NTNU

Norwegian University of
Science and Technology

Coordinated Multimarket Bidding for a Hydropower Producer using Stochastic Programming

Håkon Kongelf

Kristoffer Overrein

Industrial Economics and Technology Management
Submission date: June 2017
Supervisor: Stein-Erik Fleten, IØT

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

Problem description

Experiencing tougher competition in the Nordic power markets and higher penetration of non-dispatchable energy sources, the question arises whether reserve market opportunities should be taken into account when bidding into the day-ahead market. The coordinated bidding problem is often modeled as a multi-stage stochastic optimization problem, associated with long solution times. Hence, there is not much research currently available to provide answers to the question. This thesis considers a Norwegian hydropower producer and its participation in the day-ahead, primary reserve and balancing market.

Preface

This master thesis is written within the field of Applied Economics and Operations Research at the Department of Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU). The thesis is motivated by important trends in the Nordic power markets and the curiosity of a Norwegian hydropower producer as to whether multimarket coordinated planning can increase utilization of market opportunities.

We would like to thank our supervisors Stein-Erik Fleten and Gro Klæboe for valuable guidance and fruitful discussions. Their insight into the problem and knowledge of the field has been very helpful.

We want to express our gratitude to Ivar Nastad and the rest of the team at NTE Power Markets, for sharing of their experience, aiding the proposal of an interesting scope, and for inviting us to their HQ to introduce us to their operations and help shaping the framework of an interesting case study.

Trondheim, June 9, 2017

Håkon Kongelf

Kristoffer Overrein

Abstract

With the connection of new regions into the Nordic power markets competition intensifies. The incentive of power producers to exploit their portfolios more effectively is increasing. Expecting a higher share of non-dispatchable energy production in the years to come, supply of ancillary services (reserves) will be divided between those producers who possess flexible portfolios. The question then arises whether a hydropower producer should take the reserve markets into account when submitting bids for the day-ahead auction. This thesis considers the bidding problem of a Norwegian hydropower producer bidding into the day-ahead, taking the primary reserve and balancing market into account. A stochastic program is proposed to investigate the value of coordinating the day-ahead bid with the reserve market opportunities. Stochastic programs require good scenarios to produce high-quality solutions, and hence a great effort is put into generating scenarios that do well at predicting market prices and volumes, and representing the associated uncertainty. A comprehensive scenario generation framework that captures the dynamics of each market as well as their inter-dependencies is proposed, tested and found to perform well.

Modeling choices and assumptions are made with the goal of attaining reasonable solution times, while at the same time reflecting the actual planning procedures of the producer. Short solution times allow for a comprehensive case study to be conducted for 250 days in 2016. The case study is performed under three control variables: price deviation from water value, planning horizon granularity and the number of watercourses in the portfolio. First, the testing is done with one watercourse at disposal. Coordination is found to yield a small gain of about 1 %. Next, the profitability of coordinated bidding is further investigated with respect to price deviation from water value. The results show that the gain associated with coordination is higher when this deviation is low. The gain tends to zero when the deviation increases. Next, we investigate the effect of adding more watercourses to the planning problem. Gains decrease, but seem to stabilize around a value of 0.5 %.

Sammendrag

Med stadig flere regioners tilknytning til de nordiske kraftmarkeder tilspisser konkurransen mellom kraftprodusenter seg, og deres insentiv til å utnytte sine porteføljer på en mer effektiv måte forsterkes. Det forventes at andelen av ikke-regulerbar energiproduksjon (herunder bl.a sol- og vindkraft) skal øke i årene framover, slik at en høyere etterspørsel etter reserver vil måtte etterkommes av de produsenter som er i besittelse av regulerbar produksjonskapasitet. Et spørsmål som derfor melder seg, er om en vannkraftsprodusent bør ta mulighetene i reservemarkedene i betraktning når spot anmeldes. Denne oppgaven iøynetar en norsk vannkraftsprodusent som deltar i spot-, FNR- og RK-markedet. Et stokastisk program utvikles med det formål å undersøke lønnsomheten forbundet med å koordinere spotanmelding opp mot mulighetene i reservemarkedene. Stokastiske programmers evne til å treffe gode beslutninger avhenger av kvaliteten på scenariene som gis. Derfor legges betydelig innsats i å generere scenarier som godt predikerer markedspriser og -volumer, og som samtidig representerer usikkerheten forbundet med nevnte prediksjoner. Et integrert rammeverk som fanger dynamikken i hvert marked, så vel som sammenhengene mellom dem blir her foreslått, testet og funnet tilfredsstillende.

Modelleringsvalg og antakelser er tatt slik at vi oppnår håndterbare løsningstider, som samtidig reflekterer de virkelige planlegginsprosedyrene hos produsenten. Korte løsningstider tillater omfattende testing av beslutningsmodellen for 250 dager i 2016. Beregningsstudien utføres under tre kontrollvariabler: prisavvik fra vannverdi, oppløsning av planlegginshorisonten, og antall vassdrag i porteføljen. I første om gang kjøres testrammeverket for ett vassdrag. Det viser seg at koordinering gir en svakt økt fortjeneste på omlag 1 %. Videre undersøkes det hvordan koordinert planlegging presterer med hensyn på prisavvik til vannverdi. Denne analysen avslører at økt prosentvis fortjeneste forbundet med koordinering er betydelig når spotprisene utover et planleggingsdøgn ligger tett på vannverdien. Den positive effekten av koordinering svekkes dog når dette avviket øker. Til sist undersøkes effekten av å øke porteføljestørrelsen. Prosentvis økt fortjeneste avtar, og stabiliseres på omlag 0,5 %.

Contents

Nomenclature

Sets & Indices

$f \in \mathcal{F}$	Set of production curve points
$i \in \mathcal{I}$	Set of generators
$j \in \mathcal{J}$	Set of reservoirs
$l \in \mathcal{L}$	Set of water value reference levels
$k \in \mathcal{K}$	Set of watercourses
$p \in \mathcal{P}$	Set of pricepoints
$m \in \mathcal{M}$	Set of markets
$t \in \mathcal{T}$	Set of hours
$h \in \mathcal{H}$	Set of subperiods
$s \in \mathcal{S}$	Set of day-ahead/primary reserve scenarios
$\omega \in \Omega$	Set of balancing market scenarios

Parameters

π_s	Probability of day-ahead scenario
$\pi_{s\omega}$	Probability of balancing scenario
ρ_{mts}	Market price
$\rho_{mts\omega}$	Balancing premium
ι_m	Dummy variable for balancing price
P_p	Bid curve pricepoint
Q_i^{min}	Lower bound on production capacity
Q_i^{max}	Upper bound on production capacity
A_{if}	Production/discharge curve cut constant
B_{if}	Production/discharge curve cut slope
I_{jt}	Inflow of water
Γ_{ij}	Indicator matrix for discharge at i ending up in j
$\Lambda_{jj'}$	Indicator matrix for spill at j' ending up in j
$O_{j'ts}$	Spill
V_j^0	Initial water volume in reservoir
E_{lk}	Water value reference cut constant
W_{jk}	Water value reference cut unit value
W_k	Constant water value in well regulated one-reservoir watercourse
V_{jl}	Water value reference cut level
D_i^{min}	Minimum discharge
D_i^{max}	Maximum discharge

V_j^{min}	Minimum water level
V_j^{max}	Maximum water level
σ	Estimated balancing market share
$\nu_{mts\omega}$	Regulating volume
κ_i^{max}	Maximum droop
κ_i^{min}	Minimum droop
C_i	Cost of start-up

Variables

x_{mts}	Production commitment
$w_{ks\omega}$	Water value at end of planning period
$\Delta w_{ks\omega}$	Cost of water dispatched
z_{pt}	Day-ahead bid curve volume
$q_{mits\omega}$	Generator production
$d_{its\omega}$	Generator discharge
$v_{jts\omega}$	Reservoir level
$u_{ihs\omega}$	On/off-decision
$c_{ihs\omega}$	Start-up cost
ϕ_{pt}	Primary reserve bid curve volume

Chapter 1

Introduction

Over the past 25 years, electricity markets in the Nordic countries have been deregulated in order to improve efficiency and increase security of supply. This has shifted the operation of power markets from centralized to competitive. Ventosa et al. 2005 describe the trend towards competition in the electricity sector, and how it has led to efforts by the research community to develop decision and analysis support models adapted to the new market context. They classify the main families of approaches of energy market modeling. It is emphasized that one-firm optimization models are able to deal with difficult and detailed problems because of their better computational tractability. Good examples of such models can be found in short-term hydrothermal coordination and unit commitment, in which binary variables are required, and optimal offer curve construction under uncertainty. Optimization models can be divided in two subgroups; single-market and multi-market models. Single-market decision making softwares are widely taken advantage of of the world's power producers.

In the Nordic electricity market of today, by far the largest amount of electricity is traded in the day-ahead market. Though, the demand for reserves is expected to increase in the coming years. Due to a larger share of non-flexible renewable, intermittent, non-dispatchable energy, it is expected that the power markets' dispatchable production capacities will have a greater demand for regulating reserves to fulfill (Lorubio 2011). Contributions have been made to the literature in the last years, on whether taking smaller and subsequently cleared markets into account when bidding into a main market is profitable. An overview of the contributions made is given in Chapter 3.

We base our investigations on the market conditions of the price-area NO3 in mid-Norway. Power market bidding and the operation of hydro-power production capacity is modelled using stochastic programming (SP). We first give an introduction to the energy markets treated in this thesis in Chapter 2. The end of this chapter addresses important attributes of hydro-power; how it is traditionally planned, operated, and finally how coordinated market considerations may be profitable when possessing flexible production capacities. Thereafter, we discuss our contribution to the topic of coordinated bidding in electricity markets in context to existing literature in Chapter 3. In

chapter 4 we propose a stochastic programming model that produces bid curves for the Nord Pool Day-Ahead market, while taking into account the opportunities that may occur in subsequent reserve markets, cleared at a later point of time. To model the uncertainty in prices and dispatched volume, we put a great amount of effort into generating scenarios to the stochastic program. To begin with we perform an empirical analysis on prices and traded quantities in section 5. This is the basis of our choice of market models in section 6. In chapter 7 we propose a framework that generates scenarios to the decision model. Finally, in chapter 8 we carry out a computational case study to test the performance of coordinated planning and benchmark it against it's counterpart; sequential planning. Chapter 9 concludes the thesis.

Chapter 2

Energy markets and hydropower production

2.1 Introduction

In this section we begin by briefly presenting the Nord Pool Day-Ahead market. It is the main arena for trading power in the Nordic region. In contrary to most other European day-ahead markets, The Nord Pool day-ahead market is cleared before the reserve markets. The reserve markets that we wish to focus on in our work are the primary reserve and tertiary reserve markets. These markets are presented in Sections 2.3 and 2.4. We conclude by describing the characteristics of hydropower planning and production.

2.2 Day-ahead market

The day-ahead market (spot market) is the main arena for trading power. There is a variety of different types of bids the producer can make. This paper only considers single-hourly bids. The seller needs to decide how much he can deliver and at what price, hour by hour. A buyer needs to assess how much power it needs to meet demand the following day, and how much to pay for this volume. Nord Pool aggregates both supply and demand curves and calculates the price which balances the two; the spot price.

12:00 CET is the deadline for submitting bids for power which will be delivered the following day. Hourly prices are typically announced at 12:42 CET or later. An hourly price corresponds to a volume commitment for producers and buyers. This means that producers are obliged to deliver whatever accumulated volume they bid beneath this price. From 00:00 CET the next day, power contracts are physically delivered (meaning that the power is provided to the buyer) hour by hour according to the contracts agreed. If a producer should fail to deliver the committed volume, the TSO will buy tertiary

reserves on the behalf of the producer and make him pay for it afterwards (NordPool 2016).

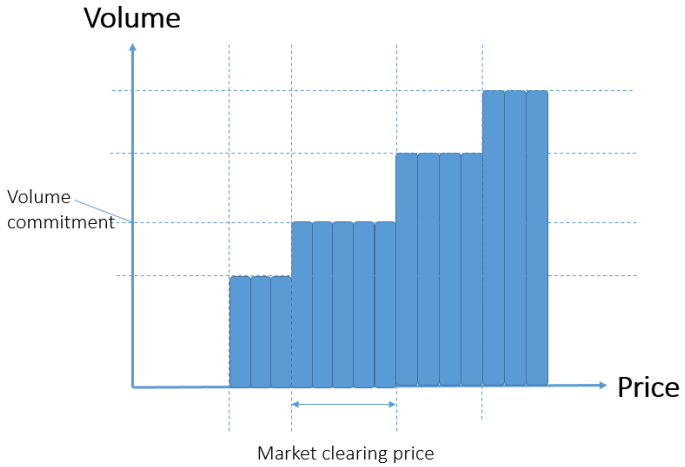

Figure 2.1: Example bid curve and relation between market price and producer commitment. Any clearing price within the interval in the figure will result in the specified commitment

2.3 Primary reserves market

When either production or consumption changes, so does the frequency in the grid. A change in frequency of +/- 0.1 Hz activates primary reserves. The regulation of these is completely automated within the respective plants, with the generator droop control restricting the degree of possible regulation. The market comprises both normal operation reserves (FCR-N) and disturbed operation reserves (FCR-D). There exist two markets for primary reserves; a day-market and a week-market. The day-market consists of daily commitments, hour by hour, and is the market considered in this paper.

Only a running plant can deliver primary reserves. The producer makes increasing bid curves for the plants he expects to run the next day, in the same manner as for the day-ahead market. He is paid for the reservation of a power band, not for an actual delivery of energy. The half-width of this band (in MW), for which the supplier is remunerated, is restricted by the equation below:

$$\min\{PROD - LB, UB - PROD, 0.1 \cdot 2 \cdot \frac{N_i}{\kappa}\}$$

where $PROD$ denotes the planned production to the day-ahead market, UB and LB upper and lower bounds on the generator capacity respectively, N_i the nominal power of the generator, and κ the droop setting. In order to secure a good distribution of primary reserves producers are obliged to set their plants' droop setting equal to or below 12 percent. The TSO will thus utilize the "residual" reserves of a producer

no matter whether he participates in the market or not. Producers that do not bid towards the market get paid by a certain rate. It is up to the TSO to decide which quantities of primary reserves to provide. The lastly committed bid sets the price for all participants, and the market is thus a marginal pricing market(Statnett AS 2016). Bids are committed before 18:00 CET on the day before delivery, and final commitments are given by the TSO at 19:00.

2.4 Tertiary reserves market

The tertiary reserves market will be referred to as the balancing- or the regulating market interchangeably in this paper. Tertiary reserves are used to reduce larger imbalances unforeseen by the day-ahead settlement, and are activated to relieve primary reserves, such that these more flexible reserves are ready for the next sudden imbalance occurring. Tertiary reserves are also activated when regional bottlenecks are present.

A producer, tentatively, places an increasing bid curve for the market within 21:30 on the evening before the day of schedule. Though, bids can be changed until 45 minutes before the hour of schedule, and bids the evening ahead is only a mere guidance for the TSO.

The activation happens after a call from the TSO to the producer, and the producer must be able to fully activate the volume agreed upon within 15 minutes. Both production and consumption capacities can be offered in the market. In the case of tertiary reserves, the producer is paid for actual energy delivery. The TSO, at all times, has a list of offered reserves, and starts off by activating the cheapest (upper) alternative whenever needed. The lastly activated bid sets the price for all market participants. In the case of upward regulation the balancing price is by market rules equal to or higher than the spot price, and in the down-regulation case the price is equal to or lower than the spot price. The absolute values of these differences (between balancing and spot price) are often referred to as balancing market premiums (Statnett 2016).

Figure 2.2 shows the time line of bidding and clearing for all three of the markets.

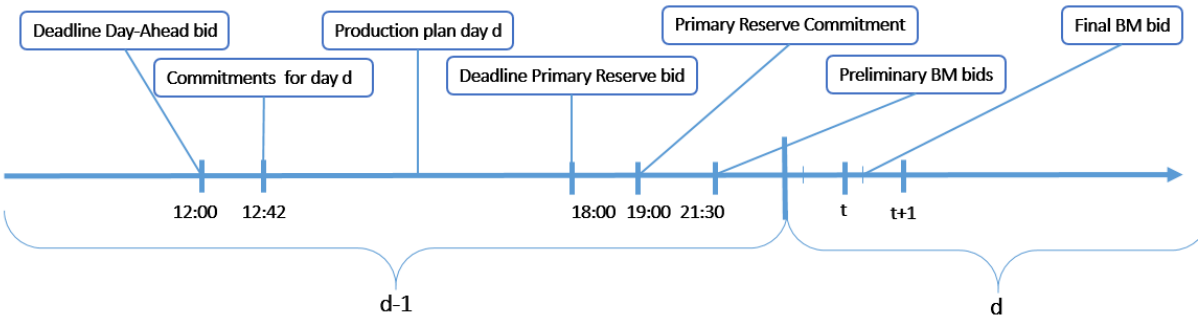

Figure 2.2: Sequential clearing of relevant markets

Table 2.1: Market characteristics

	Day-ahead	**Primary reserve**	**Balancing**
Settlement granularity	Hourly	Hourly	Hourly
Market clearing	Day-ahead	Day-ahead	Real time
Remuneration	Energy only	Reserved capacity	Energy only
Pricing scheme	Marginal pricing	Marginal pricing	Marginal pricing
Activation	Manual	Automatic	Manual
Single/dual pricing	Single	Single	Dual
Price cap/floor [NOK]	-4 500 - 27 000	>10	Day-ahead price

2.5 Hydropower production and planning

Hydro power is undoubtedly Norway's most important electric power source as it constitutes 99 % of total production (Statkraft 2016). It is a renewable energy source and possesses a different set of characteristics compared with common electricity production methods such as combustible fuels, nuclear, wind and solar.

The principle of hydropower is to utilize potential energy in elevated water reservoirs. This energy is converted to electrical energy in a turbine-generator system. The energy is transferred directly to the grid for consumption. A strength of hydropower production is its flexibility. Production volume can be regulated in a matter of seconds to accommodate for unexpected events on the power grid. This is in contrast to for example nuclear and thermal production, which require more time to adjust production. The flexibility of hydropower makes it a valuable reserve resource. In comparison with other renewable resources such as wind and solar power, hydropower is predictable. While wind and solar resources are dependent on exogenous factors in real time, hydropower has storage capacity in the reservoirs.

Water in the reservoirs comes for free. On the other hand, water is a scarce resource, and volume in the reservoirs constrain production opportunities. Hence, the producer must always consider the option of producing at a later point in time when prices are higher. This implies that the opportunity cost of producing at a later point in time is arguably the most important cost when planning hydropower production. The producer must assess the future value of each unit of stored water and produce when prices are higher than this value.

The marginal increase in future income resulting from a unit increase of water in reservoirs is commonly referred to as the water value. The water value is high when reservoir levels are low and lower when levels increase. The risk of spill increases as levels approach the maximum capacity of the reservoir, and hence the marginal water value approaches zero for full reservoirs.

It is common practice among producers to distinguish between short term, seasonal and long term planning of production. The long term planning is often done utilizing a stochastic optimization model, and water values are given by the shadow price of the reservoir volume constraints. Water values can subsequently be used as boundary conditions in short term planning models. In this paper we treat short term planning on a daily basis, and water values are treated as deterministic, given parameters throughout this paper.

When planning short term production, several markets can be considered. Some of the opportunities in the markets are mutually exclusive. For instance, if a producer commits to a large quantity in the day-ahead market, there may be no extra capacity to react to balancing opportunities. This trade-off between markets can be considered during the planning phase, by taking subsequent markets into account when bidding into the day-ahead market. We refer to this practice as coordinated bidding or planning. Making bids without taking subsequent opportunities into account, is referred to as sequential bidding or planning. Surprisingly, it is hard to determine whether a coordinated planning is worthwhile. Such models require more computations and more advanced models, but may not necessarily produce large gains. This might be because reserve markets are very small compared to the day-ahead market or because reserve markets are hard to predict in terms of prices and quantities. In this thesis we seek to investigate whether, and if so in what situations, a coordinated planning is beneficial.

Chapter 3

Coordinated bidding in the literature

Several authors have used optimization models to investigate the possible gains from coordinating bids towards different Nordic power markets. Faria and Fleten 2011 propose a two-stage stochastic model to support the bidding into the Nordic day-ahead market while taking the Elbas market into account. The impact of considering the possibility of trading in the Elbas in the day-ahead bidding decision is measured. The results imply that, for a price-taking medium-sized hydropower, considering Elbas at the time of day-ahead bidding has no significant effect on profits.

Boomsma et al. 2014 propose a multi-stage stochastic programming model for coordinated bidding into two sequential markets, the day-ahead and the balancing market. Their objective is to investigate whether there is a gain from using the coordinated model instead of a sequential model, that takes only day-ahead scenarios, and no balancing scenarios, into account when producing a day-ahead market bid curve. As multi-stage models are difficult to solve, they derive an upper bound, UB, of the objective value of the coordinated model, that can be calculated without solving the coordinated problem. They run a case study for the year 2010, and record the gains from coordinated planning, $\frac{UB-LB}{LB}$, where LB is the value of bidding sequentially; first into the day-ahead market without regarding balancing market opportunities, then to react to these later conditional on the day-ahead bid. They conclude that, under a two-price system, there is a significant gain from coordinated bidding in the day-ahead and balancing markets. However, they do not explicitly model the volume allowance in the balancing market.

Eriksrud and Braathen 2012 also consider the day-ahead and balancing market. Unlike Boomsma et al. 2014 they put an effort into modeling the risk of not being dispatched in the balancing market by giving demanded balancing volumes as a stochastic parameter to the model. They indicate that there are situations with certain relations between market prices and the cost of water, where a coordinated model is likely to outperform a sequential one. They point out, though, that more case testing has to be

done in order to gain knowledge about the profitability of accounting for the balancing market when bidding into the day-ahead market.

Not much work has been done to investigate the effects of taking the primary reserve market into account when bidding into the day-ahead market in the Nordics. Though, Kårstad and Skjong 2015 formulate a stochastic dynamic model that coordinates the day-ahead and primary reserve bidding for a thermal power producer participating in the Swiss market. The report concludes that it is profitable for the producer to participate in both markets. They forward to future work the inclusion of other markets, as it is believed that there are synergy effects from delivering capacity in different reserve markets.

To our knowledge, no research has been conducted to investigate the profitability of taking both the balancing and primary reserve market into account when bidding into the Nord Pool day-ahead market. This is our goal in this thesis. Most earlier contributors to the field have experienced long solution times associated with a coordinated decision model. We have put effort into finding the intersection between realism, computational tractability and fairness, in benchmarking a coordinated and a sequential planning regime. The decision model is presented in a rather general form in chapter 4. The achieved solution times allow for much more testing of the model, which in turn provides a great amount of results. The amount of results make a good basis for drawing conclusions on the value of coordinated bidding. In addition, we test under different circumstances: portfolio size and deviation between the day-ahead price and the value of water.

Chapter 4

Problem formulation

Participants in the electricity markets submit their bids in a sequential order. After prices clear in the day-ahead market, producers must place bids in the primary reserves market, and then the balancing market (tertiary reserves). Bids in the balancing market are only indicative and can be changed until 45 minutes before the operating hour throughout the next day. In addition, producers can sell or buy power in the intraday Elbas market opening at 2 pm the day before the operating day.

In this chapter, a stochastic mixed integer programming model (SMIP) is developed for constructing bid curves for the day-ahead market, taking the alternatives of bidding into the primary and tertiary reserves markets into account. The model has a daily planning horizon, and boundary conditions are given as parameters from a long-term model which is not treated in this thesis. The alternative of including Elbas is left for future work. As mentioned in 3.2 there exist several primary reserve markets, but we shall only include the daily normal operations reserves (FCR-N). From now on we will simply use the term primary reserves for this market.

The flow of information during the bidding process is stage-wise. A stage is a point in time when new information is revealed to the decision-maker, and he may take recourse actions to accommodate for the new information and previous actions. There are several such stages in this problem. When day-ahead prices clear, the producer obtains information about obligated production quantity, and may then allocate this quantity to available production capacity. In addition, the producer submits bids for the primary reserve market. In the next stage primary reserve prices and obligations are revealed. For the balancing market final bids are placed at latest 45 minutes before operating hour. The balancing market is then operated in real-time, and the producer may or may not be dispatched during operating hour. Hence, the problem consists of 27 stages, counting initial bid submission to the day-ahead market. Figure 4.1 illustrates this structure.

For feasibility of implementation a simplified structure is considered. The simplified structure consists of three stages, and is illustrated in Figure 4.2. In the first stage the producer places day-ahead bids. In the second stage day-ahead obligations are

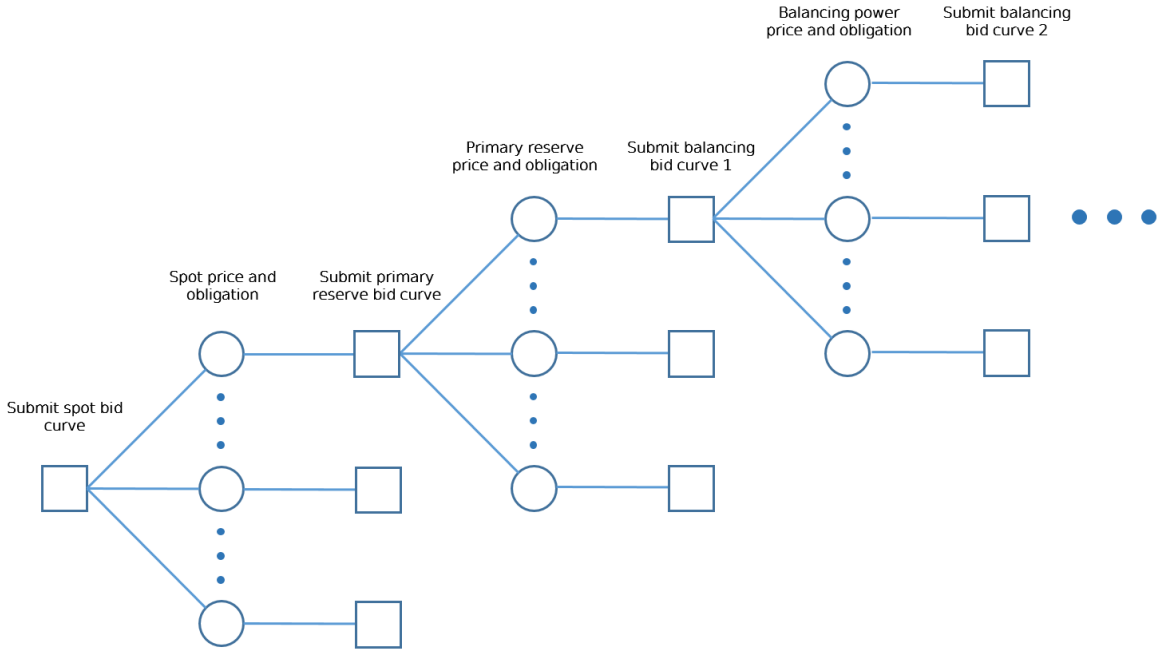

Figure 4.1: Decision tree (squares denote decisions, circles denote new information)

calculated according to the realized price. In addition, the producer receives information about the primary reserve price, and may decide on a commitment to this market. In the final stage balancing market prices are revealed for the entire day and commitments are decided upon. This structure may seem like a violation of natural non-anticipativity. However, considering marginal pricing and the price taker assumption, the producer has incentive to bid marginal cost. Only prices above marginal cost are attractive to the producer, which would be reflected in a bid curve. Therefore, only inter-hour coordination and predictability in the trade-off between reserve markets is overestimated with this structure. In addition, day-ahead commitments are much more important for inter-hour coordination than reserves because these commitments usually determine whether a generator is on or off. Value from reserve markets should nevertheless be regarded strictly as an upper bound.

Uncertainty is inevitable considering the problem at hand. The model contains parameters that are inherently stochastic; clearing prices in the three markets as well as the volumes traded in the balancing market. Because reserve markets are designed to adjust for unforeseen events, the future clearing prices and volumes are hard to predict (Klæboe et al. 2013). We therefore resolve to stochastic programming using discrete representations of uncertainty commonly known as scenario trees. The producer has a set of scenarios, $\mathcal{S}$, when constructing day-ahead bid curves. The scenarios contain day-ahead and primary reserve prices. For each $s \in \mathcal{S}$, a set of scenarios Ω^s for prices and volumes in the balancing market is given.

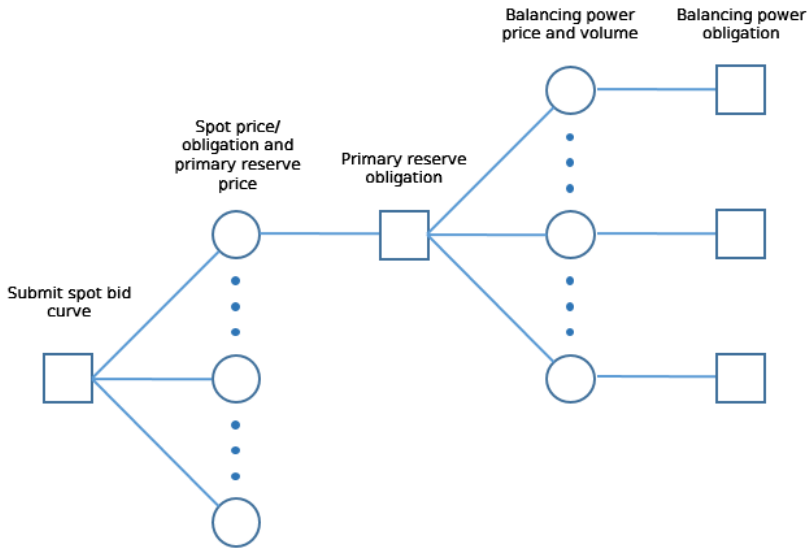

Figure 4.2: Simplified structure decision tree (squares denote decisions, circles denote new information)

The remainder of this chapter consists of three parts. In section 4.1 we introduce notation and define the mathematical model for the markets. Section 4.2 addresses modelling of production. Section 4.3 discusses input data to the model.

4.1 Modelling bidding and commitments

We seek to model four markets; day-ahead, primary reserve, up regulation and down regulation. Let the set $m \in \mathcal{M}$ denote the different markets

$$m \in \mathcal{M} = \begin{cases} 1 & \text{Day-ahead} \\ 2 & \text{Primary reserve} \\ 3 & \text{Up regulation} \\ 4 & \text{Down regulation} \end{cases} \tag{4.1.1}$$

The here-and-now decision of the producer is the bid curve to submit for the day-ahead market. Therefore, this is the only bid curve modeled explicitly. The bid curve consists of a set of non-decreasing price-volume pairs, which is handled by restriction 4.1.3. To preserve linearity of the model we assign fixed bid points and assign a volume to each. A similar approach was implemented by S.E. Fleten and Pettersen 2005. Bid points are assigned such that an equal number of price scenarios are distributed between each point. Let $p \in \mathcal{P}$ denote the set of bid points, P_p the price of the bid point and z_{pt} the bid volume of bid point p at time t. Now, because the TSO can perform linear interpolation between price-volume pairs to fill demand, commitment

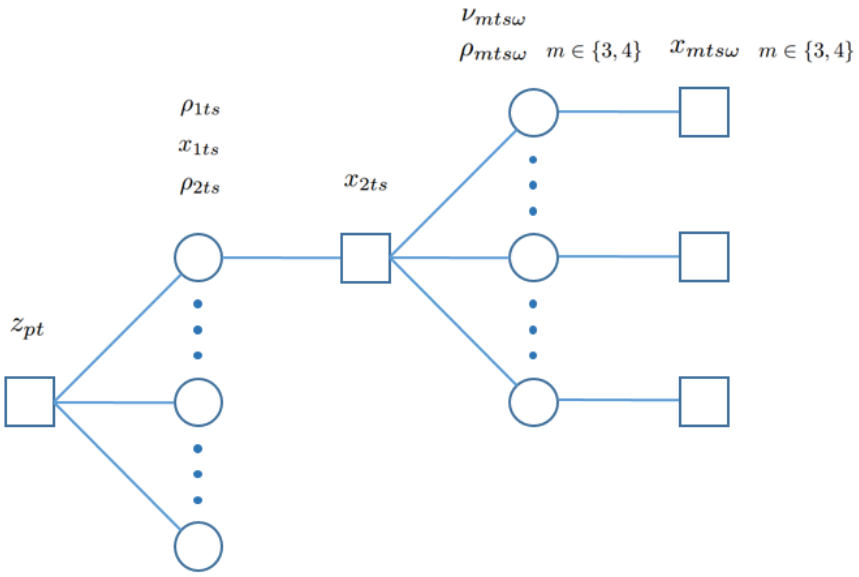

Figure 4.3: Decision tree with problem variables and parameters (squares denote decisions, circles denote new information)

will be determined by equation 4.1.2.

$$x_{1ts} = z_{pt} + (z_{p+1t} - z_{pt})\frac{\rho_{1ts} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1ts} \leq P_{(p+1)}, \quad p \in \mathcal{P}, s \in \mathcal{S}, t \in \mathcal{T} \tag{4.1.2}$$

$$z_{(p+1)t} \geq z_{pt} \qquad p \in \mathcal{P} \setminus |\mathcal{P}|, t \in \mathcal{T} \tag{4.1.3}$$

where ρ_{1ts} denotes the day-ahead price and x_{1ts} denotes the committed volume.

Next, for each price realization in $\mathcal{S}$, it has to be decided what reservation of primary reserve to offer in each hour $t \in \mathcal{T}$ for the day of planning. This reservation is allocated to the decision variable x_{2ts}.

Balancing market commitments are limited by the demand for balancing power. In addition, upward and downward balancing can not be delivered simultaneously. We restrict the commitments by an estimated market share σ multiplied by the forecasted demand in NO3 $\nu_{mts\omega}$

$$x_{mts\omega} \leq \sigma \nu_{mts\omega} \qquad m \in \{3,4\}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.1.4}$$

In terms of the parameters and decision variables defined thus far, we have the decision structure illustrated in Figure 4.3

4.2 Modelling production

Next, we seek to model the connection between commitments and production. The committed volumes in each market must be allocated between the available generators. Let $q_{1its\omega}$, $q_{2its\omega}$ and $q_{mits\omega}$ when $m \in \{3,4\}$, denote the volume or reservation allocated to generator i at time t in scenario s for day-ahead, primary reserve and balancing respectively. Then it follows that

$$\sum_{i\in\mathcal{I}} q_{1its\omega} = x_{1ts}, \qquad s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.1}$$

$$\sum_{i\in\mathcal{I}} q_{2its\omega} = x_{2ts}, \qquad s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.2}$$

$$\sum_{i\in\mathcal{I}} q_{mits\omega} = x_{mts\omega}, \qquad m \in \{3,4\}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.3}$$

Primary reserve reservation on a generator must lie within an interval, which lower and upper bounds are functions of κ_i^{max} and κ_i^{min}, the generator's maximum and minimum droop settings. Generators are grouped into two sets. Generators $\mathcal{I}'$ can be turned on and off during the planning horizon by the binary variable $u_{ihs\omega}$. This decision can be made for subperiods $h \in \mathcal{H}$ of the planning horizon. The remaning generators $\mathcal{I} \setminus \mathcal{I}'$ are producing. For generators $i \in \mathcal{I} \setminus \mathcal{I}'$ reservation is bounded by

$$\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2its\omega} \leq \frac{0.2N_i}{\kappa_i^{min}} \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.4}$$

where N_i denotes nominal production. For generators to be turned on or off by the decision model, $i \in \mathcal{I}^{'}$, the binary variable $u_{ihs\omega}$ has to be included, such that primary reserve reservation is only possible when generators are producing

$$u_{ihs\omega}\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2its\omega} \leq u_{ihs\omega}\frac{0.2N_i}{\kappa_i^{min}} \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s \tag{4.2.5}$$

Similarly, to ascertain that generators produce within possible power range, the following two restrictions are required

$$Q_i^{max} \geq q_{1its\omega} + q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.6}$$

$$Q_i^{min} \leq q_{1its\omega} - q_{2its\omega} + q_{3its\omega} - q_{4its\omega} \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.7}$$

where Q_i^{max} and Q_i^{min} denote upper and lower bounds on production respectively. Primary reserve reservation must be included in both directions, hence the change in

sign of $q_{2its\omega}$. For generators $\mathcal{I}'$ we have

$$u_{ihs\omega}Q_i^{max} \geq q_{1its\omega} + q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s \tag{4.2.8}$$

$$u_{ihs\omega}Q_i^{min} \leq q_{1its\omega} - q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s \tag{4.2.9}$$

Start-up costs are incurred whenever generators start operating after a period of standstill. Start-up cost is accounted for in the following formulation

$$c_{ihs\omega} \geq C_i(u_{ihs\omega} - u_{i(h-1)s\omega}) \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, \omega \in \Omega^s \tag{4.2.10}$$

where C_i denotes the cost of a start-up.

The quantity of power produced depends on the volume discharge through the turbine system, the efficiency of the turbine and the effective head by the physical relation

$$q = \eta(d, h_{eff})\varrho g d h_{eff}$$

where q is the power produced in watts, $\eta(d, h_{eff})$ is the system efficiency at the given discharge and head, ϱ is the density of water in kg/m^3, g is the gravitational constant, h_{eff} is the effective head in meters and d is the discharge in m^3/s. Note that the overall efficiency of the system is also dependent on other factors such as the load in the grid and adjustments on the generator. The dependence on d is stressed to reveal the non-linear relationship between discharge and production. To preserve the linearity of the model and hence utilize linear programming tools, a set of cuts $f \in \mathcal{F}$ is introduced. h_{eff} can for well regulated reservoirs i.e. little change in the height of the water surface, be treated as constant throughout a day of planning. The inclusion of the binary variable $u_{ihs\omega}$ prevents the model from dispatching water when a generator is not running, but which production cuts $\mathcal{F}$ has a negative intersection. Approximation of the concave production curve is thus done forcing

$$q_{1its\omega} + q_{3its\omega} - q_{4its\omega} \leq u_{ihs\omega}A_{if} + B_{if}d_{its\omega} \qquad h \in \mathcal{H}, i \in \mathcal{I}, f \in \mathcal{F}, s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s \tag{4.2.11}$$

where primary reserve capacity is not included because there is no net water usage in this market. There may be restrictions governing the minimum or maximum discharge of a generator e.g. considering the wildlife in a watercourse. For the same reasons there may be restrictions on minimum and maximum volume $v_{jts\omega}$ in a reservoir. Therefore, we introduce lower and upper bounds on discharge and reservoir volume

$$D_i^{min} \leq d_{its\omega} \leq D_i^{max} \qquad i \in \mathcal{I}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.12}$$

$$V_j^{min} \leq v_{jts\omega} \leq V_j^{max} \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.13}$$

where $\mathcal{J}$ is the set of reservoirs. Furthermore, one must model how the reservoir levels respond to production, and the interconnected flow between reservoirs. Since most watercourses consist of several reservoirs, production in one part of the watercourse will increase reservoir levels downstream. Let I_{jt} denote inflow, Γ_{ij} and $\Lambda_{jj'}$ the fraction discharge at i and spill at j' ending up at j, and $O_{j'ts\omega}$ the spill at j' at time t. Now we get

$$v_{jts\omega} - v_{j(t-1)s\omega} = I_{jt} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{its\omega} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'ts\omega}, \quad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s \tag{4.2.14}$$

$$v_{j1s\omega} - V_j^0 = I_{j1} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{i1s\omega} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'1s\omega}, \qquad j \in \mathcal{J}, s \in \mathcal{S}, \omega \in \Omega^s \tag{4.2.15}$$

where V_j^0 denotes the initial volume.

Next, we proceed to model the value of water in the reservoirs. The future value of a marginal volume of water is dependent on the outlook on future prices and inflow, as modelled by a long term optimization model. In addition, the marginal water value depends on the level of the reservoirs. It is high at low levels, and tends to zero at full reservoirs. Hence the water value curve is concave and must be modelled by linear cuts to gain linearity. Parameters for these cuts are estimated by the long term optimization model at different reference levels. For well regulated plants i.e. with little change in reservoir volume and head, this way of modelling is not necessary, and one could rather use a water value independent of level.

