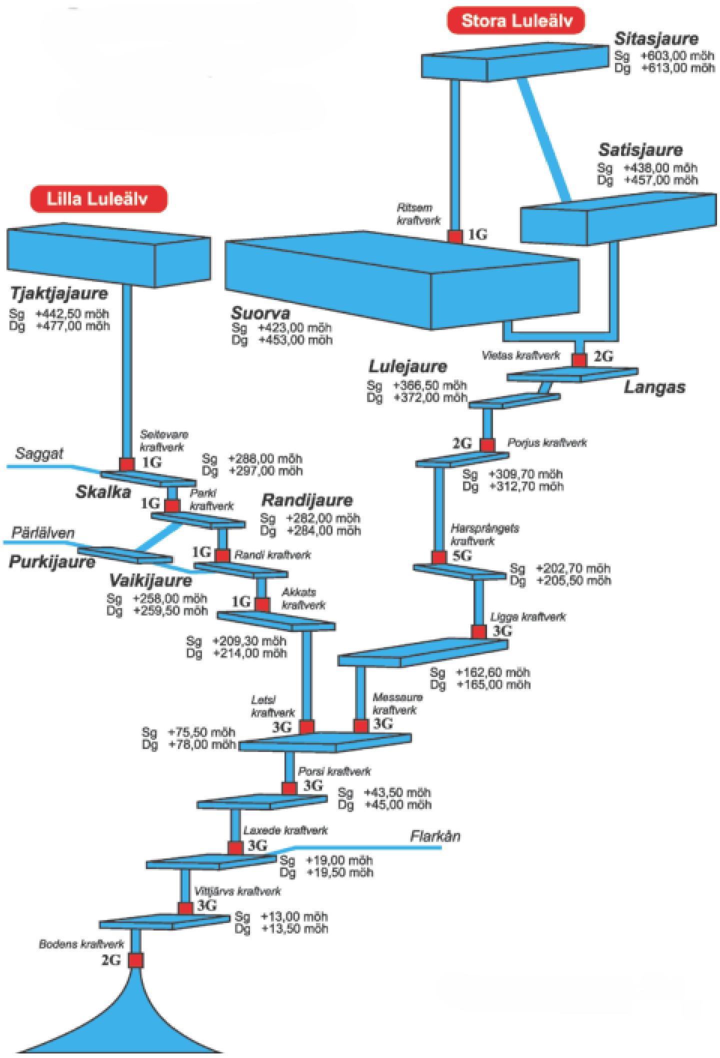

Figure 7: Example of water system with cascaded hydro plants [19]. The Lule river in price area SE1 consists of two main branches with a total of 15 hydropower plants and a total installed capacity of 4350 MW.

head can convert a small quantity of water into a large quantity of energy. Similarly, precipitation and inflow can be converted to energy according to their geographical location. Converting the precipitation amounts into inflow to reservoirs is not simple, as the topology of the catchment area needs to be considered. Hydrological processes such as soil moisture and evaporation further complicate this task. Hydrological models have been developed to accurately determine the inflow to reservoirs. One such model is the HBV (Hydrologiska Byråns Vattenbalansavdelning) model [20]. The HBV model was originally developed to help hydro producers forecast the inflow to the reservoirs by simplifying the complex hydrological environment into a small number of catchment areas. Despite its relative simplicity, the HBV model has proved useful for estimating inflow to hydro reservoirs in the Nordics [20].

Hydrological models can be used to obtain additional data besides reservoir levels and inflow. A key indicator of the hydrological situation is the hydro balance which measures the energy equivalent of the deviation from historical normal reservoir filling levels as well as the deviation from historical normal snow and groundwater levels. In early autumn when there is very little snow, most of the water in the hydrological system has ended up in the reservoir and the hydro balance corresponds almost exactly to the filling level of reservoirs. On the other hand, water is stored as snow in the winter and a surplus in snow can mean that the hydro balance is positive even if reservoirs are below normal filling level.

The availability of the hydrological data described above is somewhat limited. Typically, the reservoir level and inflow to individual reservoirs is not public information. The only publicly available data is the aggregated filling levels of reservoirs in each price area. Hydro producers must report the total energy content of each reservoir to Nord Pool on a weekly basis and the total energy content of all reservoirs within a price areas is subsequently published. There are also third party companies providing data on inflow, precipitation and hydro balance on a price area level. This data is highly useful to determine how the current hydrological situation will affect hydropower production.

The hydrological conditions follow a typical seasonal pattern during a year (Fig. 8). Inflow is low during winter, as precipitation falls as snow. During spring, the inflow increases rapidly as the snow melts. The timing of the spring flood from snow melt can vary from year to year. During summer, inflow stays at higher levels as snow is still melting from the higher mountains. Snow melt at the highest altitudes can continue into late July or early August. During autumn the inflow is usually higher due to more rainfall. Inflow can vary significantly from year to year, which can be seen in figure 9. The variation in inflow is low during winter as it is very rare to have high enough temperatures to start the snow melt before March. Inflow uncertainty becomes very large in the spring, as the timing of the spring flood changes from year to year. The uncertainty continues to be high during the summer as the amount of rain can differ significantly between years. Especially during autumns, storms can bring large amounts of rain in some years, causing spikes in the inflow.

The seasonal variations in electricity demand is very different to the seasonal variations in inflow. Demand for electricity is strongly connected to outdoor temperatures with consumption highest during winter, when the need for heating is the largest. The yearly hydro production profile follows the yearly consumption profile closely with more production in the winter compared to summer. This highlights the flexibility of hydropower and the ability to store water to produce during times of high demand. This can be observed when studying the reservoir level over the year in figure 8. The aggregate Nordic reservoir level decreases during winter when hydro production is greater than inflow and increases during spring and summer when inflow is greater than production.

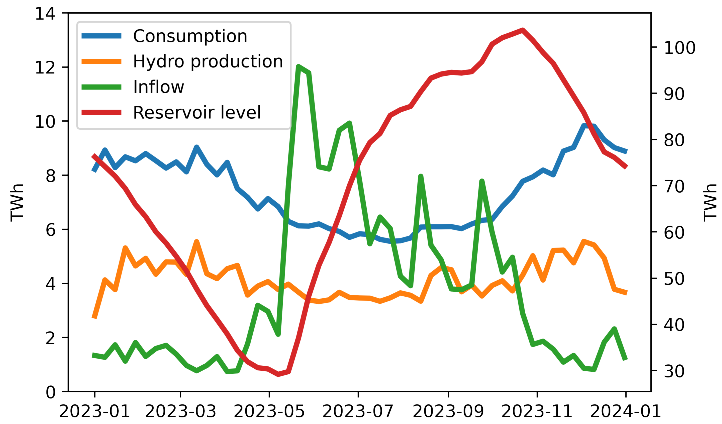

Figure 8: Total weekly consumption, hydro production and inflow (left axis) and reservoir level (right axis) during 2023 in the Nordics. Data from Volue [12].

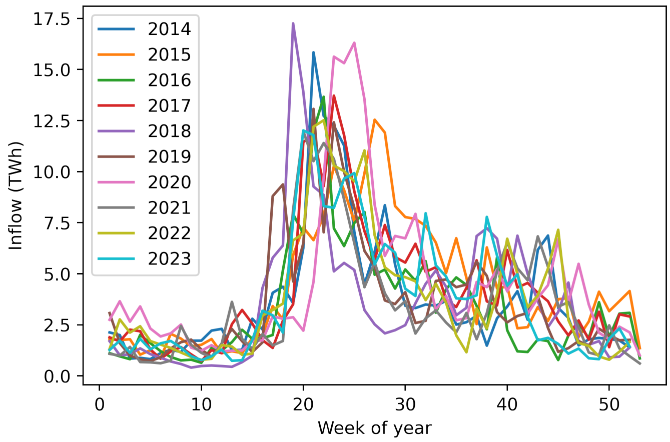

Figure 9: Weekly total inflows in the Nordics during years 2014-2023. Data from Volue [12].

3.4 Hydro supply curve and water value

In studying the bidding behaviour of hydro producers, our focus will be on the hydro supply curve, which we define as the aggregated supply curve from all hydro producers within a given price area. As the actual day-ahead bids of individual producers are not published by Nord Pool, it is impossible to reconstruct the true hydro supply curve. For clarity, the unknown hydro supply curve is referred to as the *true hydro supply curve* from here on. Information about the true hydro supply curve can be obtained by studying realized prices and realized hydro production. In areas with a large share of hydropower, most supply bids are made by hydro producers. Therefore, it can be assumed that a hydro plant is always the price setter [21]. If it is assumed that all hydropower is sold to the day-ahead market, the day-ahead price and hydro production volume for one hour will be one point on the true hydro supply curve for that hour. This is illustrated in Figure 10, where the 24 price-volume pairs of one day are plotted. The price-volume pairs of several consecutive hours will hereafter be referred to as the *empirical hydro supply curve*. Along with the empirical hydro supply curve we also plot an estimation of the corresponding true hydro supply curve in Figure 10. In an ideal case, the true hydro supply curve would be constant for multiple hours (for example all 24 hours in a day) which would result in all points on the empirical hydro supply curve being observations of the same true hydro supply curve. In the example, this is not the case as the empirical supply curve is not an increasing function. This could be because some hydropower is sold in the intraday and balancing market. Furthermore, the true hydro supply curve can change even during short time horizons violating the assumption of a constant supply curve. Even though the empirical hydro supply curve is not a perfect representation of the true hydro supply curve, it still resembles an increasing function and provides a good estimate of the true hydro supply curve by fitting an increasing function to the empirical hydro supply curve.

Based on the example in Figure 10, it can be observed that the hydro supply curve is a nonlinear function consisting of three distinct parts. At lower production volumes, the curve first rises steeply, followed by a longer flatter section and another steep rise at higher production volumes. The leftmost steep part will be referred to as the minimum production level. The minimum production level depends on the amount of unregulated run-of-river production and the steep part reflects the large increase in price when going from run-of-river to reservoir hydro. The minimum production level can only be observed during times of low hydro production, for example during times of high wind production. The flat part in the middle represents the pricing of reservoir hydro. In areas with a lot of production from reservoirs, the flat part will be longer than in areas with less reservoir production. The rightmost steep part represents the

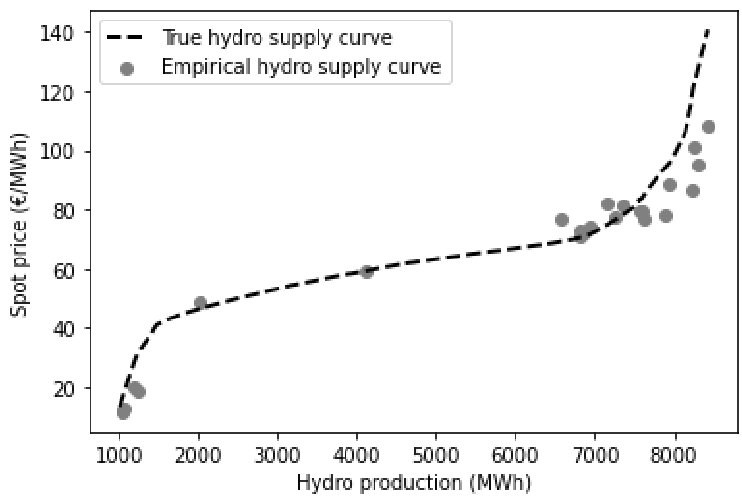

Figure 10: Illustrative example of points on the empirical hydro supply curve and true hydro supply curve.

total production capacity in the price area and will be referred to as the maximum production level. In this part, the pricing can be affected by opportunity costs in the intraday and balancing markets. This is the case if a hydro producer expects a possibility of high prices in other markets. Also, decreased turbine efficiency at maximum output can explain the sharp increase in price at the maximum production level. Another possibility is that the assumption that a hydro plant is price setting does not always hold and instead the price is set by a e.g. a thermal plant.

We now study some examples of empirical hydro supply curves in different price areas during three consecutive days in April 2024 (Fig. 11). The flexibility of the production system is an important explaining factor when studying the shape of the hydro supply curve. The price areas NO1, NO2, SE1 and SE2 were chosen as examples as these areas are dominated by hydro production and they have varying shares of run-of-river production. As previously mentioned, NO1 is dominated by run-of-river production, while NO2 has larger storage capacity and a large production volume in general. Price areas SE1 and SE2 have some storage capacity but less than NO2. The empirical hydro supply curve in NO2 is very smooth, indicating that the assumption of a constant true hydro supply curve holds. In this case it can be argued that the empirical hydro supply curve gives very good information about the true hydro supply curve. In contrast, the curve for NO1 is more scattered as there is more uncontrollable run-of-river production, indicating that the true hydro supply curve changes more from hour to hour and day to day. The empirical supply curves in SE1 and SE2 are also somewhat scattered. The price-volume pairs line up quite well during the first two days, but in the third day the curve seems to be shifted slightly higher compared to previous days. This indicates that also the true hydro supply curve may have been shifted upwards in the third day. The supply curves in all areas follow the same nonlinear shape observed in Figure 10. However, it is hard to tell exactly where the maximum production level is in areas NO1 and NO2, since there are no observations at very high prices.

The shape of the empirical hydro supply curve can change significantly over longer time horizons, which can be seen in Figure 12. Prices correlate better with production volume during the winter day than during summer days. This is likely due to the lower inflow during winter, resulting in more production from reservoir plants than run-of-river plants. Reservoirs can better optimize their production against the day-ahead price when inflow is low as it becomes easier to plan the production in complex river systems (c.f. Fig. 7). When inflow is higher during summer, it may be necessary to run relatively high hydro production, as water would be spilled otherwise. The price level is usually higher in winter which can be seen as the supply curve being shifted upwards. This is especially clear in SE2, where the flat part was much higher in the winter compared to the summer. The total hydro production level can also vary between seasons, as seen in NO1. During winter, the production from run-of-river plants is lower while it increases in summer. This can be seen as a shift of the hydro supply curve from left to right.