We introduce a set of cuts $l \in \mathcal{L}$ reflecting the different levels. Let w_{ks} denote the future income of watercourse k at the end of the period, E_{lk} the reference level future income, W_{jk} the marginal water value at the reference level, V_{jl} the reference level, and $\mathcal{J}^k$ the set of reservoirs in the watercourse. Then for all cuts it must hold that

$$w_{ks\omega} \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk}(v_{j|T|s\omega} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, s \in \mathcal{S}, \omega \in \Omega^s \tag{4.2.16}$$

In some watercourses, there is only one reservoir and one generator. That is, after water has been dispatched from the reservoir it becomes irrelevant to the hydro power producer. When in addition the reservoir is very large relative to the dispatch capacity of the generator, it is considered reasonable to assume a constant water value (W_k) over the course of 24 hours. Thus, instead of modeling the value of the remaining water in a watercourse, $w_{ks\omega}$, and adding this to the objective function to be maximized, we can model the cost of water dispatched, $\Delta w_{ks\omega}$, and subtract this from the objective function. This alternative water valuation is shown by equation 4.2.17.

$$\Delta w_{ks\omega} \geq \sum_{i \in \mathcal{I}^k} \sum_{t \in \mathcal{T}} W_k d_{its\omega} \qquad k \in \mathcal{K}, s \in \mathcal{S}, \omega \in \Omega^s \tag{4.2.17}$$

However, for the objective function presented below, we choose the general water valuation formulation from equation 4.2.16. And finally, the objective function is

$$\max \sum_{s \in \mathcal{S}} \Bigg(\pi_s \sum_{m \in \{1,2\}} \sum_{t \in \mathcal{T}} \rho_{mts} x_{mts} + \sum_{\omega \in \Omega^s} \pi_{s\omega} \Bigg(\sum_{m \in \{3,4\}} \sum_{t \in \mathcal{T}} (\rho_{mts\omega} + \iota_m \rho_{1ts}) x_{mts\omega} + \sum_{k \in \mathcal{K}} w_{ks\omega} - \sum_{i \in \mathcal{I}'} \sum_{h \in \mathcal{H}} c_{ihs\omega} \Bigg) \Bigg) \tag{4.2.18}$$

where π denotes the probability of a scenario and $\rho_{3ts\omega}$ and $\rho_{4ts\omega}$ are upward and downward balancing premium respectively. In accordance with the market mechanisms we have that $\iota_3 = 1$ and $\iota_4 = -1$. The formulation can be found in its entirety in Appendix C.

4.3 Modelling of stochastic parameters

So far we have discussed the mathematical formulation and decision variables of the problem. This section is devoted to a treatment of the parameters in the model. Before implementation is possible, all parameters must be specified. The term parameter is usually applied to deterministic quantities, and the term random variable is often used for stochastic quantities. To avoid confusion with decision variables, we will use the term parameter for all quantities (including random variables) that are input (constants) to the mathematical program. Where necessary, we will specify whether a parameter is deterministic or stochastic. Except for six parameters, all problem parameters are treated as deterministic. The deterministic parameters are generally highly specific for the watercourse and production resource considered. The deterministic parameters will therefore be treated along with the case study in Section 8.2.

The stochastic parameters are the day-ahead price, primary reserve price, balancing premium and volume in both regulating directions (ρ_{mts}, $m \in \{1,2\}$,$\rho_{mts\omega}$, $m \in \{3,4\}$, $\nu_{mts\omega}$, $m \in \{3,4\}$). Stochastic parameters are usually generated by some appropriate generation algorithm. Scenario generation is a vital part of the modelling of stochastic problems. Stochastic programming relies on the fact that modelling uncertain parameters as if they were certain, e.g. mean value deterministic problems, can be inefficient. A realistic description of uncertainty is therefore needed. However, solving an instance with continous stochastic parameters is either very hard or impossible (Conejo et al. 2010). This issue is tackled by constructing a discretization of the stochastic parameters, known as a scenario tree.

Kaut and Wallace 2007 argue that a complete problem modelling should include

a procedure for scenario generation. In order to provide valuable solutions to the decision maker, quality of the input data must be emphasized. We will provide a brief introduction to scenario generation, with special focus on electricity markets.

4.3.1 A brief review of scenario generation approaches

In the later years, the research community has put significant effort into the topic of scenario generation. Conejo et al. identify four main categories of scenario generation methodologies. These include path-based methods, moment matching, internal sampling and scenario reduction.

Path-based methods uses econometric or time series models to generate complete paths. The collection of these paths is called a fan, and clustering of this fan produces a scenario tree. Moment matching is used to generate discrete distributions that satisfy a set of statistical properties, e.g. moments, correlations or percentiles. Internal sampling comprises various procedures for sampling from the original distribution. Finally, scenario reduction is used to reduce the number of scenarios to a prescribed cardinality utilizing some probability metric.

4.3.2 Scenario generation in electricity markets

There exists no single best procedure for scenario generation in general. Conejo et al. presents a comprehensive framework for scenario generation in electricity markets. The electricity prices are modelled as time series. This can be done using ARIMA series provided that the series is homoscedastic i.e. the variance is steady in time, stationary and that residuals are normally distributed. Several transformations can be made to achieve validity of these assumptions. The important point here is that the modelling choice must be made in accordance with empirical properties. One must have in-depth knowledge of market characteristics to be able to model them efficiently, and always check the assumptions.

For multivariate distributions or time series, the modelling complexity is significantly reduced under the assumption of a Gaussian normal joint distribution. If in addition the joint behaviour does not change in time, the dependence between the variables is fully characterized by the variance-covariance matrix. In this case, correlation between time series can be modeled utilizing a Cholesky decomposition. Modelling is more complicated for other joint distributions. Once again the main point is that modeling requires profound knowledge of the empirical behaviour. To model dependence correctly, one must know the joint behaviour of the markets.

Relation to the methodology in this thesis

In subsequent chapters, we propose a methodology that combines path-based methods (quantile autoregression - QAR) and a copula based heuristic that best fits the moment

matching category. We adopt parts of the methodology by Conejo et al. and deviate where need be.

The next three chapters treat scenario generation for the problem at hand. In Chapter 5 we describe the historical characteristics of the markets, which is necessary to choose good models. In Chapter 6 we propose models for the markets that reflect the important empirical characteristics. In turn, we connect the dots and outline the entire scenario generation algorithm in Chapter 7. In addition, we test the quality of the scenario generation method.

Chapter 5

Empirical market analysis

5.1 Introduction

The goal of this section is to provide a thorough assessment of the relevant electricity markets to aid subsequent model specification. There are three markets at hand: Day-ahead market, daily primary reserve market and the balancing market. All data is from the period 2014-2017 in the price zone NO3. Data was downloaded from Nord Pool and Statnett. We will investigate the historical characteristics of these markets in the natural order; the order of market clearing.

5.2 Day-ahead price

The daily day-ahead price (spot price) in NO3 for the period 2014-2017 is plotted in Figure 5.1a. We note that the price level seems to fluctuate around a mean of 240 NOK/MWh with some periods of longer deviations from this price. Surprisingly, it is hard to identify a predictable annual seasonal pattern. The price in 2014 is relatively steady throughout the year. During 2015 and 2016 however, we observe periods in which the price differs significantly from its mean value. Mid 2015 is dominated by a regime with low prices, and 2016 starts off with days of extremely high prices and ends with a period of high prices. These variations may be explained by other fundamental factors than pure calendar effects, such as rain fall, snow melting or extreme temperature that may shift somewhat from year to year. The average daily prices are found in Figure 5.1b. We note that spot prices are generally lower during weekends, due to lower demand for power in industrial applications.

To get an impression of the hourly day-ahead prices, descriptive statistics are presented in Table 5.1. Observe that the mean price increases in the morning hours, and are highest during the working day. Demand for power is generally higher during this period. By comparing the hourly means and medians we see that probability density is asymmetrically distributed. This is further characterized by the skewness and kurtosis.

Figure 5.1: Daily and daily averaged day-ahead prices

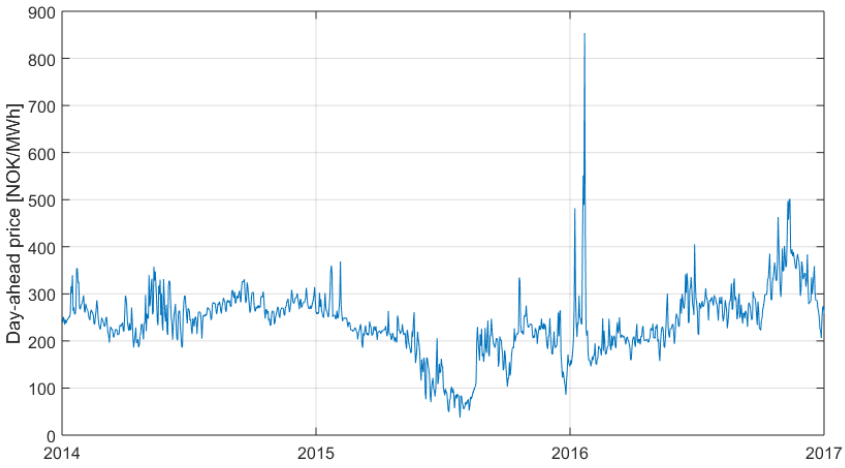

(a) 2014-2017

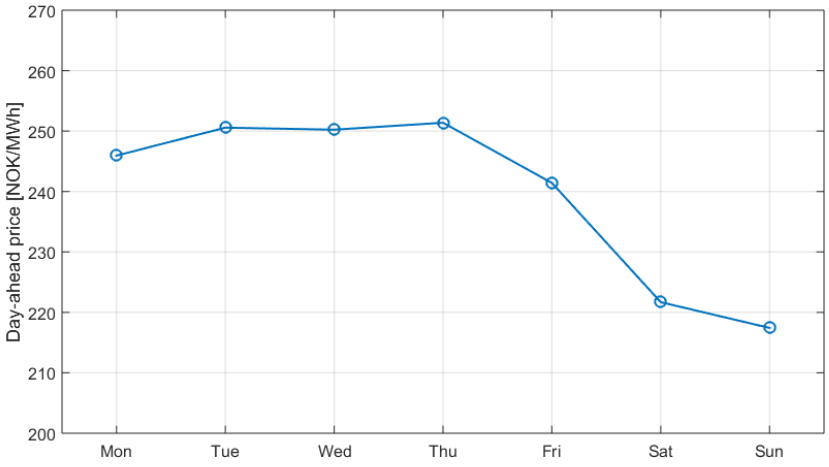

(b) Monday-Sunday

Table 5.1: Descriptive statistics for hourly day-ahead price

Hour	Mean	Median	St. dev.	Skewness	Kurtosis	Min	Max
00 - 01	209.02	219.01	53.01	-1.22	1.23	30.28	291.52
01 - 02	201.95	212.27	53.72	-1.20	1.16	17.43	286.01
02 - 03	197.69	208.39	54.50	-1.19	1.11	10.95	283.11
03 - 04	195.79	206.42	55.00	-1.18	1.06	10.86	281.98
04 - 05	198.15	210.20	55.31	-1.22	1.16	24.16	287.49
05 - 06	206.69	216.68	55.78	-1.20	1.22	23.98	300.61
06 - 07	223.59	228.32	59.93	-0.85	0.83	24.07	404.89
07 - 08	245.45	245.46	91.41	6.54	117.04	24.52	1937.81
08 - 09	261.50	255.02	107.96	5.91	85.39	25.52	2076.30
09 - 10	259.58	256.72	99.63	5.87	86.75	26.24	1938.20
10 - 11	256.57	255.81	83.39	2.49	24.39	26.96	1260.41
11 - 12	251.42	251.44	72.90	0.62	3.70	27.05	726.53
12 - 13	245.66	246.06	69.39	0.45	3.43	26.87	654.82
13 - 14	241.79	242.63	67.27	0.24	2.71	26.51	596.08
14 - 15	239.39	240.16	68.92	0.90	9.70	26.42	896.17
15 - 16	239.00	237.79	72.22	1.65	16.16	26.42	955.53
16 - 17	241.64	240.21	96.27	7.46	113.40	31.29	1937.91
17 - 18	248.23	246.85	105.90	7.88	111.46	38.86	1939.07
18 - 19	247.78	249.86	89.56	7.33	129.68	51.88	1936.94
19 - 20	243.83	247.28	65.57	0.06	3.43	52.44	672.02
20 - 21	238.04	241.61	60.19	-0.58	1.18	55.28	485.42
21 - 22	233.36	236.62	57.26	-0.84	0.74	55.46	384.42
22 - 23	226.78	232.31	55.02	-0.95	0.67	54.28	344.38
23 - 00	215.73	224.44	53.28	-1.12	0.90	38.32	303.73

Night and morning hours generally have more symmetric distributions than day hours. Note especially the extreme kurtosis that arises during the start and end of the working day, at 7-8 am and 4-7 pm respectively. Around these hours prices are generally much more volatile; hour 8-9 am and 5-6 pm are the only hours with a standard deviation above 100 NOK/MWh. Note also that prices can become quite extreme in these hours, with the highest observed price at NOK 2076.30 per MWh. We conclude that price behaviour may be very different from hour to hour.

However, hourly prices are obviously related. The information that bidders act on before market closure is the same for all hours. Hence, a forecasting error is likely to propagate throughout the day. To investigate cross-hour correlation, Figure 5.2a plots the historical correlation between all the 24 hours. The day-ahead price has a clear block structure. Night hours from 9 pm to 8 am are highly correlated. The night hours are also correlated with the hours right after noon. The early and late working day hours are cross correlated, but exhibit less correlation with the rest of the day.

Furthermore, we see from Figure 5.2b that there is significant autocorrelation for each of the hours from one day to the next. Late night and morning hours seem almost non-stationary, but the augmented Dickey-Fuller test rejects a unit root for all hours indicating stationarity in the long term. Autocorrelation is much stronger for night time hours, but weekly seasonality is more prominent for day time hours. Heteroscedasticity is obviously present in some periods. Engle's ARCH LM test show that there are in fact ARCH effects, except for in hours 8-10. Model considerations will be discussed in subsequent chapters.

Figure 5.2: Day-ahead price (a) cross-correlogram and (b) hourly autocorrelation 2014-2017 NO3

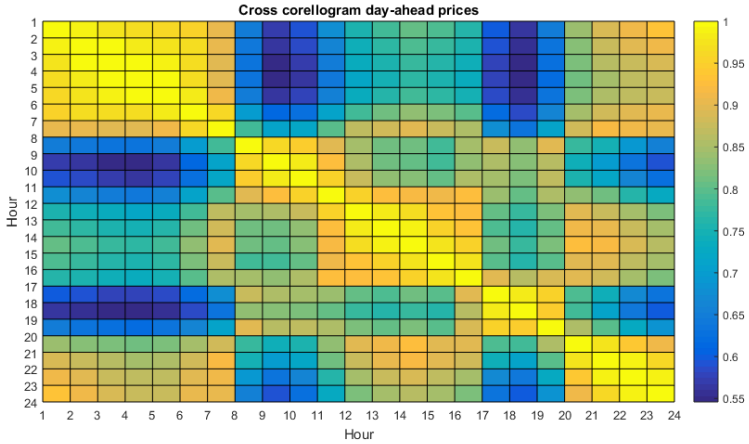

(a) Cross-correlogram

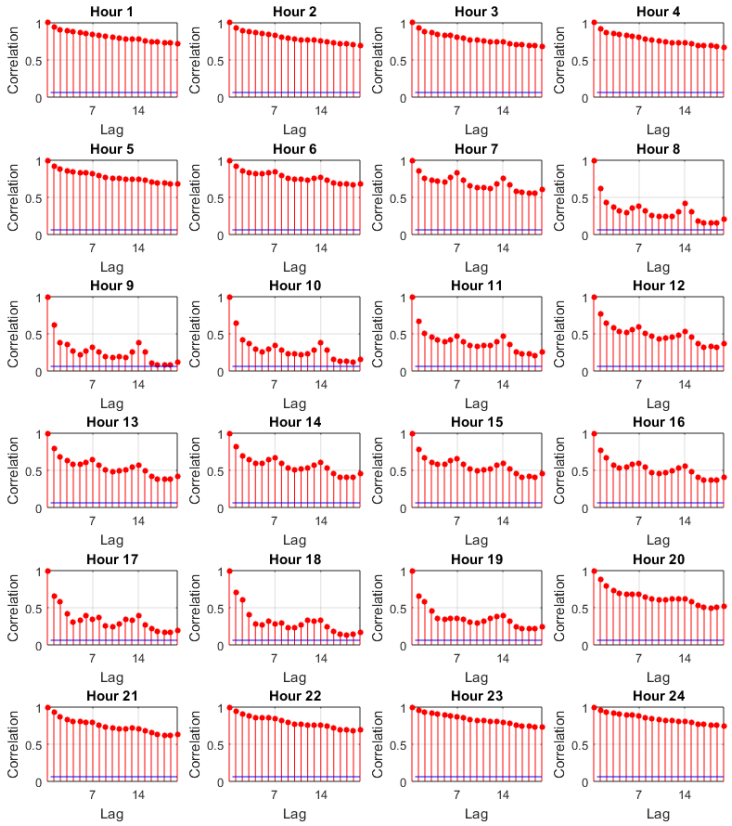

(b) Autocorrelation

5.3 Primary reserve price

The daily average of prices in NO3 for the period 2014-2017 can be found in Figure 5.3a. We note that the price behaviour is asymmetric about the mean of 40 NOK/MWh, exhibiting large spikes above the mean and more stable, lower prices below. The price is generally higher, and has more spikes during early summer. Hydrology, e.g. reservoir filling and inflow, are important factors. It seems that the day ahead price and primary reserve price are negatively correlated in this period. In general, snow melting and high rates of inflow forces high rates of discharge from the reservoirs. This gives a downward pressure on the day ahead price for energy delivery. However, capacity reservation and spinning reserves is connected to the forecasted difference between production and consumption that may be harder to estimate precisely at this time of year. The weekly prices are plotted in Figure 5.3b. We see the opposite pattern that we did for day-ahead prices, with higher reserve prices during the weekends.

Once again we would like to investigate the hourly price characteristics of the electricity market at hand. Table 5.2 shows descriptive statistics for the primary reserve price in all hours of the day. Notice the development of the mean throughout the day. Prices are high during night and morning hours, and generally lower during the rest of the day. This is in strong contrast to the behaviour of day-ahead prices. The start and end of the working day does not seem to affect the primary reserve price at all. The median lies below the mean for all hours. The distributions are clearly fat tailed and skew, which can also be seen from the skewness and kurtosis. Surprisingly, the highest skewness and kurtosis occurs in the afternoon and evening, despite the low price level. We see from the minimum and maximum prices that primary reserve prices, just as day-ahead prices, span a broad range. The highest recorded price is 700 NOK/MWh of reserved capacity, which is shockingly high because there is no actual delivery of power should there not be demand. Nevertheless, this is about three times the average of the day-ahead price, which includes power delivery. Once again we conclude that the hourly prices behave quite differently from hour to hour.

Primary reserve bids are made at the same time for all hours, just as for the day-ahead market. Hence, we expect there to be significant correlation between the hours. The correlation structure between all hours is plotted in Figure 5.4a. The primary reserve market also has a block structure correlation, but there are only two blocks: Day/evening hours and night/morning hours. This makes sense compared to the descriptive statistics that show different primary reserve price behaviour during day and night time.

To investigate the hours separately, the reader is referred to Figure 5.4b. The autocorrelation seems weaker than for day-ahead prices for all hours. The autocorrelation is generally stronger during night time. We observe a weekly seasonality in the prices. The augmented Dickey-Fuller test rejects the null hypothesis of a unit root for all hours, thus indicating stationarity. Engle's ARCH LM test shows that there are ARCH effects in about two thirds of the hours.

Figure 5.3: Daily and daily averaged primary reserve prices

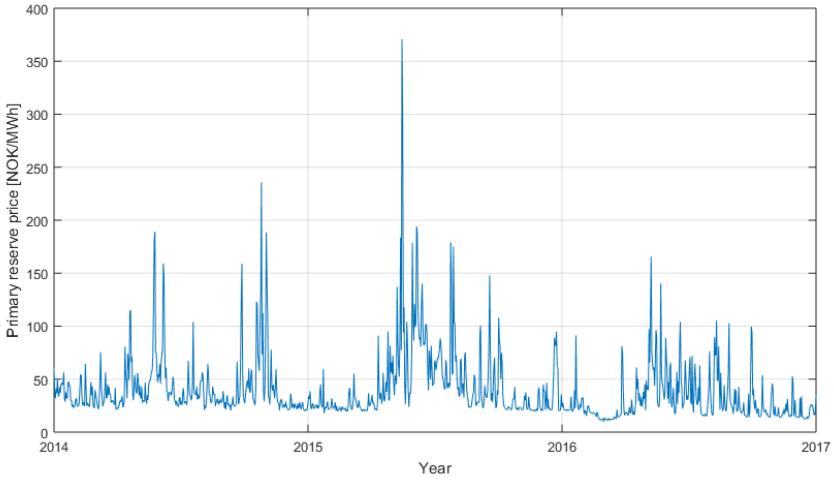

(a) 2014-2017

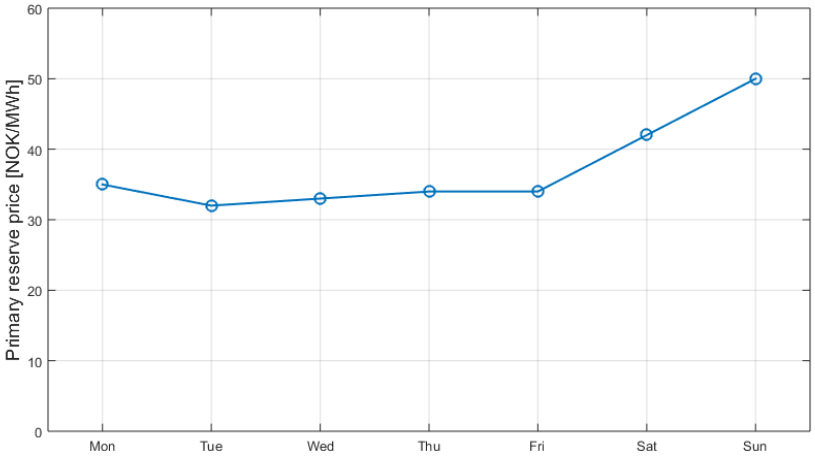

(b) Monday-Sunday

Table 5.2: Descriptive statistics for the primary reserve price

Hour	Mean	Median	St. dev.	Skewness	Kurtosis	Min	Max
00 - 01	62.04	40.00	58.39	3.24	19.72	10.00	700.00
01 - 02	63.48	40.00	59.21	2.52	8.69	10.00	450.00
02 - 03	64.95	41.00	62.21	2.58	9.21	10.00	500.00
03 - 04	64.51	44.00	60.71	2.51	8.59	10.00	450.00
04 - 05	64.18	40.00	60.64	2.48	8.64	10.00	500.00
05 - 06	55.38	35.00	52.99	2.65	9.20	10.00	400.00
06 - 07	41.61	25.00	41.99	3.58	16.75	10.00	380.00
07 - 08	36.68	25.00	36.65	4.47	27.28	10.00	380.00
08 - 09	35.32	25.00	31.51	4.31	27.21	10.00	380.00
09 - 10	33.37	24.00	28.75	4.61	32.27	10.00	380.00
10 - 11	32.29	24.00	25.93	4.27	25.69	10.00	300.00
11 - 12	31.75	24.00	27.21	5.72	51.48	10.00	380.00
12 - 13	31.78	24.00	28.28	6.85	83.27	10.00	500.00
13 - 14	31.82	24.00	28.88	6.70	77.91	10.00	500.00
14 - 15	31.55	24.00	28.34	6.89	83.65	10.00	500.00
15 - 16	31.29	24.00	28.33	6.97	84.39	10.00	500.00
16 - 17	31.24	24.00	28.43	7.13	86.34	10.00	500.00
17 - 18	31.59	24.00	28.62	6.82	81.03	10.00	500.00
18 - 19	31.12	24.00	28.43	7.61	94.34	10.00	500.00
19 - 20	31.06	24.00	25.36	5.23	41.91	10.00	350.00
20 - 21	31.69	24.00	28.59	7.77	97.70	10.00	500.00
21 - 22	33.39	24.00	31.32	6.33	65.71	10.00	500.00
22 - 23	36.95	25.00	35.05	4.91	40.03	10.00	500.00
23 - 00	47.25	29.00	46.92	2.80	9.47	10.00	380.00

Figure 5.4: Primary reserve price (a) cross-correlogram and (b) hourly autocorrelation 2014-2017 NO3

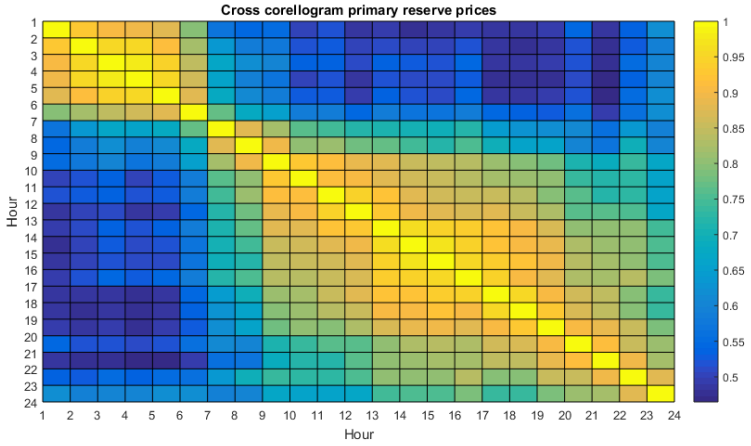

(a) Cross-correlogram

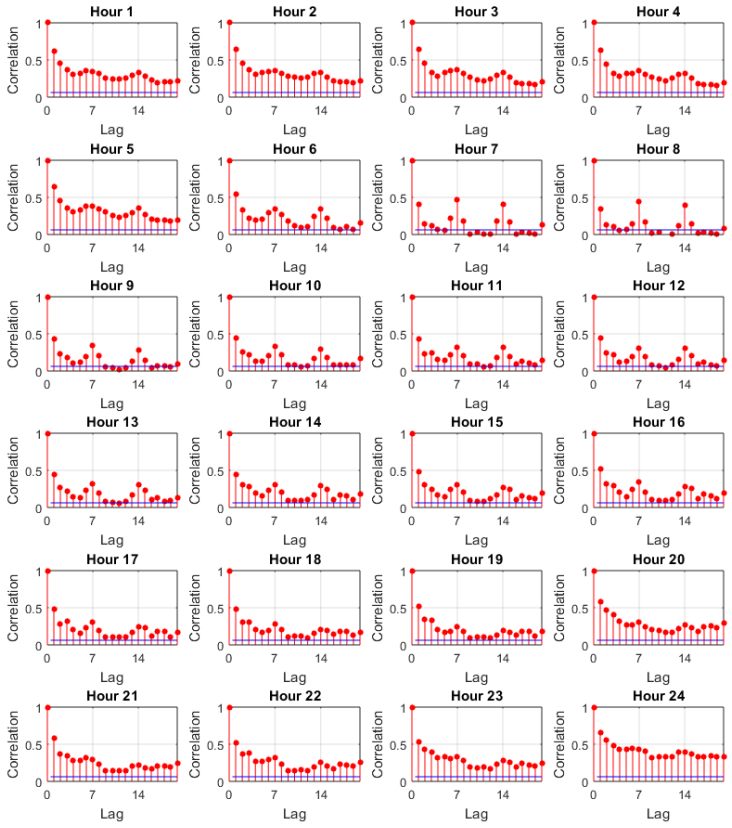

(b) Autocorrelation

5.4 Balancing market

The balancing market is fundamentally different from the two other markets in this thesis. First of all, prices are not quoted day-ahead for all hours. The balancing market is a real time market in which trading happens in hourly blocks, and bids are submitted until 45 minutes before the operating hour. Secondly, demand arises in the balancing market only when there is an imbalance between production and consumption. Primary reserve obligations on the other hand, are capacity reservations intended to normalize small and sudden frequency deviations. If the imbalance between production and consumption is larger, then balancing reserves are activated. The producer therefore faces the risk of not being dispatched. In addition, the balancing price is unknown until after operating time, because the TSO will activate reserves until regulating demand is met in real time.

In general, we expect the balancing market to be highly random and unpredictable. The reason for this is that expected power demand will be priced into the day-ahead market price at closure of this market. Hence, expected deviation between supply and demand should be zero. This discussion argues that the balancing market should be highly random and hard to forecast before day-ahead market closure. This is also supported by empirical studies on balancing market forecasting, see for example Klæboe et al. 2013.

5.4.1 Balancing market regulating states

There are three possible states in the balancing market, all mutually exclusive; we shall use the term *regulating states*. First, if there is no deviation between production and consumption, there is no need for regulation, and the market is in a *no regulation state*. If production exceeds consumption, the system is in a *downward regulation state*. The contrary holds for *upward regulation*, i.e. when consumption exceeds production and there is excess demand for power.

There are most hours of no regulation in the data set. When regulation occurs however, it usually occurs in clusters of hours. This means that consecutive hours are often dominated by the same regulation state. Intuitively this makes sense because an imbalance in production and consumption may be caused by an unforeseen event of some duration. In hours with no regulation, no traded quantity nor balancing premium exists. Hence, analysis on these parameters is complicated. The analyst must make a consistent choice of how to deal with hours of no regulation. This is necessary to assess the autocorrelation structure, which will be done as we choose model in Section 6.6.

Figure 5.5 shows the probability of the system being in some regulating state for all hours during a day. We note that no regulation is the most likely state of the system, and that this probability is almost constant throughout the day. Hence, market participants tend to forecast power demand with consistent precision even though hours towards the end of the day is further away in time at the time of bid submission. We note that

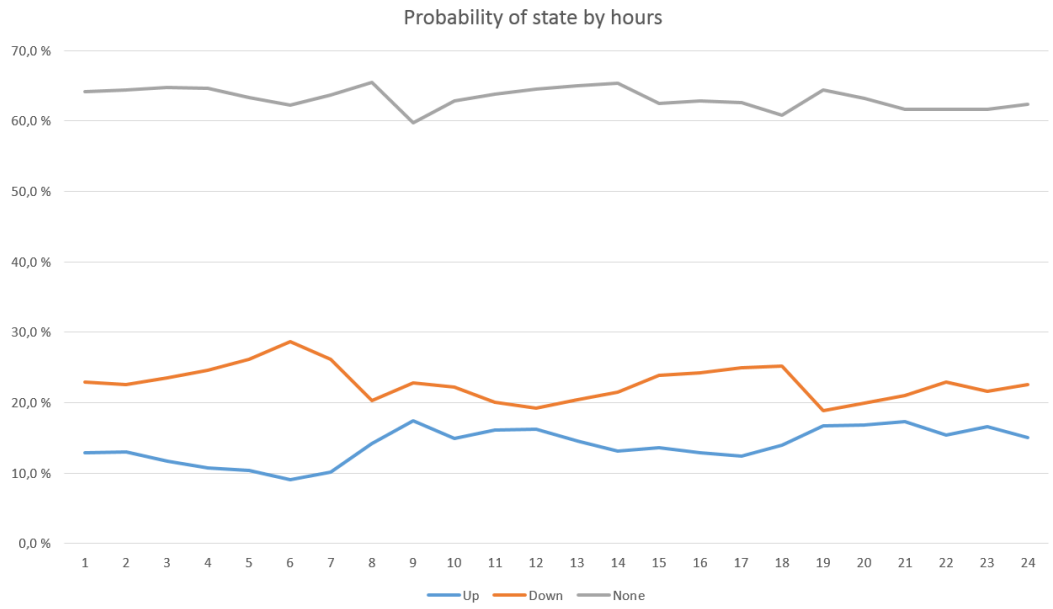

Figure 5.5: Probability of regulation states during the operating day

there is a slight dip in probability of no regulation at time 8-9 am and 5-6 pm. These hours exhibit higher volatility in demand, likely due to the start and end of the working day. Up regulation also seems to be more likely in these hours. Two main regimes are identified: midnight to around 8 am when down regulation is significantly more likely than up regulation, and the rest of the day when the difference is less prominent.

5.4.2 Balancing market volume

Regulating demand is associated with a non-zero regulating quantity in either direction. We refer to this quantity as the regulating- or balancing volume, simply because it is industry jargon. Figure 5.6b displays the balancing volume in both directions in 2014-2016. We observe that the balancing volumes are highly volatile, with tremendous spikes in the data. There is no clear annual seasonal pattern, but we note that there are periods for which balancing is more or less likely. For instance, in late 2014 there are few hours of downward regulation accompanied by low volumes, but in late 2016 the opposite is the case. Figure 5.7 shows the historical distributions of volumes and premiums. Note that the distributions are fat-tailed and highly skewed, looking somewhat like Weibull distributions for different shape parameters.

The daily distribution of balancing volumes are plotted in Figure 5.8. Investigate first the plots labeled excl. These indicate the average hourly volume in non-zero hours, i.e. it is an average over the hours that are in either upward- or downward regulation. We notice that the average size of a regulation event in both directions is approximately 60-70 MWh independent of the hour of day, except for the nightly

Figure 5.6: Historical prices and traded quantities in the balancing market

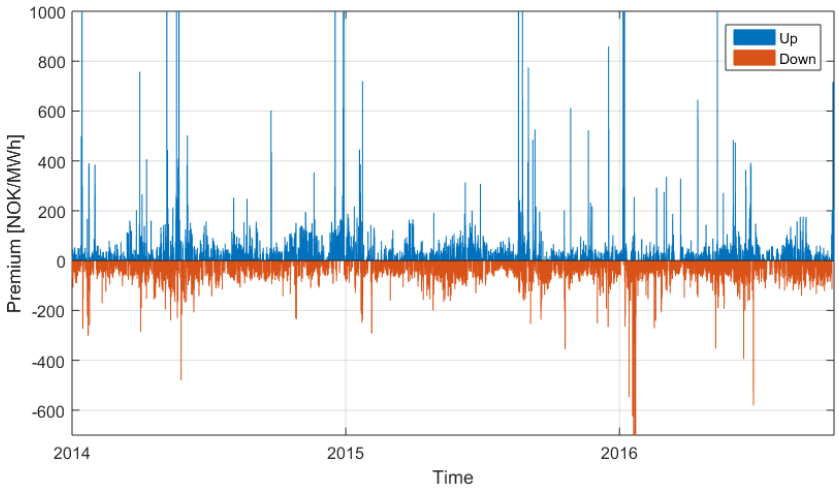

(a) Balancing premium

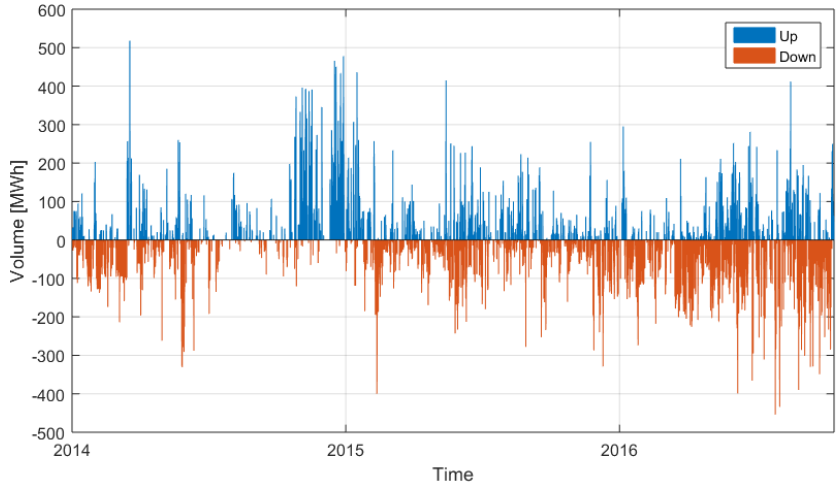

(b) Balancing volume

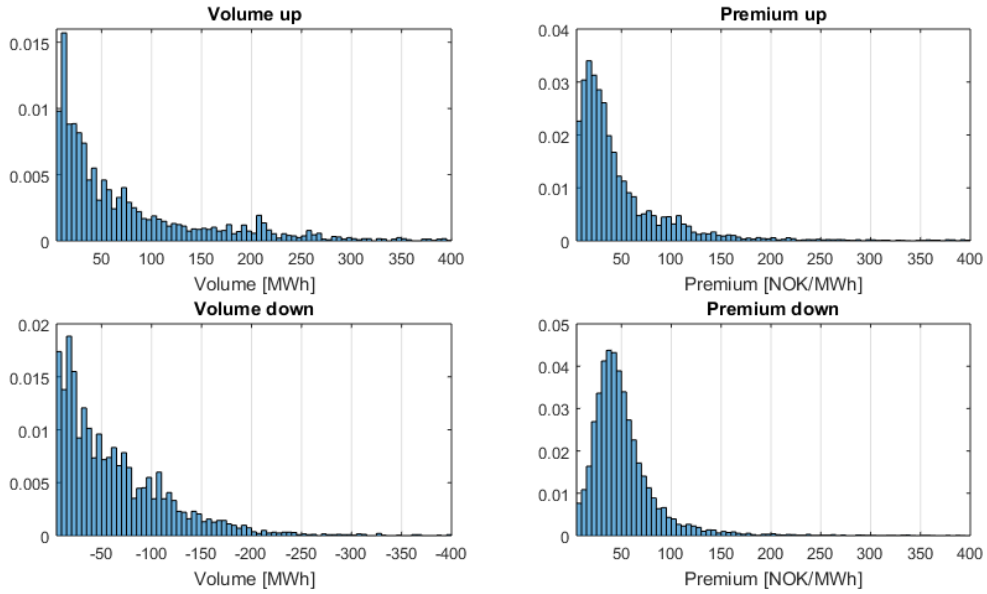

Figure 5.7: Historical volume and premium distributions

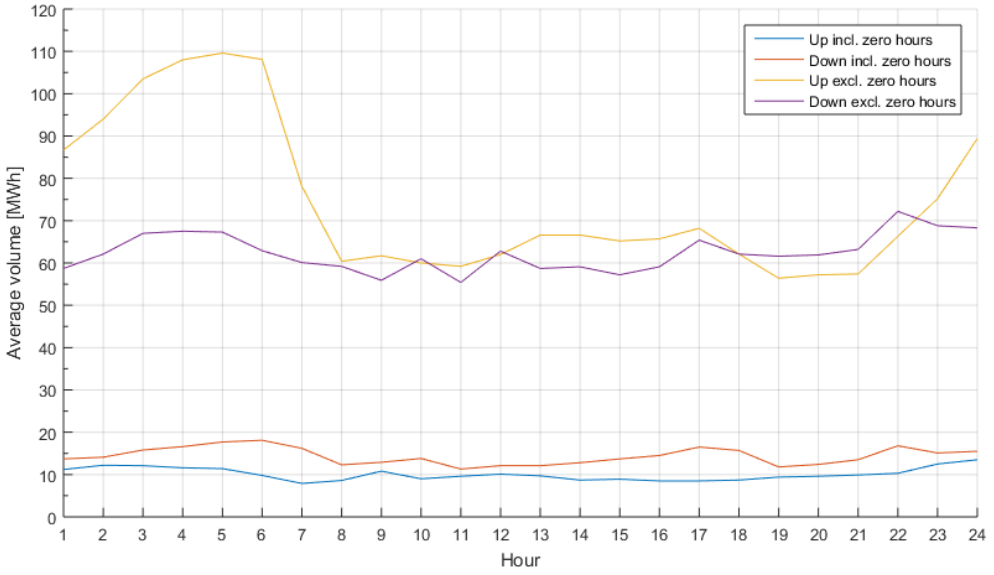

Figure 5.8: Average hourly balancing volumes

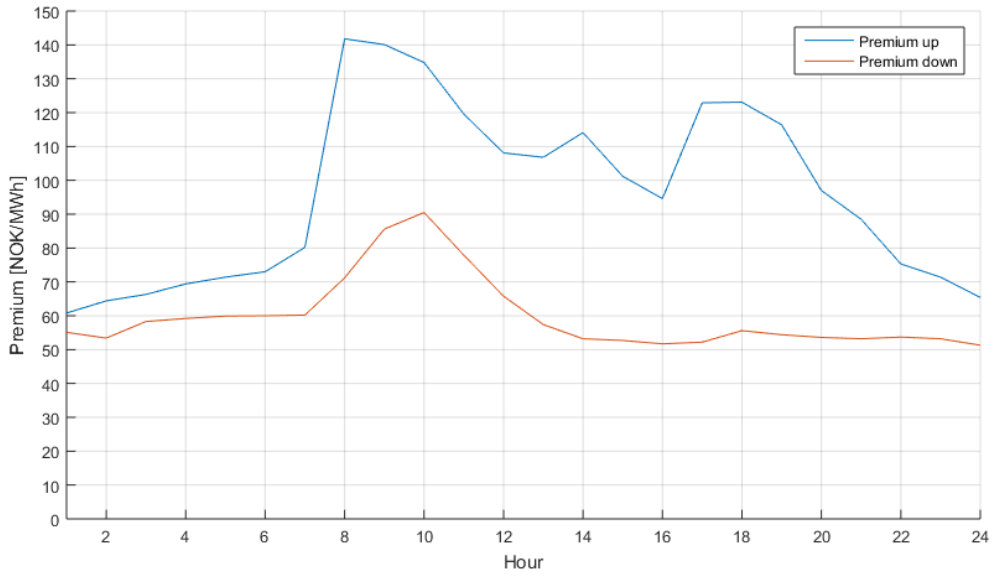

Figure 5.9: Average hourly balancing premiums

upward volumes. Between 11 pm and 7 am average upward regulation volumes are significantly higher. The plots labeled incl. show the hourly average volume including hours without regulation, i.e. it is the product of the excl. plots, and the probability of regulation occurring cf. Figure 5.5.

5.4.3 Balancing market premiums

The balancing market premiums exhibit a behaviour much alike that of the balancing volumes. However, they are in general more spiky, see for example Figure 5.6a. Once again it is hard to identify a clear annual seasonal pattern. The premiums seem to be more stable than volumes around its mean and less exposed to large changes in mean and variance in time (heteroscedasticity). The augmented Dickey-Fuller test furthermore rejects the null hypothesis of a unit root in premiums and volumes, thus indicating stationarity of the series in the long term. The hourly average premiums of hours with regulation can be seen in Figure 5.9. Once again we see that two main regimes are present: night and day hours. During the working day and the evening up regulation premiums are higher than during night time. Down regulation has a premium peak during the start of the working day.

Figure 5.10: Scatter plots for the balancing volumes and premiums

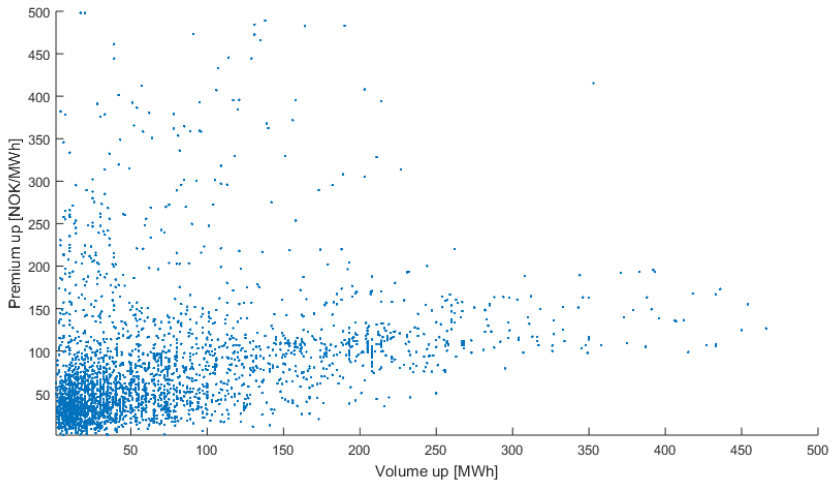

(a) Upward regulation

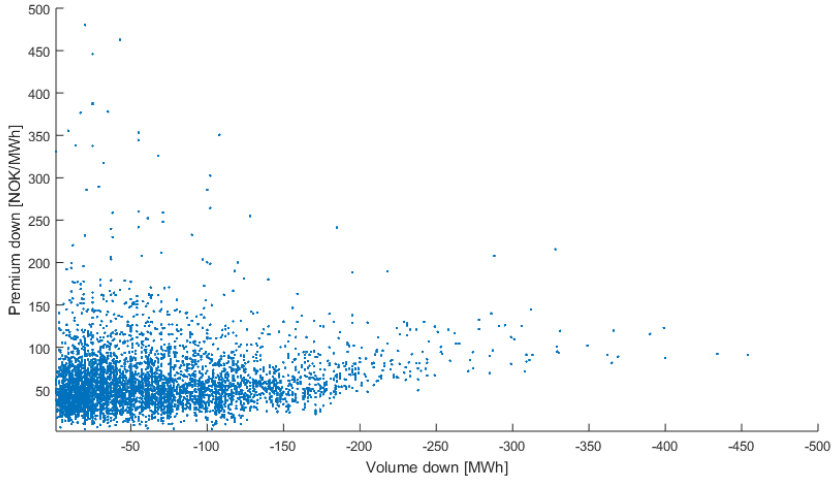

(b) Downward regulation

5.4.4 The relationship between balancing market volumes and premiums

Finally, we look at the concordance between volumes and premiums in the balancing market. The reader is asked to study Figure 5.10a and 5.10b. The figure plots data points for hours with regulation only. We notice a weak linear relationship in the data. This is confirmed by computing Pearson's correlation coefficient between premium and volume, giving 44 % and 37 % correlation for upward and downward regulation respectively. It will be an important goal to retain this relationship in subsequent modelling.

To conclude, we have found that regulating volumes and premiums are highly stochastic and that spikes occur frequently in the data. The data seems to be stationary. There is no clear annual seasonal pattern. There are two prominent regimes during the day; late night and morning hours, and day hours. State probability, volumes and premiums have different characteristics in these hours. There is a positive correlation between volumes and premiums.

5.5 Intermarket dependence

Relationship between day-ahead market and primary reserves

As pointed out previously, there is a relationship between the day-ahead and primary reserve prices. Figure 5.11a plots the hourly average prices from Table 5.1 and 5.2 in the spot and primary reserve market. The primary reserve price is high during night hours when the day-ahead price is low. During the working day the reserve price falls, while day-ahead price reaches its maximum. We also saw from Figure 5.1b and 5.3b that there is a weak negative correlation between day-ahead and primary reserve prices on a weekly basis. The scatter plot in Figure 5.11b shows the relationship between the prices for hours in the data set. The relationship is clearly non-linear. Computing Pearson's correlation coefficient for all the data gives a correlation of -26.2 %. However, if we divide the data into two sets: one set of prices for which the spot price is above 350 NOK/MWh, and one below, we get two very different results. For the above set, correlation is 43.9 %, and for the below set it is -35.1 %. To investigate whether this relationship in the prices is a function of the hour of day, we plotted scatters for the hours separately. These revealed no obvious relationship between the prices during day time, and a negative non-linear relationship during night time. A linear regression model between day-ahead and primary reserve price would therefore not be appropriate.

Relationship between day-ahead market and balancing reserves

The connection between the day-ahead price and the balancing market must be evaluated by investigating the regulation state, volumes and premiums. We start by

Figure 5.11: Relation between day-ahead and primary reserve markets

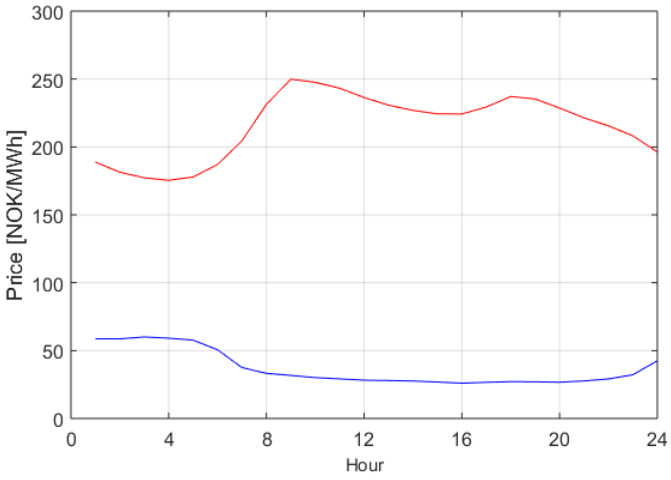

(a) Hourly spot (red) and primary reserve price (blue)

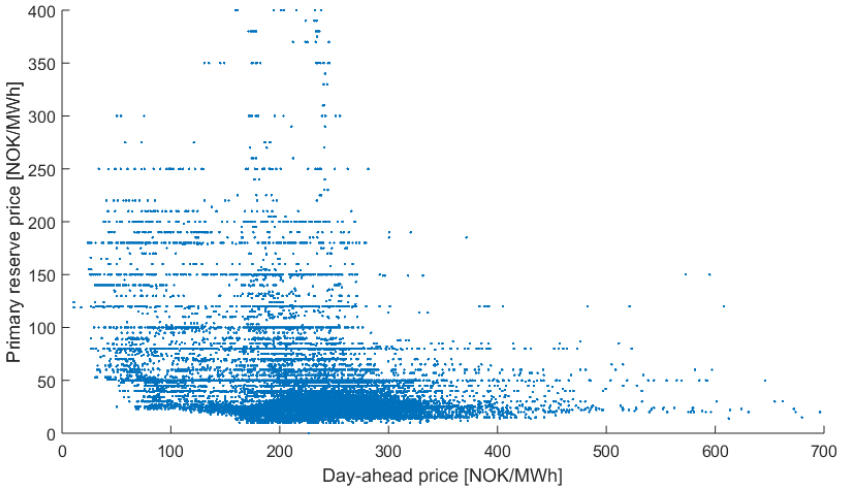

(b) Spot and primary reserve price

Table 5.3: Dependence relations among the markets

	Primary reserve Price	**Balancing** Reg. state	Reg. volume	Reg. premium
Day-ahead	Negative non-linear dependence	No explicit dependence	No dependence	Positive linear dependence
Primary reserve		Conditionally independent[2]	Conditionally independent	Conditionally independent

analysing the day-ahead price for the different regulation states. The prices are distributed very similarly for all regulation states, with a somewhat lower price level for down regulation. We therefore regard regulation state as independent of the spot price.[1]

Regulating volumes and the connection to the day-ahead price is evaluated by scatter plots, and Pearson's correlation coefficient. The correlation is computed only for hours with respective regulation. The scatter plots show no obvious relationship between the quantities. Correlation is calculated at 5.2 % for upward regulation and 1.6 % for downward regulation. Spot price and regulation volumes are therefore regarded as independent in this thesis.

Regulating premiums show correlated behaviour with the day-ahead price. Correlation is measured to be 35.9 % and 64.7 % for upward and downward regulation respectively, including only regulation hours. A scatter plot furthermore gives an impression of the linearity of the dependence. Figure 5.12 shows a scatter plot of the day-ahead price and premiums for downward regulation. Notice that the relationship exhibits weak linearity. The same is the case for upward regulating premiums.

Relationship between reserve markets

Finally, we will discuss the relationship between the reserve markets. Histograms of the primary reserve price has been plotted for all regulation states, with no significant change in distributional properties. Scatter plots and correlation coefficients for primary reserve price and regulating volumes and premiums during respective regulation, shows no relation between primary reserves and regulating reserves. The markets are thus regarded as independent. To sum up, Table 5.3 shows the dependence relations discovered in this thesis.

[1] At least explicitly - the next section will propose an hour specific Markov model for state determination, which then will result in an implicit connection to the spot price since downward balancing is more likely in the morning when spot prices are low. See section 6.6

[2] Conditional on day-ahead price

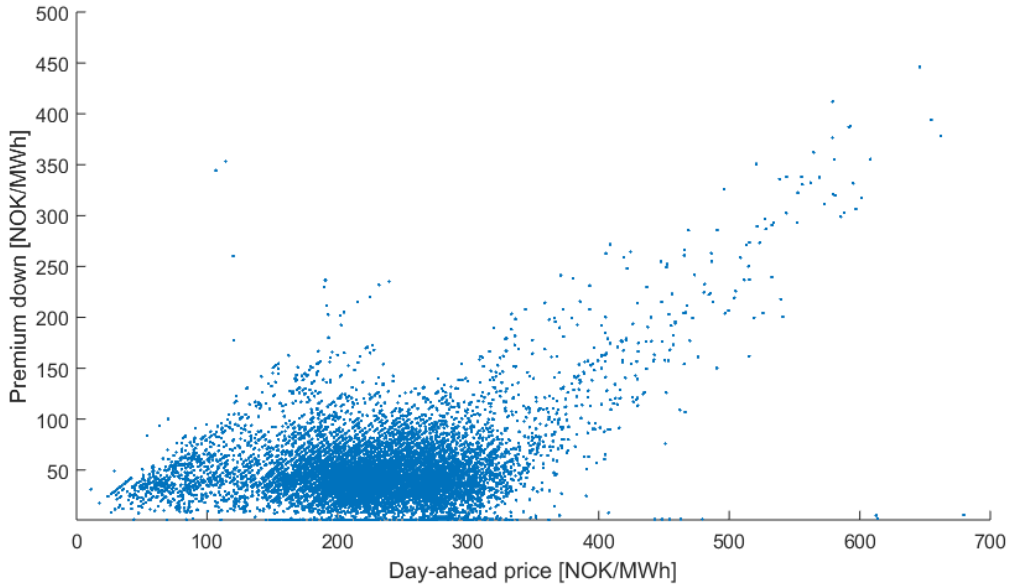

Figure 5.12: Scatter plot day-ahead price and downward premium 2014-2017

Chapter 6

Modelling the markets

6.1 Introduction

In this section the goal is to arrive at a set of models to be used in scenario generation outlined in Chapter 7. These scenarios should have a link to forecasting to be relevant for the day of planning. Good point forecasts are important to reduce uncertainty in the stochastic parameters. In addition, good probabilistic forecasts are necessary to achieve a reasonable description of the uncertainty inherent in the forecast. The quality of the forecast determines the quality of the scenarios. We therefore start off with a brief introduction to electricity price forecasting.

Most hydropower producers have in-house proprietary models for forecasting day-ahead prices. To our knowledge there is no commercial tool for forecasting primary reserve prices. We do not have access to such fundamental forecasting models. We therefore forecast prices directly with autoregressive models using lagged prices.

This chapter starts with a discussion on modelling of day-ahead and primary reserve prices which are modelled in the same way. We end up proposing a panel data model using quantile autoregression to provide probabilistic forecasts for these market prices. Non-linear dependence is modelled with a copula based heuristic. Thereafter, we discuss models for the balancing market. Regulating states are modelled with an hour-specific Markov model. We use a resampling technique to account for hours with no regulation. Thereafter we fit one-dimensional AR/ARMA models to volumes and premiums. To aid the reading of this chapter Figure 6.1 illustrates the final modelling framework we arrive at. Explanations and justifications will be given throughout this chapter.

6.2 Selection of model family

Weron 2014 makes a review of the most frequently used model families in electricity price forecasting. The author groups models into five categories: Fundamental, reduced-form, multi-agent, computational intelligence and statistical models. As many other

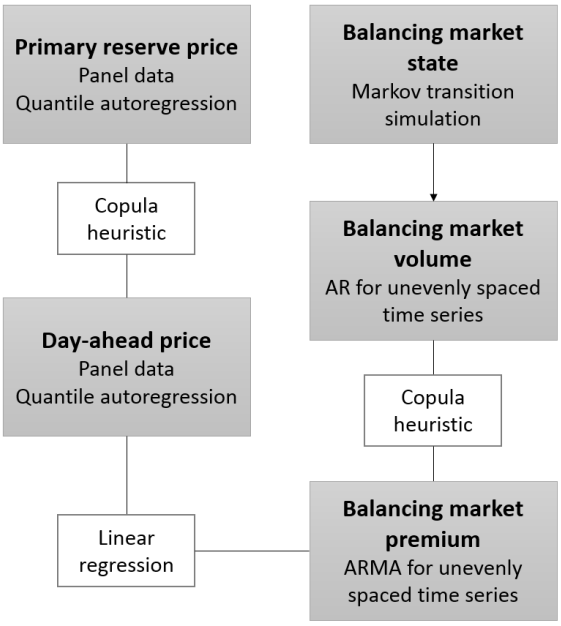

Figure 6.1: Final model overview

authors, due to short term prediction performance, we choose to proceed with the category of statistical models.

6.2.1 ARIMA-models

GARCH-models (General Auto-Regressive Conditional Heteroscedasticity) are intended to deal with heteroscedasticity, which is an occuring phenomenon in electricity market prices. However, application of various ARIMA-models (Auto-Regressive Integrated Moving Average) appears traditionally more frequently in the literature. Weron pointed out that "Although electricity prices exhibit heteroskedasticity, the general experience with GARCH-type components in electricity price forecasting models is mixed. There are cases where modeling heteroskedasticity is advantageous, but there are at least as many examples where such models perform poorly". This might hold especially for short term forecasting, because even if variance changes annually, heteroscedasticity is less prominent in the short term. We therefore choose to proceed with ARIMA models in this thesis. ARIMA will be used to model the balancing market. Boomsma et al. 2014 and Eriksrud and Braathen 2012 have previously used ARIMA models for the balancing market.

Mathematical structure of ARIMA-models

Here follows a mathematical definition of a general SARIMAX-model, which is an ARIMA-model with one or more seasonal AR-terms or MA-terms and one or more

exogenous explanatory variables. ρ_{mh} is the price in market m in hour h that is forecasted. B is the backshift operator, such that $B^k\rho_h = \rho_{h-k}$, and B^s is the seasonal backshift operator, that backwards shifts by a season instead of a unit. Let $\phi(B)$ be the autoregressive polynomial, and $\Phi(B^s)$ the seasonal autoregressive polynomial. Further, we denote the difference operators ∇^d and ∇_s^D, regular and seasonal respectively, with order d and D. Let $\theta(B)$ and $\Theta(B^s)$ represent the moving average polynomials, where the subscript s denotes seasonality as before. ϵ_t denotes the white noise error term, and $\psi(B(n_k))$ the polynomial of the exogenous variable u_{mh}, where $B(n_k)$ is the backshift operator with n_k in system dead time i.e. the delay between an input signal and the response in output. C is a constant.

$$\Phi(B^s)\phi(B)\nabla_s^D\nabla^d\rho_{mh}(t) = \Theta(B^s)\theta(B)\epsilon_{mh}(t) + \psi(B(n_k))u_{mh}(t) + C \tag{6.2.1}$$

where ∇_s^D=$(1-B^s)^D$ and $\nabla^d = (1-B)^d$.

To correctly specify which terms should be included in an ARIMA model, there are two widely applied methods. They are the Box-Jenkins methodology (Box et al. 2007) and minimization of some information criterion e.g. Bayesian or Akaike. In this thesis we follow the Box-Jenkins methodology and utilize the ACF and PACF of the prices to aid identification.

Innovation terms in standard software-packages for ARIMA estimation are often set by default to Gaussian. However, other distributions can be used, for example the Student's t-distribution. Also other tailor made estimation techniques can be used based on the distribution of the innovations.

6.2.2 QR-based models

Another relevant statistical model is quantile regression (QR) as introduced by Koenker and Bassett 1978. QR can be used in a variety of ways.

Fundamental quantile regression models

Hagfors et al. 2016 developed a fundamental quantile regression model for the UK (APX) day-ahead price. They showed that sensitivity towards fundamental factors and the lagged price changes across quantiles. Do and Molnar 2015 similarly found that the effect of fundamental factors vary substantially across quantiles for the German day-ahead electricity price. Regular regression will therefore not capture the full dynamics in the price.

Jonsson et al. 2014 propose a time-adaptive quantile regression framework to arrive at a semi-parametric description of the predictive density of the day-ahead price in Denmark. Exponential distributions are fitted to the tails due to higher uncertainty in tail estimation. The QR model is benchmarked against a GARCH model with

Gaussian innovations. Upon evaluation the parametric GARCH model seems to be unfit for probabilistic forecasts, while the QR model performs better.

Quantile regression averaging

Weron and Misiorek 2008 perform a comparison of 12 different parametric and semi-parametric models for spot point predictions and prediction intervals (PI). In terms of PI forecasting, and conditional coverage, the semi-parametric models clearly outperform the parametric models, and especially in periods with heteroscedasticity.

In a later paper Nowotarski and Weron 2013 use a quantile regression averaging (QRA) framework to calculate prediction intervals by weighting point forecasts from those same 12 models. The results indicate an enhanced performance. Maciejowska et al. 2016 further develops the model to orthogonalize input point forecasts by principal component analysis (PCA) in the QRA model, yielding the Factor QRA (FQRA). The FQRA model shows even more promising results in terms of interval forecasts.

Quantile autoregression

Quantile autoregression (QAR) captures systematic influence of previous prices on the conditional distribution of the price. It has been used for scenario generation before, see for example Tomasgard and Høeg 2005. This approach is valuable because price variance can be related to the forecasted value, e.g. high forecasts might yield a higher variance of innovations. The conditional price distribution can also be asymmetric, which is the case for electricity prices. Such dynamics cannot be modelled by ARIMA models with error terms of constant variance. In addition, the methodology is easily implemented without the need of fundamental forecasting models. We therefore choose to proceed with QAR-models in this thesis, and will evaluate such model's ability to provide probabilistic forecasts for the day-ahead and primary reserve market.

Mathematical structure of quantile autoregression models

For an introduction to quantile autoregression, the reader is referred to Koenker and Xiao 2005. We will provide details on estimation and implementation in this thesis in Section 6.4. Here we will only outline the concept of quantile autoregression to ease the reading of Section 6.3. As Koenker points out these models can capture systematic influences of conditioning variables on location, scale and shape of the conditional distribution of the response. This provides an extension to constant coefficient time series in which the the effect of conditioning is confined to a location shift.

The basic concept of quantile regression is to estimate coefficients in the autoregressive model such that, for a given quantile τ, the conditional quantile $Q_{\rho_{mh}(t)}(\tau \mid \rho_{mh}(t-1), ..., \rho_{mh}(t-p))$ has the property that $Prob(\rho_{mh}(t) \leq Q_{\rho_{mh}(t)}(\tau)) = \tau$. QR does not forecast the distribution itself, but rather each quantile directly. Hence, for a fine discretization of $\tau \in \langle 0, 1 \rangle$ $Q_{\rho_{mh}(t)}(\tau)$ approximates the conditional density of the

Table 6.1: Day-ahead and primary reserve characteristics and corresponding modelling choices

Findings from empirical study	Modelling
Different day-ahead/primary reserve hourly characteristics and information structure	Panel data
Autocorrelation and stationarity	Autoregressive models
Weekly seasonality	Seasonal term
Price variance depends on price level [1]	Quantile autoregression
Day-ahead/primary reserve non-linear dependence	Copula heuristic

price for hour h on day t. Some interpolation scheme and tail modelling will be needed, this is discussed in section 6.4. The QAR(p) model can be written (in terms of the notation in this thesis)

$$Q_{\rho_{mh}(t)}\Big(\tau|\rho_{mh}(t-1),...,\rho_{mh}(t-p)\Big) = \phi_{mh0}(\tau) + \phi_{mh1}(\tau)\rho_{mh}(t-1) + ... + \phi_{mhp}(\tau)\rho_{mh}(t-p) \tag{6.2.2}$$

6.3 Day-ahead and primary reserve price modelling

Day-ahead and primary reserve prices are very similar in several respects. Both markets are cleared day-ahead and all 24 hourly prices are quoted at the same time. In addition, the individual hours in each market have different statistical properties. Finally, both market prices are stationary, highly autocorrelated and has a similar underlying structure, see Section 5.1. We therefore aim to model day-ahead and primary reserve prices in a similar manner.

In Chapter 5 we discovered a set of market properties that we wish to retain in the modelling. Table 6.1 lists the important findings, and how each of these will be modelled. We will explain why these are important properties and provide details on modelling throughout this section.

The general concept is illustrated in Figure 6.2 and is as follows. First we use a combination of panel data and quantile autoregression to achieve predictive price densities (probabilistic forecasts) on an hourly basis for both markets. These densities (marginals) shall represent the underlying uncertainty in the forecast. Because hourly prices are correlated, innovations will be correlated between hours and markets. Our

[1] This was not treated in Chapter 5 but will be discussed in this section and in section 6.4

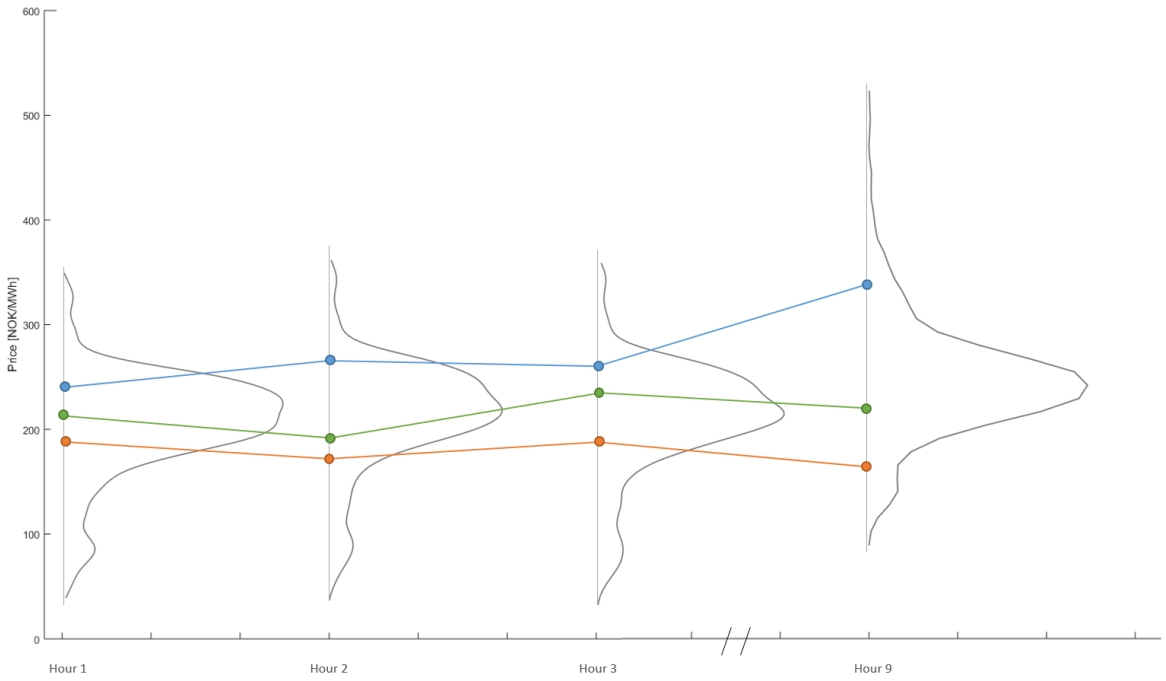

Figure 6.2: Hourly price densities constructed from quantile regression and example price paths constructed from the copula heuristic

goal is to connect daily price paths from all 48 densities, such that both intermarket and inter-temporal dependencies represent the true behaviour. For this we use a copula based scenario generation heuristic. Section 6.4 and 6.5 are devoted to a description of quantile regression and the copula based method respectively.

6.3.1 Choosing a panel data model

Autoregressive models are intended for stochastic processes in which the next observation depends on a combination of previous observations. As time progresses, the previous observations seize to influence the most recent observations gradually. Nearby observations are therefore more strongly correlated than observations further apart in time. However, both day-ahead and primary reserve prices are quoted at the same time for all hours in a day. The information set that the market participants utilize is the same for all the bidding hours. It is therefore unsound from a methodological perspective to model the sequence of day-ahead and primary reserve prices as a one dimensional hourly time series. The prices should rather be modelled as 24-dimensional panel data, with discrete time increments of one day. This makes intuitive sense if we consider the transition between two consecutive days. The price in hour 1 on day t, should not necessarily be strongly correlated with the price in hour 24 on day $t-1$. In the same manner, hour $h+1$ should not be explained by hour h because it is not

yet an actual observation at the time of forecasting. In addition, day-ahead and primary reserve prices exhibit very different characteristics in terms of variance and mean reversion in each of the hours. Individual models should therefore be fitted to each of the hours independently. In other words, the forecast in market m for hour h on day t is based on $B^k \rho_{mh}(t) = \rho_{mh}(t-k)$.

Huisman et al. 2007 proposed a panel model composed of a deterministic component and a stochastic component,

$$\rho_{mh}(t) = f_{mh}(t) + X_{mh}(t) \tag{6.3.1}$$

the deterministic component accounting for predictable regularities such as mean price and seasonal effects. Huisman builds his deterministic component on a mean price level μ_{m0}, and hourly deviations from the mean price level, μ_{mh}, to allow for differences in mean price levels over the hours. The deterministic component also allows for different price levels for different weekdays of the week. To model this, $I^d(t)$ is a dummy variable that equals 1 if the delivery day t+1 is a weekday d. β_{md} is the difference from the mean price level.

$$f_{mh}(t) = \mu_{m0} + \mu_{mh} + \sum_d \beta_{md} I^d(t) \tag{6.3.2}$$

The stochastic part accounts for the variation of the price around the deterministic component and is modelled as a mean reverting process.

$$x_{mh}(t) = -\alpha_{mh} x_{mh}(t-1) + \epsilon_{mh}(t) \tag{6.3.3}$$

The error term is assumed i.i.d., has zero mean and finite variance. The authors use Seemingly Unrelated Regression (SUR) to estimate the model. Just like Huisman et al. we choose to model prices as panel data. Apart from that, we take a different approach to model the time series. Instead of modelling predictable regularities separate from the stochastic process we include it as a constant in a quantile autoregressive model. As we shall see later, this allows for a comparison of the constant term across quantiles. In addition, we choose to model weekly seasonality by including a seasonal auto-regressive term instead of using the dummy variable $\beta_{md} I^d(t)$. We allow the model to be of higher order than plain mean reverting (AR(1)). As we will see in subsequent chapters, both day-ahead and primary reserve prices in NO3 are in fact auto-regressive processes of higher order.

6.3.2 Identifying the order of the autoregressive process

Before we proceed we must first decide on the necessary order of autoregressive terms in the QAR-model. The PACF of the primary reserve price can be found in Figure 6.3. From the PACF we see that all hours have significant spikes at lag 1. There are also significant spikes at lag 2 and 3, especially for evening hours. Some hours also

have significant spikes at higher lags, but this is generally not the case. Minimization of the Bayesian information criterion further elaborates: almost all hours are of order 3 (exceptions are hours 2-5 am). To be able to perform direct comparison of coefficients in different hours, our choice is to model all hours as the same process order which is not an unreasonable approximation for any of the hours. There are weak seasonal effects at lag 7 in some hours. We thus include the seasonal autoregressive term at lag 7. Upon investigating the residuals ACF after inclusion of 3 autoregressive terms and the seasonal term, all significant serial correlation in the residuals is gone. A similar analysis was done for the day-ahead price, but the PACF plot is left out for brevity. A model containing 3 autoregressive terms and one seasonal term at lag 7 is found appropriate for day-ahead prices as well. Hence, we arrive at the following non-parametric model for the conditional price density, which will be referred to as the QAR^{hourly}-model

$$Q_{\rho_{mh}(t)}(\tau) = \phi_{mh1}(\tau)\rho_{mh}(t-1)+\phi_{mh2}(\tau)\rho_{mh}(t-2)+\phi_{mh3}(\tau)\rho_{mh}(t-3)+\Phi_{mh7}(\tau)\rho_{mh}(t-7) + \phi_{mh0}(\tau), \qquad \tau \in \Big\langle 0,1 \Big\rangle \tag{6.3.4}$$

6.3.3 Extending the basic model

In this section we propose an extension to the QAR^{hourly}-model in 6.3.4 which turns out to provide even better probabilistic forecasts. Even though hourly prices are modelled as individual time series, they are highly correlated. Other hours than h in day $t-1$ may be good explanatory variables for hour h and could be included in the autoregressive framework. Possibly all $h \in \{1, ..., 24\}$ on day $\{t-1, ..., t-p\}$ can potentially be explanatory variables for $\rho_{mh}(t)$. For all the other hours in the panel we restrict ourselves to consider only day $t-1$. Figure 6.4 illustrates.

The hourly prices are highly correlated a given day. It is therefore likely that the true dimension of the panel is much less than 24. In that case, it would be highly impractical to model all 24 hours. We therefore performed a Principal Component Analysis (PCA) of the prices. The first principal component of the day-ahead accounts for 83 % of the variance in the price, and the first principal component of primary reserve prices accounts for 71 %. Scaling and mean centering reveals that this component is very similar to the daily mean of the prices. Figure 6.5 plots the first principal component and the mean (centered and scaled) of the daily prices. This result implies that the daily mean of the prices can be a very efficient explanatory variable. This makes intuitive sense because the daily mean price represents the underlying information set of participants upon bidding. As daily time increments progress, the information set is updated (outlooks on inflow, price, etc). New information to participants is a fundamental driver of the variance of the prices. We therefore choose to test three models: pure one-dimensional autoregression (only the same hour lagged), autoregression with yesterdays mean as explanatory variable, and 24-dimensional regression using all pre-

Figure 6.3: PACF for the primary reserve price

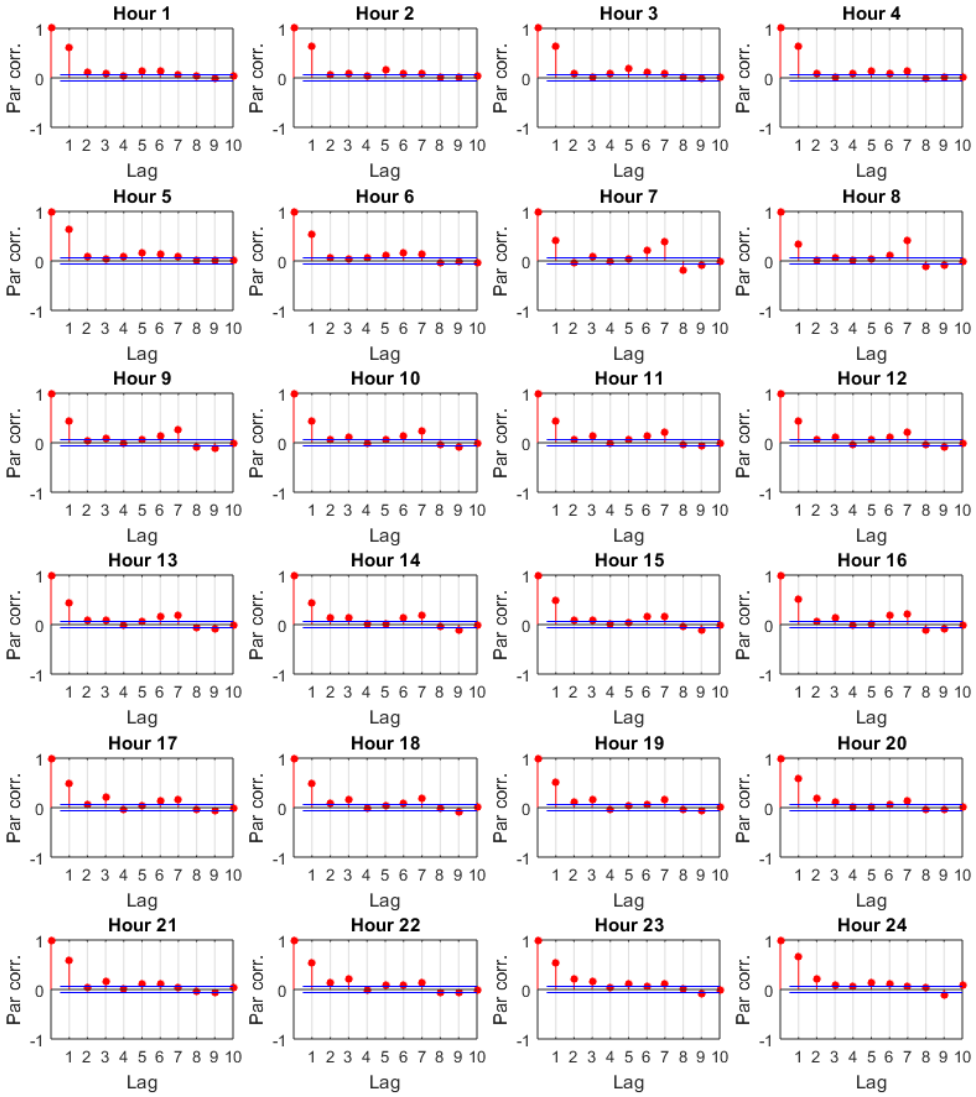

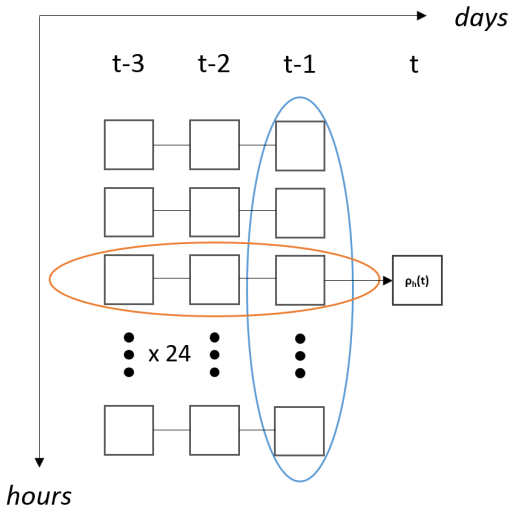

Figure 6.4: Including price information in other hours. The orange ellipse shows daily lagged prices in the same hour, while the blue ellipse shows other hours on the previous day

vious hours.

Testing has been done in two stages. First we check the explanatory power of the variables in-sample (2014-2015) by evaluating different criteria (Sum of squared error, adjusted R-sqaured and AICc), in a stepwise forward and backward regression framework. Thereafter, we evaluate out-of-sample predictive ability of the models (after ordinary AR-estimation).

All criteria for picking variables return models with similar in-sample quality for both day-ahead and primary reserve prices. Most of the hours are included, indicating explanatory power across the day. The average adjusted R-squared is 82.2 % for day-ahead prices and only 47.7 % for primary reserve. Most of the variance in day-ahead prices are well explained during night and evening. Start and end of the working day are much more problematic however, with an R-squared score as low as 65 % in hour 9. For primary reserve, the high-price morning hours score higher than afternoon hours. The adjusted R-squared drops on average 8 % for day-ahead and 13 % for primary reserve in sample when only including the lagged mean.

Out-of-sample testing tells a different story however. The models have been tested for 350 days out-of-sample in terms of root mean square error (RMSE) after ordinary AR-estimation using MLE. Figure 6.6 plots the RMSE during the test period (2016). As expected, predictive performance varies across the day. Working day hours are harder to predict for day-ahead prices, and morning hours are harder to predict for primary reserve prices. The more complex models including information about other hours on day $t-1$ perform marginally better than the pure one-dimensional. This is especially true in the high volatility hours. There is no significant difference in the performance of the model using only the lagged mean, and the model using all lagged hours. The

Figure 6.5: First principal component and daily mean (centered and scaled)

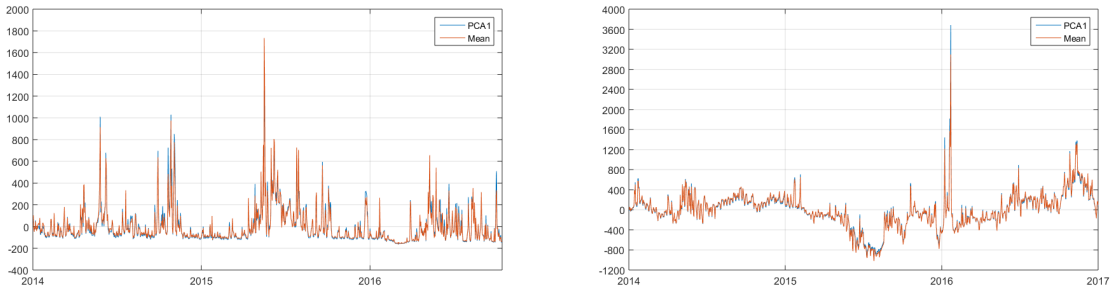

(a) Primary reserve

(b) Day-ahead

lagged-mean actually performs slightly better in hour 9 for day-ahead prices. For the sake of parsimony our choice is therefore to proceed with the lagged mean-model. As we shall see later, this model performs better than the pure one-dimensional QAR^{hourly} model in terms of prediction intervals. Hours with high variance are handled more efficiently by the lagged-mean model.

We have now arrived at the following extension of the quantile autoregressive model, which will be referred to as the QAR$^{lagged-mean}$-model

$$Q_{\rho_{mh}(t)}(\tau) = \phi_{mh1}(\tau)\rho_{mh}(t-1)+\phi_{mh2}(\tau)\rho_{mh}(t-2)+\phi_{mh3}(\tau)\rho_{mh}(t-3)+\Phi_{mh7}(\tau)\rho_{mh}(t-7)$$
$$+ \psi_{mh}(\tau)\mu_{\rho_m}(t-1) + \phi_{mh0}(\tau), \qquad \tau \in \langle 0, 1 \rangle \tag{6.3.5}$$

All coefficients are estimated for $m = \{\text{Day-Ahead}, \text{Primary reserve}\}$ in the manner explained in 6.4. The next goal will be to connect full price paths between the 24 panels of prices in each market, as well as between the markets.