We can also see that the maximum production level varies significantly during the year due to maintenance. Most maintenance is performed during summer months reducing the maximum production level. This is especially noticeable in NO2, where

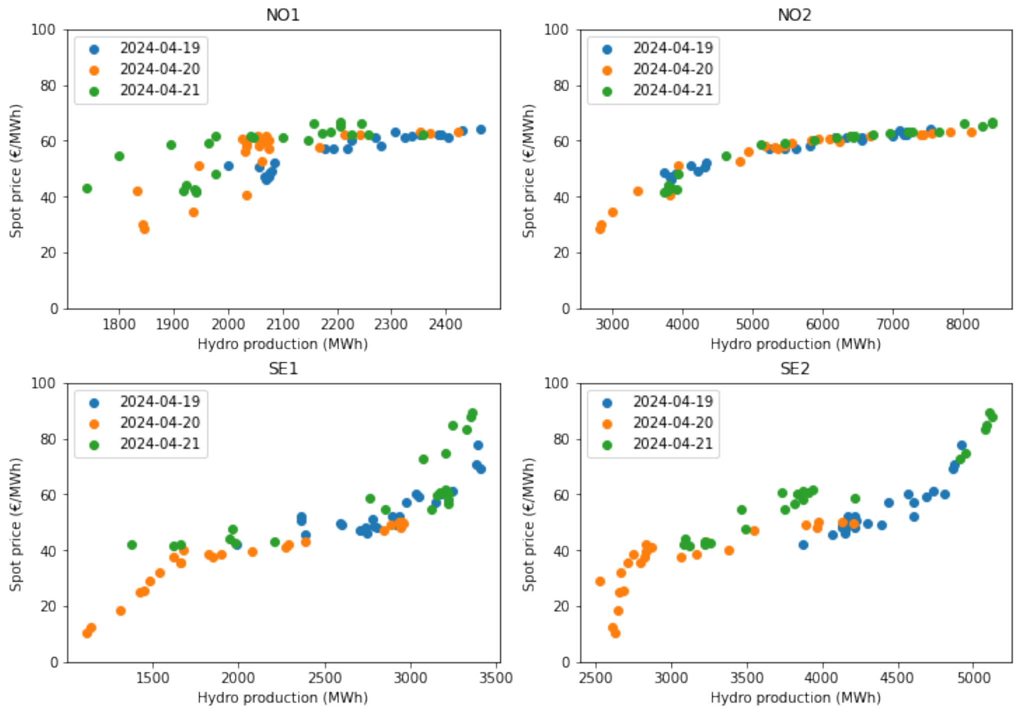

Figure 11: Empirical hydro supply curves in price areas NO1, NO2, SE1 and SE2 on three consecutive days in April 2024. Data on hydro production from Volue [12].

the maximum production level in the summer is only half compared to the winter. To summarize, it is clear that the true hydro supply curve can change significantly between seasons, as the price-volume pairs in Figure 12 are not points on the same true hydro supply curve. The shifts both in production and price levels are significant and will have a large effect on the price formation in the day-ahead market. This highlights the importance of understanding the dynamics of the hydro supply curve to model electricity prices.

The different parts of the hydro supply curve represent water values for different hydro reservoirs. The water value describes the opportunity cost associated with the use of water for electricity production and can be defined as the expected marginal value of one unit of water stored in the reservoir. If the water value is known, the production decision should be to produce at optimal efficiency if the price is above the water value, and not producing when the price is below the water value. The merit order on the hydro supply curve is decided by the water values of individual reservoirs, which in turn depend on the filling level of the particular reservoir as well as other reservoirs, thermal prices, consumption etc. Since some reservoirs are smaller, their ability to store water is limited which will lead to lower water values. The shape of the supply curve will then be determined by how much water values differ between reservoirs in the price area. If there is a large difference in water values, the supply curve will be steeper, while it will be flatter if the water values are more similar.

The hydro supply curve can be thought of as an aggregation of supply curves of individual hydropower plants. The supply curve of a hydropower plant can be described

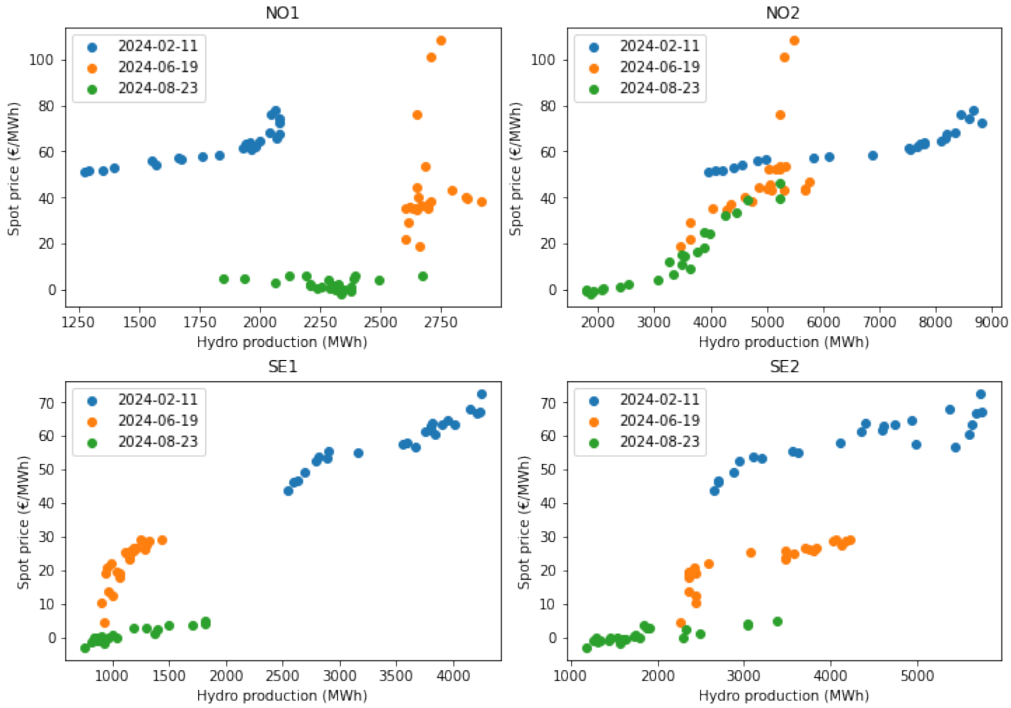

Figure 12: Empirical hydro supply curves in price areas NO1, NO2, SE1 and SE2 during different times of the year. Data on hydro production from Volue [12].

as a step function, where the step is located at the water value. This assumption can be validated by studying the hydro production of individual hydropower plants as a function of the day-ahead price. In Figure 13, the relation between production and price at Kvilldal and Tokke hydropower plants is shown for the week starting 4.2.2024. Kvilldal is the largest hydropower plant in the Nordics with 4 turbines with a maximum capacity of 410 MW each. The plant gets the water from lake Blåsjø, which is the largest reservoir in the Nordics. Based on these plant characteristics, it can be assumed

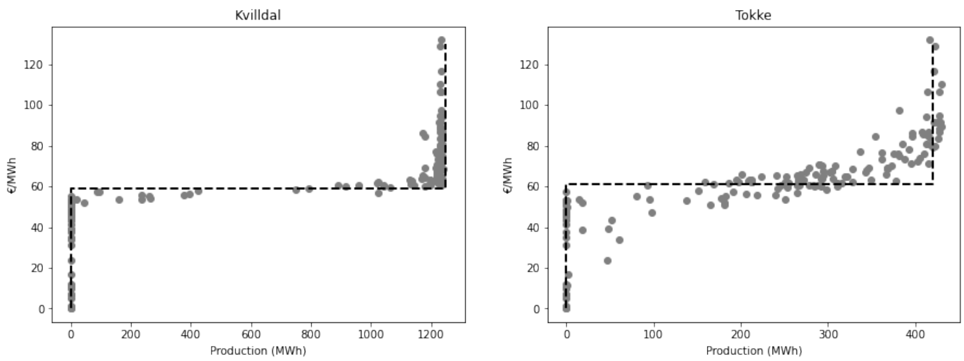

Figure 13: Scatter plot of hourly price and hourly hydro production at Kvilldal and Tokke hydropower plants in price area NO2 between 4.2.2024 and 10.2.2024. The dashed step functions represents the supply curves of these plants.

that production can be adjusted very flexibly according to the day-ahead price. Tokke hydropower plant is also connected to a large reservoir and the total capacity of the plant is 430 MW and the plant has 4 turbines, each with a capacity of 110 MW. The scatter plots indicate that the water value for the both plants was around 60 €/MWh during this time period. The production from the individual turbines is shown as a concentration of points slightly below 300 MW. It can be assumed that three of the four turbines were running at close to optimal efficiency during these hours.

Comparing the production strategies at Kvilldal and Tokke, it is clear that Kvilldal is able to better optimize the production against the day-ahead price. The efficiency curves of the turbines could explain this difference. It can, for example, be the case that the efficiency in Kvilldal reduces quickly if the turbines are not run close to optimal efficiency. Another explanation for the differing production patterns could be the ownership structure. Tokke is fully owned by Statkraft AS while Kvilldal is also owned by other smaller companies. The ownership structure can impact how the electricity generated by a given hydropower plant is sold to the market as bids are based on a company's whole production portfolio. A fully owned hydropower plant could be better utilized in the balancing market.

In conclusion, the hydro supply curve describes at what price a given amount of hydro production is run. We cannot observe the hydro supply curve directly, but instead we can observe the price-volume pairs for given hours. Based on empirical observations, the hydro supply curve does not change much over shorter time horizons, but larger changes can be observed over time. The changes can be seen both in the price level, corresponding to upwards or downwards shifts in the supply curve, or changes in production level corresponding to shifting the curve left or right. Both of these changes can happen at the same time, which can significantly alter the shape of the hydro supply curve. The underlying reason for changes in the hydro supply curve is that water values of individual hydropower plants change. To model the hydro supply curve, a good understanding of the calculation of the water value is needed.

4 Hydropower scheduling

This section presents existing methods for calculating the water value. Understanding the underlying problem from which the water value is solved is important to explain the fundamentals affecting the hydro supply curve. The water value is calculated by optimally allocating hydro production under uncertainties in inflow and other market conditions, such as fuel prices and availability of transmission lines. This problem, which will be referred to as the hydropower scheduling problem, is a stochastic dynamic optimization problem. We will study the formulation of the problem and the solution methods proposed in the literature. A vast share of the literature on the scheduling problem originates from the Nordic countries with a focus on the Nordic electricity market. We will therefore focus on the methods developed for the Nordic hydro systems, even though the problem has also been studied for hydro systems in other parts of the world, such as New Zealand [22] and Brazil [23].

The first methods for hydropower scheduling and water value calculation were developed by Stage and Larsson in the 1960's [24]. The term water value originates from their work. The aim of the water value method is to translate the long-term value of stored water into short-term operational decisions [25]. The water value method was further developed in Norway while the electricity market was still regulated. An overview of these developments as well as current scheduling methods are given in [26]. The methods described in [26] seem to be used by most hydro producers in the Nordics.

The hydropower scheduling problem is a large-scale problem which is computationally difficult to solve. Therefore, the problem is usually divided into a hierarchy of models with different time horizons and level of detail. The most common approach is to use three models: the long-term, medium-term and short-term model [27]–[29]. Splitting the scheduling problem into sub-problems reduces the computational complexity of the problem, while still allowing for long-term uncertainty in inflow and a detailed enough model of the production system in the short-term. Figure 14 shows a schematic picture of this decomposition. Long-term scheduling is done on a weekly resolution with a time horizon of 3-5 years. Medium-term scheduling is also done on a weekly resolution but with a time horizon of 12-18 months. Short-term scheduling is done on an hourly resolution for the next 1-2 weeks. Information is passed through the models from the longer to the shorter term. The long-term model gives price scenarios to the medium-term model, while the medium-term model passes water values to the short-term model. The different sub-problems make it possible to focus on different aspects of the scheduling problem. The long-term scheduling is important to optimally allocate production from large reservoirs, where the decision is whether to produce during the current year or next year. Short-term scheduling is done on a very detailed level, where specific characteristics of the hydropower plants and water systems are taken into consideration.

The problem can be modelled from the perspective of a centrally dispatched market or a deregulated market. In the case of a centrally dispatched market, the objective is to meet electricity demand while minimizing cost. In contrast, in a deregulated market the objective is to maximize profit under uncertain price and inflow. We assume that

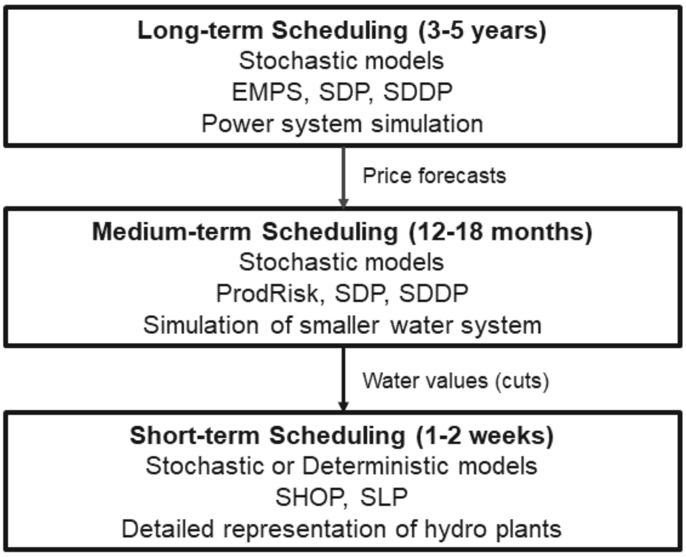

Figure 14: Flow-chart of the different sub-problems in the scheduling problem, the models used to solve them and information flow between sub-problems.

each participant in the decentralized market is a price taker, meaning that no single market participant can influence prices. If all market participants are price takers, the market is perfectly competitive and the decentralized approach should lead to the same hydro scheduling as the centralized approach. If market competitiveness cannot be assumed, oligopolistic behaviour needs to be taken into account, significantly complicating the problem. From here on, we assume a perfectly competitive market. It has been shown that the optimal dispatch schedule in the profit maximization problem converges to the cost minimization problem if the market is competitive and the hydro plants of different companies are isolated, i.e. not hydrologically connected [30].

The similar dispatch schedules in centrally dispatched and deregulated markets have been verified empirically in [2] where reservoir handling before and after deregulation is studied. They conclude that the average filling level of the reservoirs decreased by 4.6% after the deregulation. This can be explained by structural changes in the market, such as increased transmission capacities. In view of structural changes, the reservoir levels corresponded well to reservoir levels simulated by hydro scheduling models. These results indicate that optimal reservoir handling is independent on whether the problem is studied from the perspective of a central dispatcher or a producer in a decentralized market.

4.1 Long-term scheduling

The objective of long-term scheduling is to generate a set of price forecasts by simulating the whole power system over a longer time horizon. In our case, the system to be simulated includes the whole Nordic region. With the large share of hydropower, the decisions of hydropower producers will have a large influence over prices in the long-term. Therefore, it is particularly important to accurately simulate hydropower production in the long-term model. As many relevant simulation variables, such as inflow, demand, wind production and fuel costs, are unknown, the simulation must account for uncertainties. The long-term scheduling problem must therefore be modelled as a stochastic optimization problem. The long-term scheduling is formulated as a hydro-thermal scheduling problem, where the objective is to use hydro resources so that the cost of thermal generation is minimized. The complexity of the problem requires several simplifying assumptions, of which the most important is the aggregation of reservoirs [31]. In the aggregation step, all hydro reservoirs and plants, usually within the same price area, are modelled as one aggregate plant. The simulation period is divided into shorter time steps, usually one week, with the realizations of stochastic variables known one week ahead.