6.3.4 Hourly and intermarket dependence

We have now proposed how to forecast hourly conditional price distributions. Connecting the individual hours into daily price path scenarios remains. As pointed out in Section 5.5, the prices at hand are correlated across hours and markets. After fitting either of the quantile autoregressive models in this section, hourly innovations (also across markets) will still exhibit correlation and dependence cross-sectionally. The block correlation structure will be present in the residuals. The dependence structure between the day-ahead and primary reserve prices also seems to be non-linear. Thus, linear correlation is not sufficient to model the intermarket price behaviour. This dependence must be embedded into the scenario generation to achieve realistic scenarios. Kaut 2014 proposed a copula-based heuristic for scenario generation. The main concept is to provide a full description of dependence between random variables by utilizing copu-

Figure 6.6: RMSE of day-ahead and primary reserve price predictions

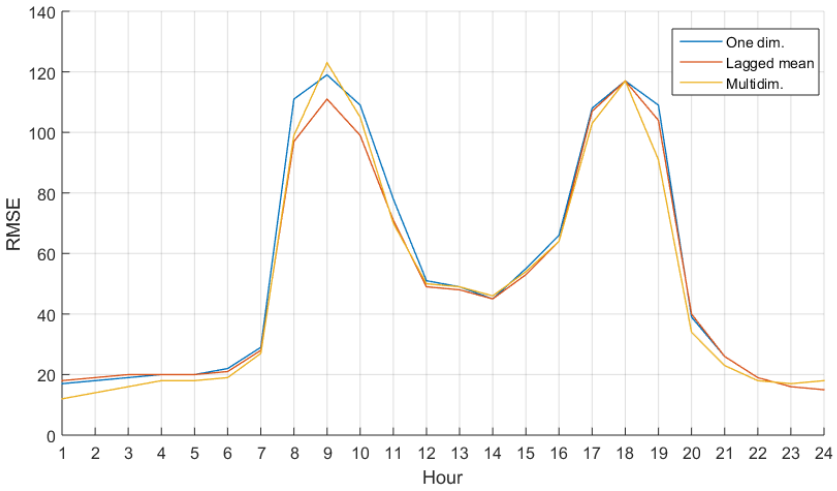

(a) Day-ahead

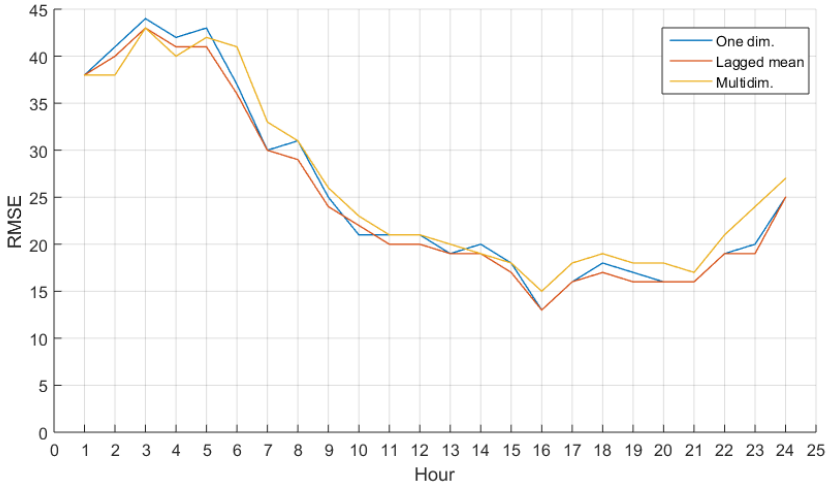

(b) Primary reserve

las. The value of the variables and their dependence are decoupled via the cumulative mapping. The construction of scenarios is done by solving an assignment problem minimizing the distance between the copula defined by the scenarios and a target copula (e.g. the empirical). The optimization problem is solved heuristically. We will use this heuristic approach to model dependence. A full explanation is provided in Section 6.5.

6.4 Modelling predictive density: Quantile regression

We use quantile regression for probabilistic forecasting of day-ahead and primary reserve prices. This section aims to describe how we have implemented the quantile autoregression framework and important considerations. For the general topic of quantile regression the reader is referred to Koenker 2005. We will briefly summarize the methodology in the following. Quantile regression can be used in various ways, but we use it to estimate optimal coefficients for a linear autoregressive forecasting model. Quantile regression can be performed multiple times to estimate coefficients for the different quantiles and hence approximate the conditional distribution. Section 6.3 presents the models to be estimated; QAR^{hourly} (6.3.4) and $\text{QAR}^{lagged-mean}$ (6.3.5). Both models are estimated for day-ahead and primary reserve prices. The section also provides an explanation of the choice of a quantile regression model. Main reasons are that it allows for a dynamic non-parametrical description of forecast uncertainty, and that conditional price variance may be related to the forecasted price.

In this section we will use a slightly modified notation corresponding to the notation in litterature. Let $f_t(\beta_0, ..., \beta_k)$ denote a linear regression function with coefficients $\beta_0 ... \beta_k \in \mathcal{R}$ in the variables $X_{t1} ... X_{tk}$, i.e. the set of explanatory variables of the stochastic variable Y_t at time t. Then, a prediction on Y_t is given by

$$\hat{Y}_t = f_t(\beta_0, ..., \beta_k) = \beta_0 + \beta_1 \tilde{X}_{t1} + ... + \beta_k \tilde{X}_{tk} \tag{6.4.1}$$

and the residual of the prediction is defined as

$$\varepsilon_t = \tilde{Y}_t - \hat{Y}_t \tag{6.4.2}$$

where ε_t denotes the error, and $\tilde{Y}_t$ denotes the observed value of Y_t.
The problem of computing the coefficients for the τ^{th} quantile $\tau \in (0, 1)$ can be formulated as

(P)

$$\min_{\beta_0, ..., \beta_k} z = \sum_{t \in \mathcal{T}} (\tau - \mathbf{1}_\varepsilon) \varepsilon_t \tag{6.4.3}$$

where

$$\mathbf{1}_{\varepsilon} = \begin{cases} 1 & \text{if } \varepsilon_t \leq 0 \\ 0 & \text{otherwise} \end{cases} \tag{6.4.4}$$

We can linearize problem (P) by introducing the binary variable δ_t. Solving (P) is then equivalent to solving (P1)

(P1)

$$\min_{\beta_0,\ldots,\beta_k} z = \sum_{t\in\mathcal{T}} \tau\varepsilon_t^+ + (1-\tau)\varepsilon_t^- \tag{6.4.5}$$

s.t.

$$\begin{aligned} \varepsilon_t &= \varepsilon_t^+ - \varepsilon_t^- && t \in \mathcal{T} \\ \varepsilon_t^+ &\leq \delta_t M && t \in \mathcal{T} \\ \varepsilon_t^- &\leq (1-\delta_t)M && t \in \mathcal{T} \\ \varepsilon_t^+ &\geq 0 && t \in \mathcal{T} \\ \varepsilon_t^- &\geq 0 && t \in \mathcal{T} \\ \delta_t &\in \{0,1\} && t \in \mathcal{T} \end{aligned}$$

Problem (P1) is easily solved in any standard software given that $\mathcal{T}$ is not a very big set. This program will give coefficients for the forecasting model to issue predictions on the τ^{th} quantile. If we compute coefficients for sufficiently large number of quantiles, we will achieve an approximation of the conditional probability distribution. In this case however, we have 48 parallel time series, one for each hour of the day in the day-ahead and primary reserve prices. This can be interpreted as a panel.

For a panel data model, that is

$$Y_{it} = \beta_0 + \boldsymbol{\beta}\boldsymbol{X}_{it} + u_{it} \qquad i \in \mathcal{N} = \{1..24\} \tag{6.4.6}$$

with possibly shared covariates between equations and covariant error terms u_{it}, the problem of estimating *pooled* quantile coefficients (Wooldridge 2010), can be formulated as

(P2)

$$\min_{\beta_0,\boldsymbol{\beta}} z = \sum_{i\in\mathcal{N}}\sum_{t\in\mathcal{T}} \tau\varepsilon_{it}^+ + (1-\tau)\varepsilon_{it}^- \tag{6.4.7}$$

s.t.

$$
\begin{aligned}
\varepsilon_t &= \varepsilon_{it}^{+} - \varepsilon_{it}^{-} && t \in \mathcal{T}, i \in \mathcal{N} \\
\varepsilon_{it}^{+} &\leq \delta_{it} M && t \in \mathcal{T}, i \in \mathcal{N} \\
\varepsilon_{it}^{-} &\leq (1 - \delta_{it}) M && t \in \mathcal{T}, i \in \mathcal{N} \\
\varepsilon_{it}^{+} &\geq 0 && t \in \mathcal{T}, i \in \mathcal{N} \\
\varepsilon_{it}^{-} &\geq 0 && t \in \mathcal{T}, i \in \mathcal{N} \\
\delta_{it} &\in \{0, 1\} && t \in \mathcal{T}, i \in \mathcal{N}
\end{aligned}
$$

Problem (P2) is harder to solve than (P1) due to huge increase in the number of binary variables. Other authors recommend to let $\boldsymbol{\beta}$ be a function of i and the constant term β_0 the same across panels, or penalizing deviations. These problems are equally hard to solve as (P2). We have to solve the problem many times to get estimates of the different quantiles. We choose to use (P1) to estimate quantiles for each hour individually. This is a reasonable choice because each of the hourly prices behave quite differently, so we regard it as a strength that the coefficients are highly specific for each hour. To illustrate this, we have plotted the quantiles using both (P1) and (P2) for September 2 2016 in Figure 6.7. Note that the price axis is shifted by 100 NOK in (b) to include the top quantile. Variance is very different from hour to hour and (P1) clearly captures this. (P2) gives a change in hourly variance only dependent on observed values; clearly not enough to capture the high variance during start and end of the working day.

One drawback of quantile regression is tail uncertainty. As we progress towards the tails of the distribution, less and less data points are available, yielding higher uncertainty. One way to cope is to include more data points in the estimation, but this increases computational effort. For scenario generation purposes a good description of distributional tails is paramount. Jonsson et al. 2014 fits exponential distributions to the tails below 5 % and above 95 %. However, upon testing this does not seem to be worthwhile. The empirical 5 % and 95 % tails of the forecast error distribution perform just equally well or better. In this thesis we therefore choose to use the empirical 5 % and 95 % tails of the forecasting error. The limit of 5 % and 95 % is set ad hoc because good results are achieved for these values, similar to Jonsson et al.

Between 5 % and 95 % we will use quantile regression to estimate the cumulative distribution of the conditional price. Ideally, a very fine discretization of the quantiles would be used. This is however very time consuming. For 5 % increments, the running time is just short of half an hour, which is deemed acceptable. We thus choose 5 % increments.

An interpolation scheme is needed to arrive at a complete cumulative distribution. We choose Matlab's built-in Piecewise Cubic Hermite Interpolating Polynomial (PCHIP). PCHIP is shape preserving, does not overshoot (does not violate monotonic-

Figure 6.7: Pooled vs. nonpooled quantile estimation Sept. 4 2016

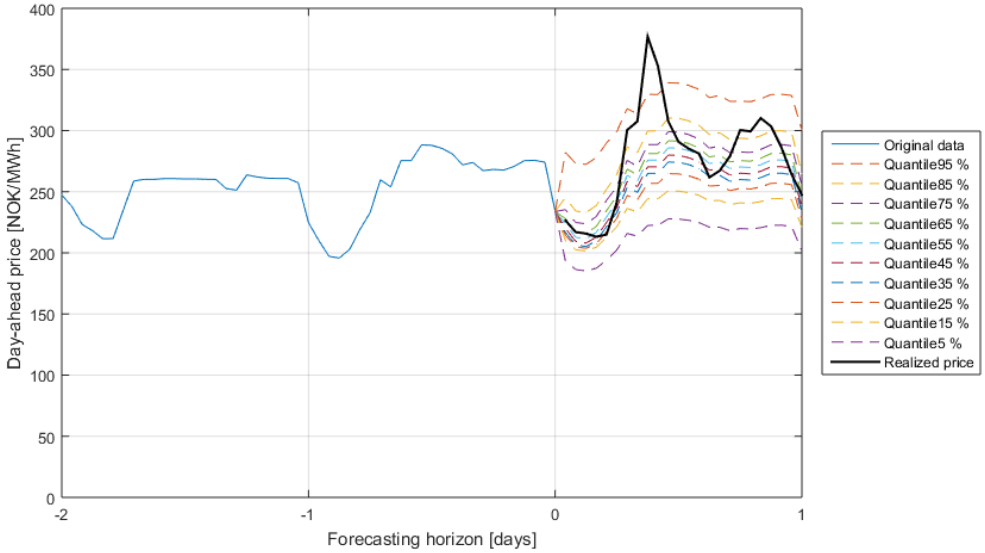

(a) Pooled (P2) estimation

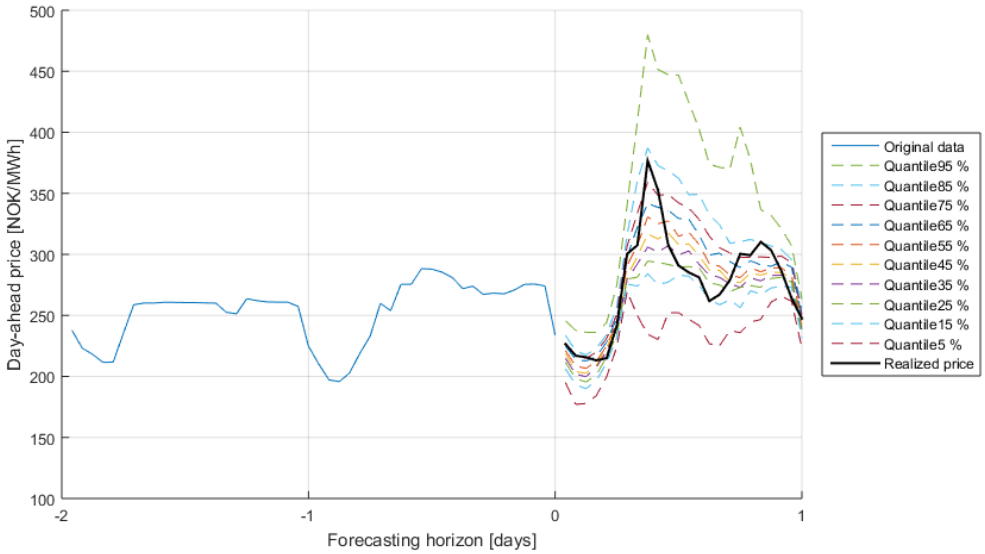

(b) Nonpooled (P1) estimation

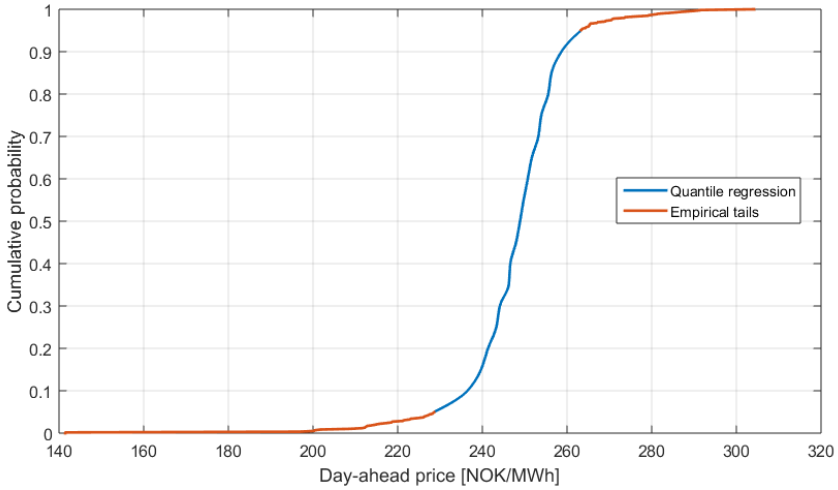

Figure 6.8: Predictive cumulative distribution for day-ahead price with empirical tails. Notice that monotonicity is maintained

ity), and works well when the underlying function is not oscillatory.

Finally, we discuss the length of the training period. Jonsson et al. choose training period based on a trial and error approach, evaluating which training length gives best fit out of sample. They arrive at longer training periods for quantiles close to the tails and shorter training periods for centre quantiles. The range is from 2000 to 210 entries in the training set. We have tested different training period sizes for our model, but we seem to get best results including all data in the in sample set, that is 2014-2015, for all quantiles. This implies less time dynamic coefficients. Figrure 6.8 shows an example of an interpolated cumulative distribution with empirical tails fitted below and above 5 % and 95 % using the methodology in this chapter. Qualitatively, the result is satisfactory. Monotonicity is clearly not violated in this case.

In order to verify our assumption that the conditional price exhibits varying sensitivity to conditioning variables across the distribution, we did an analysis on the resulting regression coefficients. These results are in themselves interesting, and the full analysis can be found in Appendix E.

6.5 Modelling dependence: Copula heuristic

Dependence between random variables should not always be modelled linearly using correlations. In our case it is necessary to model dependence between the day-ahead and primary reserve prices precisely. For instance, in the morning hours there seems to be some threshold day-ahead price, at which the primary reserve price is expected to be low, typically in the range of 15-25 NOK/MWh. If the day-ahead price is reduced from this threshold by only a few percent, the primary reserve price can be expected

to increase by many percent. On the other hand, if the day-ahead price is increased from the threshold value by a few percent, very little is expected to happen to the primary reserve price. Dependence in down-turns differs from dependence in up-turns. Correlations are not sufficient to model such dynamics.

Dependence across hours and across markets will be modelled with the copula-based framework by Kaut 2014. The model proposed in Section 6.3 combined with estimation from Section 6.4 provides hourly predictive densities for the prices. Constructing daily price paths from these densities with correct dependence properties is our next goal. In this chapter we only explain quite briefly what a copula is, and the goal of generating scenarios from a copula. A more extensive introduction to copula theory and the copula-based scenario generation heuristic used is given in Appendix D.

6.5.1 Copula-based scenario generation

A copula is the joint cumulative distribution function of any n-dimensional random vector with standard uniform margins, that is, a function C : $[0,1]^n \to [0,1]$. Sklar's theorem (Sklar, 1996) states that for any n-dimensional cumulative distribiution function F with marginal distribution functions $F_1, ..., F_n$, a copula C exists such that

$$F(x_1, ..., x_n) = C\Big(F_1(x_1), ..., F_n(x_n)\Big) \tag{6.5.1}$$

If all the marginal cumulative distribution functions F_i are continuous, then there exists only one unique C. A consequence of the theorem is that, for every $u = (u_1, ..., u_n) \in [0,1]^n$ we have that

$$C(u_1, ..., u_n) = F\Big(F_1^{-1}(u_1), ..., F_n^{-1}(u_n)\Big) \tag{6.5.2}$$

where F_i^{-1} is the generalized inverse of F_i. That means that knowing the marginal cdfs and the copula is the same as fully knowing the multivariate cdf. This stands in contrast to correlations that assume a linear dependence. Furthermore, copulas, unlike correlations, are independent from the marginal distributions. Therefore we can model the two independently. A copula only models the interdependence of two or more distributions, the information about the distributions themselves has been removed. Ergo, the copula is unchanged as long as the multivariate samples do not change order.

Obtaining the copula for a multivariate distribution is a non-trivial task. Instead we construct an empirical copula, C_T from historical samples, each historical sample containing one sample entry from each of the distributions modeled. In the case of our day-ahead and primary reserve market modeling, each sample contains 24 hourly prices for both day-ahead and primary reserve. The method takes the number of scenarios to generate, S, as input. The goal of the method is now to generate the copula scenarios for each margin $\tilde{u}_{sm}, s \in \{1..S\}, m \in \{1..n\}$, that in aggregate deviate as little as possible from the target copula. According to the definition of a copula, the scenarios $\tilde{u}_{sm}$ should be placed approximately uniformly between 0 and 1. At the same time,

Table 6.2: Balancing market characteristics and corresponding modelling choices

Findings from empirical study	Modelling
Regulating states with different probability throughout the day	Hour specific Markov
Autocorrelation and stationarity	Autoregressive models
Missing values in volumes and premiums	Unequally spaced time series
Day-ahead dependence for premiums	Linear regression
Dependence between premiums and volumes and hours	Copula heuristic

for all $s \in \{1..S\}$, $\tilde{u}_{sm}$ have to be connected across the n distributions (hours and markets), such that the dependence between the distributions resemble the one of the target copula. This is achieved by solving an assignment problem heuristically. The problem minimizes the deviation between the target copula and the copula defined by the scenarios.

After obtaining $\tilde{u}_{sm}$, we can make the transformation back to target variables. We have already estimated day-specific cumulative distributions of the prices (using quantile autoregression) for which we are making scenarios. We can now calculate the actual price scenarios using $F_m^{-1}(\tilde{u}_{sm})$. These scenarios will be correctly distributed across each margin and at the same time have correct dependence properties between hours and markets.

6.6 Balancing market modelling

This section describes the modelling of the balancing market. It is divided into four parts. First we discuss how the balancing market differs from the day-ahead and primary reserve markets. Thereafter, we present a model for regulating states. We then discuss balancing volumes, and finally premiums.

Two very important features of the balancing market will permeate the modelling in this section: regulation states and real time balancing and pricing. As pointed out in section 5.4 the balancing market is characterized by regulation states, i.e. if there is regulation or not. Should there be regulation, the regulating volume in that direction is strictly positive. Let ϑ_t denote the regulation state of the system at time t

$$\vartheta_t = \begin{cases} -1 & \text{if down regulation, i.e.} \quad \nu_{4t} > 0 \\ 0 & \text{if no regulation, i.e.} \quad \nu_{mt} = 0 \quad \forall m \\ 1 & \text{if up regulation, i.e.} \quad \nu_{3t} > 0 \end{cases} \tag{6.6.1}$$

Because of hours without regulation the data set for regulating volumes and premiums contains zeros in most hours. This must be tackled by the modelling approach because these hours will significantly bias the mean and variance of the model if included without regard to the regulation state. We therefore chose to decouple the regulation state i.e. arrival of demand, and the regulating volume, i.e. the size of the demand. A similar approach is chosen by several other authors, see for example Klæboe et al. 2013 and Olsson and Söder 2008. As we shall see later, this allows for a conceptual framework in which the volume and price processes are continuous background processes made observable by the arrival of demand. In this way we can think of the regulating volume and premium as unequally spaced- or irregularly sampled time series. That is, zero hours in the data are simply regarded as hours in which no information about the volume or price process is available. This discussion raises two questions: what is an appropriate way of modelling the arrival of regulation demand, and its duration, and how to model unequally spaced time series. Before returning to these question, we shall discuss the real time pricing mechanism of the balancing market.

Because day-ahead and primary reserve prices are quoted for all hours at the same time, and because the different hours exhibit different mean reversion and variance properties, we chose to model the prices as panel data. For the balancing market however, hourly prices are revealed consecutively throughout the day. In this structure one would expect consecutive hours to be correlated more than hours further apart. In addition, we would expect a forecasting error in one hour to propagate to the next hour, increasing forecasting uncertainty towards the end of the day. The balancing market volumes and prices can therefore be modelled as ordinary one dimensional time series on an hourly basis.

Several models for balancing market forecasting have been proposed. In an empirical study, Klæboe et al. benchmark a variety of time series based models. Klæboe test the forecast ability of four volume models, five premium models and three models for state determination. In this thesis we use the best performing model for the 12-36h forecasting horizon, with minor adjustments where we find it beneficiary.

6.6.1 Regulating states

For state determination several model options exist. We want a model to correctly reflect the probability of regulation occurring throughout the day. In addition, if a state occurs, it tends to last for some duration. That is, probability of remaining in state is higher than the probability of entering this state from another state. Klæboe et al tested duration dependent and hourly specific Markov switching models in addition to an arrival rate model from inventory control theory. For the day ahead forecasting horizon the hour specific Markov model performed best, predicting 37 % of hours correctly. This rather low fraction clearly shows the unpredictability of the balancing market. An hour specific Markov model has a transition probability matrix

$$P_t = \begin{bmatrix} P_{11t} & \dots & P_{1nt} \\ \vdots & \ddots & \vdots \\ P_{m1t} & \dots & P_{mnt} \end{bmatrix} \tag{6.6.2}$$

where P_{ijt} denotes the probability of the regulating state switching from state i to state j from hour t to hour $t+1$, and $|\cdot|$ denotes the number of entries in a set. In our case P_t is a 3×3 matrix reflecting the possible outcomes of ϑ_t. Estimators for P_{ijt} are calculated according to

$$P_{ijt} = \frac{|\vartheta_t = i \cap \vartheta_{t+1} = j|}{|\vartheta_t = i|} \tag{6.6.3}$$

using historical data. All hour specific Markov transition matrices are given in Appendix A. These can be used to simulate outcomes for the regulating state through the day of planning.

6.6.2 Balancing volumes

Because the balancing volumes are unequally spaced time series, analysis and estimation is complicated. Our first goal is to analyse the autocorrelation structure in the data. If we include zero hours in the analysis, the results will be biased. On the other hand, if we exclude zero values and artificially compress the time series, hours that are far apart in time will contribute to the calculation as if they were consecutive. Other options include linear interpolation in the data and mean value substitution, but as Erdogan et al. 2005 points out, this smoothens the data and causes bias. We cope with this in a similar way as Söder et al. Let β_{mt} denote the *regulation signal* in each regulating direction

$$\beta_{mt} = \begin{cases} 1 & \text{if} \quad \vartheta_t = 1 \text{ for } m = 3 \text{ or } \vartheta_t = -1 \text{ for } m = 4 \\ 0 & \text{otherwise} \end{cases} \tag{6.6.4}$$

β_{mt} is 1 if there is regulation in hour t and zero otherwise. Using the regulation signal we can include only regulation hours in the calculation without destroying the time structure. The autocovariance function at lag h can now be computed

$$\gamma(h) = \frac{1}{\sum_{t=1}^{T} \beta_{mt}} \sum_{t=1}^{T-h} (\nu_{mt} - \mu_\nu)(\nu_{m,t+h} - \mu_\nu)\beta_{mt}\beta_{m,t+h} \tag{6.6.5}$$

where μ_ν denotes the mean of the regulation volume. From this result the autocorrelation function $\varrho(h)$ can be computed as

$$\varrho(h) = \frac{\gamma(h)}{\gamma(0)} \tag{6.6.6}$$

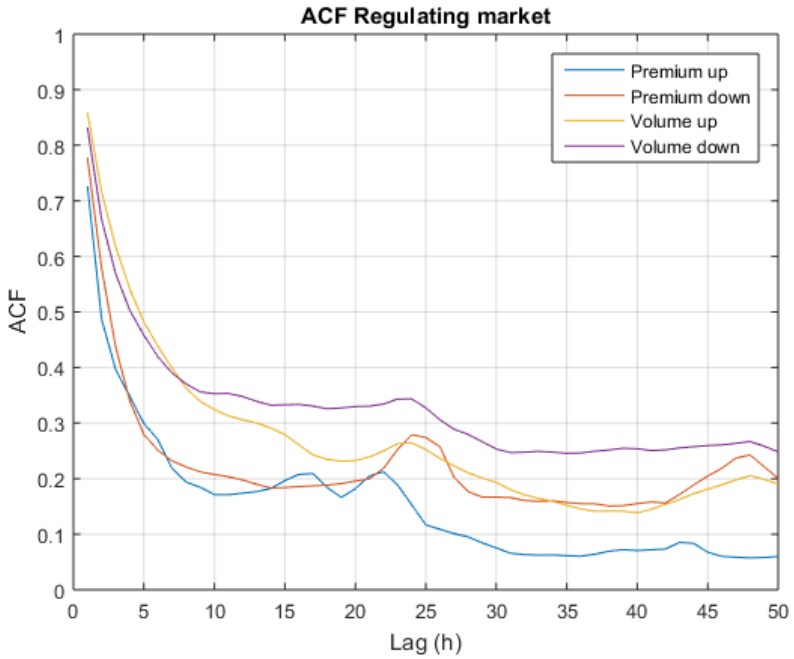

Figure 6.9: ACF for the balancing market

Figure 6.9 shows the autocorrelation computed for various lags for volumes and premiums. We note that volumes as well as premiums exhibit an autoregressive signature, with slowly decaying positive ACF. There is also some seasonality present at around 24 hour lag, especially for premiums. Taking stationarity of the series into account, ARMA models will be an appropriate modelling tool.

Erdogan et al. present a statistical model for unequally spaced time series. The core idea is to use an autoregressive process of order 1 to resample the missing values. Such a technique might be appropriate for the regulating volumes because an estimate for the PACF shows a significant spike at lag 1, leaving the other lags at essentially zero. This might however destroy some of the seasonal effects for the cases when no regulation lasts for more than a day. It should not be a problem for shorter instances because the 1 hour lag correlation is much stronger than the 24 hour lag correlation. Besides, the average duration of no regulation is 8.5 hours, well below a full day.

The details on the methodology of Erdogan can be found in Appendix B. Here follows the fundamentals of it, necessary to show why and how we pre-process the balancing volumes time series.

An autoregressive process of order 1 can be written as follows

$$X(t+1) = \theta X(t) + \sigma \epsilon_{t+1} \tag{6.6.7}$$

We start with an irregularly sampled time series $X(t_i)$ with nonzero values at times t_i. We want to "fill in the gaps" at all the times s_k when the value of the time series

is not available. In addition, we obviously want to keep all original values of the time series such that the resampled series coincide at these times, i.e. $\tilde{X}(s_k) = X(t_i)$ for $s_k = t_i$. The parameters $\hat{\theta}$ and $\hat{\sigma}$ are estimated using a least squares and maximum likelihood argument (according to the methodology in the Appendix). In general, ϵ_s cannot be estimated for hours without data. Instead we introduce an auxiliary error ε for hours without data such that

$$X(t_i + \Delta_i) = \hat{\theta}^{\Delta_i} X(t) + \varepsilon \hat{\sigma} \sum_{j=0}^{\Delta_i - 1} \hat{\theta}^j \tag{6.6.8}$$

where Δ_i denotes the time gap between the last and next hour with available data.

In this manner way we can estimate a local error term in between observations. Rearranging the terms we get a direct estimate of ε

$$\varepsilon = \frac{X(t_i + \Delta_i) - \hat{\theta}^{\Delta_i} X(t)}{\hat{\sigma} \sum_{j=0}^{\Delta_i - 1} \hat{\theta}^j} \tag{6.6.9}$$

This allows us to set recursively

$$\hat{X}(t_i + 1) = \hat{\theta} \hat{X}(t) + \hat{\sigma} \varepsilon \tag{6.6.10}$$

Note also that this prediction sets $\hat{X}(t_i) = X(t_i)$ so that the resampled values coincide with the original data, as desired. We now have all we need to resample the regulating volume time series. We first logarithmically transform the data to stabilize the variance, and fit a rational to make the residuals have zero mean.

One very important feature of the resampling procedure is that it uses only the end points of the original data to estimate the local error term ε. However, for the balancing market this will cause an issue which has not been pointed out by previous research. The balancing market is operated in real time, but quoted hourly in the data set. Balancing events tend to last for several hours when they start; averaging at 3.9, 5.0 and 8.5 hours for upward, downward and no regulation respectively. Because of this, the hours in the middle of the event are likely to be full hours. On the other hand, for the first and last hour in a regulation event, we have no guarantee that the quoted hours are full hours. Actually, the contrary seems to be the case. Figure 6.10 shows an illustrative example of what the volume data may look like. We can see that the last volume in the first event, at hour 6, is very low compared to the rest of the data. The same holds for the first hour in the next event, at hour 10. Hence, if we use hour 6 and 10 for estimation in the resampling algorithm, this will bias the mean of the series significantly. Figure 6.11 shows proof of this hypothesis. The average volumes are plotted during regulation events. It is calculated as follows: All consecutive regulation hours of duration longer than three hours are linearly stretched from 0 to 100 % of duration. For each percentile increment, the average volumes are computed over all events. From this figure we clearly see the main point: volumes are lower at

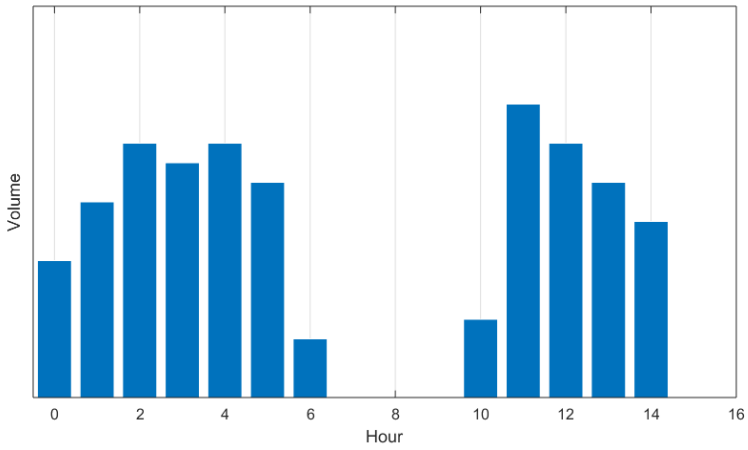

Figure 6.10: Example of volumes structure in two regulation events. End-point volumes are generally smaller than mid-event volumes

the beginning and the end of regulation events. Mid-event volumes are just short of twice the size. We therefore double the size of end-point volumes before proceeding. This is an important step to attain a reasonable mean for the resampled series. The resampling procedure has been implemented in *Matlab* according to the pseudocode in Table 6.3

Figure 6.12 shows the resampled downward volume. Qualitatively, the resampled data seems to behave much like the original data in terms of amplitude. Calculations show that the mean and variance of the series slightly decrease after resampling. This is not surprising because the data density in the unequally spaced series is much higher for periods with high volumes. The resampled series however, will have uniform density in time. We thus add more data to periods of low volumes than we do to periods of high volumes. This cannot be viewed as a weakness of the resampling technique. The variance can possibly be increased by perturbing the auxillary error term with randomness until the desired variance is achieved a posteriori. Like Erdogan et al. pointed out, it seems that the resampled data tends to underestimate the presence of spikes. Spikes are an important characteristic of the regulating market. A more in depth analysis of the balancing market would attempt to tackle this problem. For the purposes of this thesis however, the results are considered satisfactory.

We now move on to forecasting balancing volumes. AR(1) models are fitted to the resampled series for the desired training period. The AR(1) models can be simulated a large number of times for the forecast horizon to yield probabilistic forecasts. In turn, the hourly simulations can be used to estimate the predictive density for each hour. These densities will then be coupled with a copula framework to generate scenarios, see section 7. The reason we choose simulation of the AR model instead of quantile regression as previously, is twofold. Primarily, the probabilistic forecasting of the balancing market is fundamentally different from day-ahead and primary reserves. For

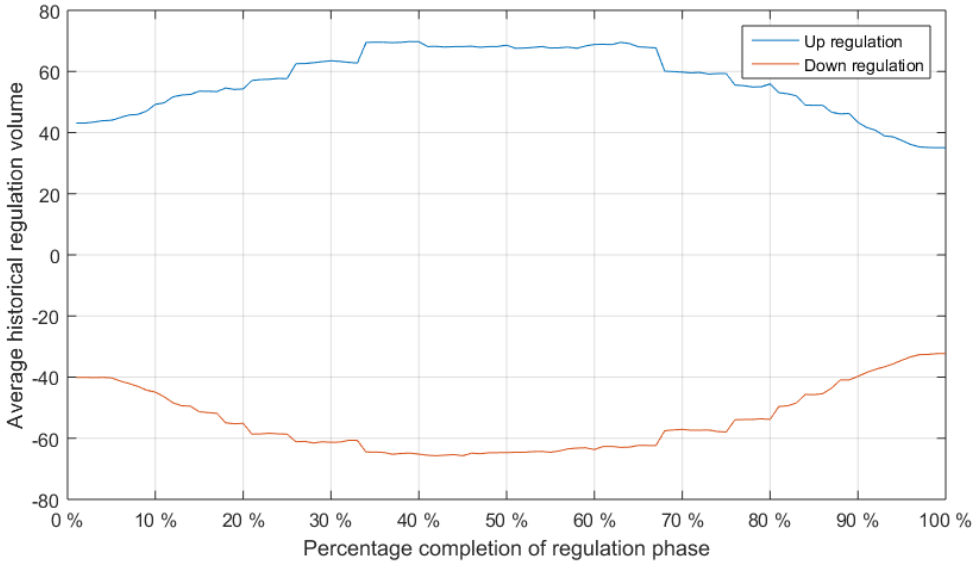

Figure 6.11: Average volume during the completion of regulation events

Table 6.3: Pseudocode for resampling unequally spaced time series

Algorithm: Resampling unequally spaced time series

$k \leftarrow 0; \quad s_k \leftarrow t_0; \quad \tilde{X}(s_0) \leftarrow X(t_0);$
For $i \in \{1...n-1\}$ **do**
    $\delta \leftarrow t_{i+1} - t_i; \quad \varepsilon \leftarrow \frac{X(t_{i+1}) - \hat{\theta}^{\delta} X(t_i)}{\hat{\sigma} \sum_{j=0}^{\delta-1} \hat{\theta}^{j}};$
    For $j \in \{1...\delta-1\}$ **do**
        $s_{k+1} \leftarrow s_k + 1;$
        $\tilde{X}(s_{k+1}) \leftarrow \hat{\theta}\tilde{X}(s_k) + \hat{\sigma}\varepsilon;$
        $k \leftarrow k + 1;$
    End for
End for

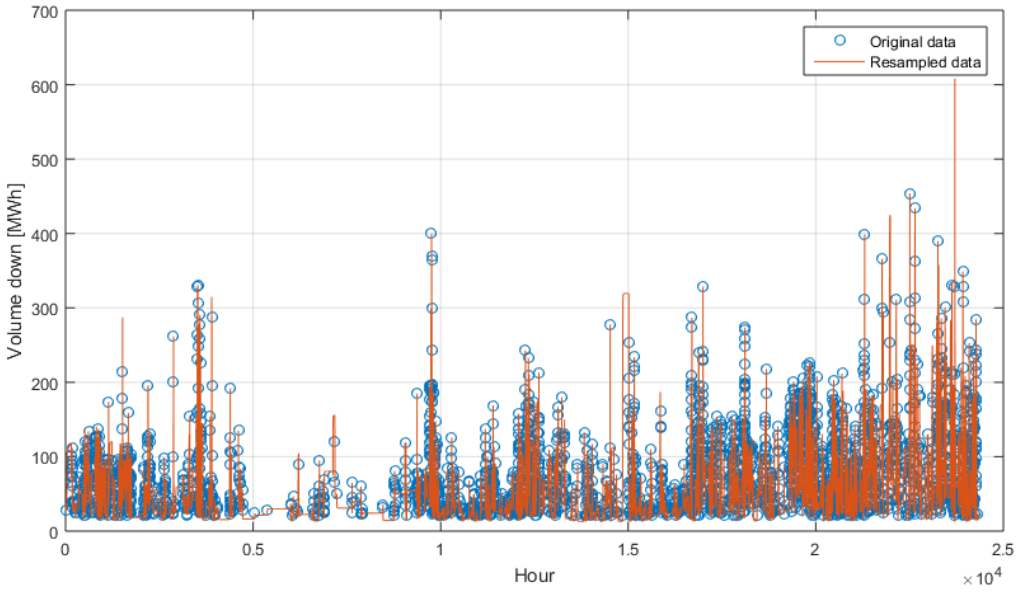

Figure 6.12: Resampled downward regulating volume

the balancing market we are to forecast 12-36 time steps ahead, while for day-ahead and primary reserve we forecasted 48 time series one step. If we were to use quantile regression for the balancing market we would have to model the propagation of forecasting uncertainty in hour h to hour $h+1$ in some way. Secondly, the balancing market is highly unpredictable, so we hypothesize that a more advanced model will not be worthwhile.

In this section we have covered a model for the balancing market volumes. To summarize, we first identify a Markov model for determination of the balancing state. We then modelled the volume time series as a continous series which is only observable when there is regulation. To obtain the regularly sampled series, we resampled the series using an AR(1) process. Probabilistic forecasts can be generated by simulating the AR model a large number of times. The next section will briefly discuss how to generate corresponding regulating premiums.