To illustrate the problem, we present a simplified mathematical model of the hydro-thermal scheduling problem, where the objective is to utilize hydropower such that the cost of electricity production from thermal plants is minimized. We assume a hydro-thermal system consisting of one thermal plant and one hydro reservoir plant. The inflow to the aggregate reservoir plant is stochastic, while the demand in all weeks is assumed known. We will further assume that the cost of thermal generation, c is constant over the whole time period. The objective is to minimize the expected total cost of thermal generation plus the future cost, Φ_T, of operating the system at the end of the scheduling horizon, (equation (1)). The future cost must be included, as the reservoir would otherwise be emptied at the end of the scheduling horizon. The expectation is taken over the stochastic inflows, which in practice means taking the average over several inflow scenarios. The decision variables in the problem are the hydro generation, q, the thermal generation g, the reservoir level w and the spillage u. The parameters in the model are demand d and the stochastic inflow variable v. All of the mentioned variables are indexed and aggregated over the time steps. The constraints in the model represent the supply and demand balance (2), the water balance in the system (3), the hydro production (4) and maximum reservoir capacity (5). All decision variables in the optimization problem must be non-negative (6).

$$\min \quad \mathbb{E}\left[c\sum_{t=1}^{T} g_t + \Phi_T\right] \tag{1}$$

$$g_t + q_t = d_t, \qquad \forall t = 1, \dots, T \tag{2}$$

$$w_{t-1} - q_t - u_t + v_t = w_t, \qquad \forall t = 1, \dots, T \tag{3}$$

$$q_t \leq q_{max}, \qquad \forall t = 1, \dots, T \tag{4}$$

$$w_t \leq w_{max}, \qquad \forall t = 1, \dots, T \tag{5}$$

$$q_t, g_t, w_t, u_t \geq 0, \qquad \forall t = 1, \dots, T \tag{6}$$

In the literature, the hydro-thermal scheduling problem has been solved using Stochastic Dynamic Programming (SDP) or Stochastic Dual Dynamic Programming (SDDP) [26]. We give a brief outline of both solution approaches. More complete descriptions can be found in [32]. Both methods are based on the concept of dynamic programming, where the problem is divided into multiple sub-problems, which are solved iteratively. At each time step, a set of state variables are chosen, which describe the state of the system. The future cost function will then be a function of the state variables. In our example, the reservoir level is the only state variable. Both SDP and SDDP rely on constructing the future cost function $\Phi_t(w_t)$ at a given time step t by solving the sub-problem (7). The sub-problem is subject to the same constraints (2)-(6) as the complete problem. The sub-problem gives the future cost at a given time and in a given state as the expected cost of thermal generation plus future cost in the next step $t+1$.

$$\Phi_t(w_t) = \min \quad \mathbb{E}\left[cg_t + \Phi_{t+1}(w_{t+1})\right] \tag{7}$$

In the SDP approach, the state variables (reservoir levels) are discretized into a finite set of states. The expected future cost is calculated iteratively for each state at each time interval, starting from the final time step T. First, the future cost in the final time step, $\Phi_T(w_T)$, has to be set manually, then the future cost function at time step $T-1$ is constructed at each discrete reservoir level. For example, the expected future cost for an empty reservoir is obtained by solving problem (7) for all inflow scenarios during time t, given an empty reservoir at time $T-1$. Repeating this process for all reservoir levels in the state space and linearly interpolating the future costs gives the complete expected future cost function. A snapshot of the future cost function at a given time step is illustrated in Figure 15. At each stage, the price of hydropower is equal to the (negative) slope of the future cost function, as this slope represents the marginal value of water. The future cost function should always be convex as the marginal value of water should decrease with increasing reservoir filling level. Finally, price scenarios for each inflow scenario can then be generated by making optimal hydro production decisions and setting the price at each stage equal to the water value.

The SDP algorithm works well if the state space, i.e. the number of reservoirs, is small. However, the number of times the sub-problem has to be solved grows exponentially when the number of reservoirs increases. Hence, the computational burden becomes too high even for a moderate number of reservoirs. The curse of dimensionality can be solved by applying SDDP [32]. Like SDP, the SDDP algorithm uses a future cost function, but its construction process is different. The state space does not need to be discretized and instead the future cost function is constructed by iteratively adding linear functions to the future cost function. In Figure 15, this can be seen as adding linear segments to the future cost function. These linear segments are also referred to as cuts. The SDDP algorithm consists of a forward and backward simulation. In the forward simulation, a set of inflow scenarios are simulated using the existing cuts. The average total cost in the forward simulation gives an upper bound for the true cost. After the forward simulation, a backward simulation is run, where the dual of the sub-problem is solved for each scenario in the forward simulation. This gives the expected future cost and the marginal value of water for this state which can be added as a cut. The total cost in the backward recursion will then represent a lower bound of the true future cost. The algorithm is iterated until the upper and lower bound converge.

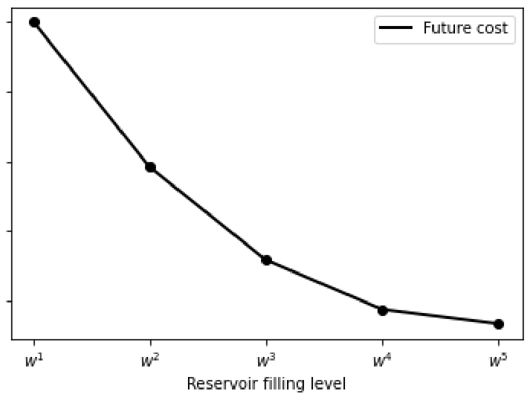

Figure 15: Illustration of the future cost function at a given time step. In the SDP algorithm, the future cost function is constructed by linear interpolation between discrete points in the state space. In the SDDP algorithm linear segments are iteratively added.

In addition to the more general SDP and SDDP approaches to the long-term scheduling problem, there are models developed specifically for the Nordic market. A popular model for long-term planning in the Nordic electricity market is the EMPS model (EFI's Multiarea Power Simulator) [26]. The EMPS model simulates the Nordic electricity market using historical weather years, with one price scenario for each weather year. The model first performs a strategy evaluation, which is similar to the SDP algorithm on an aggregate reservoir. Then, a heuristic drawdown model is used to describe the multi-reservoir characteristics of the system [26]. The long-term forecast generated by the EMPS model is not ideal for short-term price forecasting, since the model is aggregated on a weekly level and does not output hourly prices. In particular, the increase of intermittent electricity production and stronger connections to Central Europe has posed new challenges for the EMPS model, as the short-term effects from these factors are hard to handle with aggregation and disaggregation [31]. However, the aggregated model will be a good approximation if it can be assumed that no single reservoir will become full before all other reservoirs are full and no reservoir is empty before all other reservoirs are empty [26].

A challenge in the long-term scheduling is to construct inflow scenarios as inflow has a significant temporal and spatial correlation [29]. Inflow in one week is often positively correlated with the inflow of the previous week, and the inflow in one area correlates positively with inflows in neighboring areas. It has been observed that constructing inflow series using statistical methods leads to extreme scenarios being less represented [33]. This could result in the reservoirs being emptied too early before the spring flood or the reservoirs being filled too early, leading to spillage during autumn. For this reason, historical time series for inflow are preferred to statistically generated time series as the geographical and temporal correlation of the inflow is better represented in historical data. Using historical data comes at the expense of fewer available inflow series.

4.2 Medium-term scheduling

The results from the long-term model cannot be used directly for the short-term operational decisions as reservoirs are aggregated but water values for individual reservoirs are needed for operational decisions. The medium-term model aims to bridge this gap between the long-term and short-term scheduling by computing water values for individual reservoirs. The reservoir specific water values can then be used as an input to the short-term operational scheduling model [34]. In the medium-term model, the objective is changed from cost minimization to profit maximization [26]. The problem is formulated from the perspective of a hydro producer, who aims to optimally schedule a water system under uncertainties in price and inflow. The geographical area for the simulation is reduced from the whole Nordic area to individual water systems, which enables modelling on individual reservoir level without the need for their aggregation. Some complexity is added as the medium-term model must consider water flows between reservoirs. However, water systems which are not hydrologically connected can be scheduled separately, decreasing the complexity.

One model commonly used for medium-term scheduling in the Nordic region is ProdRisk [26] which applies a combination of SDP and SDDP. The mathematical formulation is described in detail in [35] and [29]. To compute water values for multiple reservoirs, SDDP is preferred as the SDP algorithm would require each of the reservoir levels to be a state variable. However, SDDP cannot be applied directly as the expected profit is not a concave function of the price level. The price scenarios from the long-term scheduling (EMPS model) is an exogenous variable and the temporal correlation in prices is also estimated from the EMPS model. Ideally, the inflow scenarios to the hydro plants should be the same inflow scenarios as in the long-term forecast. The long-term scheduling is done on price area level. It is difficult to say exactly how inflow on price area level should correlate with inflow to a given reservoir within the price area. Often a simplifying assumption is made that the local inflow and price are independent [29].

The output from the medium-term model are the future profit functions at each time step. Like in the long-term model, the future profit function is constructed iteratively by adding linear cuts in the SDDP algorithm and therefore the future profit function will be a piecewise linear approximation of a true future cost function. The slope of the linear parts can be interpreted as water values at different filling levels. This is shown in Figure 16, which shows the future profit is plotted as a function of reservoir size. The tangent to the future profit function is plotted at 25%, 50% and 75% filling levels. It can be observed that the future profit function is concave, which can be interpreted as a decrease in marginal value of water with increasing reservoir filling level.

If the water system consists of multiple reservoirs, the future profit function is a function of the filling level of all reservoirs in the system. In particular, it is important to note that the water value in a particular reservoir is therefore not only a function of the filling level of the given reservoir but also dependent on the filling level of other reservoirs. This can be interpreted so that the water value in a reservoir is lower if the other reservoirs in the same water system are well filled. The future profit function

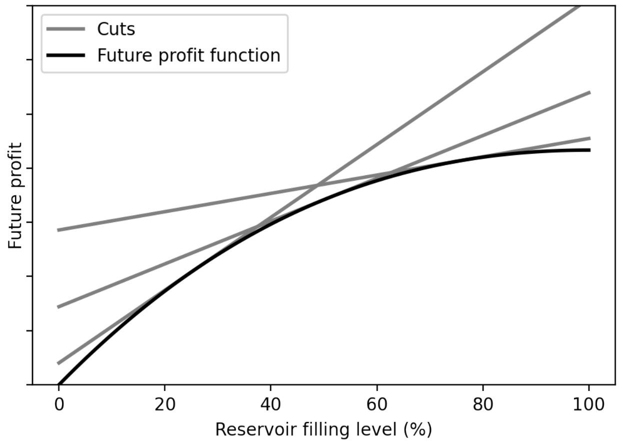

Figure 16: Illustration of the future profit function with cuts.

in Figure 16 should therefore be interpreted as a one-dimensional projection of a multi-dimensional function. The cuts would then be represented by hyperplanes and the cuts in Figure 16 are the one-dimensional projections of these hyperplanes.

4.3 Short-term scheduling

The objective of short-term scheduling is to obtain an operational plan for a single hydropower plant or a set of cascaded hydropower plants in a watercourse. A review of different approaches to short-term scheduling is found in Kong et.al. [36]. Short-term scheduling must be done on a detailed level, taking into account the efficiency curves of turbines, hydraulic losses and other plant characteristics [36]. This level of detail can increase the complexity of the model significantly, which can become a problem as the operational use of short-term scheduling requires fast computation times. The problem contains nonlinearities in the form of losses, efficiency curves and head effects [37]. These factors can even make the problem non-convex. Solution methods for the short-term problem include linear-programming, nonlinear programming and integer programming. The short-term scheduling model can be either deterministic or stochastic [38]. In a deterministic model, the prices and inflows are considered known during the scheduling horizon, while the stochastic version considers multiple price and inflow scenarios. The deterministic approach leads to a smaller computational burden, which can allow for greater detail in the modelling of the system.

The complexity of the short-term scheduling problem has also lead to the development of heuristic algorithms [39]. The short-term model requires a very detailed model of the system, considering technical restrictions on generating units and reservoirs. The use of heuristic methods makes it harder to model the hydro supply curve as the operational decisions of real hydro producers may not perfectly correspond to theoretical scheduling decisions. It is also impossible to consider all the characteristics of watercourses and individual hydropower plants that affect the operational decisions.

In the Nordics, a common program used for short-term scheduling is SHOP (Short-term Hydro Optimization Program) [28], [40], which applies successive linear programming (SLP). SLP handles nonlinearities in the problem formulation by iteratively solving a linear optimization problem. The formulation of the short-term problem is similar to the medium-term problem, where the objective is again profit maximization. Compared to the medium-term problem, more constraints are added to model the system in high detail. The short-term model is coupled to the medium-term model through the cuts representing water values from the medium-term model. The cuts at the end of the short-term horizon represent the value of storing water.

4.4 The bidding problem

In addition to the short-term scheduling problem, hydro producers face the bidding problem. The short-term scheduling gives a production plan for the upcoming days but the producer still has to decide how to optimally sell the production in different markets. The bidding problem can be viewed as an extension of the short-term scheduling, where the task is to construct the price-volume bids given to the market. As the hydro supply curve is an aggregation of the bids from individual producers, it is important to understand how individual producers bid their production to the day-ahead market in order to understand the hydro supply curve.

Most production is sold in the day-ahead market, but opportunities in other markets, such as the intraday and balancing markets should also be considered when preparing bids to the day-ahead market [41], [42]. If a hydro producer sells one unit of electricity in the day-ahead market, the ability to sell electricity in the intraday and balancing markets are limited. The prices in different markets are unknown at the time of bidding, with prices and dispatched volumes being revealed over time. The bids in the day-ahead market, and thereby also the hydro supply curve can therefore be affected by the hydro producers' price expectations in the intraday and balancing markets.

Bidding optimization methods can be divided into deterministic and stochastic. Previous studies indicate that stochastic methods have potential to increase the profits of the producer. Stochastic methods may, however, suffer from significant time-complexity which is a limiting feature as decisions need to be taken within given time limits. Fleten and Kristoffersen [43] compare the bidding strategies given by a deterministic and a stochastic formulation. They model the bidding problem as a stochastic mixed-integer program with the price being a stochastic variable. They find that a stochastic formulation yields better results compared to a deterministic model. The stochastic approach is further developed in [44], which also find that the total revenue can be improved by use of a stochastic bidding approach compared to a reference heuristic method.

There has been several bidding strategies for hydropower proposed in the literature. Three different heuristic models for hydropower bidding are presented in [45]. The models are the expected volume methods, the water value method and a method based on multi-scenario deterministic optimization. The most simplistic method is the expected volume method, where bids are optimized according to a deterministic price forecast. This method is expected to yield reasonable results as long as the price forecast is very reliable. The water value method is based on bidding according to the water value, where no bids are offered below the water value and the optimal production level is offered at the water value. This bidding behaviour yields a step function, discussed in section 3.4. The multi-scenario deterministic method generates several price scenarios and one bid curve is formed for each scenario. The bid curves from each scenario are then combined to form the final bid curve. This approach better takes uncertainty in price into consideration. Comparing the performance of the three models, the multi-scenario deterministic model was found to give the highest profits, but the difference between the models was small.