6.6.3 Balancing market premiums

The regulating premiums are obviously very dependent on the regulation state. Premiums decline to a very low level (not necessarily zero due to price zone export) whenever there is no regulation. Premiums are also dependent to a varying degree on day-ahead price and volumes. Klæboe et al. conclude that the performance in terms of forecasting MAE is similar for all four non-naive premium models they propose. Two of these models use day-ahead price and/or volume as exogenous parameters. Because of the correlation between these quantities in NO3, we want to include those same exogenous

parameters in this model.

Since there is no volume in hours without regulation, the premium is not of any relevance to the decision model in these hours. We will therefore also set the premium to zero in hours without regulation. This causes the same problem as for volumes. If we do not make any changes to the premium series data, and fit ARIMA models, the result will be biased by low premiums in hours without regulation. On the other hand, if we remove the no regulation hours and replace with zeros, the exact same problem that we had for the volumes occurs. We see two natural options: 1) regression on the volumes and 2) resampling of premium data.

The first approach will involve fitting an ARIMAX model for the premiums with volume as exogenous input in addition to the day-ahead price (see Section 5.5). This will require us to generate volume scenarios first, and then generate premium scenarios conditional on the volumes. There is one rather big problem with this approach: dependence is likely non-linear. Premiums should tend to zero when there is no regulation, i.e. $\nu_{mt} = 0$. When $\nu_{mt} > 0$ however, premiums are distributed in a wide range, especially for low volumes. The linear relationship in Figure 5.10 intensifies as volumes become larger. It is therefore likely that a linear coefficient describing the relationship between volumes and premiums is not sufficient.

The other option is to remove hours with no regulation, and resample the premium series as well. Once again, this will destroy some of the seasonality in the premiums. Keeping in mind that the average length of no regulation phases is 8.5 hours, and that seasonality is 24 hourly, we do not believe this will be a big disadvantage. Another problem is that resampling with an AR(1) process will not explicitly catch the dependence on exogenous parameters: volumes and day-ahead price. We have tested both models, and the resampling model performs better. Test results clearly show that non-linear dependence between volumes and premiums is critical. As volumes go from $\nu_{mts} = 0$ to $\nu_{mts} \geq 0$, impact on premiums is huge. The effect is especially severe when volumes are very large. In this case, corresponding premiums can be way too large. Our choice is therefore to proceed with the resampling model.

We follow the procedure in Section 6.6 for resampling the series. The resampled time series is of order 2 even after resampling. Afterwards an ARMAX(2,0,1) model is fitted to the resampled data with resampled volumes as well. Coefficients are in Table 6.4.

Generating premium scenarios is non-trivial, because for each volume scenario, a premium must be generated. In addition, we seek to reduce sampling error. The obvious option is to simulate a number of premiums for each volume scenario. This will yield a realistic predictive distribution for premiums. What remains is to choose full premium paths to pair with volumes. For this we will once again use to copula heuristic with the empirical copula directly as target, see Chapter 7.

Table 6.4: Coefficients for premium up model

Parameter	Value	Standard error	t statistic
Constant	0.0368	1.95	0.0188
AR{1}	0.0873	0.00517	16.9
AR{2}	0.755	0.00417	181
MA{1}	0.774	0.00533	145
ψ^{vol}	0.0911	0.00608	15.0
ψ^{spot}	0.0300	0.00530	5.66
Variance	2584	1.38	1870

6.6.4 Out-of-sample testing

We have tested the point forecasting ability of the models in this section out of sample. Tests were done for the period October 2015 to September 2016 and benchmarked against a naive[2] model. The optimal training period (in terms of lowest forecasting error) was found to be 80 days. We have chosen to quote forecasting error in terms of Monthly-weighted Mean Absolute Error (MMAE). This is a modified MAPE (Mean Absolute Percentage Error) that avoids the problem of volumes close to zero. Weron and Misiorek 2008 used a similar measure on a weekly basis for day-ahead prices. The MMAE is defined in terms of the predictions and realization of the variables e.g. volumes, as

$$MMAE = \frac{1}{N \cdot \bar{\nu}_N} \sum_{i=1}^{N} |\nu_{mi} - \hat{\nu}_{mi}| \tag{6.6.11}$$

where N is the number of hours in the month. Results are found in Table 6.5. The results are rather depressing. As expected, the balancing market is extremely hard to predict before day-ahead market closure. However, in contrary to what Klæboe et al. 2013 found, the naive model performs equally well as the more sophisticated model. The best performance has bold font in the table. Note especially that premiums and volumes for upward regulating are problematic. Surprisingly, the naive model has an error averaging well below 100 %, which is much better than what Klæboe et al. found in NO2.

Even though these results are bad in terms of point forecasting, the models may be valuable for probabilistic forecasting. Lack of a good point forecast simply implies more uncertainty in the model parameter. In turn, this must be handled by scenarios that correctly portray the uncertainty. These scenarios may span a broad range. There is not necessarily a strong connection between point forecasting ability and probabilistic forecasting ability. Hence, it may be possible to construct a probabilistic forecasting model that works very well even though the underlying point forecasts are highly uncer-

[2]Using yesterday's prices as forecast

Table 6.5: Results in MMAE from out-of-sample testing Oct 2015-Sept 2016

	Prem. down		**Prem. up**		**Vol. down**		**Vol. up**	
Period	Naive	ARMA	Naive	ARMA	Naive	AR	Naive	AR
Oct	**36 %**	42 %	99 %	**90 %**	51 %	**44 %**	**17 %**	23 %
Nov	29 %	**29 %**	**73 %**	74 %	44 %	**40 %**	**30 %**	40 %
Dec	44 %	**39 %**	99 %	**65 %**	54 %	**50 %**	**26 %**	42 %
Jan	99 %	**84 %**	102 %	**78 %**	48 %	**44 %**	**23 %**	33 %
Feb	36 %	**36 %**	**21 %**	42 %	56 %	**51 %**	**14 %**	24 %
Mar	**26 %**	31 %	**35 %**	43 %	**37 %**	50 %	**25 %**	30 %
Apr	43 %	**30 %**	**44 %**	64 %	64 %	**52 %**	**16 %**	21 %
May	46 %	**37 %**	95 %	**67 %**	57 %	**44 %**	**47 %**	50 %
Jun	**39 %**	42 %	89 %	**73 %**	57 %	**53 %**	64 %	**60 %**
Jul	53 %	**41 %**	**50 %**	55 %	64 %	**57 %**	**34 %**	52 %
Aug	**31 %**	36 %	**50 %**	52 %	61 %	**59 %**	**56 %**	57 %
Sep	53 %	**43 %**	**45 %**	48 %	74 %	**53 %**	52 %	**38 %**

tain. Probabilistic forecasting ability of all models proposed thus far will be evaluated in Section 6.7.

6.7 Evaluation of probabilistic forecasting ability

This section is devoted to testing of the probabilistic forecasting ability of the proposed models. Good probabilistic forecasts is a necessity to provide a description of uncertainty. These forecasts are used directly in scenario generation by choosing scenarios that represent the underlying density as effectively as possible. In contrast to single-valued or point forecasts, probabilistic forecasting amounts to assigning a probability to each of the possible outcomes, i.e. a probability density. In this thesis we evaluate probabilistic performance in terms of interval forecasts. An interval forecast consists of a lower and upper bound $L_{t|t-1}(p)$ and $U_{t|t-1}(p)$ respectively, and a probability p that the random variable will take a value between these bounds. By evaluating interval forecasts for different probabilities p, we can assess how well the probabilistic forecasts performs. Other options to evaluate probabilistic forecasts include for example the Continous Ranked Probability Score (CRPS). In this thesis we take the approach proposed by Christoffersen 1998 to evaluate a set of interval forecasts.

Interval forecasts can be evaluated in terms of unconditional and conditional coverage. Put informally, unconditional coverage is achieved if a fraction p of the observed values were between the bounds for some test period. If in addition the probability that the random variable takes a value in the interval remains constantly equal to p throughout the test period, we have conditional coverage. This is especially important for time series exhibiting heteroscedasticity, such as electricity prices, in which the probabilistic

forecast may perform well in some periods and worse in other.

Some notation is necessary. Let the indicator variable I_t define whether or not the realization of the random variable y_t is in the interval.

$$I_t = \begin{cases} 1 & \text{if } y_t \in [L_{t|t-1}(p), U_{t|t-1}(p)] \\ 0 & \text{otherwise} \end{cases} \tag{6.7.1}$$

With the indicator variable obtained from some test period $t \in \mathcal{T}^{test}$, we can perform analysis on conditional and unconditional coverage without regard to the underlying forecasting model nor any distributional assumptions. The next subsections establish briefly the testing criterion on I_t. Unconditional coverage is treated first, and then we move on to conditional coverage. Finally, we review the results from testing.

6.7.1 Unconditional coverage

Given a sequence of indicators from the test period $\{I_t\}_{t=1}^{t=T}$, we wish to test whether $\mathbb{E}(I_t) = p$ against the alternative hypothesis that $\mathbb{E}(I_t) \neq p$. Christoffersen proposes a log likelihood ratio test

$$LR_{uc} = -2\log\left[\frac{L(p; I_1, I_2, ..., I_T)}{L(\hat{\pi}; I_1, I_2, ..., I_T)}\right] \overset{asy}{\sim} \chi^2(1) \tag{6.7.2}$$

where $\hat{\pi} = n_1/(n_0 + n_1)$ is the maximum likelihood estimator for π, and n_0 and n_1 are the number of I_t equal to 0 and 1 respectively. The LR_{uc}-statistic is plotted and compared with the 5 % and 1 % significance level for a chi squared distribution with one degree of freedom ($\chi^2_{\alpha=0.05}(1)$ and $\chi^2_{\alpha=0.01}(1)$). $\chi^2_\alpha(1)$ denotes the value for which there is a probability of α of observing $LR_{uc} \geq \chi^2_\alpha(1)$ under the assumption that the hypothesis $\mathbb{E}(I_t) = p$ is true. That is, if the LR-statistic lies beneath the significance limit, then the forecast passes the test.

This procedure tests for overall coverage of the interval, but does not test whether the zeros and ones come clustered in a time dependent manner.

6.7.2 Conditional coverage

Conditional coverage is a measure of coverage that also checks whether indicator ones and zeros come clustered or independently distributed in time. We want to test if $\mathbb{E}(I_t|I_{t-1}, ..., I_1) = p$ for the entire test period. This is similar to test whether I_t is identically and independently Bernoulli distributed with parameter p, $\{I_t\} \overset{iid}{\sim} Bern(p)$. The statistic is

$$LR_{cc} = -2\log\left[\frac{L(p; I_1, I_2, ..., I_T)}{L(\hat{\Pi}_1; I_1, I_2, ..., I_T)}\right] \overset{asy}{\sim} \chi^2(2) \tag{6.7.3}$$

Table 6.6: Unconditional coverage day-ahead and primary reserve

	Day-ahead		**Primary reserve**	
Target coverage	QAR^{hourly}	$\text{QAR}^{lagged-mean}$	QAR^{hourly}	$\text{QAR}^{lagged-mean}$
95 %	95 %	95 %	95 %	95 %
90 %	90 %	89 %	88 %	90 %
75 %	74 %	72 %	67 %	72 %
50 %	50 %	47 %	41 %	42 %

with $\hat{\Pi}_1$ a Markov transition probability matrix estimated by

$$\hat{\Pi}_1 = \begin{bmatrix} \frac{n_{00}}{n_{00}+n_{01}} & \frac{n_{01}}{n_{00}+n_{01}} \\ \frac{n_{10}}{n_{10}+n_{11}} & \frac{n_{11}}{n_{10}+n_{11}} \end{bmatrix} \tag{6.7.4}$$

This statistic suffices to check for both independence and coverage simultaneously, and will be plotted against 5 % and 1 % significance level for a chi squared distribution with two degrees of freedom.

6.7.3 Results

The models for day-ahead and primary reserve prices have been tested for the 50 first days (1200 hours) in 2016. Testing is done by computing forecast intervals for the 95 %, 90 %, 75 % and 50 % coverages with quantile autoregression. Thereafter, we check whether the realization of the price is within the respective intervals. We test both the QAR^{hourly}-model (6.3.4) and the $\text{QAR}^{lagged-mean}$-model (6.3.5) for both markets. To give an intuitive impression of coverage, we first present the coverage in Table 6.6. The widest intervals seem to perform very well. The 75 % and 50 % intervals seem to be a bit to narrow, especially for primary reserve. It is not clear whether the lagged mean model gives better probabilistic forecasts at this point.

Next, we perform the LR-test by Christoffersen on an hourly basis. This will provide us with information on how the model performs in each of the 24 hours, which is interesting because the hours have such different characteristics. The results are plotted in Figure 6.13 and 6.14. Investigate first only the upper plots in each figure, labeled unconditional. Results are very good for both markets. Almost all hours and intervals pass the test at the 1 % significance level. The primary reserve prices are somewhat problematic in the early hours of the day, but the inclusion of the lagged mean seems to relieve this. Inclusion of the lagged mean in the day-ahead model does not seem to give an advantage in terms of unconditional coverage.

In terms of conditional coverage, results are also very satisfying. The QAR^{hourly} day-ahead model passes the test for most hours, but has some fails across several hours through the day, with no obvious structure. The $\text{QAR}^{lagged-mean}$ model outperforms the pure hourly model. Only two test instances fail the test at the 1 % significance

Table 6.7: Unconditional coverage balancing market

Target coverage	Prem. up	Prem. down	Vol. up	Vol. down
95 %	92 %	93 %	95 %	95 %
90 %	88 %	89 %	92 %	90 %
75 %	76 %	80 %	84 %	73 %
50 %	56 %	60 %	58 %	52 %

level. We thus conclude that the lagged mean model performs best, and we choose this model for scenario generation in Chapter 7 & 8.

For primary reserve prices, the $\text{QAR}^{lagged-mean}$ model provides significantly better conditional coverage than the pure one-dimensional hourly model. This is especially true in the highly volatile early morning hours. Some hours still do not pass the test, but this is generally the more narrow intervals. We are not aware of any probabilistic forecasting on primary reserve prices in Norway. Still, comparing with day-ahead coverage, results are very good. We once again conclude that the lagged mean model performs best, and thus choose it for use in subsequent scenario generation.

As an additional benchmark, we compare the mean width of the prediction interval for the two models. Generally, if a model has correct probabilistic properties (calibration), we want the mean width of the prediction interval to be as narrow as possible (sharpness). Figure 6.15 plots the mean PI width across hours for both markets and models. For day-ahead prices we see that the lagged mean model generally provides tighter bounds, except for hour 18-19. For primary reserve the lagged mean model is significantly tighter during early morning hours. Taking conditional coverage into account, there is little doubt that the lagged mean model has best performance for both markets.

A similar test was done for the balancing market. The test was carried out for 350 days (8400 hours), from October 2015 to September 2016. Aggregated results are in Table 6.7. Results are very similar to Klæboe et al. 2013. The widest intervals perform pretty well, while the narrowest tend to include too many realizations. The results are further elaborated by computing Christoffersen's LR-test for unconditional and conditional coverage. Figure 6.16 plots the results. Unconditional coverage is acceptable for wide intervals, but narrow intervals generally perform worse. The wider intervals perform badly for downward premium in morning hours. Klæboe et al. did not check conditional coverage for the proposed models, but encouraged future researchers to do so before implementing. Conditional coverage is generally bad, especially for narrow intervals. This is not unexpected reminding ourselves of the inherently unpredictable nature of the balancing market. It is not reasonable to expect good forecasts in a market which is designed to respond to unforeseen events.

Figure 6.13: LR-test for unconditional (top row) and conditional (bottom row) coverage day-ahead price

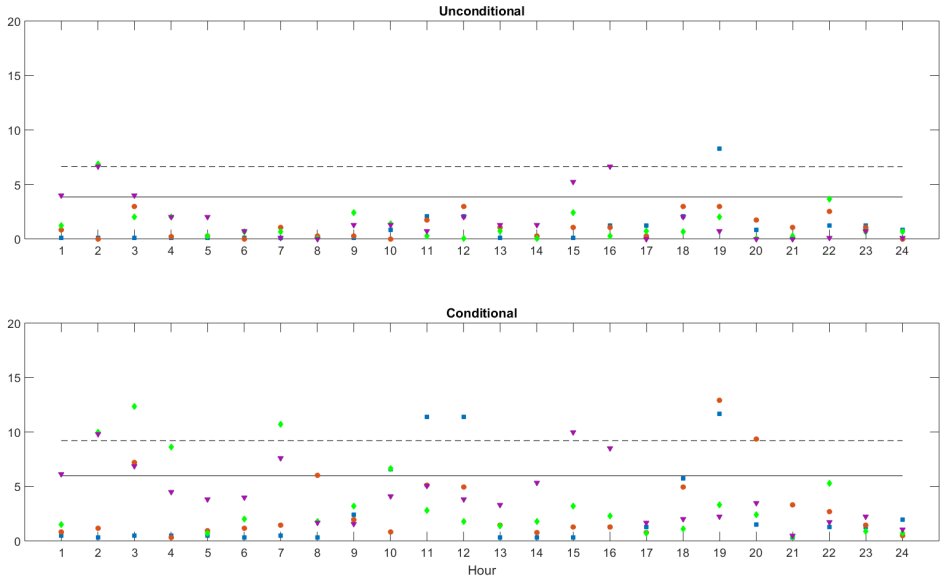

(a) Day-ahead QAR^{hourly} model

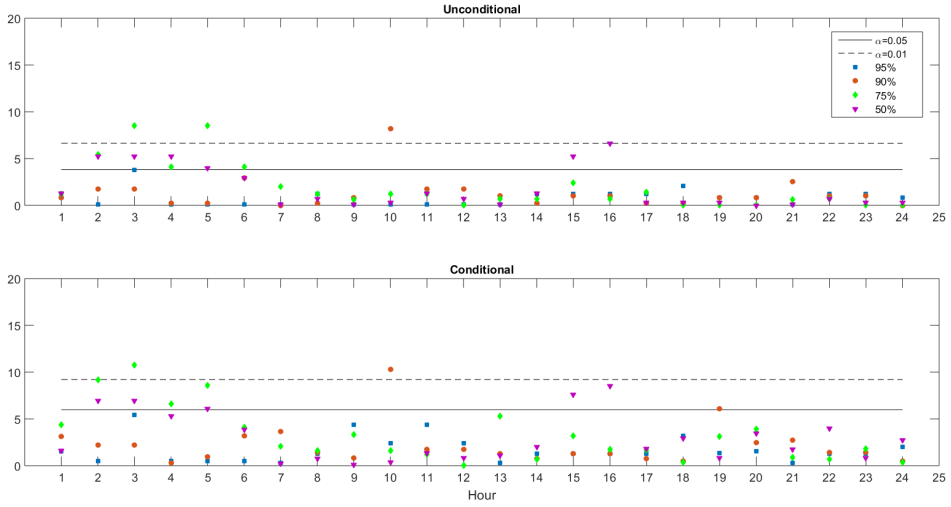

(b) Day-ahead QAR$^{lagged-mean}$ model

Figure 6.14: LR-test for unconditional (top row) and conditional (bottom row) coverage primary reserve price

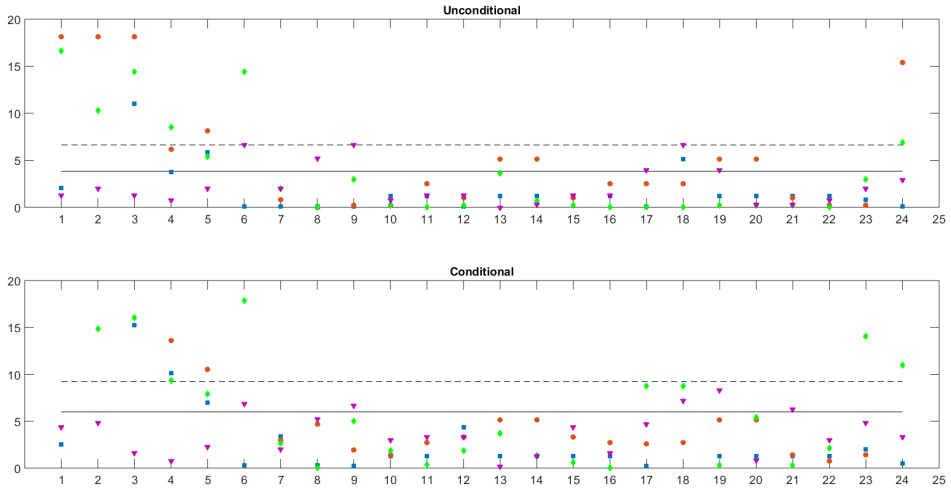

(a) Primary reserve QAR^{hourly} model

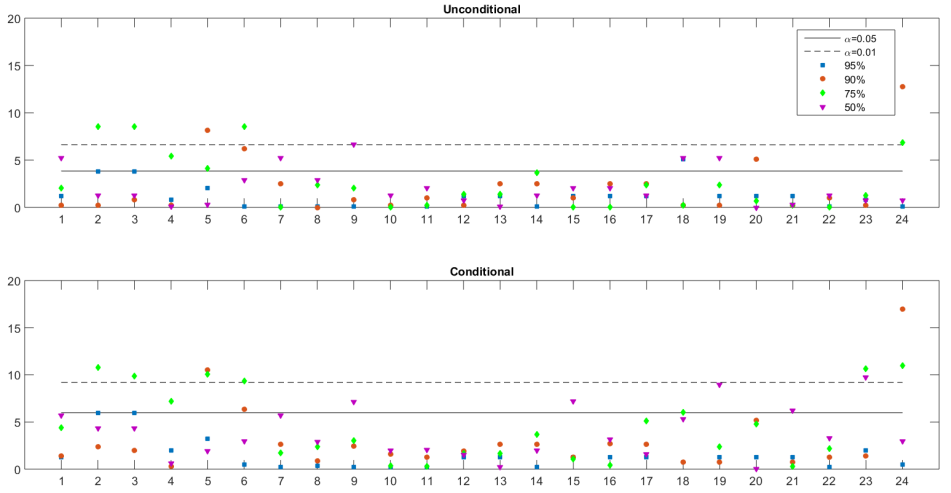

(b) Primary reserve QAR$^{lagged-mean}$ model

Figure 6.15: Mean PI width - solid line represents QAR^{hourly} model, dashed line represents QAR$^{lagged-mean}$ model

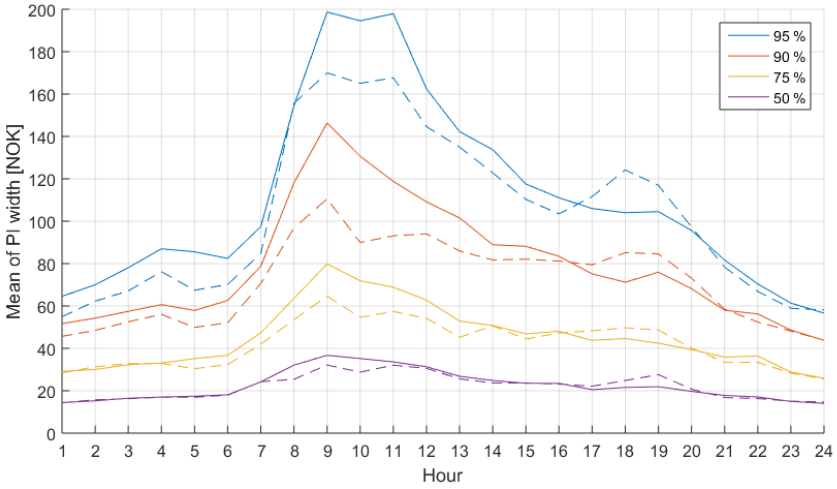

(a) Day-ahead

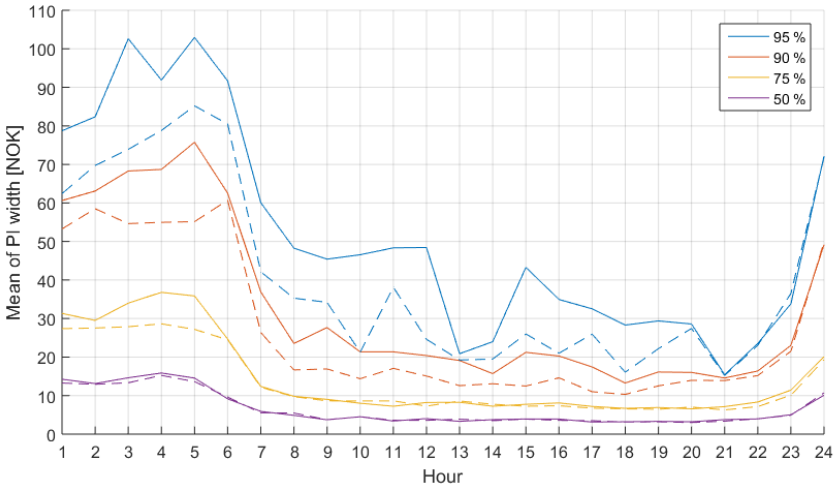

(b) Primary reserve

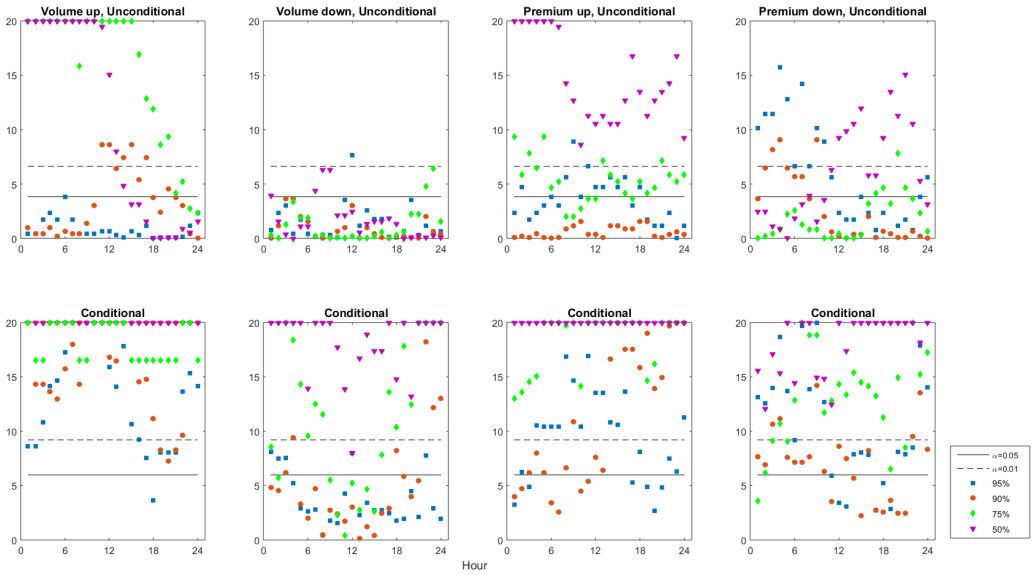

Figure 6.16: LR-test for unconditional (top row) and conditional (bottom row) coverage balancing market

Chapter 7

Scenario generation and evaluation

7.1 Introduction

Chapter 6 presented the necessary tools for generating probabilistic forecasts and modelling dependence between hours and markets. This chapter aims to connect the dots, and outline the scenario generation algorithm in its entirety. In addition, we test important properties of the generated scenarios.

An important characteristic of good scenarios is that the scenario distribution reflects the uncertainty in the forecast method. As we saw in section 6.7, the probabilistic forecasts perform well. However, upon generating scenarios, only a few discrete values must be picked to represent the probabilistic forecast. In addition, the scenarios must represent the underlying connection between hours and markets.

We test the scenarios in two ways. First, we qualitatively test dependence properties between markets and hours by evaluating scatter plots with scenarios and real prices. Thereafter we subject several scenario sets of different size to the optimization problem. The resulting objective value can be used to check for in-sample and out-of-sample stability. Results are used to determine the necessary number of scenarios to be used in the case study in Chapter 8.

7.2 Scenario generation algorithm

The reader is reminded of the problem structure in Figure 4.3, consisting of three stages. The here-and-now decision is the bid curve to submit in the day-ahead market. No decision is made between day-ahead price realization and primary reserve price realization. Hence there is no need to create additional branching between these events. If for each day-ahead price we had a branch of primary reserve price scenarios, there would be redundant information (duplicates) in day-ahead prices. We rather model pairs of day-ahead and primary reserve prices in the same scenario set $s \in \mathcal{S}$. The copula based heuristic is used to model dependence in the scenarios. In the second

stage, a primary reserve commitment is decided upon with the opportunity cost of lost balancing market opportunities. Hence, for each s there must be a set of possible balancing market realizations $\omega \in \Omega^s$. The balancing market premiums are dependent on day-ahead price, and is therefore generated conditional on each day-ahead price scenario. Dependence internally in the balancing market volumes and premiums is also modelled with the copula-based heuristic.

Figure 6.2 outlines the concept. We use either quantile regression (day-ahead and primary reserve) or simulation (balancing premium and volume) to produce hourly predictive densities. Thereafter we treat each density as the marginal distribution of a multivariate joint distribution. That is, for day-ahead and primary reserve, we have 48 marginals (one for each hour in each market). Our objective is to pick $|\mathcal{S}|$ scenarios of the 48 variables that reflect their joint behaviour as effectively as possible. For this we use the copula heuristic.

The scenario generation algorithm is implemented partly in Matlab and partly in C++. In the following we will outline a high-level pseudocode for the algorithm. Our goal is clarify how the different building blocks from Chapter 6 fit into the procedure. Figure 7.1 illustrates the program.

1. Hourly predictive densities/cumulative functions are created with quantile regression on a hourly basis for day-ahead and primary reserve prices

2. The copula heuristic is used to create the desired number of day-ahead and primary reserve scenarios, with respect to each hourly density, the connection between hours, and the connection between markets. We use the empirical copula of the training period directly as target to model dependence

3. Balancing market data is resampled

4. Balancing market dependence is modelled using the empirical copula, and copula scenarios are generated before simulation

5. Appropriate AR/ARMA models are fitted to the log-transformed balancing market volumes and premiums and used to simulate a large number N times. Premiums have day-ahead price scenarios as explanatory variable, and in this way the linear dependence between day-ahead price and premiums is respected. Based on the simulations, we estimate the empirical cumulative distribution for volumes and premiums on hourly basis for the 12-36 hour forecasting horizon

6. Cumulative distributions for the balancing market are matched with copula scenarios to yield scenarios in terms of target variables. Back-transformation is performed; half the estimated variance is added to cope with bias from log transformation

7. Regulation state is simulated with the hour specific Markov model in each branch. Volumes and premiums in hours without relevant regulation are set to zero

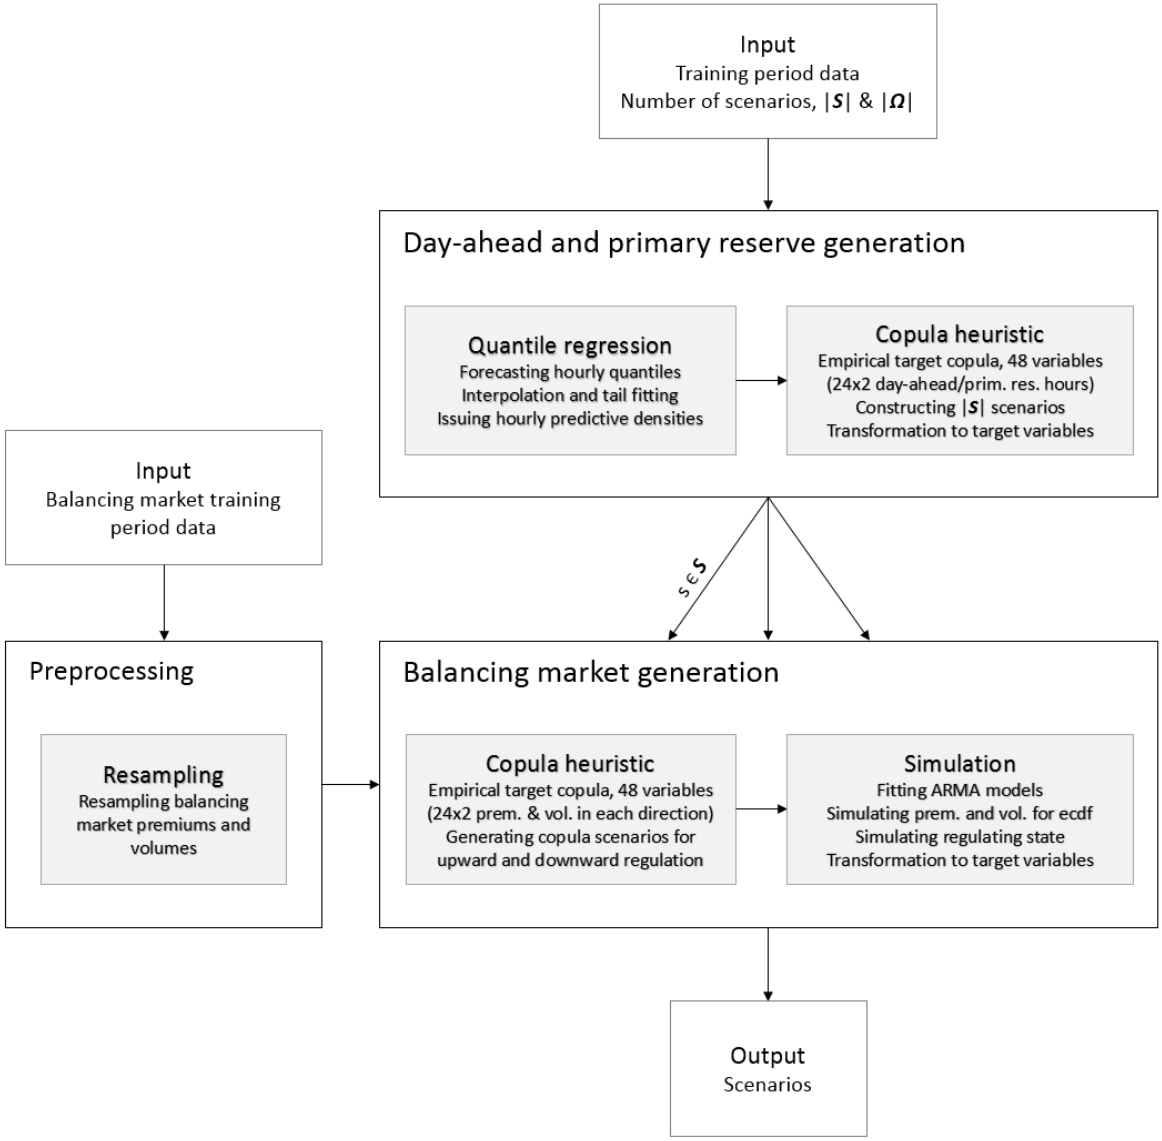

Figure 7.1: Scenario generation overview

7.3 Evaluation of scenarios

This section is devoted to testing of the scenario generation procedure. From Section 6.7 we already know that the predictive densities for each hour perform well. These densities are treated as marginals of a multivariate joint distribution. The copula heuristic ensures that each marginal is represented by $|\mathcal{S}|$ points uniformly distributed along the cumulative mapping. Hence, the only loss of information in the marginals stems from the discretization into scenarios, and not from the heuristic itself. We

therefore assume that the scenarios constitute an effective description of uncertainty in each individual variable given the number of scenarios. In 7.3.2 we determine the necessary number of scenarios by solving the planning problem multiple times with different scenario sets. It remains to assess dependence in scenarios in the multivariate setting (between hours and markets). Assessing dependence is interesting because it determines the daily price path and the trade-off between markets. The price path or price development between hours is of interest because it determines how to coordinate production inter-temporally. Start-up costs are inflicted when production restarts after stand-still.

7.3.1 Evaluating inter-temporal and inter-market dependence

We want to find out if the scenarios represent plausible price paths. There are multiple reasons for this. The copula heuristic solves an assignment problem to minimize deviation between the target copula and the copula defined by the scenarios. This is done heuristically by matching margins as in the bivariate case, and then subsequently add margins until all margins are added. No higher order dependencies are taken into account. In addition, the assignment problem is solved by a greedy algorithm. The loss of quality, i.e. the deviation from the target copula, may result in an altered description of dependence in the scenarios.

Visual evaluations in this section should only be regarded as a rough estimate of performance, and would reveal nothing but severe flaws in the scenario generation. Conclusions are drawn accordingly. If one is primarily concerned with the performance of scenarios subject to a stochastic problem, one can run stability tests and compare to some large reference scenario set. This follows in subsection 7.3.2.

To evaluate the suggested scenario generation method's ability to create inter-hourly price paths, the following methodology is used. 40 daily scenarios are generated out-of-sample for a period of 50 days in 2016. Daily price means are calculated for all scenarios and real prices, μ_{ms} and $\tilde{\mu}_m$ for each day in the test period. The means are subtracted from each of the hourly prices in the respective scenarios. The same is done for the realizations, the realized daily average price is subtracted from the hourly price, see Figure 7.2. Thereafter, we plot empirical distributions for the resulting variables $\rho_{mts} - \mu_{ms}$ and $\tilde{\rho}_{mt} - \tilde{\mu}_m$ for all test days. The resulting distributions should coincide if scenarios have similar fluctuations around daily means as real prices. Figure 7.3 and 7.4 show the results. The results are generally satisfactory. Notice that distributions seem to have approximately equal support and probability density. The primary reserve realized prices are unevenly distributed, and more days should have been included in the test period. However, primary reserve prices often take values that are multiples of ten because of bidding behaviour. This might cause some of the spikes in the figure.

Next, we seek to evaluate price development from one hour to the next. How much the price fluctuates between consecutive hours, implicitly determines risk of unplanned start-up costs. We therefore need to verify that the scenarios have similar properties

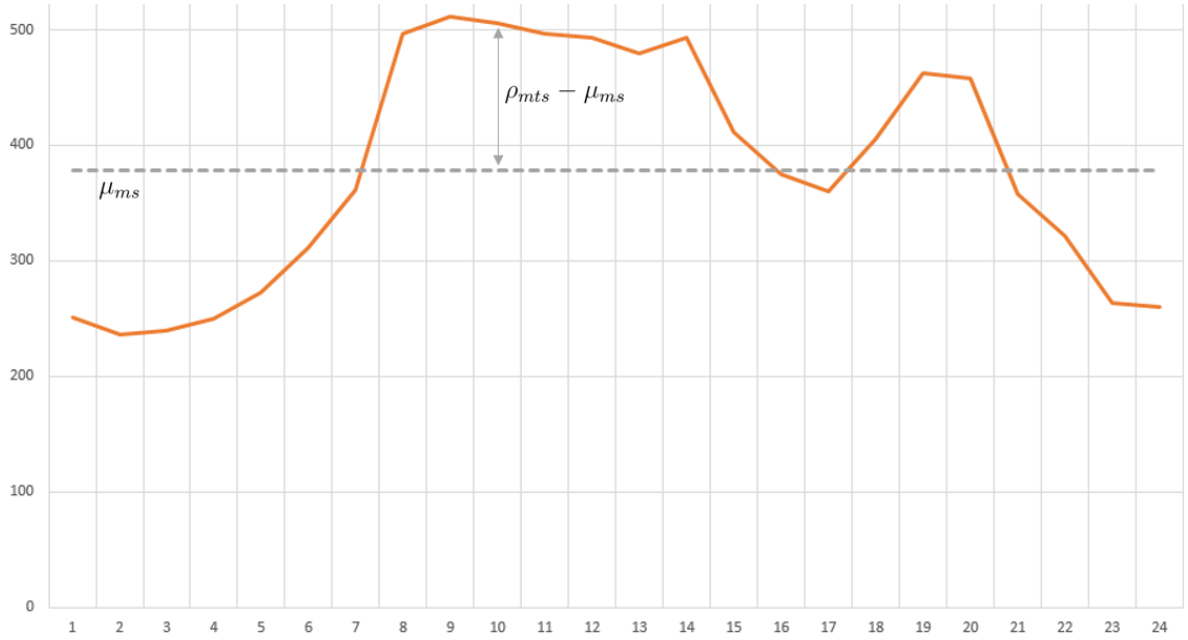

Figure 7.2: Price fluctuations around daily mean

compared to real prices. This is done by plotting scatters for ρ_{mts} against $\rho_{m,t+1,s}$ and comparing with realized prices. Figure 7.5 shows the results. For day-ahead prices results are very good. The scenarios seem to exhibit the same behaviour as the real prices; the variance at different price levels coincide. For primary reserve prices, there might be a tendency to produce scenarios with too extreme fluctuations from one hour to the next. Due to lack of realized prices, it is hard to conclude. Primary reserve has enourmous spikes in some periods, and are very stable in other periods. For additional reference, we compared the scenarios to the whole set of historical prices. In that case, the results were very similar.

Finally, we visually check the relationship between the day-ahead and primary reserve market. Figure 7.6 shows the generated scenarios plotted against realized prices from the period (50 days in 2016). Notice how the non-linear relationship between the prices is portrayed well by the scenarios. The supports of the real prices and the scenarios seem to coincide. In sum, visual inspections indicate that non-linear dependence seems to be modelled successfully. This is a property we would not have been able to capture using only correlations. Similar scatter plots were made to assess the dependence between balancing market volumes and premiums, and can be found in Figure 7.7. There seems to be some tendency to underestimate the presence of spikes.

7.3.2 Evaluating stability

The quality of a scenario tree should be evaluated by how good decision support it provides rather than by statistical properties alone. Performance of a scenario generation

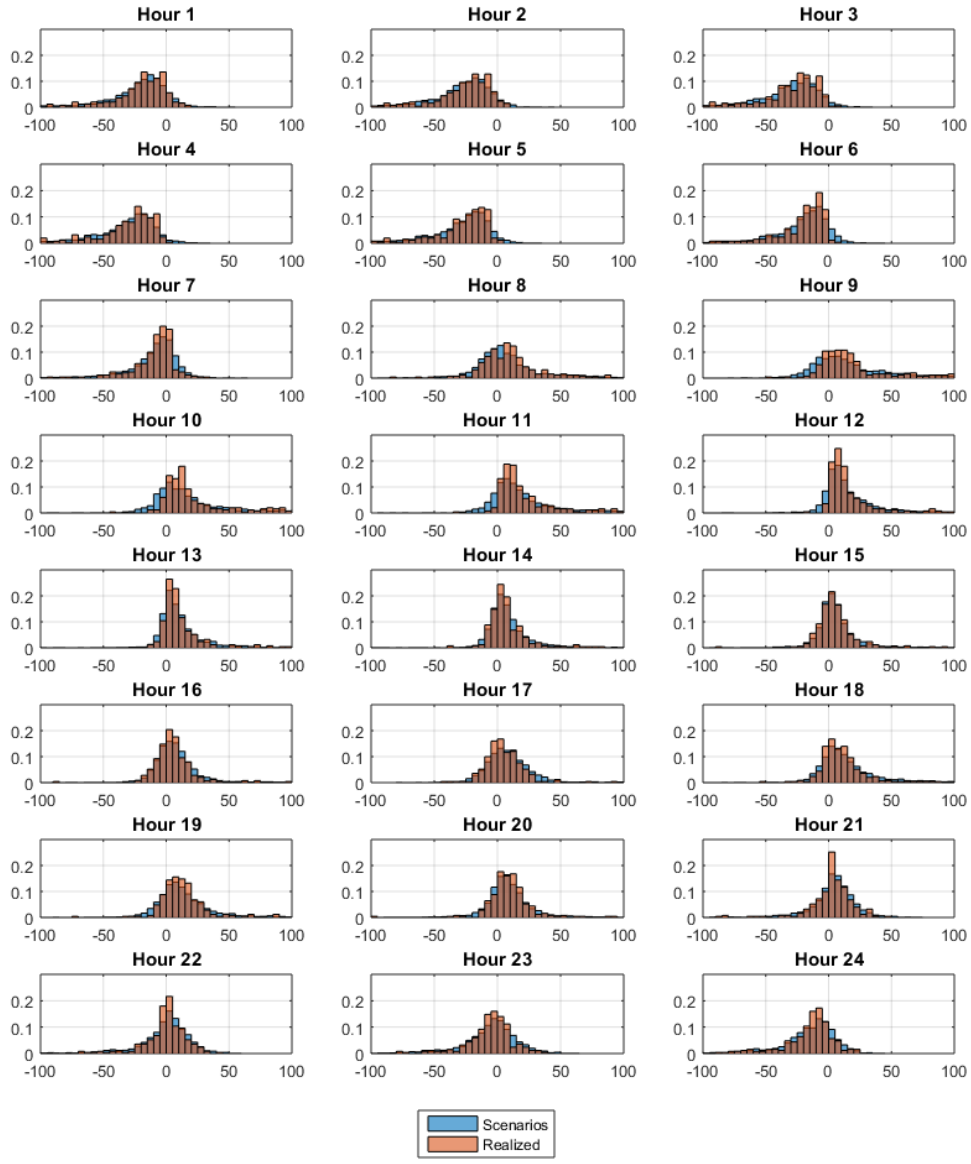

Figure 7.3: Fluctuations of day-ahead prices around daily mean for scenarios and real prices

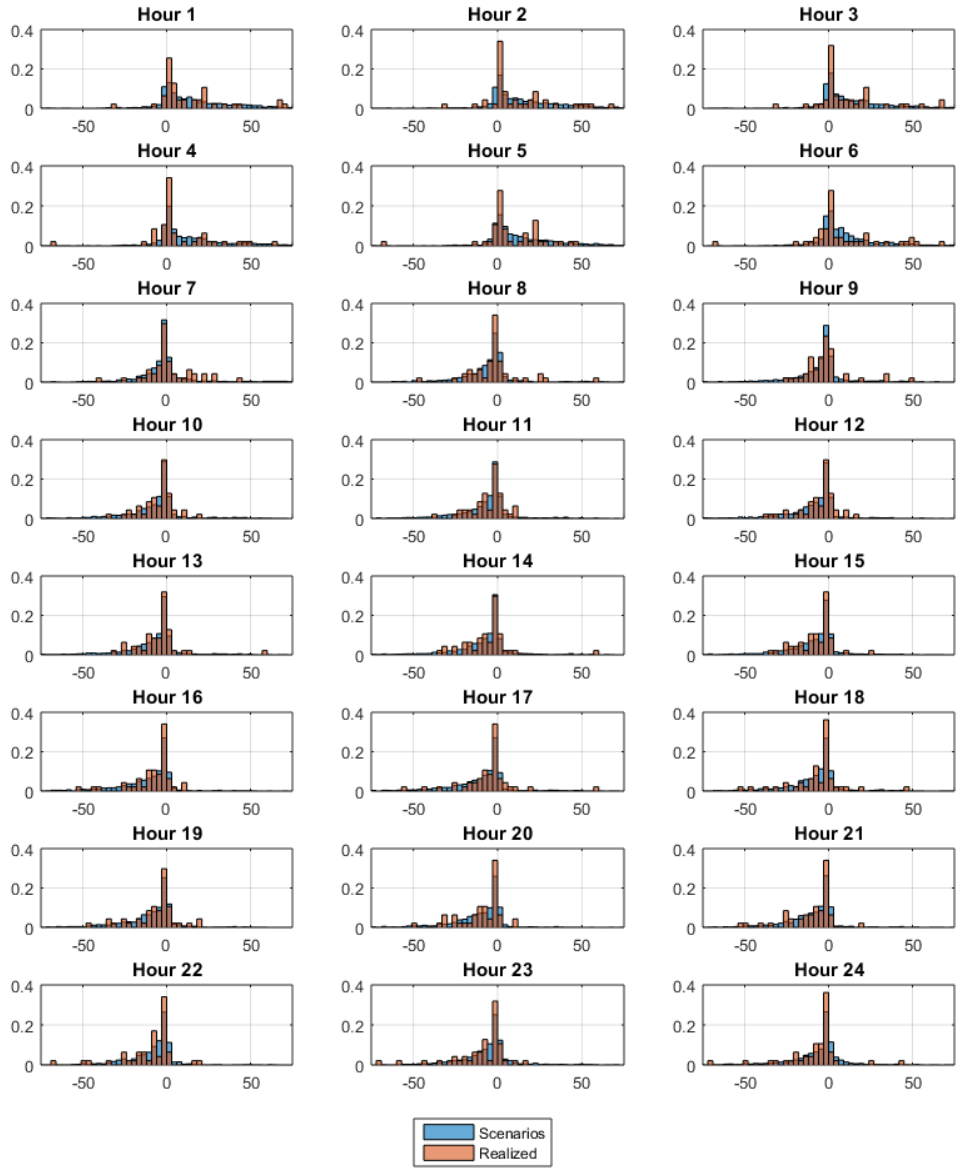

Figure 7.4: Fluctuations of primary reserve prices around daily mean for scenarios and real prices

Figure 7.5: Hour to hour price fluctuations. Blue dots denote scenarios, red circles denote realized prices

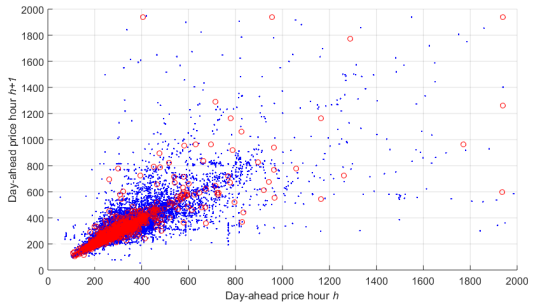

(a) Day-ahead price

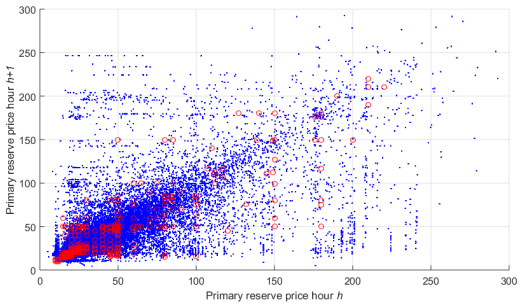

(b) Primary reserve price

Figure 7.6: Generated day-ahead and primary reserve scenarios. Blue dots denote scenarios, red circles denote realized prices

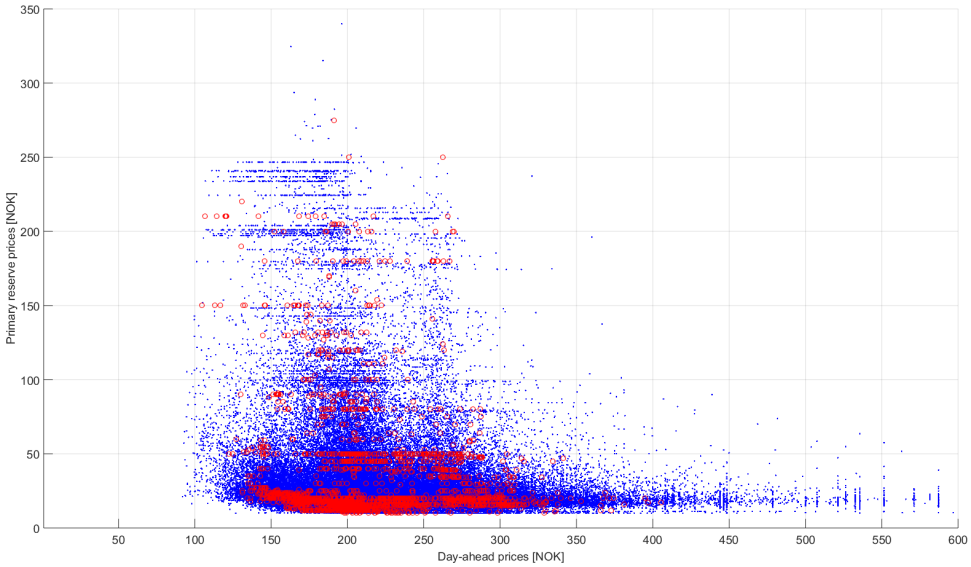

Figure 7.7: Balancing market premiums and volumes. Blue dots denote scenarios, red circles denote realized prices and volumes

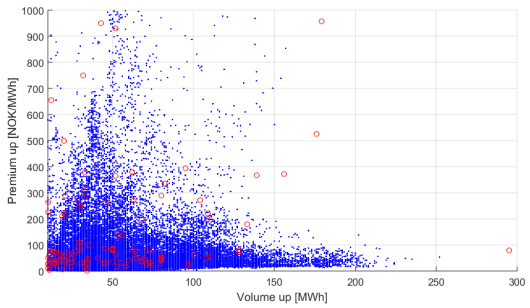

(a) Upward balancing

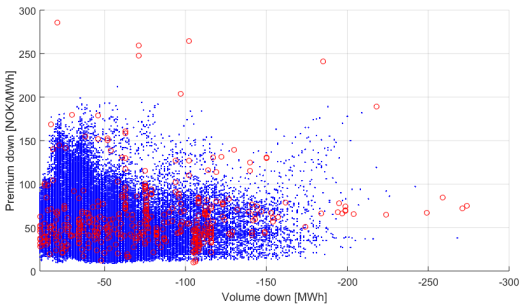

(b) Downward balancing

method is therefore often evaluated by subjecting it to the optimization problem at hand. To evaluate the quality of a scenario generation method we can run an in-sample stability test, as defined in the work of Kaut and Wallace 2007, in which we require the following relationship between objective functions on different scenario trees

$$\max_{x \in X} \hat{F}_a(x) \approx \max_{x \in X} \hat{F}_b(x) \tag{7.3.1}$$

where $\hat{F}_a(x)$ denotes the objective function evaluated on scenario tree a. In practice this means that any two scenario trees representing the same planning period should yield an approximately equal objective value. If there is randomness in the scenario generation method, we require in-sample stability between different trees of the same size. If there is not much or no randomness in the scenario generation method, we can still compare objective values when changing the size of the scenario tree.

Another form of stability testing is the out-of-sample test, also defined by Kaut and Wallace, which is carried out by evaluating the relationship described by

$$\hat{F}_a(\hat{x_b}^*) \approx \hat{F}_b(\hat{x_a}^*) \tag{7.3.2}$$

A solution $\hat{x_b}^*$ from running the model with tree b is imposed on tree a, and vice versa. Generally, an out-of-sample test indicates the performance of a solution given a different realization of stochastic parameters. Solutions with high out-of-sample stability tend to be robust, but then again, if the scenario generation method is consistent, there is little variation between the scenario trees. Hence a solution from one tree would be expected to perform well on another tree. There is no simple relationship between in-sample and out-of-sample stability, but both are required in a good scenario tree method (Kaut and Wallace 2007).

In-sample stability

We have tested the scenario generation method in this thesis in-sample and out-of-sample. In-sample testing is done by running the model several times with different scenario trees for January 8, 2016. The decision model from Chapter 4 has been run with different scenario trees of different sizes; 20 times for each size. Note that we use the simplified modelling for the value of water according to 4.2.17, yielding net daily profit in the objective. Not all scenario tree size combinations are tractable; for large numbers of day-ahead/primary reserve scenarios, the number of possible balancing scenarios per branch is gradually limited. The mean objective value from the 20 runs is plotted in Figure 7.8 as a function of the number of scenarios. Notice that the number of day-ahead/primary reserve scenarios affects the objective value significantly. As the number of day-ahead/primary reserve scenarios approaches 100, the objective value seems to stabilize perfectly with respect to the number of scenarios. Notice also that the number of balancing market scenarios has negligible effect on the mean of the objective value. This might be due to the small size of the balancing market compared to the day-ahead market.

The objective value is extremely stable between trees of the same size. Figure 7.9 plots the standard deviation in the objective for the 20 runs. The standard deviation does not exceed 0.4 % of the objective value. The standard deviation decreases with respect to the number of day-ahead/primary reserve scenarios and balancing market scenarios, but this increased stability is likely not very important. Stability from run to run is due to the low inherent randomness in the generation method [1]. The Markov model is a source of randomness. However, the important point here seems to be that the objective value changes as we increase the number of day-ahead/primary reserve scenarios. No reference scenario tree is available, but the 200 day-ahead/primary reserve tree is the largest tree we solved for. Using this objective value as reference, other objective values vary between -2.5 and 10.0 % as the number of scenarios increase.

Out-of-sample stability

Next, we check out-of-sample stability. This has been done by imposing the first stage solutions (bid curves for day-ahead) from the smaller scenario trees to a large tree of 200 $\times$ 2 scenarios, and resolving. We want to assess what number of scenarios is necessary to provide a good solution relative to the larger reference tree. Figure 7.10 shows the results. The variance in performance among the 20 trees of the same size is negligible. Changing the number of balancing market scenarios per branch neither affects out-of-sample performance. Any number of day-ahead and primary reserve scenarios above $\sim$20 yields good solutions, deviating less than -0.5 % from optimum.

Considering in-sample and out-of-sample stability our choice is to use 40 day-ahead and primary reserve scenarios and 4 balancing market scenarios per branch for the case

[1] Only when the deviation from the target copula is equal for different rank assignments, one rank will be chosen randomly. See Appendix D

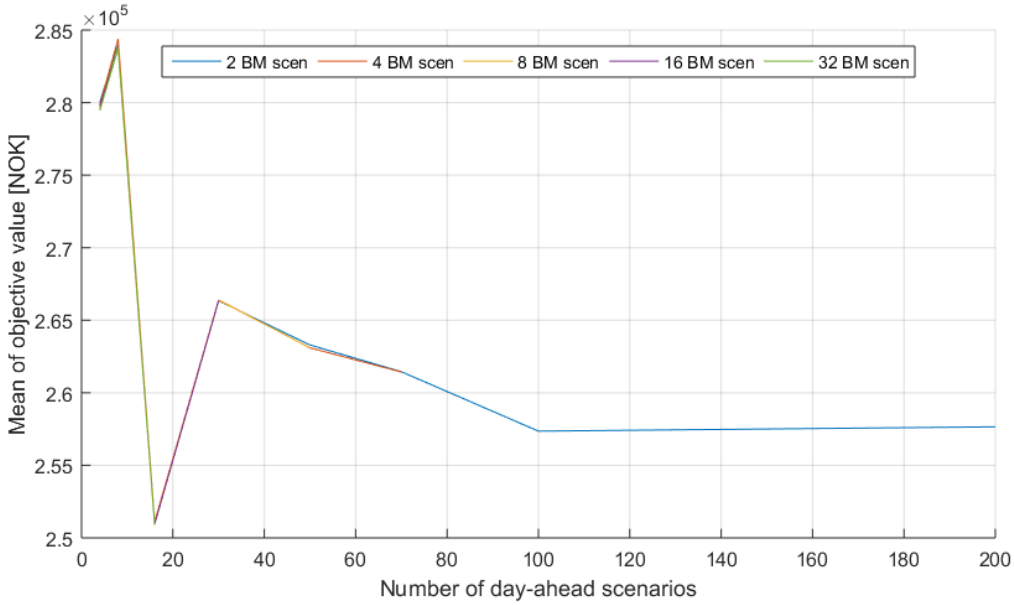

Figure 7.8: Mean of objective value for different scenario tree sizes. The different colors represent the number of balancing market scenarios per branch

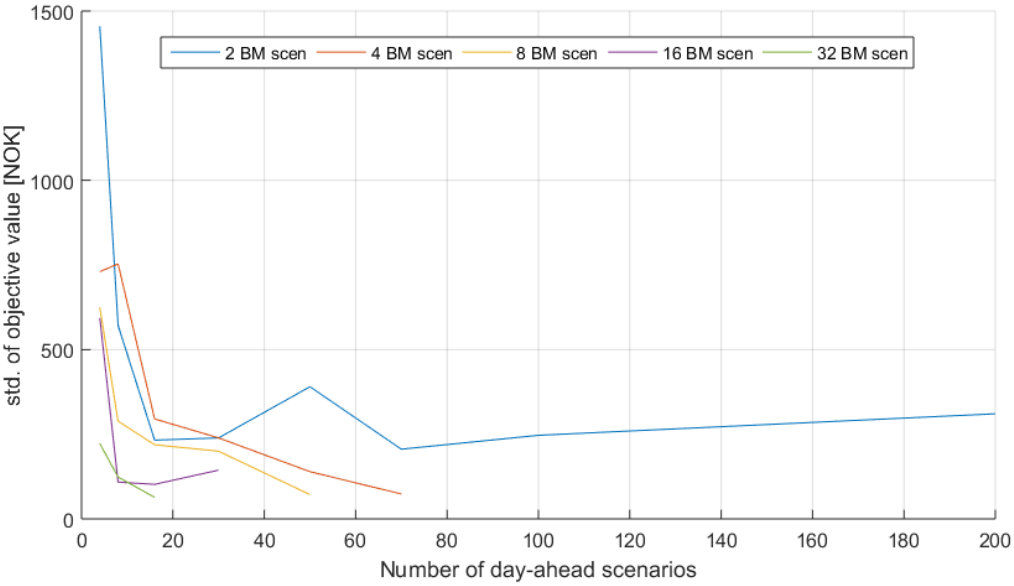

Figure 7.9: Stability (standard deviation of each set of 20 runs) of the objective value for different scenario tree sizes

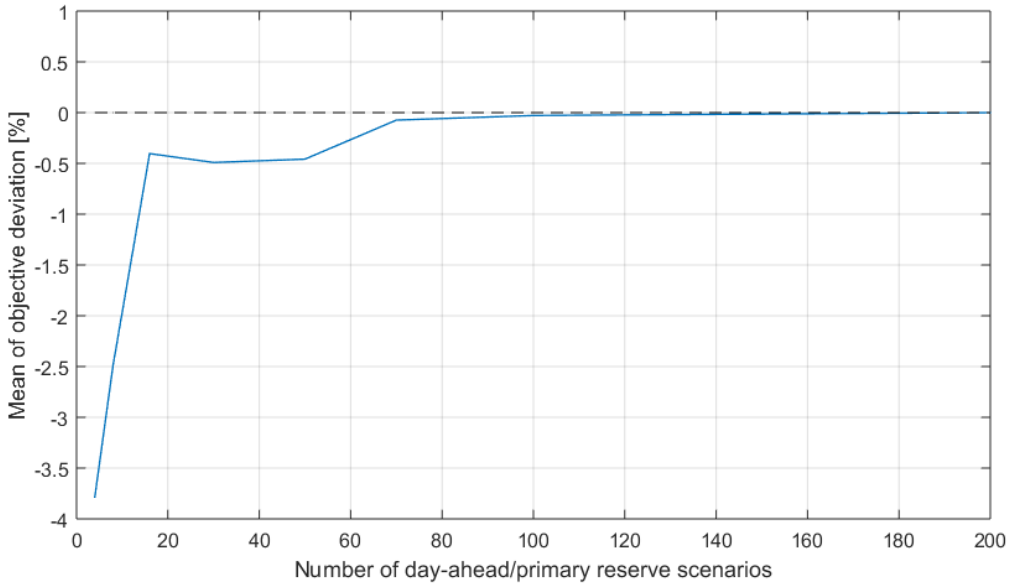

Figure 7.10: Out of sample stability. Bid curves constructed for different scenario tree sizes are subjected to a reference tree with 200 day-ahead/primary reserve scenarios and 400 balancing scenarios (200×2)

study in Chapter 8. The problem can be efficiently solved with this tree size, and still provide very good solutions.

Chapter 8

Computational case study

The goal of this chapter is to assess whether the hydpropower producer should take primary reserve and balancing market opportunities into account upon bidding into the day-ahead market. Section 8.1 describes two planning regimes, that henceforth are referred to as *coordinated* and *sequential* planning. Coordinated planning takes subsequent market opportunities into account when bidding, while sequential planning does not. The possible gains from coordinated planning will be investigated under three control variables, all introduced in the next section. In subsection 8.1.1 we propose a measure for deviation between water value and the day-ahead price, price-water-value squared deviation (PWSD). The PWSD is used to identify interesting sets of historical days for testing. This measure is actively referred to throughout the case study. In subsection 8.1.2 we present and discuss the granularity of the planning horizon. We discuss two settings; hourly and sub-periods. A sub-period is a set of hours for which a generator is planned either on or off the entire sub-period. This practice reduces the number of binary variables and hence decreases solution time significantly. Finally, in subsection 8.1.3 we discuss the effects of an increasing production portfolio size. Section 8.2 addresses assumptions and problem parameters that are independent of planning framework; boundary conditions, production functions, market shares etc. The results from the case study are presented in section 8.3. All decision models involved have been implemented in Xpress-IVE with Xpress-Mosel and solved by Xpress-Optimizer. All test results were obtained on 3.20 GHz Intel Core PCs with 16 GB RAM.

8.1 Case study framework

Coordinated bidding

Figure 8.1 illustrates the case study framework. The model from Chapter 4 (Model 1) is used to calculate a *coordinated* day-ahead bid curve. Next, this curve is fed into another model (Model 2) along with the realized day-ahead price, resulting in a production commitment. Furthermore, Model 2 receives primary reserve and balancing market

scenarios[1] and calculates a bid curve for the primary reserve market. Finally, in a third model (Model 3) we reveal the primary reserve prices, resulting in a commitment. Also, balancing market volumes and premiums are revealed. The producer simply sells the balancing reserves that are found profitable and feasible under the incurred commitments, and the profit of the coordinated framework is output.

Sequential bidding

The profit from coordinated bidding is at all times compared to the equivalent resulting from a *sequential framework*. The sequential models calculate bid curves based on scenarios for the next upcoming market clearing. No scenarios for subsequent markets are used. This way we get a benchmark to assess the value of coordinated planning. All involved models can be found for inspection in Appendix C.

The next sections describe the control variables of the case study. These control variables will be used to determine when possible gains from a coordinated planning can be expected.

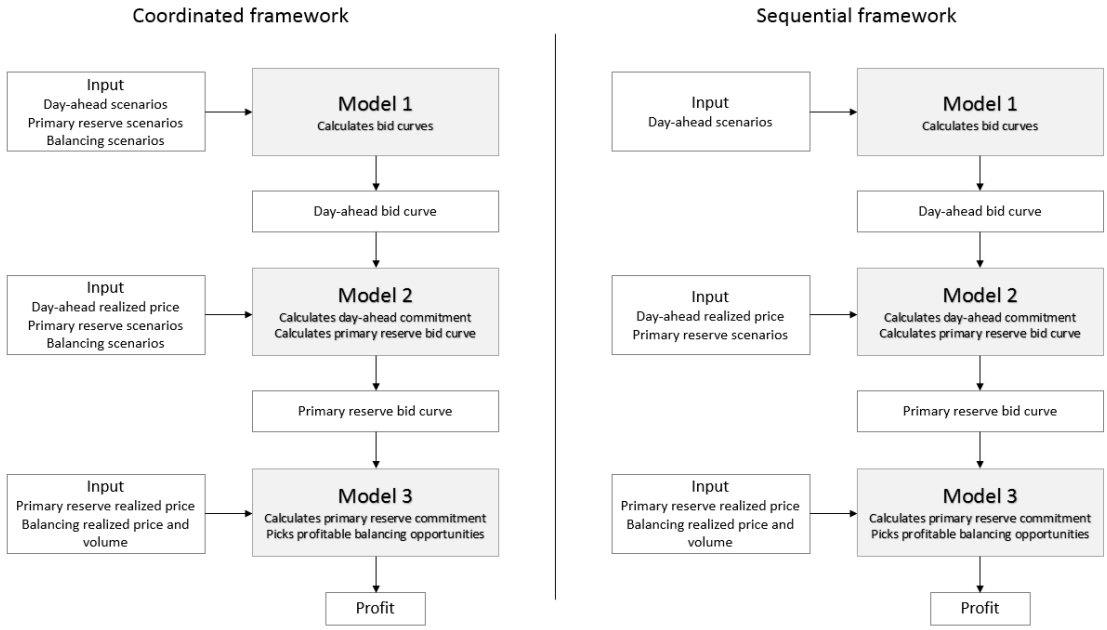

Figure 8.1: Case study framework

[1]The same primary reserve scenarios are used again, although the prediction should be somewhat less uncertain knowing the day-ahead price. One should in practice generate new scenarios conditional on the realized day-ahead price. This would require a thorough study of post-spot reserve price forecasting, which is not treated in this thesis.

8.1.1 Effect of price-water-value-deviation

Eriksrud and Braathen 2012 identify different situations in which coordinated bidding may be valuable. Relations between water value, day-ahead prices, reserve market prices and volumes determine these situations. Eriksrud and Braathen only consider the balancing market, but a similar argument can be made for primary reserves. The water value is here represented by V [NOK/MWh]. The first four relate day-ahead price (ρ_1) and downward balancing. The downward regulation premium is denoted by (ρ_4).

(a) $\rho_1 \ll V$: It is not profitable to bid into the day-ahead market, and no reserve opportunities can compensate. No bids will be placed.

(b) $\rho_1 < V$: It is profitable to bid into both markets if one expects high volumes/premiums in BM, and not profitable to bid into any market if one expects low volumes/premiums in BM.

(c) $\rho_1 > V$ and $\rho_4 > (\rho_1 - V)$: It is profitable to bid into the day-ahead market and one should bid into BM whenever possible.

(d) $\rho_1 > V$ and $\rho_4 < (\rho_1 - V)$: It is profitable to bid into the day-ahead market and not into BM, because it yields smaller profits than the pure day-ahead profit.

These four relations are possible in the case of upward regulation. The upward regulation premium is denoted by (ρ_3).

(e) $\rho_1 < V$ and $\rho_1 + \rho_3 < V$: It is not profitable to bid into any of the two markets.

(f) $\rho_1 < V$ and $\rho_1 + \rho_3 > V$: It is not profitable to bid into the day-ahead market, and it is profitable to bid into BM.

(g) $\rho_1 > V$ and $\rho_1 + \rho_3 > V$: It is profitable to bid into both markets. If one anticipates high volumes in BM, one might be willing to reduce the day-ahead volume to achieve higher volumes in BM.

(h) $\rho_1 \gg V$ and $\rho_1 + \rho_3 \gg V$: It is very profitable to bid into the day-ahead market, and the producer will not risk reducing its total commitments in order to obtain flexibility in BM.

From this discussion it is clear that coordinated bidding theoretically should yield higher gains whenever the day-ahead price does not deviate much from the water value. We therefore seek to quantify such a deviation. Because we are considering several reserve markets, we only use the day-ahead price to quantify price-water-value-deviation.

We define a simple measure for the purpose of choosing test days, and we refer to it as the PWSD, absolute price-water-value squared deviation:

$$PWSD = \sqrt{\frac{1}{24}\sum_{t=1}^{24}(\rho_{1t} - V)^2} \tag{8.1.1}$$

In addition, we define a slightly modified measure that accounts for whether the deviation is mainly negative or positive; net price-water-value-deviation NPWD.

$$NPWD = \begin{cases} + & \text{if } \frac{1}{24}\sum_{t=1}^{24}(\rho_{1t} - V) > 0 \\ - & \text{otherwise} \end{cases} \tag{8.1.2}$$

The motivation for defining these measures is to identify which days that might yield higher or lower gains from a coordinated framework. The terms are squared to penalize large deviations in some hours, compared to the alternative of smaller deviations throughout the day. Generally, smaller deviations throughout the day is more interesting in terms of coordinated bidding. We utilize the NPWD to determine whether we expect profits to be positive or zero/negative, i.e. whether the price is generally higher or lower than water value.

8.1.2 Effect of planning horizon granularity

We do not include block bids in this thesis, but model only single hourly bids. However, we allow for partitioning of the planning horizon into sub-periods, $h \in \mathcal{H}$ with corresponding operating hours $t \in \mathcal{T}^h$. The generator on/off-decision variable, $u_{ighs\omega}$ therefore requires that a generator in a given scenario is either on or off during the entire sub-period. The motivation for this choice is twofold. As shown in Chapter 5, all markets are governed by different price regimes during the day. Day-ahead prices are generally low during night hours, while primary reserve prices are high, downward regulation is more likely in the morning etc. This also has practical implications for most producers. If day-ahead prices are too low to be profitable in night hours, the producer is probably not committed in any of the night hours. Start-up costs play an important role, because profitability in some hours is not enough to overcome additional start-up costs. This is often the case in practice; generators are either on or off for entire sub-periods. Planning under this assumption drastically reduces the number of binary third stage variables, and thus solution time. This approximation is necessary to be able to solve the models with increased portfolio sizes. If a generator runs in a certain sub-period in a scenario, the bid curve will be made such that a commitment is granted in all hours of the sub-period. This has an important consequence. Given a realized price in any hour, the sub-period based solution (if production is planned) will always have equal or higher probability of being committed compared to the hourly model solution.

Production is planned in sub-periods. However, realized production may look very different. Sub-periods are only used in the planning phase of Model 1. When the real prices are revealed, any commitment is possible, and the next models (Model 2 and Model 3) allow for production in any hour according to the commitment, without regard to sub-periods.

To conclude, the introduction of sub-periods represents a restriction to the original problem, but at the same time it may have a positive effect on unplanned start-up costs. Our main motivation is still that it allows us to solve the planning problem to optimality within reasonable time also for 2 or 3 watercourses, so that we can investigate the influence of portfolio effects on coordinated bidding. Before that we also run tests with sub-periods when the portfolio only consists of one watercourse. The results are presented in chapter 8.3.1, in succession to those obtained from running with one watercourse and full hourly granularity. The results are compared qualitatively in order to evaluate the effect of the sub-period restriction before moving on to two and three watercourses.

8.1.3 Effect of increasing production portfolio size

We start off by investigating how the coordinated framework performs relative to the sequential framework when we only have one watercourse at disposal. The framework is run with both hourly granularity and with sub-periods. Thereafter, we wish to investigate possible gains from coordinated bidding when we add more production resources (watercourses) to the planning problem. It is only possible to solve these models to optimality within reasonable time ($\sim$20 min per run) when the models are run with sub-periods.

Our motivation is to investigate how the value of coordinated bidding changes when recourse flexibility changes. Theoretically, coordinated planning outperforms sequential planning when reserve opportunities are predictable, and little flexibility is at hand to respond to a reserve opportunity not accounted for in the first-stage decision.

Obviously, predictability of reserve opportunities is not a function of the portfolio size. On the other hand, when more production resources are added, possible reallocation (and thus recourse options) of committed production increases. For instance, a day-ahead obligation can be redistributed to some of the generators while one generator can be turned off to respond to downward balancing demand. Another example is primary reserve commitments. If there is only one generator in the portfolio, a primary reserve obligation restricts all large downward balancing events. Because the producer is committed, there is no way to turn the generator off in the case of large downward balancing volumes. With two or three generators, primary reserve obligations can possibly be reallocated, and flexibility increases.

To sum up, flexibility of recourse actions increases when additional production resources are added to the portfolio. Note that flexibility increases in both frameworks. Subsequent models can react to realizations and redistribute production in the coordi-

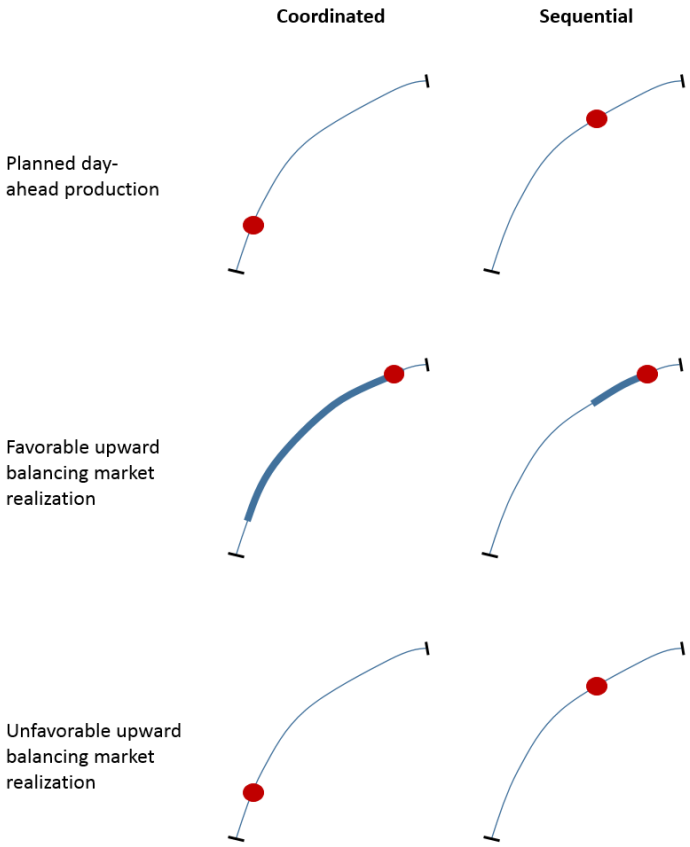

Figure 8.2: Single watercourse recourse options illustrated on a generator production/discharge-curve

nated as well as the sequential framework. Our question should be which framework benefits the most from increased recourse options. To simplify the discussion we consider only the day-ahead and the upward balancing market first. We will also distinguish between the cases when the reserve market realization turns out to be as planned, and when it does not. We will refer to these situations as favorable and unfavorable realizations, respectively. First we consider a portfolio size of 1, that is only one generator. Figure 8.2 illustrates. The day-ahead price is weakly profitable, so day-ahead production is planned in both frameworks. However, large volumes are expected in the balancing market, so the coordinated model retains some capacity to respond to these opportunities. If the expectation turns out to be correct, the coordinated model will perform better than the sequential model, earning a premium on the capacity sold to the balancing market. On the other hand, if there is no balancing volume, the coordinated model has no option but to deliver the obligated day-ahead production at low efficiency.

We now move on to the case in which we have three generators available. Everything else is the same; day-ahead prices are barely profitable, and one expects large amounts

of upward regulation. The sequential framework will plan production until marginal cost on all generators. The coordinated framework might withhold some capacity to respond to balancing market opportunities. If there are in fact large balancing volumes, the coordinated framework will benefit compared to the pure sequential. On the other hand, if there is no upward regulation, we have more recourse options than for the single watercourse. If there is enough capacity on two of the generators, one generator can be turned off to produce at higher total efficiency.

In terms of increased flexibility and recourse, there are two points we want to address: I) increased recourse increases ability to react to profitable predictable (but still unplanned in the sequential framework) realizations, and II) increased recourse increases the ability to react to costly unpredictable (obviously unplanned in both frameworks) realizations. The practical implications are twofold:

- When the sequential model constructs bad bid curves compared to the coordinated model (due to favorable reserve realizations), the sequential framework has more recourse (with a large portfolio) to compensate for the bad bid curve, and catch up with coordinated performance. That is, the expected gain from having planned coordinated instead of sequentially decreases in the case of a favorable reserve realization.

- When the coordinated model constructs a bad bid curve (due to unpredictability), the coordinated model has more recourse to mitigate unplanned costs. Put more formally, the expected cost of making a bad coordinated decision decreases.

To conclude, it is not *a priori* clear whether gains from coordinated bidding will increase or decrease with increasing portfolio size; it depends on the specific situation and which of the two aforementioned effects has more power. It is an interesting question, and results from the case study will aid a conclusion.

8.2 Case description

This section describes all problem specific parameters that are input to the decision model in chapter 4, and all assumptions made.

Restrictions on traded volumes

We assume a price-taking role in all of the modeled markets. In the day-ahead market the producer is committed to produce whatever volume that is accumulated for bid curve price points below the realized price. This is a reasonable assumption because the total demand in the day-ahead market is very large. In the primary reserve market the producer can commit to any volume below the realized price. Primary reserve commitments are however limited strictly by the minimum static setting on each generator. In these two markets the demand is rather predictable and continuous.