Understanding and modelling the bidding behaviour of hydro producers is difficult since some producers may use heuristic methods instead of advanced optimization tools [45]. It is hard to know exactly what methods are being applied in the industry, as the bidding strategies are central to the operations of hydro producers and thus not public information. The bidding behaviour of three hydropower producers in Norway is studied in [46]. They conclude that the bidding behaviour is sometimes irrational. There may be many factors affecting the bidding behaviour such that the producers price expectations in the intraday and balancing markets. It is impossible to take all factors affecting the bidding behaviour into account when trying to model the hydro supply curve. Thus, we assume that all producers bids in a rationally.

In summary, the bids from hydropower producers should in theory reflect the water value calculated from the scheduling models. The most important factors determining the water value in the long-term and medium-term models are uncertainty in inflow and cost of thermal production. However, in the short-term, the bidding process is more complicated than the simple water value approach, where the whole volume is bid at the water value. Characteristics of watercourses as well as portfolio optimization can affect the bidding behaviour.

5 Methods

In this section we propose a method to forecast the hydro supply curve for a short-term time horizon of six weeks. Predicting the hydro supply curve is very difficult as the current pricing of hydropower is already based on current price and inflow expectations. Known changes in market fundamentals should therefore already be priced into the current water values. Furthermore, changes in the supply curve over longer time horizons are strongly influenced by weather, which cannot be accurately predicted over longer time horizons. We first describe the data used and review methods which has previously been used to model hydro supply curves in electricity markets. A theoretical motivation for the proposed method is then given followed by the mathematical formulation of the model. Finally, properties of the model are studied followed by an overview of the model inputs.

5.1 Data

The lack of publicly available data is a major challenge when analyzing hydro supply curves. In this section, we describe the data which will be used in the modelling part of this thesis. The hourly hydropower production in each price area is public information as well as the day-ahead price. We will use the hydropower production data reported by Volue and the day-ahead price reported by Nord Pool. Assuming that hydropower is the price setting technology, this data can be used to estimate one point on the supply curve for each hour. A potential problem is that the realized hydro production contains volumes sold in other markets than the day-ahead market. Some share of the production may be sold, for example, in the intraday market, in the balancing market or using bilateral contracts. The traded volumes in the intraday and balancing markets have increased in recent years, due to the increase of weather dependent electricity production from renewable sources [47].

It is hard to compensate for the hydropower traded in the intraday market. Nord Pool reports the buy and sell volume from the intraday market but it is not possible to know what type of production is traded. If a unit of energy is sold in the intraday market a hydro producer may be the seller or the buyer or both. It is also possible that none of the parties in the trade are hydro producers. For example, if a wind power producer has underestimated their production in their day-ahead bids due to a poor wind forecast and a retailer has underestimated the consumption of their customers, this wind producer may sell their extra production to the retailer in the intraday market. For these reasons, volumes traded in the intraday market were not used to adjust the hydro production.

The volumes traded in the balancing market are easier to account for, as it can be assumed that all volumes traded in the balancing market are hydropower. This assumption is reasonable at least in Norway and northern Sweden where hydro production has a dominant share of production capacity. Furthermore, hydropower is one of the only production technologies capable of making quick and large production changes. Price areas Finland and SE3 are the only areas with large hydro production where this assumption does not necessarily hold. To adjust hydropower volumes, the buy volumes in the balancing market were added to the actual production volume, while the sell volume was subtracted. To test whether this gives a better representation of the volume sold in the day-ahead market, the Spearman's rank correlation between production volumes and day-ahead price was studied. It was observed that subtracting volumes from the balancing market generally increased Spearman's rank correlation of price and hydro production, indicating that it may better represent the true hydro supply curve. From here on, the hydro production will refer to the adjusted hydro production, where volumes traded in the balancing markets are subtracted.

Other sources of data which were also considered are the sell volumes and aggregated bid curves published by Nord Pool. The bid curves are aggregated on country level in the Nordics. However, the aggregated bid curves are only available starting from July 2022. The aggregated bid curves contain bids from all producers and production technologies and it is impossible to know which bids are made by hydro producers. The data from the aggregated bid curves is nonetheless interesting as the majority of bids come from hydro producers in some areas. The shapes of the empirical supply curve and the aggregated bid curve can be compared in these cases. All production is not necessarily sold to the day-ahead market which makes it more challenging to use the true bid curves for analysis. Furthermore, the aggregated bid curves are only available on country level and not price area level. Because of these challenges, the aggregated bid curves were not used in this thesis.

5.2 Supply curve forecasting

As seen in section 4, the literature on optimal scheduling of hydropower is well developed. However, it is unclear how to apply the scheduling models on a macroscopic level, with hundreds of hydropower plants. The aggregation approach used in long-term scheduling is not sufficient for short-term price forecasting as the aggregation only gives one water value for the whole aggregated price area and can therefore not model short-term price fluctuations. Creating a detailed model of the hydro system and calculating water values for individual hydro plants would in theory be an ideal way to reconstruct the hydro supply curve. This is unfortunately not feasible due to the large number of hydro plants and reservoirs in the Nordics. It is hard to find data on the different hydro systems and the filling levels of individual reservoirs are unknown. One must therefore use some approximate method to model the hydro supply curve, which should replicate the global behaviour of the hydro producers as well as possible.

There is very little literature on the problem of forecasting the hydropower supply curves, as noted in [48]. We have only identified four papers focusing on modelling the hydro supply curve [21], [49]–[51]. Most of these focus only on modelling the shape of the supply curve, with no attempt made to forecast the supply curve. The lack of previous studies on the subject highlights the extreme difficulty in modelling and forecasting the hydro supply curve.

The first challenge encountered when modelling the hydro supply curve is how to estimate the true hydro supply curve from the empirical hydro supply curve. The shape of the supply curve can differ significantly in different areas and during different times of the year. A possible solution is to fit a parametric function to the points on the empirical hydro supply curve, which is explored by Dueholm and Ravn [49]. They compare the fit of three different parametric models. Their choice of parametric models to be tested is motivated by the dynamics of the water values. Their work is mostly concerned with finding a functional form that describes the hydro supply curve and no attempt is made to use the model to generate predictions of the hydro supply curve in the future.

An interesting fundamental approach to model the supply curve is to estimate the water values of individual plants. This is done by Nycander and Söder [50], who study the aggregated hydro supply curve in Sweden. Their objective is to explain the short-term volatility in electricity prices by modelling the hydro supply curve. Models for hydro-dominated electricity markets often aggregate the reservoirs and struggle to reproduce the short-term volatility of power prices, as hydropower can be used to flatten out large price spikes. They find that differing water values for different hydro plants give rise to the hydro supply curve, similar to our analysis in section 3.4. They propose a method to model the water value of individual hydropower plants by studying their runtime, describing the fraction of time the plant is producing. Hydropower plants with higher runtime should have a lower water value as these plants must also produce when the price is lower. It is hard to apply their idea to model the full hydro supply curve as production data is only available from the largest hydropower plants and published with a delay of a couple of days.

Regression models is another option to model the hydro supply curve. This approach is used by Jahns et al. [21]. In their model, the hydro supply curve is modelled as a linear function, estimated from the empirical hydro supply curve. If the supply curve is modelled as a linear function, the slope and intercept have a clear meaning and it is possible to use these variables as the target variables in the regression model. The influence on the hydro supply curve based on the the aggregated filling level of reservoirs and SRMC of thermal generation is studied. Four hypothesis about these relations are subsequently proposed and verified by studying historical data. They conclude that modelling the hydro supply curve can help market participants understand the market fundamentals better, which can be used to improve price forecasts.

The hydro supply curve could also be modelled using machine learning, as done by Tolonen [51]. The approach is based on simultaneously forecasting the price and production and then constructing the hydro supply curve from these forecasts. This allows for a very flexible model of the supply curve, which can also model the nonlinearities in the supply curve. Machine learning can use a large number of features to model complex relationships. However, it can also suffer from overfitting if too many features are used. Furthermore, the availability of historical data to train a model is limited as the dynamics of the electricity markets have changed a lot in recent years, with price volatility increasing. Data older than a couple of years should therefore not be used to train a machine learning model.

5.3 Changes in the hydro supply curve

We proceed to study which factors influence the water values and thereby the hydro supply curve. The analysis will be based on the theoretical models discussed in section 4. This theoretical part will focus on some of the properties that are desirable for a model to forecast changes in the hydro supply curve.

As a base case, the empirical hydro supply curve from recent days should be a good estimate for the hydro supply curve in the future. Since water values are calculated based on future expectations, most predictable changes in market conditions should already be priced into the hydro supply curve. There are however times during the year when changes in the hydro supply curve could be expected. During spring flooding when inflows increase, the production from run-of-river plants increase rapidly which should move the hydro supply curve laterally to the right. The water values of reservoirs are also likely to drop in spring, as the risk of emptying the reservoirs in the near future decreases and the next time reservoirs risk running empty is during next winter. Similarly water values should increase in late fall when inflow levels decrease and there is no more a risk of spillage.

We seek to forecast the hydro supply curve during a relatively short horizon and therefore we are mostly concerned with the theoretical properties of the short-term and medium-term models. We assume the short-term scheduling is done for one week. The connection between the short-term model and the medium-term model can be illustrated with the following example. We consider a simple case where the production system consists of only one reservoir. We assume availability of accurate price and inflow forecasts for the next week and water values from a medium-term model. The water values from the medium-term model describe the future profit at each possible reservoir level in the end of the week. Let the price forecast during each hour of the week be p_t and the inflow forecast be v_t. The future profit as a function of reservoir level is $\alpha(w)$ and the reservoir level at the start of the week be w_0. The production during each hour of the week is denoted q_t, with the upper limit on production in each hour being q_{max}. We use capital letters to denote aggregated values over the week. The total inflow during the week is denoted V. We denote the total production during the week Q and similarly Q_{max} denotes the total produced electricity during the week if the plant is run at maximum production during each hour. The storage capacity in the reservoir is assumed to be large such that there is no possibility of the reservoir becoming full or empty during one week.

Deciding whether to produce hydropower is based on the trade-off between immediate profits and expected future profit. The immediate profit is the profit during the first week, while the expected future profit is the expected profit made after the first week. We define both the immediate and future profit functions as functions of the reservoir volume at the end of the week. The concept is described in [32], with the difference that the objective is to maximize profit instead of minimizing operating costs. Constructing the immediate profit as a function of end reservoir volume is straightforward when the price for the coming week is assumed known. First, the immediate profit is calculated as a function of the hydro production, which can be done by ordering the hours from the most expensive to the cheapest. If production is only run for one hour, the immediate profit is the price in the most expensive hour. If production is run for two hours, the immediate profit is the sum of the prices in the two most expensive hours. Repeating this process, the full immediate profit function can be constructed for all total production volumes between 0 and Q_{max}. This function can then be translated into a function of the final reservoir level. The reservoir volume at the end of the week can be computed from the total production as $w_{end} = w_0 - Q + V$. The highest possible final reservoir volume is then $w_0 + V$, corresponding to no production during the week. Similarly, the minimum final reservoir volume is $w_0 + V - Q_{max}$ corresponding to full production in each hour during the week.

We illustrate the immediate and future profit functions in Figures 17a and 17b. If more water is used in the coming week, the end reservoir level will be lower and the immediate profit from production will be higher. However, since the final reservoir level will be lower, the expected future profit will be smaller. The optimal final storage level will be the reservoir level where the derivatives of the functions are the same, with opposite signs. At this reservoir level, the marginal value of using water in the next week is the same as the marginal value of storing water for later use. There exists at most one such point, since both the immediate profit function and the future profit functions are concave functions of the final reservoir level. This is illustrated in Figure 17c, where the sum of the immediate and future profit functions are plotted. The maximum can be found where the derivative of this function is zero. It should be noted that scale of the future profit function is different compared to the immediate profit function.

The shape of the future profit function in the example is almost linear, which is realistic if the reservoir is large. When the expected future profit function is almost always linear, the expected future profit function is much more important for computing the water value. This follows from the fact that the water value must be the slope at some point on the future profit function. Hence, the decisions made within the horizon of the short-term model will have small effects on the storage level, compared with the total size of storage. The water values of larger reservoirs should therefore change slower over time as the reservoir level must change significantly to cause a large change in the water value from the future profit function. On the other hand, the production decision will increase or decrease the reservoir level quicker in a smaller reservoir, which can lead to quicker changes in the water values.

Based on the example, we can identify how different factors should influence the water values. In this simple model, changes in water values can happen for two reasons, either the short-term price forecast changes or the inflow forecast changes. An increase in the price forecast in the near future will increase the slope of the immediate profit function and should lead to higher water values. It should also result in more water being allocated for immediate production. The increase in the price forecast could be due to a number of factors, such as increased consumption, lower production from renewables or changes in transmission capacities. A change in inflow will shift the whole immediate profit function to the left or right as the final reservoir level will change. The water value will then change as the equilibrium final reservoir level will also be moved left or right along the future profit function. Increasing inflow will decrease the water value, as the final reservoir level will increase and the concave nature of the future profit function will lead to a smaller slope of the future profit function and similarly decreasing inflow will increase the water value.

The example cannot only be used to study what factors influence the water values, but also some sensitivity analysis of the water values can be done. The sensitivity of the water value to changes in market and hydrological conditions will depend on the curvatures of the immediate profit and future profit functions. The curvature of the immediate profit function will depend on the short term volatility of prices. Therefore, volatility of the water values for smaller reservoirs should be highly influenced by short-term price volatility. For larger reservoirs, the volatility of the water value will depend more on the curvature of the future profit function. The slope of the future profit function changes more quickly when the reservoir is almost empty which will make the water value sensitive to changes in the immediate profit function. The dynamics of the water value should therefore be different during different times of the year, with the potential for large changes being highest during autumn and spring, when the reservoirs are almost full or almost empty.

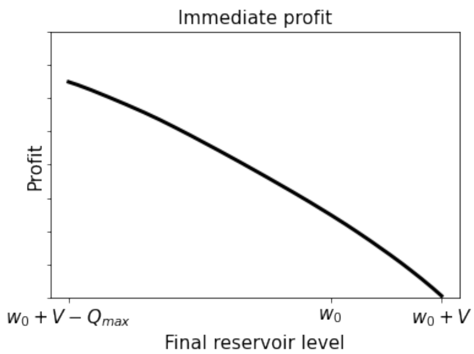

(a) An illustration of the immediate profit as a function of final reservoir level.

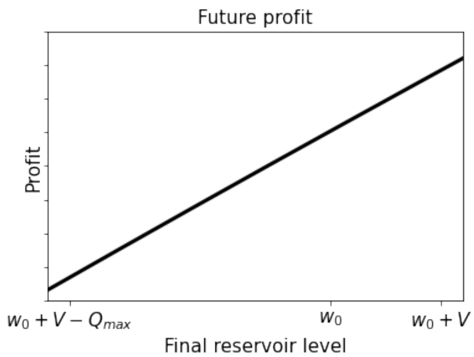

(b) An illustration of the future profit as a function of final reservoir level.