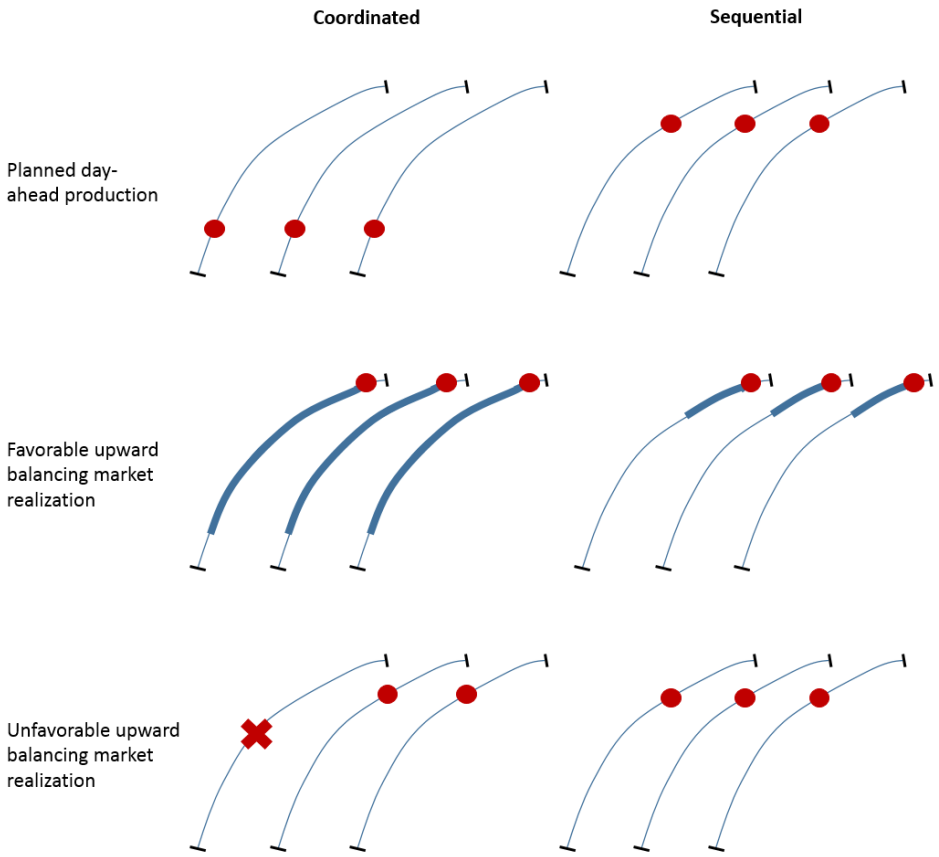

Figure 8.3: Portfolio recourse options illustrated on a generator production/discharge-curve

In the balancing market on the other hand, demand is inherently unpredictable and clustered, and varies from low to high volumes. In the case of small volumes demanded, one producer can very well sell the entire volume. In this case the producer would act as price maker. When volumes are larger it is likely that several producers contribute to delivering this volume. ν_{mts} forecasts the entire demand in hour t in scenario s in NO3. It is therefore unsound to let one producer deliver the entire volume. This would imply that no other producers placed bids beneath the market price, and it would violate the price taker assumption. Our approach is simply to restrict the balancing volume sold by the decision model to a market share multiplied by the total volume traded in a balancing hour. NTE delivered roughly 15 % of total production in NO3 in 2016. We therefore set the market share to 15 % in the balancing market as well. This might be an overly restrictive constraint, and we will discuss its implications later on.

Bid curve price points

Bid curves have to be decided explicitly for the day-ahead (Model 1) and primary reserve market (Model 2). To retain linearity of the model, bid points are set as constants before the model is run. Price levels vary with respect to test day so we update the bid points for every planning day of the test period. This is done by placing $|P|$ bid points such that the price scenarios are distributed approximately uniformly between the bins defined by neighbouring bid points (according to market rules). We choose $|P| = \frac{|S|-2}{2}$, which gives the highest number of bid points, with which non-anticipativity of the bid curve with respect to the scenarios is still ensured (Löhndorf et al. 2013). If interpolation causes too large or too small commitments, production is set to the corresponding largest or smallest feasible value.

Watercourses

In order to be able to solve the optimization problems to optimality within reasonable time, very "simple" watercourses from NTE's portfolio have been wisely chosen. By simple we mean that the watercourses each consist of one very large reservoir and one generator downstreams, after which the water is not any longer of relevance to the producer. The three watercourses are not interconnected in terms of water flow. Hence, each of the watercourses Bogna, Mosvik and Ormsetfoss are modeled as shown in figure 8.4. This despite that there is in fact a restriction on the water level change in the lake (Snåsavatnet) to which water dispatched at Bogna flows. We nevertheless defend our assumption with the fact that Snåsavatnet is the sixth largest lake in Norway, and has a water level that is hard to influence. More importantly, this assumption does not discriminate the performance of the coordinated framework compared to the sequential framework.

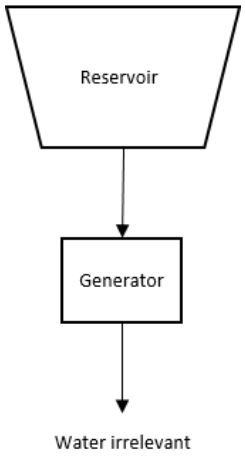

Figure 8.4: Simple watercourses

The value of water

Further, the generators' discharge are considered small relative to the respective reservoir volumes. Hence, the head cannot change drastically over the time of a daily planning horizon. According to NTE, it is therefore reasonable to assume that the water values can be treated as constants over the 24 hours of planning, independent of inflow and discharge. Therefore, instead of modeling the value of the water remaining in a watercourse with the following restriction

$$w_{ks\omega} \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk}(v_{j|T|s\omega} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, s \in \mathcal{S}, \omega \in \Omega^s \tag{8.2.1}$$

and adding the summation of w_{ks} to the objective function that is to be maximized, we require

$$\Delta w_{ks\omega} \geq \sum_{i \in \mathcal{I}^k} \sum_{t \in \mathcal{T}} W_k d_{its\omega} \qquad k \in \mathcal{K}, s \in \mathcal{S}, \omega \in \Omega^s \tag{8.2.2}$$

to be satisfied, and we subtract $\Delta w_{ks\omega}$ from the objective function.

Generators

The relation between water discharge and production in each generator is modeled using six linear cuts, assuming an average water level in the respective reservoirs. The production cuts for all involved power plants are displayed in figure 8.5. The set of production cuts is concave for all generators. Some of the cuts do not intersect below zero on the discharge axis. This is handled by introduction of the binary decision variable $u_{ihs\omega}$, see Chapter 4. NTE has estimated the cost of a start-up to 4000 NOK. This cost is incurred each time one of the generators starts running after stand-still.

Estimating each markets' contribution to profits

In addition to total profit, we estimate the profit contribution from each market. Production efficiency (amount of water used to produce a certain quantity of power) is a function of all decisions made by the producer. We therefore calculate hourly water cost per megawatt after final allocation of production. We use this number to estimate water cost in all markets with actual power delivery.

Water cost per megawatt is calculated as

$$V_t = \frac{\sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{I}^k} W_k d_{it}}{x_{1t} + x_{3t} - x_{4t}} \qquad t \in \mathcal{T} \tag{8.2.3}$$

or $V_t = \rho_{1t}$ when the denominator is zero and there is downward regulation. Primary reserve commitment is not included because net water usage is zero for this market. As an example, profit from upward regulation can now be calculated according to

$$\pi_t^{BMup} = (\rho_{1t} + \rho_{3t})x_{3t} - V_t x_{3t} \qquad t \in \mathcal{T} \tag{8.2.4}$$

Planning horizon granularity

When running with a coarser planning granularity than an hourly resolution, we always use the same three sub-periods; $\mathcal{T}^1 = \{1..7\}$, $\mathcal{T}^2 = \{8..19\}$ and $\mathcal{T}^3 = \{20..24\}$. These sub-periods are all carefully chosen. $\mathcal{T}^1$ is chosen due the day-ahead and primary reserve price levels mainly. In these hours, day-ahead prices are generally low while primary reserve prices are high. This may effectively increase the price of each day-ahead unit sold, rendering production profitable. Besides, as can be seen in Figure 5.5, the probability of down-regulating is higher in the morning hours. Ergo, planning in a coordinated manner could provide production in these hours, that can be profitably reduced in case of downward regulating volumes. The day-ahead price (Table 5.1) often peaks in the eight hour and then stays quite stable above the daily mean until somewhere around the nineteenth hour. Profitable production is thus likely in the second sub-period, hence the choice of $\mathcal{T}^2$. The day-ahead price drops towards late evening and night hours, so the third sub-period $\mathcal{T}^3$ contains these hours.

Table 8.1: Number of test days

PWSD	<20	<20	20-40	20-40	40-60	>60
NPWD	+	-	+	-	+	+
Bogna	43	34	59	42	28	36
Mosvik	29	39	46	58	35	24

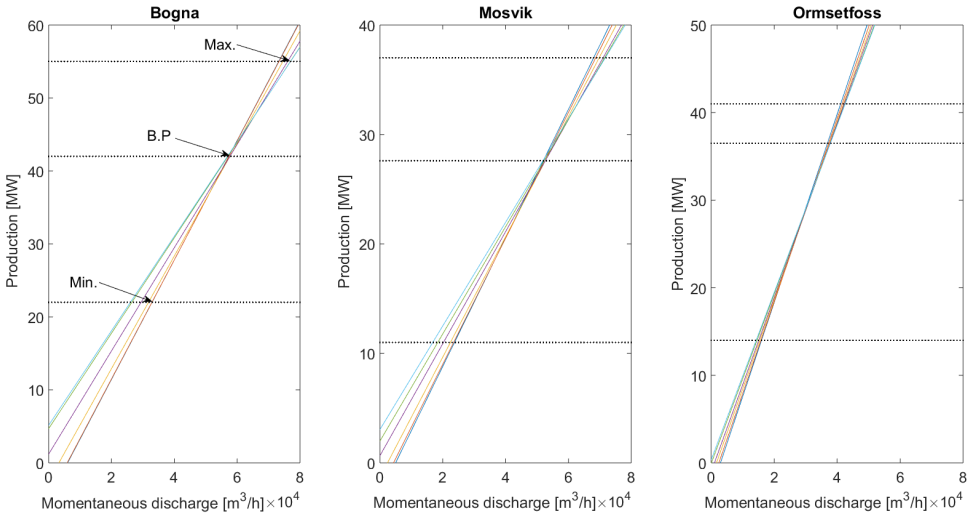

Figure 8.5: Production cuts

Test days

Historical days used for testing are the 250 first days of 2016. The testing days are grouped into sets according to the PWSD defined in 8.1.1. The PWSD is watercourse-specific, and all testing days are chosen on the basis of the PWSDs of Bogna and Mosvik. Note therefore that one calendar day might correspond to two test days.

Analyzing profit at various PWSD-levels does not make sense for a two-watercourse portfolio. The water values are consistently different at each watercourse, so they deviate differently from the day-ahead price. We would therefore not be able to relate the coordinated gain to a certain PWSD interval.

All runs with portfolios of either one or two generators are conducted with the watercourses of Bogna and Mosvik. To further investigate the effect of portfolio size, we also add the third, simple watercourse to the portfolio. We test for the same days as for the two-watercourse portfolio. The number of test days can be seen in Table 8.1.

Table 8.2: Accumulated profit at Bogna for all 250 test days distributed across markets. Full hourly granularity. Up and Down denote upward and downward balancing market profits respectively

Bogna	**Profit**	**Spot**	**Prim. res.**	**Up**	**Down**	**Start-up**
Sequential	8,535,911	9,070,850	92,825	136,935	375,319	-1,140,018
Coordinated	8,611,677	8,999,220	157,046	167,435	455,974	-1,167,998
% change	0.89 %	-0.79 %	69.19 %	22.27 %	21.49 %	2.45 %

Table 8.3: Accumulated profit at Mosvik for all 250 test days distributed across markets. Full hourly granularity. Up and Down denote upward and downward balancing market profits respectively

Mosvik	**Profit**	**Spot**	**Prim. res.**	**Up**	**Down**	**Start-up**
Sequential	4,947,357	5,509,812	71,086	161,979	284,487	-1,080,007
Coordinated	5,004,864	5,514,083	141,416	174,696	366,183	-1,191,514
% change	1.16 %	0.11 %	98.9 %	7.85 %	28.72 %	10.32 %

8.3 Case results

In this section we present the results from the case study. We first investigate the performance of coordinated bidding relative to sequential when we only have a one-watercourse portfolio at a time, and hourly planning in all models. We then move on and see what happens when the day of planning is divided into sub-periods. This means that a first-stage decision is based on the restriction that a generator is either on or off for all of the hours $t \in \mathcal{T}^h$ in a scenario. The use of sub-periods reduces the computational burden, and allows for reasonable solution times when more watercourses are added to the portfolio. The results with two- and three-watercourse portfolios are presented in 8.3.2 and 8.3.3 respectively. Finally, in section 8.4, we highlight and discuss the important findings.

8.3.1 Results: One watercourse

Total gains from coordinated planning

The accumulated results for all 250 test days can be found in Table 8.2 and 8.3. Coordinated planning increases the total profits by approximately one percent. At Bogna we observe a light reduction in profit from the day-ahead market compared to the sequential planning. On the other hand, profits from all reserve markets increase significantly. Total start-up costs also increase marginally. At Mosvik profits from the day-ahead market actually increase slightly. Once again, reserve markets account for a greater portion of the profits.

Effect of price-water-value-deviation

Next, we move on to investigate coordinated performance at various PWSD levels. A summary of the results can be found in Table 8.4. Let π_d^c and π_d^s denote daily profits for some test day d using the coordinated and sequential framework respectively. The row labeled *gain (%)* is calculated as $\sum_{d\in D} \pi_d^c / \sum_{d\in D} \pi_d^s - 1$.

We observe that profits increase as PWSD increases. Total profit contribution is much larger for high PWSD. The row labeled gain should be interpreted with some caution and always be compared to the sum of profits. When profits are close to zero, the gain percentage is very sensitive to small differences. We observe that the gain from coordinated bidding is generally higher for prices close to water value. As expected, gain tends to zero when the price is much higher than water value.

Table 8.4: Full hourly granularity profits in NOK. Negative NPWD are excluded for large PWSD because there is very little production on these days

	PWSD	<20	20-40	40-60	>60	**Totals**
	NPWD	+ & -	+	+	+	
Bogna	mean of π_d^c	3,518	25,381	48,580	155,786	42,905
	std. of π_d^c	8,160	9,178	10,825	163,852	-
	$\sum_{d\in D} \pi_d^c$	270,896	1,497,482	1,360,246	5,452,504	8,581,128
	gain (%)	7.1	0.61	2.0	0.19	0.75
Mosvik	mean of π_d^c	1,211	19,409	31,749	118 479	29,180
	std. of π_d^c	6,841	8,357	9,672	121,529	-
	$\sum_{d\in D} \pi_d^c$	82,331	892,798	1,111,218	2,961,965	5,048,312
	gain (%)	77	2.6	0.01	0.13	1.2

Profits in each hour of the day

To understand how the coordinated framework benefits compared to the sequential, we further investigate the test days when PWSD<20 at Bogna. The average profit contribution from each market throughout the day is plotted in Figure 8.6. The coordinated model tends to submit unprofitable bids into the day-ahead market in the morning. However, upon delivering primary reserve and downward balancing, total profit exceeds that of the sequential model. Throughout the rest of the day, the coordinated model sells slightly more reserves than the sequential, but the difference is smaller. Reserve market opportunities in the morning (when day-ahead prices are low) seem to be the main source of increased profits under coordinated planning.

Effect from planning with sub-periods

We then move on to investigate the results from planning with sub-periods instead of hourly granularity. We only include the results at Bogna. Our goal is to assess whether

Figure 8.6: Hourly average profit contribution from the markets

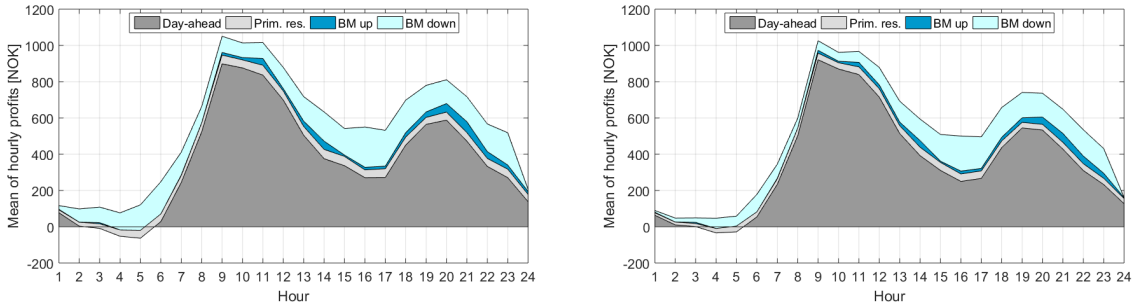

(a) Coordinated framework

(b) Sequential framework

Table 8.5: Accumulated profit at Bogna for all 250 test days distributed across markets. Sub-periods planning. Up and Down denote upward and downward balancing market profits respectively

Bogna	**Profit**	**Spot**	**Prim. res.**	**Up**	**Down**	**Start-up**
Sequential	7,802,044	8,006,040	121,965	168,253	453,809	-948,023
Coordinated	7,856,833	7,925,567	186,712	198,362	526,197	-980,005
% change	0.70 %	-1.01 %	53.09 %	17.90 %	15.95 %	3.37 %

sub-period planning changes the profit structure and more importantly if coordinated gains remain the same as for hourly planning. The reader is referred to table 8.5 for accumulated results. We first note that profits are comparable with the profits from hourly planning, although we see a light decline. The percentage change between sequential and coordinated planning is very similar to what we saw for hourly planning. Results for different PWSD-levels are in Table 8.6. We observe the same decline in gain as the PWSD increases. The total gain at Bogna is approximately equal to the gain with hourly planning. This increases our confidence that sub-period planning will provide results that correctly reflect the potential of coordinated bidding. This planning regime will be used for two and three watercourses in the portfolio.

Table 8.6: Sub-period profits at Bogna in NOK. Negative NPWD are excluded for large PWSD because there is very little production on these days

	PWSD	<20	20-40	40-60	>60	**Totals**
	NPWD	+ & -	+	+	+	
Bogna	mean of π_d^c	3,937	25,779	49,007	138,381	39,261
	std. of π_d^c	8,223	8,587	9,084	105,930	-
	$\sum_{d \in D} \pi_d^c$	303,168	1,520,973	1,323,196	4,704,946	7,852,283
	gain (%)	7.2	1.5	0.84	0.22	0.83

Table 8.7: Accumulated profit at Bogna & Mosvik for all 250 test days distributed across markets. Sub-periods. Up and Down denote upward and downward balancing market profits respectively

Portfolio	Profit	Spot	Prim. res.	Up	Down	Start-up
Sequential	23,802,574	25,211,480	365,611	444,083	957,400	-3,176,000
Coordinated	23,923,208	25,094,139	501,099	491,389	1,044,580	-3,207,999
% change	0.51 %	-0.47 %	37.06 %	10.65 %	9.11 %	1.01 %

8.3.2 Results: Two watercourses

As explained in section 8.1.3 it is not a priori clear whether an additional production capacity will increase or dilute the gains from coordinated bidding. We move on to investigate accumulated results with a two-watercourse portfolio presented in Table 8.7. In general, results are very similar to the one-watercourse portfolios. Wee see a light decline in coordinated gain down to approximately one half percent. Coordinated day-ahead profits decline slightly compared to sequential, while profits from reserve markets increase. Note however that reserve market profits do not increase as much as they did for one watercourse.

8.3.3 Results: Three watercourses

We move on to the case of three watercourses in the portfolio. The plant at Ormsetfoss has been added to the planning problem, and the models have been solved for the same days as for two watercourses. Accumulated results are in Table 8.8. As for two watercourses, there is a small gain from coordinated bidding. Once again there is a marginal decline in profit from the day-ahead market under coordinated planning. Accordingly, reserve markets account for a larger portion of total profits. Compared to the case of two watercourses, there seems to be a light convergence between the two planning regimes. The coordinated and sequential frameworks exploit the markets more equally. Put differently, the difference in day-ahead and reserve profits between the planning regimes is less than it was before. Obviously, there is a limited demand for balancing reserves. The acquired amount of balancing volume must be distributed to the generators. It is thus natural that the balancing market gain from coordinated bidding declines when we add more watercourses to the planning problem.

8.4 Discussion of results

One watercourse

For one watercourse a gain of about 1 % can be expected, thus showing that coordinated planning has a moderate potential. Compared to previous research we have tested

Table 8.8: Accumulated profit at Bogna, Mosvik & Ormsetfoss for all 250 test days distributed across markets. Sub-periods. Up and Down denote upward and downward balancing market profits respectively

Portfolio	**Profit**	**Spot**	**Prim. res.**	**Up**	**Down**	**Start-up**
Sequential	31,251,845	33,552,921	435,944	567,621	1,041,595	-4,348,000
Coordinated	31,399,066	33,456,550	565,826	611,124	1,119,835	-4,356,000
% change	0.47 %	-0.29 %	29.79 %	7.66 %	7.51 %	0.18 %

for a large number of days, and therefore obtained reliable results. The coordinated planning reduces total profits from the day-ahead market and increases profits from reserve markets even more. Obviously, this increases the importance of assumptions and modelling of the reserve markets. There may be additional risks associated with a shift to reserve markets. Especially, we want to address the assumptions on the balancing market. In the testing framework, all balancing volumes and premiums are revealed at the same time. The model may subsequently pick the balancing opportunities that are profitable (compared to water value) and feasible (capacity). This is a rather reasonable approximation because I) there is no/little opportunity cost associated with balancing that is not reflected in water value (other intraday markets such as Elbas are rather illiquid), and II) inter-temporal coordination is already very restricted by day-ahead and primary reserve commitments. Therefore, any balancing opportunity above marginal cost would be reflected in a bid curve for the balancing market. In addition, inter-temporal coordination is to a large extent already determined by day-ahead and primary reserve commitments. Only decisions about turning a generator on or off would affect this. Such situations are likely to be so valuable that the producer would seize these opportunities anyway. To conclude, this modelling choice may slightly overestimate the value of the balancing market.

In this thesis we use a market share to restrict obligations in the balancing market. This might be an excessively restrictive modelling, and will probably limit the value of the balancing market. A producer could possibly deliver the entire regulation demand in one hour, but would in this case act as price maker. Future research should look into the sensitivity in gains with respect to the market share, or look into ways to model the limited balancing demand. To conclude, this modelling choice may slightly underestimate the value of the balancing market.

We conclude that higher gains can be expected for days when the day-ahead price does not deviate much from water value. This is in line with our hypothesis in section 8.1.1. In these cases a producer can enhance profit by taking subsequent reserve markets into account. However, the total value of such days is very small compared to days when the day-ahead price is higher.

Coordinated planning often outperforms sequential planning because reserve markets can be better exploited in morning hours. Day-ahead prices are low in the morning, but primary reserve prices are generally high. In addition, the probability of downward

balancing is higher in the morning. A sequential planning fails to account for this. Coordinated planning allows for unprofitable day-ahead bids in these hours, and subsequently creates opportunities to deliver profitable reserves.

Planning with sub-periods does not seem to significantly deteriorate the quality of bidding decisions. More importantly, it cannot be seen to favour either of the planning frameworks before the other. This is an important finding because it allows us to solve the models for larger portfolios, and compare gains from coordinated bidding. We must still exhibit some caution interpreting the results. The sub-periods are chosen based on average price behaviour. It will fail to reflect hourly variations within each sub-period. Hence, when we test a different set of days, we cannot guarantee that sub-period based planning performs equally well as hourly.

Two watercourses

When planning for two watercourses simultaneously, we get very similar results as we did for only one watercourse. The gain from coordinated bidding decreases to about 0.5 %. In addition, we observe that the trade-off between markets differs less between coordinated and sequential planning. Day-ahead profits are more similar between the two planning regimes, and so are the reserve market profits. As discussed in Section 8.1.3, recourse options increase in both frameworks when additional watercourses are added. This implies that a sequential planning now has more options to react to reserve opportunities. Therefore, we would expect the sequential framework to close some of the gap to coordinated performance, which it does.

Three watercourses

Moving on to three watercourses, we once again see a light decline in gain from coordinated bidding. However, this decline is extremely small. The profit contribution from each market converges even more for the two planning regimes. This may have a natural interpretation. In addition to increased recourse, the convergence may stem from the limited balancing demand. Even though we add more production capacity, the balancing demand is exactly the same. In this way the sequential framework may catch up with coordinated performance simply because there is no more balancing demand that coordinated planning can utilize.

Chapter 9

Conclusion

In this thesis we have assessed the value of coordinated bidding for a hydropower producer. All work is put in context to the sparse literature on the topic. We develop a stochastic mixed integer program for constructing bid curves to the day-ahead market, taking subsequent market opportunities into account. To develop a scenario generation methodology, we start off with an empirical analysis of all relevant markets. We propose quantile autoregression model to yield probabilistic forecasts for the day-ahead and primary reserve markets. We further develop existing models for the balancing market, and utilize a copula based method to model dependence. Finally we evaluate the generated scenarios before carrying out an extensive case study. We compare the performance of coordinated planning to that of pure sequential under three relevant control variables, and conclude that coordinated planning has a moderate potential for increasing total profits of the producer.

Scenario generation

The extended QAR model proposed in Chapter 6 provides very good probabilistic forecasts for day-ahead and primary reserve prices. Forecasts are shown to have good unconditional and conditional coverage. A similar model has to our knowledge never been used for probabilistic forecasting of day-ahead and primary reserve prices. A producer with good forecast models could instead model the deviation between the forecast and the realized price using a similar methodology.

Our work on balancing market prices and volumes extends the current literature. Despite the effort, simulation of the ARMA models provide probabilistic forecasts with rather poor conditional coverage. Future research should look into ways to enhance conditional coverage. In addition, the balancing market is characterized by price and volume spikes. ARMA models have been shown to perform poorly under the presence of spikes. We encourage future researchers to look into mixture models for the balancing market.

We use a copula based method to model the non-linear relationship between stochastic parameters. The method seems to work well, generating scenarios of high quality.

We investigate the stability of scenario trees generated, and conclude that the method provides good decision support even for a rather low number of scenarios.

Case study

Our initial problem description aimed to evaluate the possible gains from coordinated bidding. We have benchmarked a coordinated planning framework against a pure sequential under three different control variables. In light of the results from Chapter 8 we are now ready to make a conclusion. Coordinated bidding seems to have a very moderate potential for increased profits. When planning with only one watercourse, gains of about 1 % can be expected. The gain is substantially higher for days when the day-ahead price lies in near proximity to water value. However, the total value of these days is low compared to days with higher prices. A coordinated planning increases the producer's ability to exploit profitable reserve market opportunities. This is especially true in morning hours when day-ahead prices are low and reserve prices and demand are higher. Day-ahead commitments below water value turn out to be profitable when accounting for reserve opportunities.

When planning for more watercourses simultaneously, the value of coordination decreases to about 0.5 %. More generators increase recourse options, and it is likely that this dilutes the value of coordinated planning. The producer can simply reallocate production to respond to reserve market opportunities.

Future research

There are however questions that remain unanswered. It is unclear what happens to coordinated gains when planning for very large portfolios. Our results indicate that gains should decline, but there is also a tendency of stabilization. When moving from one to two watercourses, gains decline. Some of this decline may be explained by the binary nature of decisions on the production resources. Generators are either on or off, and if primary reserves are offered, the producer cannot turn the generator off. When moving from two to three watercourses, gains decrease only marginally. It might very well be that case that coordinated bidding has *some* value even for very large portfolios. For instance, downward balancing and offering primary reserve require operating generators. Hence, it does not matter how many generators are in the portfolio if none of them are running. Such opportunities can never be exploited, and a sequential planning will fail to capture this.

One can also raise the question what will happen to coordinated gains when considering more complex watercourse structures. In this problem all watercourses have only one generator, and are physically disconnected. In more complex structures decisions at one generator will affect possible decisions at other generators (due to restrictions on reservoir levels etc). This will limit the portfolio flexibility, and possible recourse options. It might very well be the case that coordinated planning is even more valuable under such circumstances.

To our knowledge, we are the first to evaluate value of coordination considering multiple reserve markets. This may increase value of coordination because of diversification of reserve opportunities. It is more unlikely that all reserve markets have unfavorable realizations simultaneously. There are in addition other markets and bid types that could be considered. Weekly primary reserves and the intraday market Elbas are examples of such markets. Block bids is another option concerning bid types. Value from coordinated bidding may very well look differently considering these situations instead.

Obviously, the value from coordinated planning is very dependent on the quality of scenarios. This thesis considers pre-spot modelling, and estimation of the value of reserve opportunities at the time of day-ahead bidding. The day-ahead bid is the major decision of the day for the planner. However the bids submitted to subsequent markets also have an effect on final profits, and the gain associated with coordinated bidding. In our testing framework we need to model the post-spot decisions taken by the producer to assess the value of coordination. We cannot distinguish profit gains stemming from the first stage decision and post-spot decisions. A specialized post-spot reserve forecasting model may very well increase the potential of coordinated bidding estimated in our work, facilitating even better post-spot decisions.

Bibliography

Boomsma, T.K, N. Juul, and S Fleten (2014). "Bidding in sequential markets: The Nordic case". In: *European Journal of Operational Research* 238.3, pp. 797–809.

Box, G.E.P., G.M. Jenkins, and G.C. Reinsel (2007). *Time Series Analysis: Forecasting and Control*.

Christoffersen, P.F. (1998). "Evaluating Interval Forecasts". In: *International Economic Review* 39.4, pp. 841–862.

Conejo, A.J., M. Carrión, and J.M. Morales (2010). *Decision Making Under Uncertainty in Electricity Markets*.

Do, L.P.C. and P. Molnar (2015). "Day ahead electricity price modelling using quantile regression". In: *Master thesis, NTNU*.

Erdogan, E., S. Ma, A. Beygelzimer, and I. Rish (2005). "Statistical Models for Unequally Spaced Time Series". In: *IBM Research Division*.

Eriksrud, A.L. and J. Braathen (2012). "Spot Market Bidding Taking the Balancing Power Market Into Account". In: *Specialization Project, NTNU*, pp. 1–42.

Faria, E. and S Fleten (2011). "Day-ahead market bidding for a Nordic hydropower producer: taking the Elbas market into account". In: *Computational Management Science* 8.1, pp. 75–101.

Fleten, S.E. and E. Pettersen (2005). "Constructing Bidding Curves for a Price-Taking Retailer in the Norwegian Electricity Market". In: *IEEE Transactions on Power Systems*, pp. 701–708.

Hagfors, L.I., D. Bunn, E. Kristoffersen, T.T. Staver, and S. Westgaard (2016). "Modeling the UK electricity price distributions using quantile regression". In: *Energy* 102, pp. 231–243.

Huisman, R., C. Huurman, and R. Mahieu (2007). "Hourly electricity prices in day-ahead markets". In: *Energy Economics*, pp. 240–248.

Jonsson, T., P. Pinson, H. Madsen, and H.A. Nielsen (2014). "Predictive densities for Day-Ahead Electricity Prices Using Time-Adaptive Quantile Regression". In: *Energies*, pp. 5523–5547.

Kårstad, I. and C. Skjong (2015). "Capacity Allocation of Electricity Generation Between a Primary Reserve Market and a Day-Ahead Market". In: *Specialization Project, NTNU*, pp. 1–42.

Kaut, M. (2014). "A copula-based heuristic for scenario generation". In: *Computational Management Science*, pp. 503–516.

Kaut, M. and S. Wallace (2007). "Evaluation of scenario-generation methods for stochastic programming". In: *Pacific Journal of Optimization*, pp. 342–351.

Klæboe, G., A.L. Eriksen, and S.E. Fleten (2013). "Benchmarking time series based forecasting models for electricity balancing market prices". In: *Energy Systems, Springer Verlag*, pp. 43–61.

Koenker, R. (2005). *Quantile Regression*.

Koenker, R. and G. Bassett (1978). *Regression Quantiles*.

Koenker, R. and Z. Xiao (2005). "Quantile autoregression". In: *Journal of the American Statistical Association* 101.475, pp. 980–1006.

Löhndorf, Nils, David Wozabal, and Stefan Minner (2013). "Optimizing Trading Decisions for Hydro Storage Systems Using Approximate Dual Dynamic Programming". In: *Operations Research* 61.4, pp. 810–823.

Lorubio, G. (2011). "Flexible generation: Backing up renewables." In: *Union of the Electricity Industry - EURELECTRIC*, p. 26.

Maciejowska, K., J. Nowotarski, and R. Weron (2016). "Probabilistic forecasting of electricity spot prices using Factor Quantile Regression Averaging". In: *International Journal of Forecasting*, pp. 957–965.

NordPool (2016). *Day-ahead market*. NordPool. URL: http://www.nordpoolspot.com/How-does-it-work/Day-ahead-market-Elspot-/ (visited on 11/07/2016).

Nowotarski, J. and R. Weron (2013). "Computing electricity spot price prediction intervals using quantile regression and forecast averaging". In: *Comput Stat - Springer*, pp. 791–803.

Olsson, M. and L. Söder (2008). "Modeling Real-Time Balancing Power Market Prices Using Combined SARIMA and Markov Processes". In: *IEEE Transactions on Power Systems*, pp. 443–450.

Statkraft (2016). *Hydropower*. Statkraft Energi AS. URL: http://www.statkraft.com/energy-sources/hydropower/ (visited on 11/20/2016).

Statnett (2016). *Om RK*. Statnett. URL: http://statnett.no/Kraftsystemet/Markedsinformasjon/RKOM1/Om-regulerkraftmarkedet-RKM/ (visited on 11/07/2016).

Statnett AS (2016). *Om primærreserver*. Statnett AS. URL: http://statnett.no/Kraftsystemet/Markedsinformasjon/Primarreserver/ (visited on 11/07/2016).

Tomasgard, A. and E. Høeg (2005). "A Supply Chain Optimization Model for the Norwegian Meat Cooperative". In: *Applications of Stochastic Programming*, pp. 253–276.

Ventosa, M., A. Baíllo, A. Ramos, and M. Rivier (2005). "Electricity market modeling trends". In: *Energy Policy*, pp. 897–913.

Weron, R. (2014). "Electricity price forecasting: A review of the state-of-the-art with a look into the future". In: *International Journal of Forecasting*, pp. 1030–1081.

Weron, R. and A. Misiorek (2008). "Forecasting spot electricity prices: A comparison of parametric and semiparametric time series models". In: *International Journal of Forecasting*, pp. 744–763.

Wooldridge, J.M. (2010). *Econometric Analysis of Cross Section and Panel Data*.

Appendix A

Markov transition matrix

$$\mathrm{P}_1 = \begin{pmatrix} 0,76 & 0,03 & 0,21 \\ 0,01 & 0,78 & 0,21 \\ 0,05 & 0,07 & 0,88 \end{pmatrix} \; P_2 = \begin{pmatrix} 0,79 & 0,02 & 0,19 \\ 0,00 & 0,88 & 0,11 \\ 0,02 & 0,05 & 0,93 \end{pmatrix} \; P_3 = \begin{pmatrix} 0,83 & 0,03 & 0,14 \\ 0,01 & 0,92 & 0,07 \\ 0,01 & 0,04 & 0,95 \end{pmatrix}$$

$$P_4 = \begin{pmatrix} 0,87 & 0,03 & 0,10 \\ 0,00 & 0,93 & 0,07 \\ 0,02 & 0,05 & 0,94 \end{pmatrix} \; P_5 = \begin{pmatrix} 0,75 & 0,03 & 0,23 \\ 0,01 & 0,89 & 0,10 \\ 0,02 & 0,08 & 0,90 \end{pmatrix} \; P_6 = \begin{pmatrix} 0,66 & 0,01 & 0,33 \\ 0,02 & 0,74 & 0,24 \\ 0,06 & 0,08 & 0,87 \end{pmatrix}$$

$$P_7 = \begin{pmatrix} 0,64 & 0,01 & 0,35 \\ 0,03 & 0,57 & 0,39 \\ 0,11 & 0,08 & 0,81 \end{pmatrix} \; P_8 = \begin{pmatrix} 0,80 & 0,01 & 0,19 \\ 0,02 & 0,80 & 0,18 \\ 0,09 & 0,10 & 0,82 \end{pmatrix} \; P_9 = \begin{pmatrix} 0,69 & 0,03 & 0,28 \\ 0,02 & 0,75 & 0,24 \\ 0,04 & 0,08 & 0,88 \end{pmatrix}$$

$$P_{10} = \begin{pmatrix} 0,72 & 0,03 & 0,25 \\ 0,02 & 0,77 & 0,21 \\ 0,08 & 0,04 & 0,88 \end{pmatrix} \; P_{11} = \begin{pmatrix} 0,81 & 0,05 & 0,14 \\ 0,00 & 0,79 & 0,21 \\ 0,05 & 0,04 & 0,91 \end{pmatrix} \; P_{12} = \begin{pmatrix} 0,73 & 0,02 & 0,24 \\ 0,00 & 0,86 & 0,14 \\ 0,04 & 0,05 & 0,90 \end{pmatrix}$$

$$P_{13} = \begin{pmatrix} 0,72 & 0,03 & 0,26 \\ 0,01 & 0,82 & 0,17 \\ 0,04 & 0,07 & 0,90 \end{pmatrix} \; P_{14} = \begin{pmatrix} 0,80 & 0,00 & 0,20 \\ 0,00 & 0,88 & 0,11 \\ 0,05 & 0,08 & 0,88 \end{pmatrix} \; P_{15} = \begin{pmatrix} 0,77 & 0,01 & 0,22 \\ 0,01 & 0,82 & 0,16 \\ 0,03 & 0,07 & 0,89 \end{pmatrix}$$

$$P_{16} = \begin{pmatrix} 0,76 & 0,03 & 0,21 \\ 0,01 & 0,80 & 0,19 \\ 0,04 & 0,08 & 0,88 \end{pmatrix} \; P_{17} = \begin{pmatrix} 0,75 & 0,01 & 0,24 \\ 0,01 & 0,85 & 0,14 \\ 0,07 & 0,06 & 0,87 \end{pmatrix} \; P_{18} = \begin{pmatrix} 0,82 & 0,02 & 0,16 \\ 0,05 & 0,63 & 0,32 \\ 0,07 & 0,04 & 0,89 \end{pmatrix}$$

$$P_{19} = \begin{pmatrix} 0,79 & 0,04 & 0,17 \\ 0,01 & 0,82 & 0,17 \\ 0,05 & 0,06 & 0,89 \end{pmatrix} \; P_{20} = \begin{pmatrix} 0,77 & 0,01 & 0,23 \\ 0,04 & 0,81 & 0,15 \\ 0,06 & 0,07 & 0,87 \end{pmatrix} \; P_{21} = \begin{pmatrix} 0,69 & 0,02 & 0,30 \\ 0,01 & 0,85 & 0,14 \\ 0,05 & 0,08 & 0,87 \end{pmatrix}$$

$$P_{22} = \begin{pmatrix} 0,75 & 0,02 & 0,24 \\ 0,05 & 0,78 & 0,18 \\ 0,07 & 0,06 & 0,88 \end{pmatrix} \; P_{23} = \begin{pmatrix} 0,69 & 0,04 & 0,27 \\ 0,02 & 0,80 & 0,18 \\ 0,05 & 0,07 & 0,88 \end{pmatrix} \; P_{24} = \begin{pmatrix} 0,46 & 0,06 & 0,48 \\ 0,03 & 0,71 & 0,26 \\ 0,08 & 0,10 & 0,82 \end{pmatrix}$$

Appendix B

Resampling unequally spaced time series

An autoregressive process of order 1 can be written as

$$X(t+1) = \theta X(t) + \sigma\epsilon_{t+1} \tag{B.0.1}$$

where $X(t)$ are observations of the random variable, and ϵ_t is a white noise error term distributed $\epsilon_t \sim \mathcal{N}(0,1)$. Extending the notation recursively an observation at arbitrary lag h can be written as

$$X(t+h) = \theta^h X(t) + \sigma \sum_{j=0}^{h-1} \theta^{h-1-j}\epsilon_{t+1+j} \tag{B.0.2}$$

Noting that the last term is distributed as $\sigma \sum_{j=0}^{h-1} \theta^{h-1-j}\epsilon_{t+1+j} \sim \mathcal{N}\Big(0, \sigma^2(\frac{1-\theta^{2h}}{1-\theta^2})\Big)$ we make the substitution

$$X(t+\Delta) = \theta^\Delta X(t) + \sigma_\Delta \epsilon_{t+\Delta} \tag{B.0.3}$$

for $\Delta \geq 0$. Estimation of the two parameters θ and σ remains. Because the error terms are assumed to have zero mean, and because there is no intertemporal correlation of the error terms, least squares can be used to estimate θ without an estimate for σ. An estimator $\hat{\theta}$ is found by solving the convex non-linear problem

$$\hat{\theta} = \underset{\theta\in(-1,1)}{\arg\min} \sum_{i=1}^{n-1} \Big(x(t_{i+1}) - \theta^{\Delta_i} x(t_i)\Big)^2 \tag{B.0.4}$$

with $\Delta_i = t_{i+1} - t_i$ and assuming $\Delta_i \geq 1$, or alternatively rescaling the problem. The problem is very effectively solved using convex optimization. With an estimator for θ available, an estimator for σ can be found using maximum likelihood. First, if we define $z_i = x(t_{i+1}) - \hat{\theta}^{\Delta_i} x(t_i)$, we can see that $z_i \sim \mathcal{N}\Big(0, \sigma^2_{\Delta_i}\Big)$. Furthermore, we have that

$$\sigma^2_{\Delta_i} = \sigma^2\left(\frac{1-\hat{\theta}^{2\Delta_i}}{1-\hat{\theta}^2}\right) \quad \Leftrightarrow \quad \sigma^2 = \frac{\sigma^2_{\Delta_i}}{\left(\frac{1-\hat{\theta}^{2\Delta_i}}{1-\hat{\theta}^2}\right)} \tag{B.0.5}$$

The Gaussian likelihood of the residuals is maximized by calculating the mean of this normalized variance according to

$$\hat{\sigma} = \sqrt{\frac{1}{n}\sum_{i=0}^{n-1}\frac{z_i^2}{\left(\frac{1-\hat{\theta}^{2\Delta_i}}{1-\hat{\theta}^2}\right)}} \tag{B.0.6}$$

Now that we have estimators for the coefficients of the series, we proceed with the resampling process. We start with an irregularly sampled time series $X(t_i)$ with nonzero values at times t_i. We want to "fill in the gaps" at all the times s_k where the value of the time series are not available. In addition, we obviously want to keep all original values of the time series such that the resampled series coincide at these times, i.e. $\tilde{X}(s_k) = X(t_i)$ for $s_k = t_i$. However, we miss one important thing. In general, ϵ_s cannot be estimated for hours without data. Instead we introduce an auxiliary error ε such that

$$X(t_i + \Delta_i) = \hat{\theta}^{\Delta_i} X(t) + \varepsilon\hat{\sigma}\sum_{j=0}^{\Delta_i - 1}\hat{\theta}^j \tag{B.0.7}$$

This way we can estimate a local error term in between observations. Rearranging the terms we get a direct estimate of ε

$$\varepsilon = \frac{X(t_i + \Delta_i) - \hat{\theta}^{\Delta_i} X(t)}{\hat{\sigma}\sum_{j=0}^{\Delta_i - 1}\hat{\theta}^j} \tag{B.0.8}$$

This allows us to set recursively

$$\hat{X}(t_i + 1) = \hat{\theta}\hat{X}(t) + \hat{\sigma}\varepsilon \tag{B.0.9}$$

Note also that this prediction sets $\hat{X}(t_i) = X(t_i)$ so that the resampled values coincide with the original data, as desired. We now have all we need to resample the unequally spaced time series.