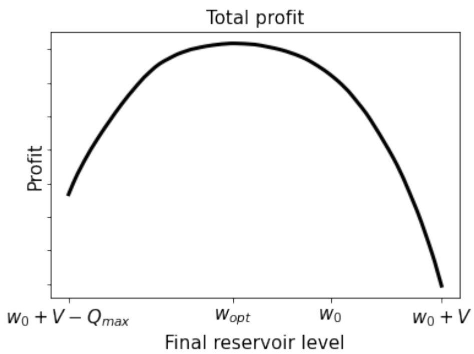

(c) The total profit as the sum of the immediate profit and future profit.

Figure 17: Visualisation of immediate and future profit functions. The optimal reservoir level at the end of the short-term scheduling horizon is the maximum of the total profit function, where the derivatives of the immediate and future profit function are equal with opposite sign.

5.4 A model to forecast the hydro supply curve

Based on the fundamental analysis in the previous section, a model to forecast the hydro supply curve will be presented. The model will be described in two steps, where we first present a simple model which is then refined. We have established that hydro producers aim to optimize their production against the day-ahead price, which will be the starting point of our approach to model the hydro supply curve. In the model, each price area will be considered as one aggregated hydropower plant with one aggregated hydro reservoir. The objective is to optimally allocate the production from this aggregated hydro plant over the next week. The inputs to the model will be a prior price forecast, minimum and maximum hourly production for the aggregated plant and a target production level, describing the average hourly production during the week. The prior price forecast is generated by a fundamental market equilibrium model, which will be explained in section 5.5.

The hydro supply curve will be estimated using a linear programming approach. The model resembles a simplified version of the short-term scheduling problem faced by a hydro producer, which was discussed in section 4.3. We first describe a simplified version of the model to illustrate how the hydro production is optimally allocated. Two key objects in the optimization problem are the price-duration and load-duration curves. In the price-duration curve, the prices are ordered from the most expensive to the cheapest, and correspondingly in the load-duration curve, hourly hydro production is ordered from highest to lowest. In the optimization model, a price-duration curve is constructed from the price forecast, and the objective is to find the optimal load-duration curve for the coming week. The resulting optimization problem is described in equations (8)-(11). The decision variables in the optimization problem are the points on the load-duration curve q_i, where q_1 represents the highest production during the week and q_{168} represents the lowest production during the week. The forecasted price-duration curve is denoted by p_i, where p_1 is the most expensive hour during the week and p_{168} is the cheapest hour during the week. It may be emphasized that the index i is not representing the hour within a week indexed by time but rather the values ordered from largest to smallest. The objective function (8) represents the total profit during the week. Constraint (9) ensures that the average production during the week is equal to the target production q_{target}, while constraints (10) and (11) restrict the production in each hour to be within the specified minimum and maximum production.

$$\max \quad \sum_{i=1}^{168} p_i q_i \tag{8}$$

$$\text{s.t.} \quad \frac{1}{168}\sum_{i=1}^{168} q_i = q_{target} \tag{9}$$

$$q_i \geq q_{min}, \qquad \forall i = 1, \ldots, 168 \tag{10}$$

$$q_i \leq q_{max}, \qquad \forall i = 1, \ldots, 168 \tag{11}$$

We illustrate the properties of the simple model in Figure 18. The example given is based on price area NO5 during the week between 4.3.2024 and 10.3.2024. The minimum and maximum production in the example were set to 1500 MWh and 6000 MWh, respectively, and the target production was set to 4250 MWh. The price forecast used is not necessarily a perfect price forecast but rather used for illustrative purposes.

From the example, we see that the simple optimization model leads to a trivial solution. Since the price forecast is assumed to be perfect, hydro production will be run at maximum capacity during the most expensive hours and run at minimum capacity in the cheapest hours. The number of hours with maximum production can be determined by dividing the target production with the maximum production level, which in the example leads to 119 hours of maximum production. The water value can be found as the value of the price-duration curve at the number of production hours. Production is run at maximum level when the price is above the water value. When the price is below the water value, production is run at minimum. The hydro supply curve can be constructed from the solution of the optimization problem by combining the price duration forecast with the optimized load duration curve. This leads to 168 price-volume pairs forming the supply curve. In this simple model, the hydro supply curve would then be a step function at the water value.

It is clear that the simple model does not explain the shape of the hydro supply curve very well. In reality, hydro producers cannot perfectly predict prices and the production systems are not flexible enough to perfectly adjust production according to day-ahead price. To better take the flexibility of the production system into consideration, we propose a revised model, where we impose a restriction on the load-duration curve in the optimization problem. The load-duration curve for a given price area should reflect the flexibility of hydro production in the price area. If the production is more flexible, the production can better be adjusted according to price and hydro production

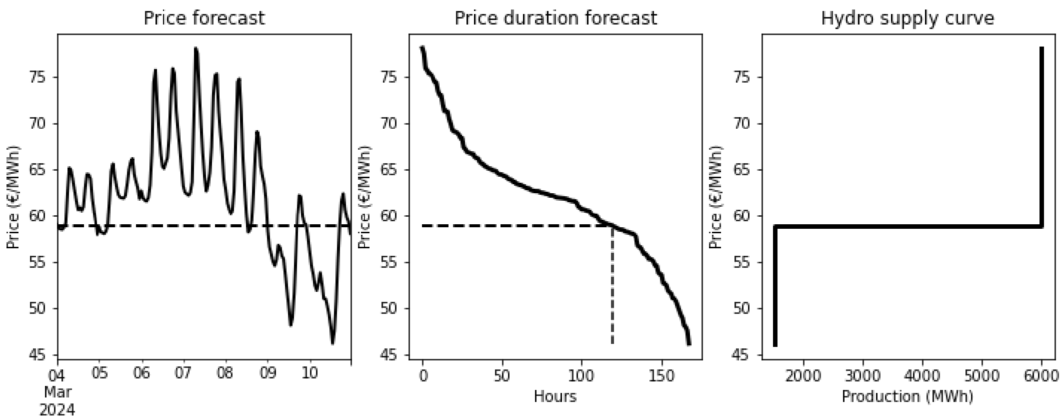

Figure 18: Example price forecast (left), corresponding price-duration forecast (middle) and generated hydro supply curve (right). The water value is represented by the dashed line and determined from the price-duration forecast as the value of the price-duration curve at the given number of production hours. Maximum hydro production is run when the price is above the water value.

can vary more significantly within a week. To restrict the model, a constraint was introduced where the load-duration curve must be a linear combination of historical load-duration curves. Constructing the load-duration curve in this way makes sure that the distribution of hydro production during the week will be reasonable and production is not run only at minimum or maximum production, as observed in the simple model. To account for seasonality effects, the historical load-duration curves are taken from the same week in previous years. It is hard to give an exact theoretical justification for this constraint, but it was observed to work well in practice. In short, the new model can be summarized as follows: Given a deterministic price forecast and a target hydro production level, find the optimal load-duration curve satisfying the constraints and giving a total production equal to the target production.

The mathematical formulation of the revised model is otherwise identical to the simple model (8)-(11), except for the addition of constraints (12)-(14). The variable l_{ij} represents the element i on the load-duration curve in historical year j, with J being the set of historical years. The coefficients a_j represent the weights given to the load-duration curve from historical year j and the coefficient b shifts the overall production level in the load-duration curve up or down. The constant b is needed to make the problem feasible. Otherwise, there would be a risk that a very large or very small target production would make the problem infeasible, even if it is larger than the minimum production and smaller than the maximum production. The coefficients a_j must be positive to ensure that the linear combination of the load-duration curves is always increasing. Furthermore, the coefficients a_j were restricted to be smaller than 1 to reduce the flexibility of the model. Larger coefficients would allow for steeper load-duration curves, which was not desired.

$$q_i = \sum_{j \in J} a_j l_{ij} + b, \qquad \forall i = 1, \dots, 168 \tag{12}$$

$$a_j \geq 0, \qquad \forall j \in J \tag{13}$$

$$a_j \leq 1, \qquad \forall j \in J \tag{14}$$

We illustrate the new model using the same example as for the simple model. The example uses the same inputs and the only difference is the addition of constraints (12)-(14). In the example, the historical load-duration curves from years 2015 to 2023 were used to predict the supply curve in 2024. The historical load-duration curves are presented in Figure 19 along with the optimal load-duration curve generated by the model. It can be observed that the shape of the optimal load-duration curve follows the shape of the historical load-duration curve. When testing the model, it was observed that the linear combination usually results in most of the coefficients being zero, with usually only one or two years having non-zero coefficients. The model can therefore be interpreted as choosing a load-duration curve from the past to match with the price-duration forecast.

The hydro supply curve in the new model is again constructed from the solution of the optimization problem by combining the price-duration forecast with the optimized load-duration curve. The resulting hydro supply curve from the new model is presented in Figure 19. Comparing the generated hydro supply curve from the new model with the hydro supply curve from the simple model, we see that the new hydro supply curve looks much more realistic. The supply curve is no longer a step function but an arbitrary increasing function, which resembles the shape of the previously seen empirical supply curves.

An attractive feature of this model is that it captures the dynamics of the hydro supply curve well. There are two main inputs, the price forecast and target production level, which strongly influence the generated supply curve. If the price forecast contains many expensive hours, it will result in an upwards shift of the supply curve. Also the volatility of the prices in the forecast will be reflected in the generated supply curve, with more volatile prices giving a steeper supply curve. Even if the price forecast is unchanged, the supply curve can change based on the given target production. The theory is that water values will be lowered if an area must produce more during the week and increased if production can be lower. Theoretically, the dynamics of the modelled hydro supply curve should reflect the real dynamics, where water values increase if the price forecast increases and decrease when the price forecast decreases. The model can also account for changes in the hydrological situation, where lower inflow should mean less hydro production and higher inflow should mean higher hydro production.

To illustrate the effect of changing the target production, we construct the supply curve from the example using two alternative values for the target production. In the high production scenario, the target production is increased by 500 MWh to 4750 MWh and in the low production scenario the production is decreased by 500 MWh to 3750 MWh. The alternative supply curves are presented in Figure 20, along with the supply curve from Figure 19. An increase in the target production shifts the whole curve to the right and also lowers the middle part of the curve. The production level in the lower priced hours is increased, representing the pressure on producers with limited storage to produce at lower prices to avoid spilling water. Similarly, the whole curve is shifted to the left and the middle part is slightly increased if the target production is decreased. The production during higher priced hours is decreased, representing the increased flexibility as producers can receive better prices when less water needs to be used. The shifts in the supply curve are somewhat similar to those observed in the empirical supply curves, where the supply curve can simultaneously shift up and

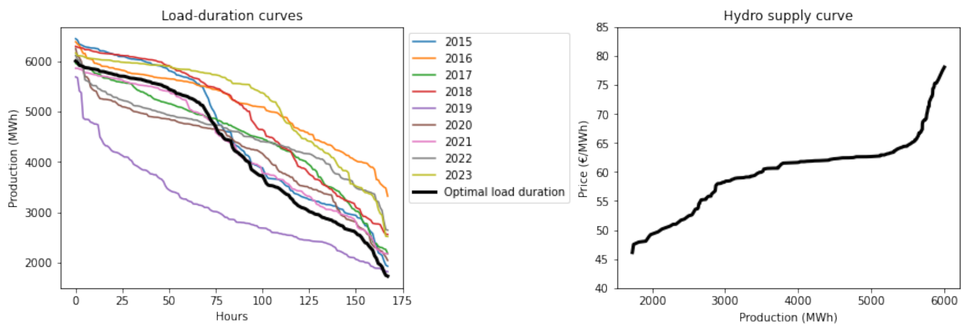

Figure 19: The historical load-duration curves (left) and the optimized hydro supply curve (right) in the refined model.

down in price level as well as production level.

The proposed model has some disadvantages. A fundamental problem with forecasting hydro supply curves is that the supply of hydropower depends on the price expectations. On the other hand, the price expectation depends on the current valuation of hydropower. Since the objective is to use the generated hydro supply curve for price forecasting, the model can be seen as a way to update a price forecast where potential changes in water values are included. Using a price forecast to predict the hydro supply curve can be motivated by the complexity of calculating water values. In the current electricity market, there are numerous factors other than weather which influence the hydro supply curves. Such factors are SRMC of thermal production, nuclear outages, transmission outages as well as prices in neighboring areas. It is very hard to model the individual influence of all these factors but they are all reflected in the price forecast.

Another disadvantage of the model is that the hydrological conditions in the historical years for the load-duration curves are not considered. In dry years, hydro production is more flexible as the utilization is lower and production can better be optimized against the day-ahead price. In the proposed model, it is possible that the optimal load-duration closely follows the load-duration from a dry year even if the current hydrological situation is not dry, leading to a supply curve that is not ideal. A typical error in this case is that the generated supply curve becomes too flat.

The generated supply curve is heavily influenced by the price-duration forecast and therefore the price-duration forecast should be as realistic as possible. However, constructing a realistic price duration forecast is no simple task. In particular, using only one deterministic price forecast does not necessarily describe the uncertainty in the price in the next week that well. The high volatility of electricity prices means that there should always be a possibility of either very low or very high prices, which can be difficult to replicate in the price forecasting process. A possible improvement which

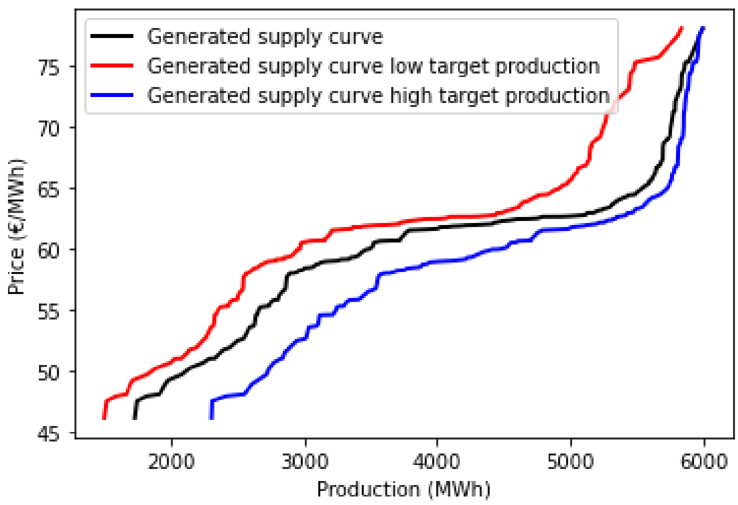

Figure 20: Hydro supply curves obtained by decreasing (red) or increasing (blue) 500 MWh from the original target production.

could be made is to use multiple price forecasts and construct a price-duration curve from each forecast. The expected price duration curve would then be obtained by averaging the price duration curves from each scenario. The multiple price scenarios should increase the probability that very high or very low scenarios are represented, which could result in a greater slope of the expected price-duration curve.

5.5 Model inputs

The optimization model for generating the hydro supply curve was described in the previous section but we did not explain how the inputs are generated. The most important input to the model presented in section 5.4 is the prior price forecast, which is generated based on fundamental modelling of the market. The fundamental model uses forecasts of important drivers of the electricity prices and gives a price forecast by replicating the EUPHEMIA algorithm. The modelling setup is summarized in Figure 21. The price forecast is converted to price-duration curves for each week in the forecasting horizon. The price-duration curves, historical load-duration curves, target hydro production and maximum and minimum hydro production levels are then used to generate a prediction for the hydro supply curve for the next week.

To generate the prior price forecast an initial hydro supply curve must be generated. The initial hydro supply curve can be generated based on the empirical hydro supply curve in the most recent days. It is challenging to come up with a method to fit an increasing function to the observed points. As there is only one observation for each hour, it must be assumed that the supply curve is constant over some period of time. Using a longer time span increases the likelihood of observing hours with very low

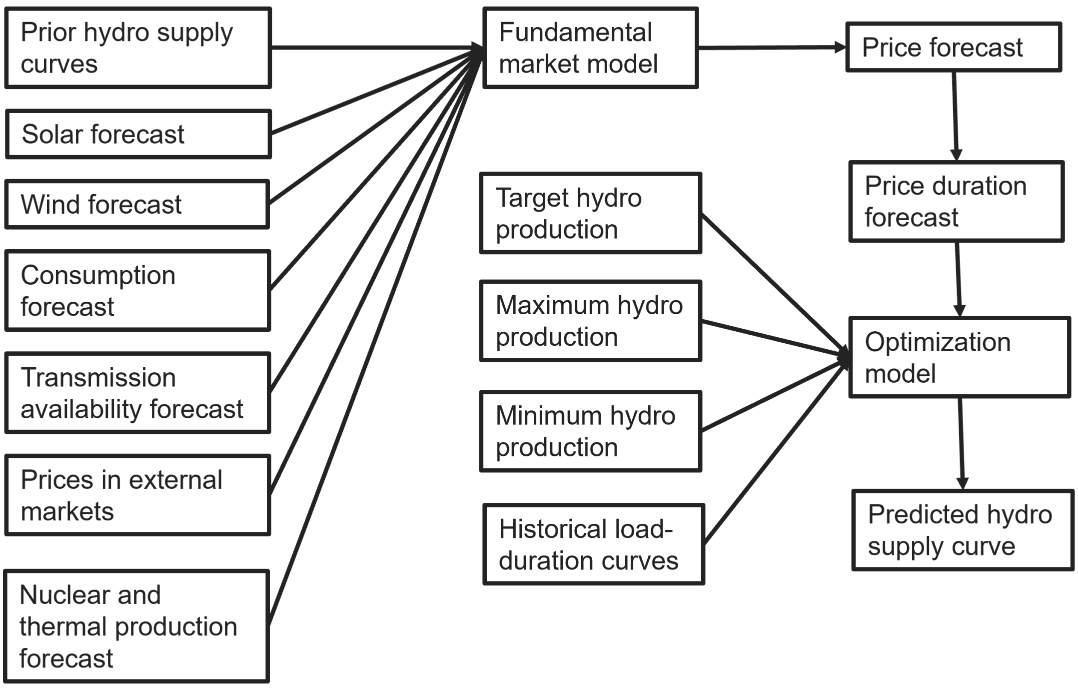

Figure 21: Flow chart of the hydro supply curve prediction model.

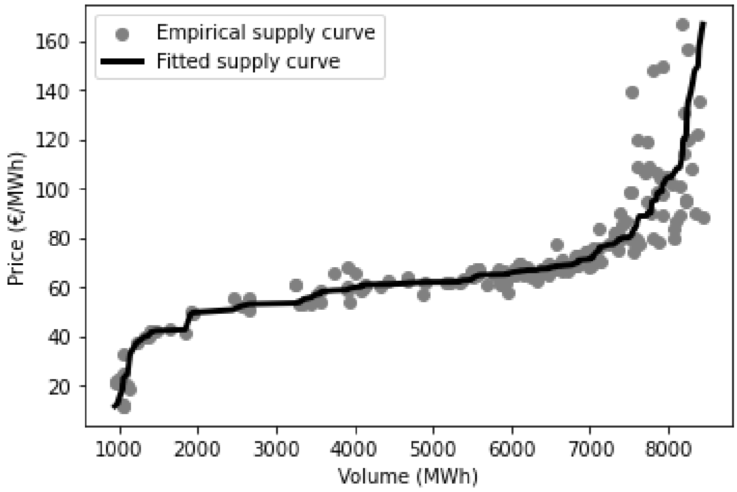

Figure 22: Example of a fitted supply curve using the weekly price-duration and load-duration curves.

or very high hydro production. Thus, a longer time span will give more information about the whole curve. On the other hand, information about changes in the hydro supply curve over this time span is lost.

The initial hydro supply curve could be fitted using a parametric function, e.g. a polynomial or linear function. This approach could suffer from outliers in the data and therefore no parametric model was deemed sufficient to accurately model the shape of the hydro supply curve. Instead, a non-parametric approach was chosen, allowing for greater flexibility in the modelled supply curve. The non-parametric model uses data from one week to generate one supply curve and it is hence assumed that the supply curve stays constant for one week. The weekly supply curve is constructed as follows. We take all price-volume observations from the last seven days. The prices and production volumes are then sorted such that the highest price correspond to the highest volume, the second highest price to the second highest volume and so on. This method was chosen as it is robust to outlying observations. An example of a fitted hydro supply curve can be seen in Figure 22. We see that the price-volume pairs in the empirical supply curve line up well for low and medium production volumes but the maximum production volumes appears to have changed during the week. The fitted supply curve gives a reasonable approximation of the true unobserved hydro supply curve, passing through the middle of the observed points. A problem with the chosen approach is that a given week may not contain any hours with very low or very high production, even if such extreme production volumes are possible. In these cases a part of the supply curve remains unobserved. This was implemented by setting the price at the minimum production level to 3 €/MWh lower than the lowest observed price and the price at the maximum production level to 3€/MWh above the highest observed price. The supply curve was then linearly interpolated up to these new artificially introduced maximum and minimum points.

One of the most important inputs to the model is the target production. The target production should reflect the average hourly hydro production which should be produced in a week. However, the target production should not necessarily be interpreted as the expected hydro production for an area. There may be situations where an area would want to produce more due to high inflow but production cannot be increased due to low demand. In such situations water values may be lowered even though production is not increased. Using different target production levels could be a good idea to generate different scenarios for the hydro supply curve. Adding more production would correspond to a scenario with higher inflow, where more hydro production is needed in order not to fill up the reservoirs, while in a drier scenario hydro should produce less to not empty the reservoirs.

The other inputs to the model are simpler to obtain compared to the prior supply curve and target production. The forecasts for solar and wind production are based on weather forecasts and the assumption that these production technologies are bid to the day-ahead market at zero or negative prices. The consumption forecast is also affected by weather in the form of temperature and also the time of day and holidays. Transmission availability can be estimated based on historical data as well as market information about upcoming outages. The prices in external markets refer to the prices in countries neighboring the Nordic countries. These markets also influence the prices in Nordic countries through trade. The nuclear production forecast is based on the maintenance schedules of nuclear power plants and the assumption that nuclear power is run at full capacity as long as the price is positive. Thermal power production is forecasted based on the SRMC.

The model requires a maximum and minimum limit for the production in each price area. These limits can be set based on historical data. The minimum level of hydro production is affected by the amount of run-of-river production as well as volumes of down regulation sold in reserve markets. The maximum production level is mostly affected by maintenance which is usually performed in summer. During summer, the prices are usually lower and producers can therefore minimize the revenue loss due to maintenance. The maximum available capacity can be significantly reduced during these times of the year. The total hydropower capacity available can be estimated from Urgent Market Messages (UMM) publicly available from Nord Pool. Producers are required to inform other market participants about the availability of their plants. Planned maintenance schedules are published in advance which is useful when estimating maximum production.

6 Results

The model proposed in section 5.4 was tested over a six week test period in the summer of 2024. The test period started on July 8th and ended on August 18th. The weather during this period was quite wet, which resulted in decreasing water values throughout the testing period. Hence, there were many weeks where the hydro supply curve shifted downwards. The model was tested over this period to see how well it was able to replicate this decrease in water values. The tested areas included all areas in the Nordics with significant hydro production, which were FI, SE1, SE2, SE3, NO1, NO2, NO3, NO4 and NO5. To construct the price forecast, the market equilibrium model was given the realized values for transmission capacity, demand, nuclear power production, wind power production and solar power production. This will give an optimistic view on the performance of the model as these values are not known in reality. This choice was made to remove effects from errors in these inputs and only focus on errors arising from the model itself. Furthermore, this modelling setup can be thought of as a benchmark on how well the model can generate scenarios for the hydro supply curve given a weather scenario.

The target production in the model was set to the realized production. It is impossible to know the realized hydro production ahead of time and the choice of the realized production may also lead to too optimistic forecasts. It is however unclear how to choose the target production in an optimal way and we leave this question open for future research. At least the realized production should give reasonable values and act as a benchmark. The minimum and maximum production levels were set manually based on historical values. Sometimes the minimum and maximum values were also modified based on the target production. This was especially the case in run-of-river areas, where the minimum and maximum production was increased when the general production level was high.

Two testing setups were studied, with different prior supply curves. First, we study how well the model can predict the supply curve one week ahead. The prior supply curves are constructed using the empirical supply curves from the week prior to the test week. Thus, we obtain six independent test weeks. This test is done to see how well the supply curve forecasting model can predict the movement in the supply curve in the very short-term. With this shorter time horizon it is more reasonable to assume e.g. wind production and consumption to be known.

In the second testing setup, we try to forecast the hydro supply curve multiple weeks ahead. In this testing setup, the prior supply curves are set based on the empirical supply curves from before the start of the six week test period, July 1st to July 7th. This setup was used to study how well the model can predict the supply curve multiple weeks into the future, using only information about the current hydro supply curves. The other inputs to the market equilibrium model still contain the realized values. In reality, the inputs to the model cannot be predicted for such a long time horizon. However, it is interesting to see how well the model can replicate the hydro supply curve given the realized inputs.

In both of the testing setups the prior supply curves were smoothed using a Savitzky-Golay filter [52]. The Savitzky-Golay filter smooths a function by fitting polynomial functions to a subset of points. The parameters used for the filter was a window length of 75 and polynomials of order 1. The smoothing was applied to remove sudden jumps and increase the slope of the middle of the supply curve. When the filter was not applied, it was observed that the hourly prices in the prior forecast often converged to a given price level, which is undesirable as the predicted supply curve would become too flat. The effect of the filter was however not that significant and the model could likely be used without it.

Measuring the quality of the generated supply curves is difficult as the true hydro supply curve, which is to be estimated, is unknown. Therefore we will mostly focus on visual inspection of the generated supply curves and their qualitative properties rather than doing a more quantitative analysis. One way to qualitatively study the results is to compare the generated supply curve with the realized empirical supply curves for some of the test weeks.

We attempt to quantify the accuracy by studying the error in price at the realized production volume. This is an interesting measure of the performance, as hydropower is usually price setting, and comparing the prices in the supply curves could be a good approximation of the error of a price forecast generated based on the predicted supply curves. The error metric chosen was Root Mean Squared Error (RMSE). Another way to think of this error measure is to take every point on the empirical supply curve and to compute the mean vertical distance to the predicted supply curve. We compare the RMSE of the modelled supply curves with the RMSE of the prior supply curves. This comparison gives an indication whether the generated supply curve would be better than the prior supply curve. Using RMSE can be heavily influenced by outliers and therefore the results should be used somewhat cautiously and the supply curves should be visually inspected before drawing any stronger conclusions.

6.1 Week ahead forecasts

We first study how well the model predicts the hydro supply curve for the upcoming week. The objective is that the predicted hydro supply curve would be better than the prior supply curve based on observations from the previous week. The prior supply curve already contains a lot of information about the current hydrological situation and is therefore a very good benchmark which may be hard to beat. The RMSE values of the prior and predicted hydro supply curves are compared in Table 1. The differences in RMSE values were generally small. The total number of tests where the predicted supply curve outperformed the prior supply curve were 19/36, indicating that the predicted supply curves did not offer that much improvement. When the predicted supply curve is worse than the prior curve, it is often not by a large margin but when the predicted curve performs better it is sometimes by a significant margin. Comparing the performance on individual weeks, the model performed worst on the week starting 15th of July while the performance on other weeks were quite similar.

We look into the potential causes for the errors by visually examining the prior and predicted supply curves for two of the test weeks, the week starting 15th of July (Fig. 23) and the week starting 5th of August (Fig. 24). These weeks were chosen to compare a week with relatively good performance and a week with worse performance.

		8.7.	15.7.	22.7.	29.7.	5.8.	12.8.
FI	Predicted	10.61	8.79	9.44	8.44	6.63	20.97
	Prior	13.15	8.36	17.97	8.57	8.87	22.00
NO1	Predicted	8.91	9.63	8.68	6.37	11.71	9.50
	Prior	13.13	14.42	11.64	15.45	11.45	5.71
NO2	Predicted	5.19	4.11	9.60	20.31	10.49	10.92
	Prior	5.34	2.92	9.07	20.32	12.87	10.74
NO3	Predicted	5.08	4.74	2.60	1.91	8.33	7.70
	Prior	5.37	4.01	2.34	2.06	8.66	7.48
NO4	Predicted	4.76	4.74	2.76	1.76	7.84	7.28
	Prior	4.64	4.97	3.28	1.24	9.31	7.30
NO5	Predicted	2.28	8.50	3.53	2.67	8.33	8.74
	Prior	2.78	8.27	3.94	2.94	9.97	10.09
SE1	Predicted	7.47	5.75	6.83	4.28	4.79	6.44
	Prior	12.50	8.51	6.76	7.55	5.41	7.00
SE2	Predicted	4.71	7.45	6.08	4.39	8.83	3.51
	Prior	4.28	6.62	7.23	4.15	8.69	5.66
SE3	Predicted	11.94	10.21	7.61	10.19	10.64	6.54
	Prior	18.78	15.93	10.32	12.04	15.57	9.14

Table 1: RMSE values for price for the week ahead forecasts. Cells marked with green represent the weeks and areas where the predicted supply curve outperformed the prior supply curve.

The relatively smooth shape of the prior and generated supply curve is due to the application of the Savitzky-Golay filter. Also the effect of extending the supply curves can be seen, where for example a large part of the supply curve in NO4 is defined based on the linear extension at higher production volumes.