Appendix C

Decision models

C.1 Coordinated framework: Model 1

$$\max \sum_{s\in\mathcal{S}}\Bigg(\pi_s \sum_{m\in\{1,2\}}\sum_{t\in\mathcal{T}}\rho_{mts}x_{mts} + \sum_{\omega\in\Omega^s}\Bigg(\pi_{s\omega}\sum_{m\in\{3,4\}}\sum_{t\in\mathcal{T}}(\rho_{mts\omega}+\iota_m\rho_{1ts\omega})x_{mts\omega} + \sum_{k\in\mathcal{K}} w_{ks\omega} - \sum_{i\in\mathcal{I}'}\sum_{h\in\mathcal{H}} c_{ihs\omega}\Bigg)\Bigg)$$

$$x_{1ts} = z_{pt} + (z_{p+1t} - z_{pt})\frac{\rho_{1ts} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1ts} \leq P_{(p+1)},\ \ p \in \mathcal{P}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$z_{(p+1)t} \geq z_{pt} \qquad p \in \mathcal{P} \setminus |\mathcal{P}|, t \in \mathcal{T}$$

$$\sum_{i\in\mathcal{I}} q_{1its\omega} = x_{1ts}, \qquad s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$\sum_{i\in\mathcal{I}} q_{2its\omega} = x_{2ts}, \qquad s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$\sum_{i\in\mathcal{I}} q_{mits\omega} = x_{mts\omega}, \qquad m \in \{3,4\}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$x_{mts\omega} \leq \sigma\nu_{mts\omega} \qquad m \in \{3,4\}, s \in S, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ihs\omega}\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2its\omega} \leq u_{ihs\omega}\frac{0.2N_i}{\kappa_i^{min}} \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2its\omega} \leq \frac{0.2N_i}{\kappa_i^{min}} \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ihs\omega}Q_i^{max} \geq q_{1its\omega} + q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$Q_i^{max} \geq q_{1its\omega} + q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ihs\omega}Q_i^{min} \leq q_{1its\omega} - q_{2its\omega} + q_{3its\omega} - q_{4its\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$Q_i^{min} \leq q_{1its\omega} - q_{2its\omega} + q_{3its\omega} - q_{4its\omega} \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$c_{ihs\omega} \geq C_i(u_{ihs\omega} - u_{i(h-1)s\omega}) \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, \omega \in \Omega^s$$

$$q_{1its\omega} + q_{3its\omega} - q_{4its\omega} \leq A_{if} + B_{if}d_{its\omega} \qquad i \in \mathcal{I}, f \in \mathcal{F}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{jts\omega} - v_{j(t-1)s\omega} = I_{jt} + \sum_{i \in \mathcal{I}} \Gamma_{ij}d_{its\omega} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'}O_{j'ts\omega}, \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{j1s\omega} - V_j^0 = I_{j1} + \sum_{i \in \mathcal{I}} \Gamma_{ij}d_{i1s\omega} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'}O_{j'1s\omega}, \qquad j \in \mathcal{J}, s \in \mathcal{S}, \omega \in \Omega^s$$

$$w_{ks\omega} \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk}(v_{j|T|s\omega} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, s \in \mathcal{S}, \omega \in \Omega^s$$

$$D_i^{min} \leq d_{its\omega} \leq D_i^{max} \qquad i \in \mathcal{I}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$V_j^{min} \leq v_{jts\omega} \leq V_j^{max} \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$x_{mts} \geq 0 \qquad m \in \{1, 2\}, s \in S, t \in \mathcal{T}$$

$$x_{mts\omega} \geq 0 \qquad m \in \{3, 4\}, s \in S, t \in \mathcal{T}, \omega \in \Omega^s$$

$$w_{ks\omega} \geq 0 \qquad k \in \mathcal{K}, s \in S, \omega \in \Omega^s$$

$$z_{pts} \geq 0 \qquad p \in \mathcal{P}, s \in S, t \in \mathcal{T}$$

$$q_{mits\omega} \geq 0 \qquad m \in \mathcal{M}, i \in \mathcal{I}, s \in S, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ihs\omega} \in \{0,1\} \qquad i \in \mathcal{I}', h \in \mathcal{H}, s \in \mathcal{S}, \omega \in \Omega^s$$

$$d_{its\omega} \geq 0 \qquad i \in \mathcal{I}, s \in S, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{jts\omega} \geq 0 \qquad j \in \mathcal{J}, s \in S, t \in \mathcal{T}, \omega \in \Omega^s$$

$$c_{ihs\omega} \geq 0 \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in S, \omega \in \Omega^s$$

C.2 Coordinated framework: Model 2

[1]

$$\max \sum_{\omega \in \Omega^s} \pi_\omega \left(\sum_{t \in \mathcal{T}} \rho_{2t\omega} x_{2t\omega} + \sum_{m \in \{3,4\}} \sum_{t \in \mathcal{T}} (\rho_{mt\omega} + \iota_m \rho_{1t\omega}) x_{mt\omega} + \sum_{k \in \mathcal{K}} w_{k\omega} - \sum_{i \in \mathcal{I}'} \sum_{h \in \mathcal{H}} c_{ih\omega} \right)$$

$$x_{1t} = z^*_{pt} + (z^*_{p+1t} - z^*_{pt}) \frac{\rho_{1t} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1ts} \leq P_{(p+1)}, \ \ p \in \mathcal{P}, t \in \mathcal{T}$$

$$x_{2t\omega} = \phi_{pt} + (\phi_{p+1,t} - \phi_{pt}) \frac{\rho_{2t\omega} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{2ts} \leq P_{(p+1)}, \ \ p \in \mathcal{P}, \omega \in \Omega^s, t \in \mathcal{T}$$

$$\phi_{(p+1)t} \geq \phi_{pt} \qquad p \in \mathcal{P} \setminus |\mathcal{P}|, t \in \mathcal{T}$$

$$\sum_{i \in \mathcal{I}} q_{1it\omega} = x_{1t}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$\sum_{i \in \mathcal{I}} q_{2it\omega} = x_{2t\omega}, \qquad t \in \mathcal{T}, \omega \in \Omega^s$$

[1] Exception from Chapter 4: In this model both primary reserve and balancing market scenarios are indexed with ω.

$$\sum_{i\in\mathcal{I}} q_{mit\omega} = x_{mt\omega}, \qquad m \in \{3,4\}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$x_{mt\omega} \leq \sigma \nu_{mt\omega} \qquad m \in \{3,4\}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ih\omega}\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2it\omega} \leq u_{ih\omega}\frac{0.2N_i}{\kappa_i^{min}} \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2it\omega} \leq \frac{0.2N_i}{\kappa_i^{min}} \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ih\omega}Q_i^{max} \geq q_{1it\omega} + q_{2it\omega} + q_{3it\omega} - q_{4it\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$Q_i^{max} \geq q_{1it\omega} + q_{2it\omega} + q_{3it\omega} - q_{4it\omega}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ih\omega}Q_i^{min} \leq q_{1it\omega} - q_{2its\omega} + q_{3it\omega} - q_{4it\omega}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h, \omega \in \Omega^s$$

$$Q_i^{min} \leq q_{1it\omega} - q_{2it\omega} + q_{3it\omega} - q_{4it\omega} \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}, \omega \in \Omega^s$$

$$c_{ih\omega} \geq C_i(u_{ih\omega} - u_{i(h-1)\omega}) \qquad h \in \mathcal{H}, i \in \mathcal{I}', \omega \in \Omega^s$$

$$q_{1it\omega} + q_{3it\omega} - q_{4it\omega} \leq A_{if} + B_{if}d_{it\omega} \qquad i \in \mathcal{I}, f \in \mathcal{F}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{jt\omega} - v_{j(t-1)\omega} = I_{jt} + \sum_{i\in\mathcal{I}} \Gamma_{ij}d_{it\omega} + \sum_{j'\in\mathcal{J}} \Lambda_{jj'}O_{j't\omega}, \qquad j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{j1\omega} - V_j^0 = I_{j1} + \sum_{i\in\mathcal{I}} \Gamma_{ij}d_{i1\omega} + \sum_{j'\in\mathcal{J}} \Lambda_{jj'}O_{j'1\omega}, \qquad j \in \mathcal{J}, \omega \in \Omega^s$$

$$w_{k\omega} \leq E_{lk} + \sum_{j\in\mathcal{J}^k} W_{jk}(v_{j|T|\omega} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, \omega \in \Omega^s$$

$$D_i^{min} \leq d_{it\omega} \leq D_i^{max} \qquad i \in \mathcal{I}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$V_j^{min} \leq v_{jt\omega} \leq V_j^{max} \qquad j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$x_{mt\omega} \geq 0 \qquad m \in \{2,3,4\}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$w_{k\omega} \geq 0 \qquad k \in \mathcal{K}, \omega \in \Omega^s$$

$$\phi_{pt} \geq 0 \qquad p \in \mathcal{P}, t \in \mathcal{T}$$

$$q_{mit\omega} \geq 0 \qquad m \in \mathcal{M}, i \in \mathcal{I}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$u_{ih\omega} \in \{0,1\} \qquad i \in \mathcal{I}', h \in \mathcal{H}, \omega \in \Omega^s$$

$$d_{it\omega} \geq 0 \qquad i \in \mathcal{I}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$v_{jt\omega} \geq 0 \qquad j \in \mathcal{J}, t \in \mathcal{T}, \omega \in \Omega^s$$

$$c_{ih\omega} \geq 0 \qquad h \in \mathcal{H}, i \in \mathcal{I}', \omega \in \Omega^s$$

C.3 Coordinated/Sequential framework: Model 3

$$\max \sum_{m\in\{3,4\}} \sum_{t\in\mathcal{T}} (\rho_{mt} + \iota_m \rho_{1t}) x_{mt} + \sum_{k\in\mathcal{K}} w_k - \sum_{i\in\mathcal{I}'} \sum_{h\in\mathcal{H}} c_{ih}$$

$$x_{2t} = \phi^*_{pt} + (\phi^*_{p+1t} - \phi^*_{pt}) \frac{\rho_{2t} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1ts} \leq P_{(p+1)}, \;\; p \in \mathcal{P}, t \in \mathcal{T}$$

$$\sum_{i\in\mathcal{I}} q_{1it} = x^*_{1t}, \qquad t \in \mathcal{T}$$

$$\sum_{i\in\mathcal{I}} q_{2it} = x_{2t}, \qquad t \in \mathcal{T}$$

$$\sum_{i\in\mathcal{I}} q_{mit} = x_{mt}, \qquad m \in \{3,4\}, t \in \mathcal{T}$$

$$x_{mt} \leq \sigma \nu_{mt} \qquad m \in \{3,4\}, t \in \mathcal{T}$$

$$u_{ih} \frac{0.2N_i}{\kappa_i^{max}} \leq q_{2it} \leq u_{ih} \frac{0.2N_i}{\kappa_i^{min}} \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h$$

$$\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2it} \leq \frac{0.2N_i}{\kappa_i^{min}} \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}$$

$$u_{ih}Q_i^{max} \geq q_{1it} + q_{2it} + q_{3it} - q_{4it}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h,$$

$$Q_i^{max} \geq q_{1it} + q_{2it} + q_{3it} - q_{4it}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}$$

$$u_{ih}Q_i^{min} \leq q_{1it} - q_{2it} + q_{3it} - q_{4it}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h$$

$$Q_i^{min} \leq q_{1it} - q_{2it} + q_{3it} - q_{4it} \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}$$

$$c_{ih} \geq C_i(u_{ih} - u_{i(h-1)}) \qquad h \in \mathcal{H}, i \in \mathcal{I}'$$

$$q_{1it} + q_{3it} - q_{4it} \leq A_{if} + B_{if}d_{it} \qquad i \in \mathcal{I}, f \in \mathcal{F}, t \in \mathcal{T}$$

$$v_{jt} - v_{j(t-1)} = I_{jt} + \sum_{i \in \mathcal{I}} \Gamma_{ij}d_{it} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'}O_{j't}, \qquad j \in \mathcal{J}, t \in \mathcal{T}$$

$$v_{j1s} - V_j^0 = I_{j1} + \sum_{i \in \mathcal{I}} \Gamma_{ij}d_{i1} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'}O_{j'1}, \qquad j \in \mathcal{J}$$

$$w_k \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk}(v_{j|T|\omega} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}$$

$$D_i^{min} \leq d_{it} \leq D_i^{max} \qquad i \in \mathcal{I}, t \in \mathcal{T}$$

$$V_j^{min} \leq v_{jt} \leq V_j^{max} \qquad j \in \mathcal{J}, t \in \mathcal{T}$$

$$x_{mt} \geq 0 \qquad m \in \{3, 4\}, t \in \mathcal{T}$$

$$w_k \geq 0 \qquad k \in \mathcal{K}$$

$$q_{mit} \geq 0 \qquad m \in \mathcal{M}, i \in \mathcal{I}, t \in \mathcal{T}$$

$$u_{ih} \in \{0, 1\} \qquad i \in \mathcal{I}', h \in \mathcal{H}$$

$$d_{it} \geq 0 \qquad i \in \mathcal{I}, t \in \mathcal{T}$$

$$v_{jt} \geq 0 \qquad j \in \mathcal{J}, t \in \mathcal{T}$$

$$c_{ih} \geq 0 \qquad h \in \mathcal{H}, i \in \mathcal{I}'$$

C.4 Sequential framework: Model 1

$$\max \sum_{s \in \mathcal{S}} \pi_s \left(\sum_{t \in \mathcal{T}} \rho_{1ts} x_{1ts} + \sum_{k \in \mathcal{K}} w_{ks} - \sum_{i \in \mathcal{I}'} \sum_{h \in \mathcal{H}} c_{ihs} \right)$$

$$x_{1ts} = z_{pt} + (z_{p+1t} - z_{pt}) \frac{\rho_{1ts} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1ts} \leq P_{(p+1)}, \;\; p \in \mathcal{P}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$z_{(p+1)t} \geq z_{pt} \qquad p \in \mathcal{P} \setminus |\mathcal{P}|, t \in \mathcal{T}$$

$$\sum_{i \in \mathcal{I}} q_{1its} = x_{1ts} \qquad s \in \mathcal{S}, t \in \mathcal{T}$$

$$u_{ihs} Q_i^{max} \geq q_{1its}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h$$

$$Q_i^{max} \geq q_{1its}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}$$

$$u_{ihs} Q_i^{min} \leq q_{1its}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h$$

$$Q_i^{min} \leq q_{1its} \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}$$

$$c_{ihs} \geq C_i (u_{ihs} - u_{i(h-1)s}) \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}$$

$$q_{1its} \leq A_{if} + B_{if} d_{its} \qquad i \in \mathcal{I}, f \in \mathcal{F}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$v_{jts} - v_{j(t-1)s} = I_{jt} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{its} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'ts}, \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$v_{j1s} - V_j^0 = I_{j1} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{i1s} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'1s}, \qquad j \in \mathcal{J}, s \in \mathcal{S}$$

$$w_{ks} \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk} (v_{j|T|s} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, s \in \mathcal{S}$$

$$D_i^{min} \leq d_{its} \leq D_i^{max} \qquad i \in \mathcal{I}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$V_j^{min} \leq v_{jts} \leq V_j^{max} \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$x_{1ts} \geq 0 \qquad s \in \mathcal{S}, t \in \mathcal{T}$$

$$w_{ks} \geq 0 \qquad k \in \mathcal{K}, s \in \mathcal{S}$$

$$q_{1its} \geq 0 \qquad i \in \mathcal{I}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$u_{ihs} \in \{0,1\} \qquad i \in \mathcal{I}', h \in \mathcal{H}, s \in \mathcal{S}$$

$$d_{its} \geq 0 \qquad i \in \mathcal{I}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$v_{jts} \geq 0 \qquad j \in \mathcal{J}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$c_{ihs} \geq 0 \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}$$

C.5 Sequential framework: Model 2

$$\max \sum_{s \in \mathcal{S}} \pi_s \left(\sum_{t \in \mathcal{T}} \rho_{2ts} x_{2ts} + \sum_{k \in \mathcal{K}} w_{ks} - \sum_{i \in \mathcal{I}'} \sum_{h \in \mathcal{H}} c_{ihs} \right)$$

$$x_{1t} = z^*_{pt} + (z^*_{p+1t} - z^*_{pt}) \frac{\rho_{1t} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{1t} \leq P_{(p+1)}, \ \ p \in \mathcal{P}, t \in \mathcal{T}$$

$$x_{2ts} = \phi_{pt} + (\phi_{p+1,t} - \phi_{pt}) \frac{\rho_{2ts} - P_p}{P_{p+1} - P_p} \text{ if } P_p \leq \rho_{2ts} \leq P_{(p+1)}, \ \ p \in \mathcal{P}, s \in \mathcal{S}, t \in \mathcal{T}$$

$$\phi_{(p+1)t} \geq \phi_{pt} \qquad p \in \mathcal{P} \setminus |\mathcal{P}|, t \in \mathcal{T}$$

$$\sum_{i \in \mathcal{I}} q_{1its} = x_{1t}, \qquad t \in \mathcal{T}, s \in \mathcal{S}$$

$$\sum_{i \in \mathcal{I}} q_{2its} = x_{2ts}, \qquad t \in \mathcal{T}, s \in \mathcal{S}$$

$$u_{ihs} \frac{0.2 N_i}{\kappa_i^{max}} \leq q_{2its} \leq u_{ihs} \frac{0.2 N_i}{\kappa_i^{min}} \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h, s \in \mathcal{S}$$

$$\frac{0.2N_i}{\kappa_i^{max}} \leq q_{2its} \leq \frac{0.2N_i}{\kappa_i^{min}} \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}, s \in \mathcal{S}$$

$$u_{ihs} Q_i^{max} \geq q_{1its} + q_{2its}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}^h$$

$$Q_i^{max} \geq q_{1its} + q_{2its}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', s \in \mathcal{S}, t \in \mathcal{T}$$

$$u_{ihs} Q_i^{min} \leq q_{1its} - q_{2its}, \qquad h \in \mathcal{H}, i \in \mathcal{I}', t \in \mathcal{T}^h, s \in \mathcal{S}$$

$$Q_i^{min} \leq q_{1its} - q_{2its}, \qquad i \in \mathcal{I} \setminus \mathcal{I}', t \in \mathcal{T}, s \in \mathcal{S}$$

$$c_{ihs} \geq C_i(u_{ihs} - u_{i(h-1)s}), \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}$$

$$q_{1its} \leq A_{if} + B_{if} d_{its} \qquad i \in \mathcal{I}, f \in \mathcal{F}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$v_{jts} - v_{j(t-1)s} = I_{jt} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{its} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'ts}, \qquad j \in \mathcal{J}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$v_{j1s} - V_j^0 = I_{j1} + \sum_{i \in \mathcal{I}} \Gamma_{ij} d_{i1s} + \sum_{j' \in \mathcal{J}} \Lambda_{jj'} O_{j'1s}, \qquad j \in \mathcal{J}, s \in \mathcal{S}$$

$$w_{ks} \leq E_{lk} + \sum_{j \in \mathcal{J}^k} W_{jk}(v_{j|T|s} - V_{jl}) \qquad k \in \mathcal{K}, l \in \mathcal{L}, s \in \mathcal{S}$$

$$D_i^{min} \leq d_{its} \leq D_i^{max} \qquad i \in \mathcal{I}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$V_j^{min} \leq v_{jts} \leq V_j^{max} \qquad j \in \mathcal{J}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$x_{2ts} \geq 0 \qquad t \in \mathcal{T}, s \in \mathcal{S}$$

$$w_{ks} \geq 0 \qquad k \in \mathcal{K}, s \in \mathcal{S}$$

$$\phi_{pt} \geq 0 \qquad p \in \mathcal{P}, t \in \mathcal{T}$$

$$q_{mits} \geq 0 \qquad m \in \{1, 2\}, i \in \mathcal{I}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$u_{ihs} \in \{0,1\} \qquad i \in \mathcal{I}', h \in \mathcal{H}, s \in \mathcal{S}$$

$$d_{its} \geq 0 \qquad i \in \mathcal{I}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$v_{jts} \geq 0 \qquad j \in \mathcal{J}, t \in \mathcal{T}, s \in \mathcal{S}$$

$$c_{ihs} \geq 0 \qquad h \in \mathcal{H}, i \in \mathcal{I}', s \in \mathcal{S}$$

Appendix D

Modelling dependence: Copula heuristic

D.1 Introduction

Dependence in the prices across hours and across markets is modelled with the copula-based framework by Kaut 2014. The model proposed in Section 6.3 combined with estimation from Section 6.4 provides hourly predictive densities for the prices. Constructing price paths from these densities with correct dependence properties is our next goal. The target copula, i.e. the target dependence characteristics, is determined from historical data from some training period. Dependence across the domain of the cumulative mapping is independent of price level. In other words, predictive density and dependence is decoupled.

The following extracts the fundamentals of Kaut's work. The principles of copula theory are first introduced, for explaining why and how we use copulas in our scenario generation method. Aiming to ease the explanation of this non-trivial subject, we first introduce some basics and notation, before turning to discuss the case of bivariate scenario generation, which can be carried out solving a MIP. Next, a heuristic to replace the MIP is presented. Finally, the logic from this heuristic is extended, such that a heuristic for the general multivariate scenario generation can be derived.

A copula is the joint cumulative distribution function of any n-dimensional random vector with standard uniform margins, that is, a function $\mathrm{C} : [0,1]^n \to [0,1]$. Sklar's theorem (Sklar, 1996) states that for any n-dimensional cumulative distribiution function F with marginal distribution functions $F_1, ..., F_n$, a copula C exists such that

$$F(x_1, ..., x_n) = C\Big(F_1(x_1), ..., F_n(x_n)\Big) \tag{D.1.1}$$

If all the marginal cumulative distribution functions F_i are continuous, then there exists only one unique C. A consequence of the theorem is that, for every $u = (u_1, ..., u_n) \in [0,1]^n$ we have that

$$C(u_1, ..., u_n) = F\Big(F_1^{-1}(u_1), ..., F_n^{-1}(u_n)\Big) \tag{D.1.2}$$

where F_i^{-1} is the generalized inverse of F_i. That means that knowing the marginal cfds and the copula is the same as fully knowing the multivariate cdf. This stands in contrast to correlations that assume a linear dependence. Furthermore, copulas, unlike correlations, are independent from the marginal distributions. Therefore we can model the two independently.

A copula only models the interdependence of two or more distributions, the information about the distributions themselves has been removed. Ergo, the copula is unchanged as long as the multivariate samples do not change order. One natural option is to change the values to the ranks of the values in the sample, 1 denoting the minimum and S the maximum. The copula sample is thus equivalent to an assignment between the ranks of the margins.

$$\mathcal{C} = \{r = (r_1, ..., r_n) : 1 \leq r_i \leq S, \forall i \leq n\} \tag{D.1.3}$$

$$\mathcal{C}_r(r) = \mathcal{C}(\frac{r_1}{S}, ..., \frac{r_n}{S}) = Pr(\Big[0, \frac{r_1}{S}\Big] \times ... \times \Big[0, \frac{r_1}{S}\Big]) \tag{D.1.4}$$

Whilst transforming a copula sample to a multivariate scenario is trivial keeping track of the accordance of values and ranks, the difficult part is to generate copula scenarios. The most straight-forward method is sampling. The problem with this however, is that we need many samples in order to reach a satisfactory resemblance of the distribution. And again, we stress our wish to keep the number of input scenarios to our decision model as low as possible. The objective of the MIP presented in the next subsection is to minimize the average deviation between the to be generated copula sample, consisting of S scenarios, and the target copula. This, subject to the constraint that a rank from both margins (in the bivariate case) must be allocated to all of the scenarios S. A margin rank shall only be allocated once. In reality, obtaining the copulas is a non-trivial task, similar to estimating distribution functions by sampling. We will use the empirical copula directly as target.

D.2 MIP for the bivariate problem

The assignment of ranks to make a scenario is modelled by binary variables x_{ij} that are equal to one if the j-th rank of the second margin is assigned to the i-th rank of the first margin. Put differently, if there is a scenario with $r = (i, j)$, the C_r function is equal to

$$C_r(i, j) = \frac{1}{S} \sum_{k=1}^{i} \sum_{l=1}^{j} x_{kl} \tag{D.2.1}$$

An assignment's deviation from the target copula can thus be written

$$dev(i,j) = C_r(i,j) - C_r^*(i,j) \tag{D.2.2}$$

Which can decomposed into its positive and negative parts,

$$dev(i,j) = \frac{1}{S}(y_{ij}^+ - y_{ij}^-) \tag{D.2.3}$$

further giving

$$dev(i,j) = \sum_{k=1}^{i}\sum_{l=1}^{j} x_{kl} - SC_r^* = y_{ij}^+ - y_{ij}^- \tag{D.2.4}$$

Consequently we have that,

$$S|dev(i,j)| = y_{ij}^+ - y_{ij}^- \tag{D.2.5}$$

The problem of minimizing the average absolute deviation d_{avg} can be formulated in the following way. The scaling factor $\frac{1}{S}$ is omitted.

$$\min \sum_{i,j=1}^{S} (y_{ij}^+ + y_{ij}^-) \tag{D.2.6}$$

$$\text{s.t} \sum_{i=1}^{S} x_{ij} = 1 \qquad \forall j \in \{1...S\} \tag{D.2.7}$$

$$\sum_{j=1}^{S} x_{ij} = 1 \qquad \forall j \in \{1...S\} \tag{D.2.8}$$

$$\sum_{k=1}^{i}\sum_{l=1}^{j} x_{kl} - y_{ij}^+ + y_{ij}^- = SC_r^*(i,j) \qquad \forall i,j \in \{1...S\} \tag{D.2.9}$$

$$y_{ij}^+ \geq 0, y_{ij}^- \geq 0, x_{ij} \in \{0,1\} \qquad \forall i,j \in \{1...S\} \tag{D.2.10}$$

D.3 Copula-based scenario generation heuristic

In order to reduce solution times, a greedy heuristic is proposed. The algorithm for the bivariate case matches every rank j from the first distribution with a rank i from the second. This is done through the evaluation of the expression

$$\delta_r(i,j) = \sum_{l=1}^{S} |dev(l,j)| = \sum_{l=1}^{S} |C_r(l,j) - C_r^*(l,j)|, \tag{D.3.1}$$

greedily choosing the rank i resulting in the lowest deviation between copula sample

Table D.1: Heuristic for bivariate scenario generation

Algorithm: Bivariate heuristic

1: $S \leftarrow \{1...S\}; \quad r_j \leftarrow i^*$
2: **For** $j \in \{1...S\}$ **do**
3: **For** $i \in I$ **do**
4: calculate the deviation $\delta_r(i,j)$
5: **If** $\delta_r(i,j) < \delta_r^*$ **then**
6: $i^* \leftarrow i; \delta_r^* \leftarrow \delta_r(i,j)$
7: **End if**
8: **End for**
9: $r_j \leftarrow i^*; I \leftarrow I \setminus i^*$
10: **End for**

and target copula. A rank i that is chosen to be paired with a rank j is removed from the set of possible ranks to match the rest of the ranks j with.

Next, we extend the heuristic from a bivariate to a general multivariate case. This is done by starting with two margins and then adding one margin at a time. The procedure can be decribed in an inductive manner: assume that we have already generated values for m margins and want to add a new margin m+1, using the bivariate copulas of variable pairs $(1, m+1), ..., (m, m+1)$ as targets. Unlike the bivariate case, we cannot simply connect rows and columns, as m margins have already been connected. Instead, we assign the ranks of the new margin to scenarios. Hence, r_s^k becomes the rank of variable k assigned to scenario s, C_r^k and C_r^{*k} denote respectively the sample and target cdfs of the bivariate copula of variables k and (m+1), δ_r^k the deviation function of this copula, $\delta_r^k = C_r^k - C_r^{*k}$. The heuristic is open for the possibility that several scenarios can have the same deviation. Therefore all the best scenarios found at line 8 are stored, before one is chosen randomly at line 12. The output is finally S scenarios with length m, such that each margin $k \in \{1..m\}$ has a rank assigned to the scenario s. The rank values of the scenarios can then be transformed according to D.3.2 such that $C^s = [0,1]^k$. Giving these values to the respective inverse cumulative density functions of the margins k, F_k^{-1}, we finally reach the point that we have meaningful scenarios.

$$C_r(i_1, i_2, .., i_n) = \frac{1}{S} \sum_{k_1=1}^{i_1} \sum_{k_2=1}^{i_2} ... \sum_{k_n=1}^{i_n} x_{k_1,k_2..,k_n}^s \tag{D.3.2}$$

Table D.2: Heuristic for multivariate scenario generation

Algorithm: Multivariate heuristic

$\mathcal{S} \leftarrow \{1...S\}; \quad \delta_r^* \leftarrow \infty$
For $j \in \{1...S\}$ **do**
    For $s \in \mathcal{S}$ **do**
        For $k \in \{1...m\}$ **do**
            calculate the deviation $\delta_r^k(r_s^k, j)$
        End for
        $\delta_{rs} \leftarrow \sum_{i=1}^{m} \delta_r^k(r_s^k, j)$
        If $\delta_{rs} < \delta_r^*$ **then**
            $s^* \leftarrow s; \delta_r^* \leftarrow \delta_{rs}$
        End if
    End for
    $r_j^{m+1} \leftarrow s^*; \mathcal{S} \leftarrow \mathcal{S} \setminus s^*$
End for

Appendix E

Regression coefficients

It remains to verify that day-ahead and primary reserve prices indeed exhibit varying sensitivity to conditioning variables across quantiles and hours. The estimation of the regression model in equation 6.3.5 has been performed using (P1) from section 6.4. Estimation has been done through 2014 and 2015 for both day-ahead and primary reserve prices. The resulting model coefficients are plotted in Figure E.1 and E.2.

Observe that the constant term in each market varies across quantiles and hours. In addition, sensitivity to the lagged prices changes significantly throughout quantiles and hours. For the day-ahead price, sensitivity to yesterday's price in the same hour is high for low quantiles in the morning and high for high quantile prices later in the day. This is qualitatively very similar to the results found in Hagfors et al. 2016. Lag 2 and lag 3 have less dynamic effects across the quantiles. The weekly lag is interesting because it varies greatly across quantiles for hours around the start of the working day. This might indicate strong explanatory power of the previous week prices for the conditional distribution in these hours. Suprisingly, the lagged mean has the opposite effect of lag 1.

For primary reserve prices the constant term varies across the quantiles, but are ordered opposite in terms of rank compared to the day-ahead constants. Low quantiles have high constants, and high quantiles have low constants. The lag 1 coefficients are less dependent on quantile compared to day-ahead. For lag 2 and 3 all coefficients are very close to zero except for the 95 % quantile throughout the day. The 95 % quantile coefficients exhibit a rather extreme behaviour compared to other quantiles, see lag 7 and the lagged mean.

To conclude, we have verified that indeed day-ahead and primary reserve prices exhibit varying sensitivity across quantiles. In addition, we see that sensitivity to previous prices changes throughout the day. These results substantiate our modelling approach.

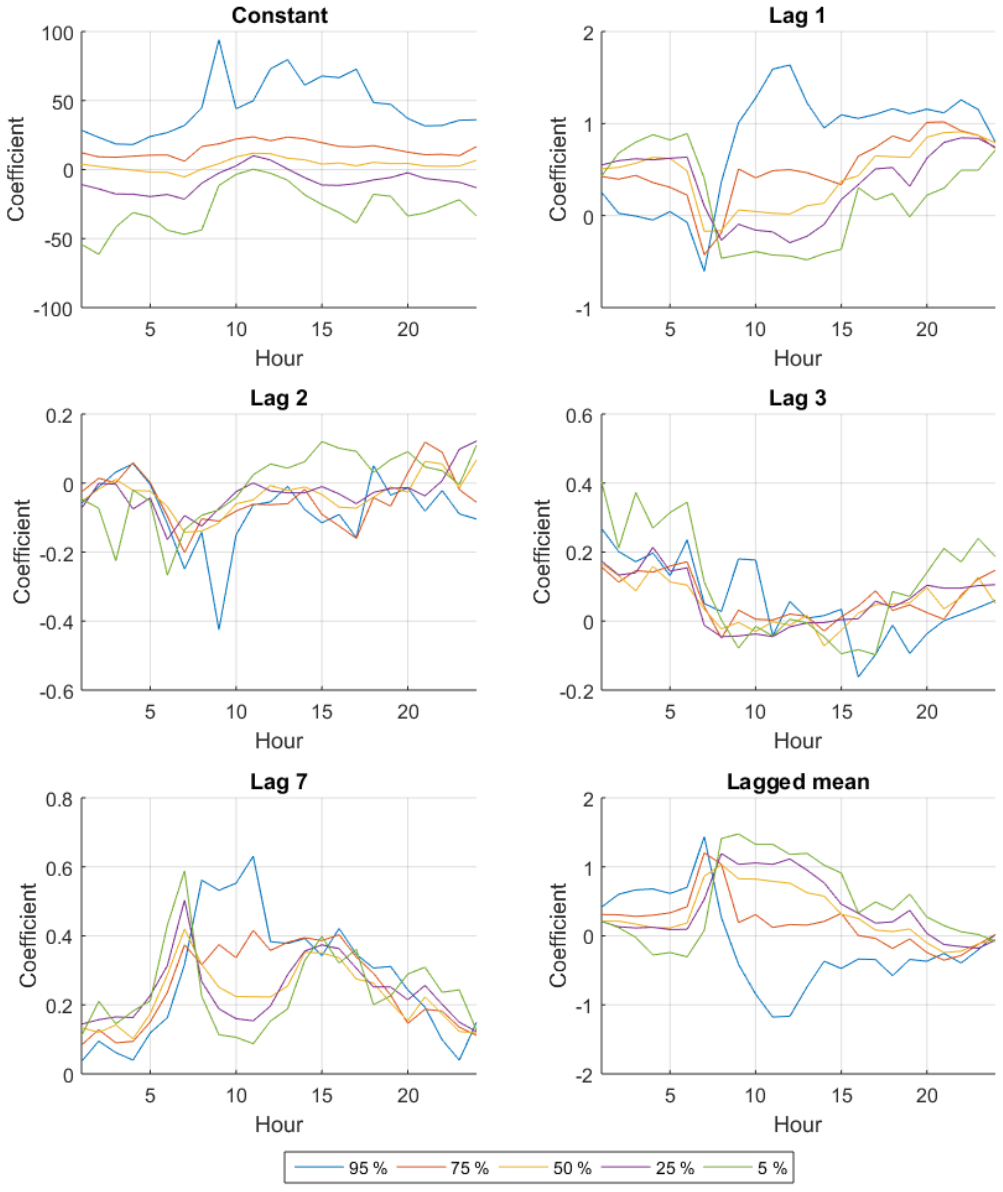

Figure E.1: Coefficients from day-ahead quantile regression model (2014-2015)

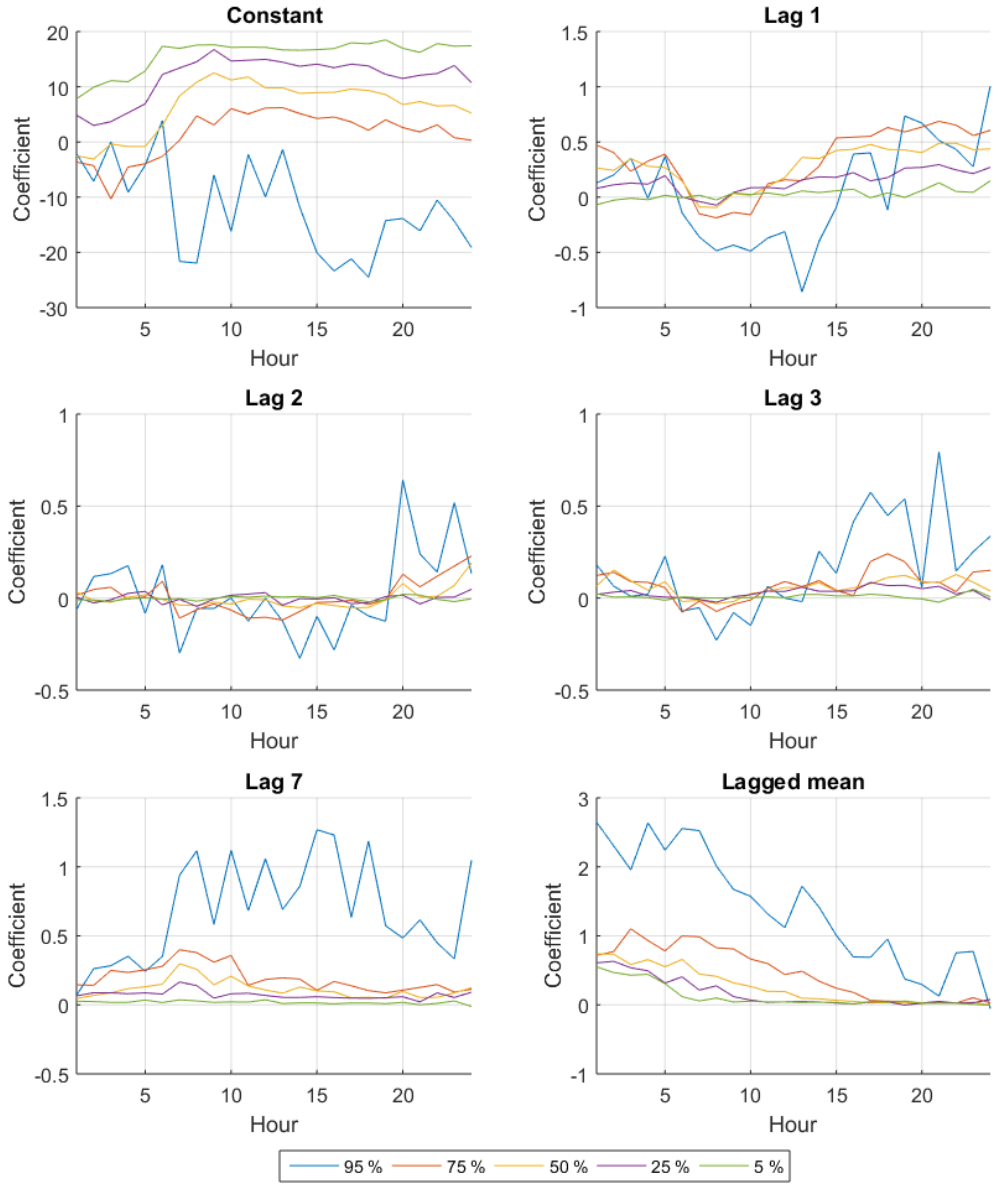

Figure E.2: Coefficients from primary reserve quantile regression model (2014-2015)