In the week starting 15th of July the model had the worst performance judging by the RMSE values. Looking at the generated supply curves in Figure 23, no area stands out with a particularly poor performance. The model predicted a too high supply curve in areas SE2, NO3 and NO4 which explains the poorer RMSE values. This could be due to the unusually low hydro production levels seen in this area during summer of 2024 as a result of low reservoir filling levels. Therefore the hydro production profile may not be well described by the historical load-duration curves. Another explanation could be the use of a deterministic price forecast. If the deterministic price forecast does not contain any hours with very low or zero prices, the model will keep the leftmost part of the supply curve at a higher level, resulting in a higher supply curve. In contrast, the model performed well in price areas with more run-of-river production and lower storage capacity (NO1, SE3 and FI). This is logical as the realized hydro production was used as the target hydro production level. The model is therefore able to correctly predict shifts in production volume, corresponding to a left or right shift in the supply curve. Hence, if one were able to correctly predict the weekly hydro production levels in these areas, the model would be able to improve the supply curve

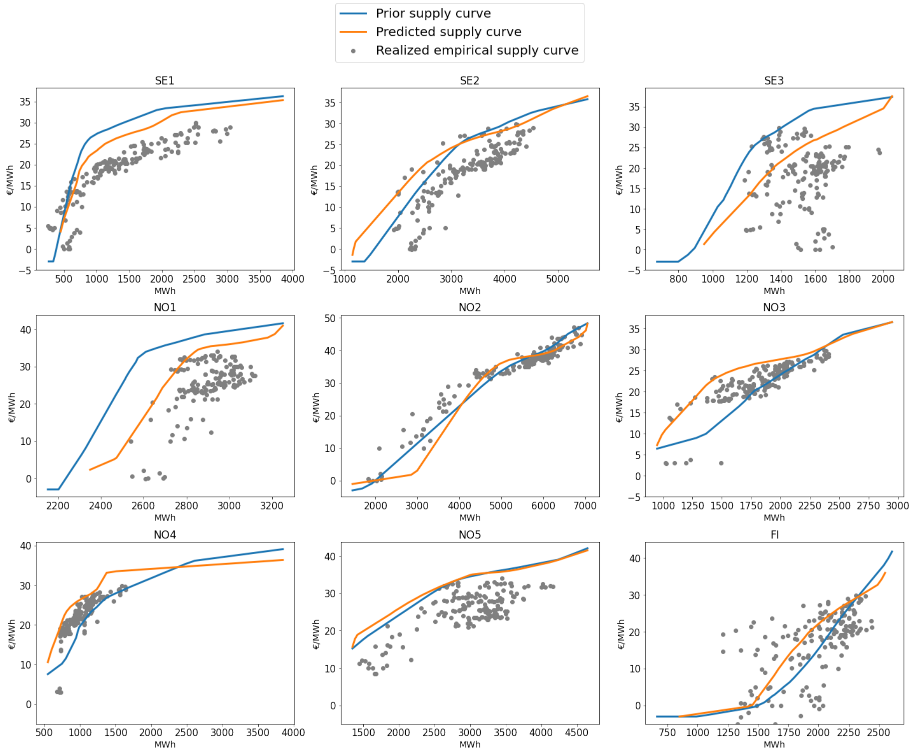

Figure 23: Week-ahead supply curve forecasts and prior supply curves for the week starting 15th of July.

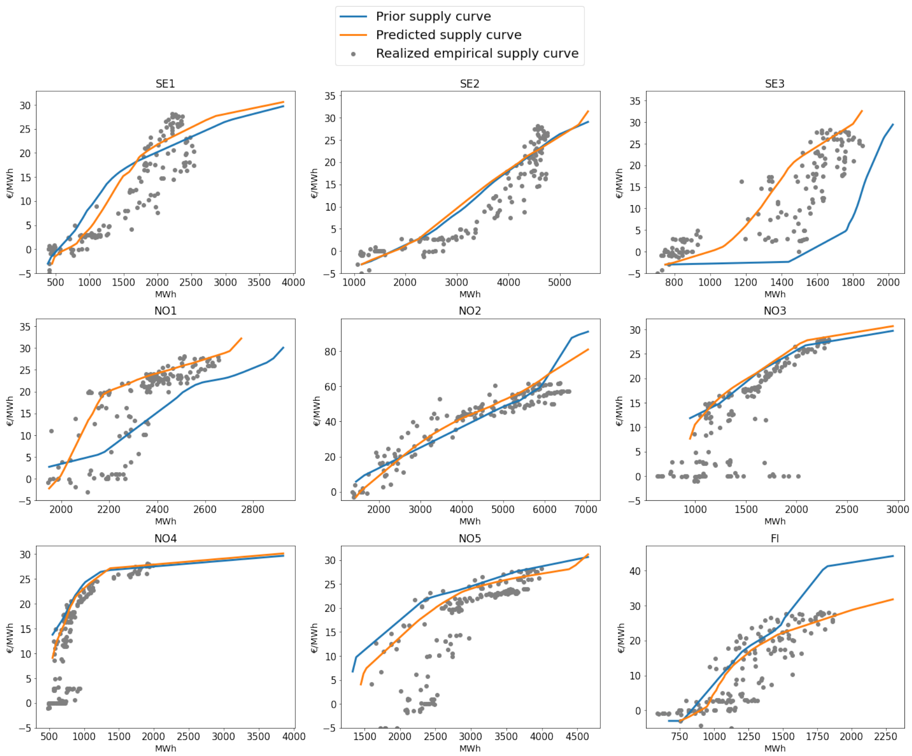

Figure 24: Week-ahead supply curve forecasts and prior supply curves for the week starting 5th of August.

prediction compared to using a constant supply curve. In areas with more run-of-river production, it may not be valid to assume that the supply curve stays constant over one week as changes in inflow can rapidly shift the curve to the left or to the right.

Looking at the generated supply curves in the week starting August 5th, we again see that the performance was good in run-of-river dominated areas. Areas SE3 and NO1 saw shifts in production volumes which were correctly predicted. The model also improved the supply curve in Finland, where the price at higher production levels where correctly lowered. This improvement in Finland was however most likely due to outliers in the prior supply curve, where some rare high priced hours had an effect on the prior supply curve. The model did, however, not predict the changes in the supply curves in areas SE1 and SE2. In both of the northern Swedish areas, the predicted supply curve followed the prior supply curve closely, instead of predicting lower prices or higher production. Compared to other weeks, this week stands out as performance in NO3 and NO4 was quite good. However, neither the prior or the predicted supply curve managed to get down to the very low prices which were observed during this week.

6.2 Forecasting multiple weeks

Next, we study how the model performed when the prior supply curves were set based on the observations from the first week in July. Using these prior supply curves, the price forecast becomes less accurate. In particular, it is possible that the general price level is off in the forecast, as the price level is often set by hydro production. It can therefore be expected that the errors will now be larger compared to forecasting only one week ahead. The performance of the prior supply curves will also naturally decrease and therefore the comparison between prior and predicted curves is still interesting. If the predicted supply curves can offer some improvement compared to the prior supply curves, it indicates that the model can also be used over longer time horizons to predict in which direction the hydro supply curve will change.

We again compare the weekly RMSE values for the price in the predicted and prior supply curves (Table 2). The errors in the first test week are exactly the same as before, since the same prior supply curves were used this week. The performance difference between the prior and predicted supply curves are again mixed, with the predicted supply curve again having better performance in 19/36 tests. The weeks and areas where the predicted supply curves outperformed the prior supply curves however changed. The errors increase when moving further into the future both for the predicted and prior supply curves but no clear trend can be seen when it comes to the performance of the predicted supply curve relative to the prior supply curve.

We study the generated supply curves in the same weeks as in the previous sections in Figures 25 and 26. The observations are very similar to the week ahead forecast. The predicted supply curves rarely differ significantly from the prior supply curves and the shape of the predicted supply curves are often reasonable.

During the week starting 15th of July the predicted supply curves again performed worse compared to other weeks, as was the case for the week-ahead forecasts. The generated supply curves look very similar compared to the ones generated in the week-ahead forecast. This can be attributed to the fact that the prior supply curves are only from two weeks before the test week and there was not much time for the supply curves to change. The model again performed well in areas dominated by run-of-river production, but it did not predict the decrease in water values in areas SE2 and NO5.

In the week starting 5th of August, we can see that in almost all areas, the predicted supply curve moves in the right direction compared to the prior supply curves. The only exception is NO3, where the predicted supply curve was shifted upwards while the prior supply curve would have been an almost perfect prediction. In NO4 the predicted supply curve is correctly shifted upwards but the magnitude of the shift is too large. The same holds in the Finnish price area. In all other areas, the predicted supply curve moves nicely in the correct direction of the realization without making too large changes. Thus, the model can also be used to model the supply curve further into the future, but the magnitude of changes are likely underestimated.

6.3 Interpreting the results

When interpreting the results, it is important to consider that realized values were used for many inputs to the price forecast. It should, however, be possible to obtain satisfactory results without knowing these inputs, because only the price duration curve really matters when generating the supply curves. The price forecast can be

		8.7.	15.7.	22.7.	29.7.	5.8.	12.8.
FI	Predicted	10.61	9.57	9.58	8.60	10.88	19.63
	Prior	13.15	12.20	13.97	15.87	10.12	26.34
NO1	Predicted	8.91	10.26	13.04	12.45	15.99	20.02
	Prior	13.13	17.68	23.96	18.47	19.95	20.02
NO2	Predicted	5.19	5.38	7.44	20.52	12.09	14.83
	Prior	5.34	5.77	7.16	21.26	16.14	16.92
NO3	Predicted	5.08	4.95	6.68	3.14	10.23	13.35
	Prior	5.37	3.80	4.18	2.05	7.77	10.62
NO4	Predicted	4.76	4.86	7.01	3.19	7.78	12.72
	Prior	4.64	7.88	9.55	9.72	6.65	4.98
NO5	Predicted	2.28	9.16	11.41	10.93	14.86	19.60
	Prior	2.78	8.11	10.37	8.88	17.16	20.86
SE1	Predicted	7.47	6.16	10.20	8.50	9.50	12.84
	Prior	12.50	8.34	9.24	10.74	13.69	14.14
SE2	Predicted	4.71	7.42	9.69	8.47	12.02	12.54
	Prior	4.28	6.31	11.21	11.63	15.19	18.05
SE3	Predicted	11.94	10.48	10.58	15.82	14.79	14.81
	Prior	18.78	24.91	25.64	32.22	23.53	24.19

Table 2: RMSE values for price when using original hydro supply curves. Cells marked with green represent the weeks and areas where the predicted supply curve outperformed the prior supply curve.

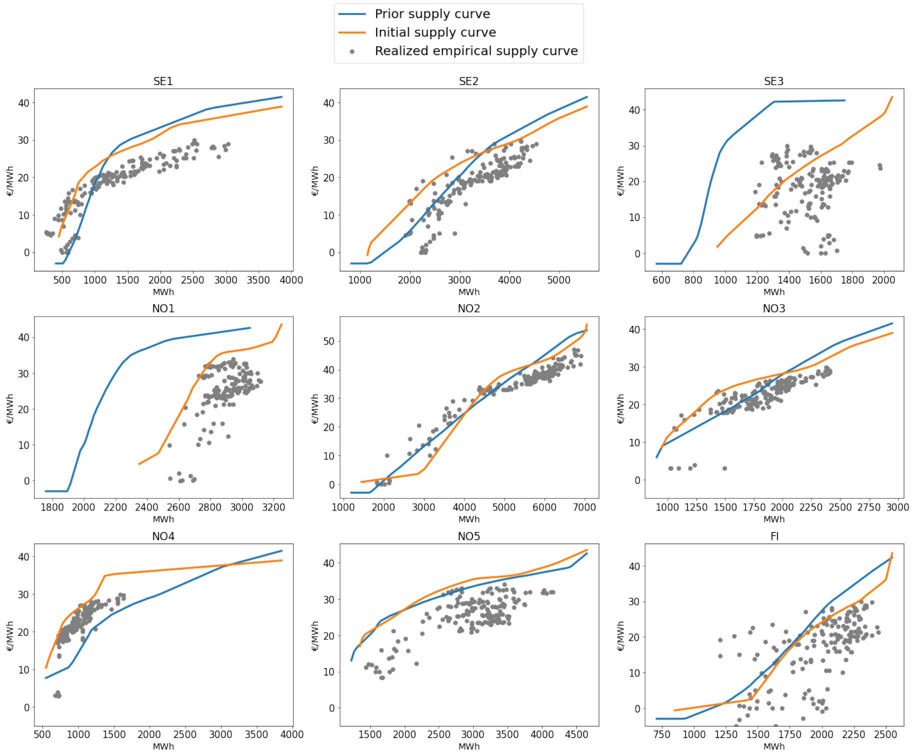

Figure 25: Supply curve forecast for the week starting July 15th using prior supply curves from July 1st to July 8th.

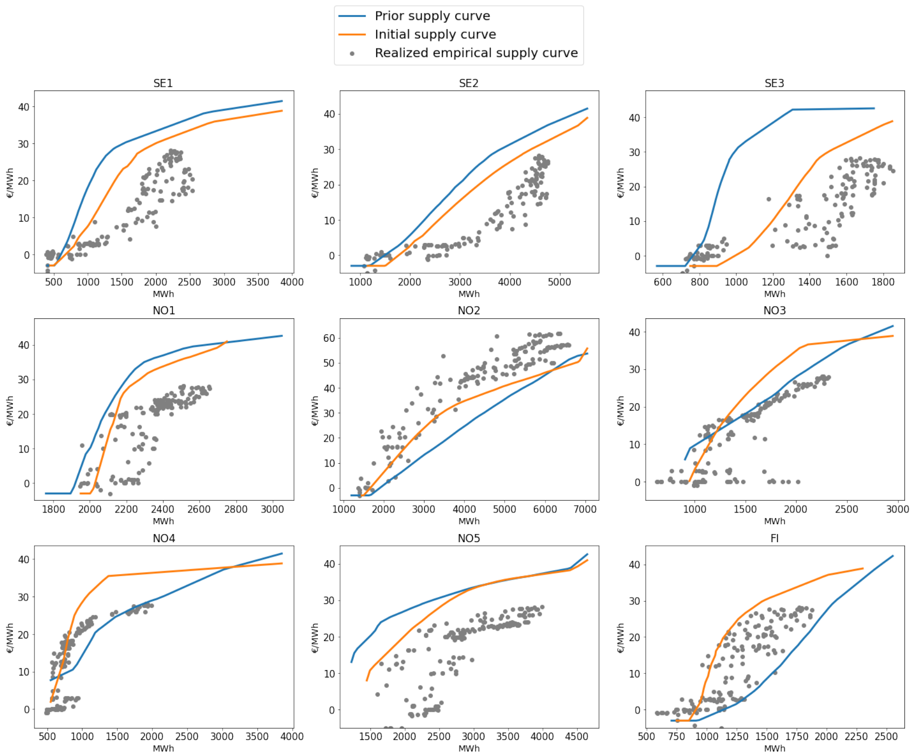

Figure 26: Supply curve forecast for the week starting August 5th using prior supply curves from July 1st to July 8th

quite inaccurate as long as it contains a realistic number of cheap and expensive hours which would result in a realistic price duration curve. This can be achieved as long as the variation in inputs resembles reality, while the exact values in the forecast are not important.

The most likely reason for poor performance in some weeks is the choice of target production. If inflow is high, the reservoir at the end of the week will increase even if hydro production is running at full capacity all the time. This increased reservoir level should decrease water values, but in the model the total hydro production during the week is the only factor representing the hydrological conditions. The model should be adjusted so that the target production is not set according to the expected realized production, but rather the expected production adjusted by the deviation from normal inflow. If inflow is higher than normal, the target production should be increased to reflect this increase in inflow, even if the actual hydro production does not increase as a consequence of the increased inflow.

The model performance generally varied across different areas. The areas NO3 and NO4 stand out from the rest, as the model consistently performed worse in these areas. Furthermore, In price area NO2, the model did not predict the shape of the supply curve very accurately. One explanation may be the large influence of low prices from Central Europe during times of large renewable output. As the transmission capacity between NO2 and Central Europe has increased only in recent years, the historical load-duration curves might not be flexible enough to reduce the production during cheaper priced hours. For the Northern Swedish areas SE1 and SE2, the model performance was mixed with some good predicted supply curves in certain weeks while the results were worse in others. The model had the best performance in areas NO1, SE3 and FI where the storage capacity is relatively small. The differences between different areas can be attributed to differences in storage capacity. It was expected that the model would perform worse on areas with larger storage capacities, since the model performs the optimization on a weekly granularity. In areas with larger storage capacity, the valuation from the long-term model will have a larger influence on the water values and considering only one week at a time will not capture the long-term effects.

7 Conclusions

The water value is one of the most important drivers of electricity prices in the Nordics. Computing the water value is very complex and requires solving a large-scale stochastic dynamic programming problem. Solving this problem to predict the hydro supply curve would suffer from high computational burden, high complexity and lack of available data, highlighting the need for simpler models, which can quickly give results with acceptable accuracy. Forecasting the hydro supply curve has become even more complex in recent years with increased short-term price volatility, arising from renewable energy sources. The increased volatility has changed the dynamics of the hydro supply curve, with hydro supply curves changing quicker, especially in areas with limited storage capacity.

Literature on forecasting the hydro supply curve is very scarce. Most of the literature focus on pricing of hydro resources from a producer's perspective rather than from a price forecaster's perspective. The literature on hydro scheduling is however very useful to understand the water values and can be used as a theoretical background. An explaining factor of the lack of literature could be the lack of available data in the form of supply curves for individual regions as well as hydro production of individual plants. The lack of literature also highlights the difficulty in predicting the hydro supply curve and it may be questionable whether it is possible to create a satisfactory model to forecast the hydro supply curve.

In this thesis, we presented a model to forecast the hydro supply curves in the Nordic electricity market over a short-term time horizon. Even if the proposed model did not manage to explain the movements in the supply curve in all areas, it still provides a useful framework for modelling the hydro supply curve. The model is highly explainable as it is essentially only a function of price forecasts and hydro production. Interpreting the model results should therefore be easy and provide more understanding of the dynamics of the hydro supply curve. The model could be used as a decision support tool, where altering the inputs gives a supply curve corresponding to the inputs. Especially altering the target production can give a good idea of how the hydro supply curve should behave in a wet or dry hydrological scenario. Using the model along with expert opinion can potentially provide insights into what direction the hydro supply curve could move.

Several improvements could enhance the model. The deterministic price forecast does not describe the electricity price in a perfect way as uncertainty is high and sudden price spikes are common. Electricity producers should be pricing in the small possibility of very high prices in the future, which is not necessarily captured in a deterministic price forecast. One way to address this would be to use an ensemble of price forecasts, generated from different weather scenarios. The scenarios could then be averaged to form an expected price-duration curve, which could be used as input to the model. This approach would better account for the uncertainty in future prices but would be more computationally expensive because of the multiple price scenarios which need to be generated. Another possible improvement could be to use different time horizons for different price areas. In price areas with larger storage capacity, the optimization could be carried out over e.g. a month instead of a week.

References

[1] *History | nord pool*, https://www.nordpoolgroup.com/en/About-us/History/, (Accessed on 08/02/2024).

[2] O. Wolfgang, A. Haugstad, B. Mo, A. Gjelsvik, I. Wangensteen, and G. Doorman, "Hydro reservoir handling in norway before and after deregulation", *Energy*, vol. 34, no. 10, pp. 1642–1651, 2009.

[3] T. Bye and E. Hope, "Deregulation of electricity markets: The Norwegian experience", *Economic and Political Weekly*, pp. 5269–5278, 2005.

[4] V. Ahlqvist, P. Holmberg, and T. Tangerås, "A survey comparing centralized and decentralized electricity markets", *Energy Strategy Reviews*, vol. 40, p. 100 812, 2022.

[5] *Download center | nord pool*, https://www.nordpoolgroup.com/en/download-center/, (Accessed on 07/10/2024).

[6] *Nord pool | day-ahead prices*, https://data.nordpoolgroup.com, (Accessed on 08/02/2024).

[7] A. Khodadadi, L. Herre, P. Shinde, R. Eriksson, L. Söder, and M. Amelin, "Nordic balancing markets: Overview of market rules", in *2020 17th International Conference on the European Energy Market (EEM)*, IEEE, 2020, pp. 1–6.

[8] E. Näsäkkälä and J. Keppo, "Hydropower with financial information", *Applied Mathematical Finance*, vol. 15, no. 5-6, pp. 503–529, 2008.

[9] N. Committee *et al.*, "Euphemia public description: Single price coupling algorithm", *URL: https://www. epexspot. com/sites/default/files/2020-02/Euphemia_Public% 20Description_Single% 20Price% 20Coupling% 20Algorithm_190410. pdf*, 2019.

[10] *Intraday trading | nord pool*, https://www.nordpoolgroup.com/en/trading/intraday-trading/, (Accessed on 08/02/2024).

[11] L. Šćekić, S. Mujović, and V. Radulović, "Pumped hydroelectric energy storage as a facilitator of renewable energy in liberalized electricity market", *Energies*, vol. 13, no. 22, p. 6076, 2020.

[12] *Volue insight*, https://app.volueinsight.com, (Accessed on 08/01/2024).

[13] T. Tellefsen, J. Van Putten, and O. Gjerde, "Norwegian hydropower: Connecting to Continental Europe", *IEEE Power and Energy Magazine*, vol. 18, no. 5, pp. 27–35, 2020.

[14] J. Jääskeläinen, K. Huhta, and S. Syri, "The anatomy of unaffordable electricity in Northern Europe in 2021", *Energies*, vol. 15, no. 20, p. 7504, 2022.

[15] J. D. Jenkins, Z. Zhou, R. Ponciroli, *et al.*, "The benefits of nuclear flexibility in power system operations with renewable energy", *Applied Energy*, vol. 222, pp. 872–884, 2018.

[16] C. S. Kaunda, C. Z. Kimambo, and T. K. Nielsen, “Hydropower in the context of sustainable energy supply: A review of technologies and challenges”, *International Scholarly Research Notices*, vol. 2012, no. 1, p. 730 631, 2012.

[17] *Hydropower modelling – new database complementing pecd*, https://eepublicdownloads.entsoe.eu/clean-documents/sdc-documents/MAF/2020/Hydropower_Modelling_New_database_and_methodology_V1_0.pdf, (Accessed on 10/14/2024).

[18] *Nve - norges vassdrags- og energidirektorat*, https://www.nve.no/, (Accessed on 07/10/2024).

[19] *Välkommen till viss*, https://viss.lansstyrelsen.se/, (Accessed on 09/18/2024).

[20] J. Seibert and S. Bergström, “A retrospective on hydrological catchment modelling based on half a century with the HBV model”, *Hydrology and Earth System Sciences*, vol. 26, no. 5, pp. 1371–1388, 2022.

[21] C. Jahns, C. Podewski, and C. Weber, “Supply curves for hydro reservoirs–estimation and usage in large-scale electricity market models”, *Energy Economics*, vol. 87, p. 104 696, 2020.

[22] G. Pritchard, A. B. Philpott, and P. J. Neame, “Hydroelectric reservoir optimization in a pool market”, *Mathematical Programming*, vol. 103, pp. 445–461, 2005.

[23] M. V. Pereira, “Optimal stochastic operations scheduling of large hydroelectric systems”, *International Journal of Electrical Power & Energy Systems*, vol. 11, no. 3, pp. 161–169, 1989.

[24] S. Stage and Y. Larsson, “Incremental cost of water power”, *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems*, vol. 80, no. 3, pp. 361–364, 1961.

[25] L. E. Schäffer, A. Helseth, and M. Korpås, “A stochastic dynamic programming model for hydropower scheduling with state-dependent maximum discharge constraints”, *Renewable Energy*, vol. 194, pp. 571–581, 2022.

[26] A. Helseth and A. C. Geber de Melo, “Scheduling toolchains in hydro-dominated systems: Evolution, current status and future challenges for Norway and Brazil”, *SINTEF Rapport*, 2020.

[27] O. B. Fosso, A. Gjelsvik, A. Haugstad, B. Mo, and I. Wangensteen, “Generation scheduling in a deregulated system. the Norwegian case”, *IEEE Transactions on Power Systems*, vol. 14, no. 1, pp. 75–81, 1999.

[28] O. B. Fosso and M. M. Belsnes, “Short-term hydro scheduling in a liberalized power system”, in *2004 International Conference on Power System Technology, 2004. PowerCon 2004.*, IEEE, vol. 2, 2004, pp. 1321–1326.

[29] A. Gjelsvik, B. Mo, and A. Haugstad, "Long-and medium-term operations planning and stochastic modelling in hydro-dominated power systems based on stochastic dual dynamic programming", *Handbook of Power Systems I*, pp. 33–55, 2010.

[30] P. Lino, L. A. N. Barroso, M. V. Pereira, R. Kelman, and M. H. Fampa, "Bid-based dispatch of hydrothermal systems in competitive markets", *Annals of Operations Research*, vol. 120, no. 1, pp. 81–97, 2003.

[31] B. Mo, O. M. Hansen, and L. E. Schäffer, "Methods of aggregation and disaggregation-project results", *SINTEF Rapport*, 2020.

[32] M. Kelman, M. Pereira, and N. Campodónico, "Long-term hydro scheduling based on stochastic models", *EPSOM*, vol. 98, pp. 23–25, 1998.

[33] K. S. Gjerden, A. Helseth, B. Mo, and G. Warland, "Hydrothermal scheduling in norway using stochastic dual dynamic programming; a large-scale case study", in *2015 IEEE Eindhoven Powertech*, IEEE, 2015, pp. 1–6.

[34] N. Flatabø, A. Haugstad, B. Mo, and O. B. Fosso, "Short-term and medium-term generation scheduling in the Norwegian hydro system under a competitive power market structure", in *EPSOM'98 (International Conference on Electrical Power System Operation and Management), Switzerland*, 1998.

[35] A. Gjelsvik, M. M. Belsnes, and A. Haugstad, "An algorithm for stochastic medium-term hydrothermal scheduling under spot price uncertainty", in *Proceedings of 13th Power Systems Computation Conference*, 1999.

[36] J. Kong, H. I. Skjelbred, and O. B. Fosso, "An overview on formulations and optimization methods for the unit-based short-term hydro scheduling problem", *Electric Power Systems Research*, vol. 178, p. 106 027, 2020.

[37] H. I. Skjelbred, J. Kong, and O. B. Fosso, "Dynamic incorporation of nonlinearity into MILP formulation for short-term hydro scheduling", *International Journal of Electrical Power & Energy Systems*, vol. 116, p. 105 530, 2020.

[38] S.-E. Fleten and T. K. Kristoffersen, "Short-term hydropower production planning by stochastic programming", *Computers & Operations Research*, vol. 35, no. 8, pp. 2656–2671, 2008.

[39] M. Nazari-Heris, B. Mohammadi-Ivatloo, and G. Gharehpetian, "Short-term scheduling of hydro-based power plants considering application of heuristic algorithms: A comprehensive review", *Renewable and Sustainable Energy Reviews*, vol. 74, pp. 116–129, 2017.

[40] M. M. Belsnes, O. Wolfgang, T. Follestad, and E. K. Aasgård, "Applying successive linear programming for stochastic short-term hydropower optimization", *Electric Power Systems Research*, vol. 130, pp. 167–180, 2016.

[41] E. K. Aasgård, S.-E. Fleten, M. Kaut, K. Midthun, and G. A. Perez-Valdes, "Hydropower bidding in a multi-market setting", *Energy Systems*, vol. 10, pp. 543–565, 2019.

[42] T. K. Boomsma, N. Juul, and S.-E. Fleten, “Bidding in sequential electricity markets: The Nordic case”, *European Journal of Operational Research*, vol. 238, no. 3, pp. 797–809, 2014.

[43] S.-E. Fleten and T. K. Kristoffersen, “Stochastic programming for optimizing bidding strategies of a Nordic hydropower producer”, *European Journal of Operational Research*, vol. 181, no. 2, pp. 916–928, 2007.

[44] E. K. Aasgård, G. S. Andersen, S.-E. Fleten, and D. Haugstvedt, “Evaluating a stochastic-programming-based bidding model for a multireservoir system”, *IEEE Transactions on Power Systems*, vol. 29, no. 4, pp. 1748–1757, 2014.

[45] E. K. Aasgård, H. I. Skjelbred, and F. Solbakk, “Comparing bidding methods for hydropower”, *Energy Procedia*, vol. 87, pp. 181–188, 2016.

[46] E. N. Alnæs, R. B. Grøndahl, S.-E. Fleten, and T. K. Boomsma, “Insights from actual day-ahead bidding of hydropower”, *International Journal of Sustainable Energy Planning and Management*, vol. 7, pp. 34–54, 2015.

[47] P. Shinde, I. Kouveliotis-Lysikatos, and M. Amelin, “Analysing trading trends in continuous intraday electricity markets”, in *2021 56th International Universities Power Engineering Conference (UPEC)*, IEEE, 2021, pp. 1–6.

[48] M. Birkedal and T. F. Bolkesjø, “Determinants of regulated hydropower supply in Norway”, *Energy Procedia*, vol. 87, pp. 11–18, 2016.

[49] L. Dueholm and H. F. Ravn, “Modelling of short term electricity prices, hydro inflow and water values in the Norwegian hydro system”, in *Proceedings of the 6th IAEE European Conference, Zürich*, 2004.

[50] E. Nycander and L. Söder, “Modelling prices in hydro dominated electricity markets”, in *2022 18th International Conference on the European Energy Market (EEM)*, IEEE, 2022, pp. 1–6.

[51] A. Tolonen, “Forecasting short-term valuation of regional hydropower generation with machine learning”, Master’s Thesis, Aalto University, 2022.

[52] A. Savitzky and M. J. Golay, “Smoothing and differentiation of data by simplified least squares procedures.”, *Analytical Chemistry*, vol. 36, no. 8, pp. 1627–1639, 1964.